On the Design of Interim Cognitive Diagnostic Computerized Adaptive Testing (CD-CAT) in
Learning Context

Chun Wang
University of Washington

Correspondence concerning this manuscript should be addressed to Chun Wang at:

312E Miller Hall
Measurement & Statistics
College of Education
Box 353600
Seattle, WA 98195-3600
e-mail: wang4066@uw.edu
phone: 206-616-6306

Acknowledgement: This project is supported by NSF SES-1659328 and University of Washington Royalty Research Fund A143697

On the Design of Interim Cognitive Diagnostic Computerized Adaptive Testing (CD-CAT) in Learning Context

Abstract

Interim assessment occurs throughout instruction to provide feedback about what students know and have achieved (Popham, 2008; Shepard, 2009). Different from the current available cognitive diagnostic computerized adaptive testing (CD-CAT) design that focuses on assessment at a single time point, we discuss several designs of *interim* CD-CAT that are suitable in the learning context. The interim CD-CAT differs from the current available CD-CAT designs primarily because students' mastery profile (i.e., skills mastery) changes due to learning, and new attributes are added periodically. Moreover, hierarchies exist among attributes taught sequentially and such information could be used during item selection. Two specific designs are considered: the first one is when new attributes are taught in stage II but the student mastery status of the previously taught attributes stay the same. The second design is when both new attributes are taught and previously taught attributes can be further learned or forgotten in stage II. For both designs, we propose an individual prior, which takes into account a person's learning history and population learning model, to start an interim CD-CAT. Simulation results show that the new method outperforms the current method that treats each interim CD-CAT independently. The generalized diagnostic index (GDI) is extended to accommodate item hierarchies, and analytic results are provided to further illustrate the types of items that are most popular during item selection.

Key words: cognitive diagnostic computerized adaptive testing, learning, item selection

Every Student Success Act (ESSA) strongly promotes using technology to its maximum extent to develop, administer, and score assessment results, and CAT is on the technology forefront. To fully embody the potential of CAT for facilitating individualized learning on a mass scale (Quellmalz & Pellegrino, 2009), a CAT should have built-in diagnostic features. Kingsbury (2009) has dubbed adaptive tests geared toward cognitive diagnosis as “*Idiosyncratic CAT*” (ICAT) and has found promising applications in providing teachers with information for targeted instruction. CAT based on cognitive diagnostic models (CDMs), namely CD-CAT, has become a psychometrically sound option. CDMs provide a level of control in scaling, linking, and item banking that is unavailable with other simpler subscore methods.

A large body of simulation work has been done to explore different item selection methods. For example, Xu, Wang, and Shang (2016) provided a non-asymptotic theory-based approach to guide initial item selection in CD-CAT, which considerably reduces the test length. After the initial-stage selection, well-established information-based criteria including the posterior weighted Kullback-Leibler (KL) information index (Cheng, 2009, 2010), the mutual information index (Wang, 2013), or the generalized diagnostic index (GDI, Kaplan, de la Torre, & Barrada, 2015) could be used subsequently. However, these available CD-CAT methods mostly focus on assessment at a single time point (e.g., beginning or end-of semester). As a result, the test often starts randomly and the items usually cover a broad range of attributes with the goal of maximizing the precision of the entire cognitive profile. In contrast, item selections for interim CD-CAT should take into account the general learning model as well as an individual student’s learning history. Developing CD-CAT for such a purpose is essential to close the learning-assessment feedback loop.

The interim CD-CAT differs from the current available CD-CAT designs primarily because students' mastery profile (denoted as α) changes due to learning, and new attributes are added periodically. Therefore, at the $(t)^{th}$ learning stage such as the $(t)^{th}$ week of the instructional period, the interim assessment will select items that require skills taught during the corresponding period and relevant prerequisite skills taught prior to stage t . Furthermore, interim CD-CAT has two additional unique features that should be considered in its design. First, the number of measured attributes could grow, yet test length needs to be kept desirably low. Hence, challenge remains as to how to accurately recover a potentially high-dimensional α with a limited number of items. Second, hierarchies likely exist among attributes taught sequentially (e.g., curriculum design) and such information could aid the item selection design.

In this paper, we propose viable interim CD-CAT designs that reflect above mentioned features. In what follows, we will first briefly introduce CDMs and GDI, followed by the interim CD-CAT designs as well as relevant analytic results. Then a series of simulation studies are presented to demonstrate the performance of interim CD-CAT under different scenarios.

Cognitive Diagnostic Models (CDM)

To support these next-generation assessments aimed at providing fine-grained feedback for students and teachers (Leighton & Gierl, 2007; Templin & Bradshaw, 2014), CDMs have arisen as advanced psychometric models in the past few decades. In essence, CDMs are restrictive latent class models that uncover the skills/attributes a student possesses at the time of assessment. Each latent profile constitutes one latent class. Denote the mastery profile of a person by $\alpha = (\alpha_1, \dots, \alpha_K)$, where K is the total number of attributes measured by a test. Similarly, the item to attribute mapping could be summarized in a Q-matrix (Embretson, 1984).

The q -vector of item j is denoted by $\mathbf{q}_j = (q_{j1}, \dots, q_{jK})$, which is a binary vector that implies the attributes measured by item j . Recent advancement in CDMs also allows polytomous attributes (Chen & de la Torre, 2018), but we restrain the discussion to binary attributes in this paper.

To model the item responses as a function of item and person characteristics, de la Torre (2011) proposed a generalized DINA (GDINA) model, which is one of the most flexible CDMs. The GDINA model specifies the probability that person i answers item j correctly as follows

$$P(y_{ij} = 1 | \boldsymbol{\alpha}_i) = \delta_{j,0} + \sum_{k=1}^{K_j} \delta_{j,k} q_{jk} \alpha_{ik} + \sum_{k=1}^{K_j} \sum_{k' > k} \delta_{j,kk'} q_{jk} q_{jk'} \alpha_{ik} \alpha_{ik'} + \dots + \delta_{j,12\dots K_j} \prod_{k=1}^{K_j} q_{jk} \alpha_{ik} \quad (1)$$

where $\delta_{j,0}$ is the intercept, $\delta_{j,k}$ is the main effect due to α_k , $\delta_{j,kk'}$ is the two-way interaction, and $\delta_{j,12\dots K_j}$ is the K_j -way interaction (de la Torre, 2009), K_j is the number of relevant attributes measured by item j . From Equation 1, it is shown that the GDINA model is essentially a generalized linear model (GLM) with binary latent predictors. If the identity link in Equation 1 is replaced by a logit link, the model becomes the log-linear CDM (LCDM; Henson, Templin, & Willse, 2009). Assume K is the total number of attributes measured in a test, then without imposing any additional constraints, both GDINA and LCDM assume an individual profile, $\boldsymbol{\alpha}_i$, could be any one of the possible profiles, with the class mixing proportions denoted as π_c .

Generalized Diagnostic Index (GDI)

The GDI is considered as the basis item selection method in this paper because it is computationally fast and it performs as well as or even sometimes better than the other computationally more intensive method, such as the mutual information method (Wang, 2013) or the modified PWKL (Kaplan, et al., 2015; Xu et al., 2016). The GDI measures the weighted variance of the conditional success probabilities of an item. The definition is as follows. Let $\boldsymbol{\alpha}_{jc}^*$ denote a reduced attribute vector consisting of K_j attributes measured by item j . The superscript

“*” is added to show that α_{jc}^* is not of the same length as α , and the length of α_{jc}^* varies across different items. Denote $P(y_{ij} = 1|\alpha_{jc}^*)$ as the correct response probability of item j given α_{jc}^* , and denote $\pi_i^n(\alpha_{jc}^*)$ as the posterior probability of the reduced attribute vector $\alpha_i^* = \alpha_{jc}^*$ after n administered items in CAT. Then the GDI for item j is

$$\eta_j = \sum_{c=1}^{2^{K_j}} \pi_i^n(\alpha_{jc}^*) [P(y_{ij} = 1|\alpha_{jc}^*) - \bar{P}_j]^2, \quad (2)$$

where $\bar{P}_j = \sum_{c=1}^{2^{K_j}} \pi_i^n(\alpha_{jc}^*) P(y_{ij} = 1|\alpha_{jc}^*)$ is the average success probability. In essence, the GDI measures the discrimination power of an item in differentiating different relevant reduced attribute vectors. It also attaches higher weight to α_{jc}^* if its posterior density is high, implying that α_{jc}^* is more likely to be the true reduced attribute vector for person i . One feature that makes GDI compelling for interim CD-CAT is that it uses reduced vector α_{jc}^* in its computation. Hence, regardless how large K is in a test, as long as the number of attributes measured by a single item is relatively small, selecting items using GDI is always efficient.

Interim CD-CAT Design

In this section, several design aspects for interim CD-CAT will be discussed in sequel, which include: (1) efficiently synthesizing information from prior stages to update α when the size of α keeps growing over time; (2) selection of items in the presence of attribute hierarchies; and (3) updating prior information to take into account learning.

Specifically, at the $(t+1)$ th learning stage such as the $(t+1)$ th week of the instructional period, if a set of skills denoted by $s^{(t+1)}$ is taught which could contain one or multiple skills, then the interim assessment will select items that require skills in $s^{(t+1)}$ and skills taught prior to stage $(t+1)$. The goal of item selection is to evaluate whether skills in $s^{(t+1)}$ are learned, and if

not, whether any prerequisites are missing. Items that can provide maximum information to recover mastery status of skills in $s^{(t+1)}$ will be selected.

Update of α . First of all, during a CD-CAT when the size of α is fixed, then to save computation time for sequential update of α using the maximum a posteriori (MAP, Huebner & Wang, 2011), we recommend to use the posterior density from the previous step as the current prior. That is, suppose the cardinality of $s^{(t)}$ is k_t , then α is a vector of length K_t , which includes all attributes taught up to time t . By definition, we have

$$\pi_i^n(\alpha | \mathbf{y}_i^n) = \frac{L(\mathbf{y}_i^n | \alpha) \pi_0(\alpha)}{\sum_{c=1}^{2^{K_t}} L(\mathbf{y}_i^n | \alpha_c) \pi_0(\alpha_c)}, \quad (3)$$

where $\pi_0(\alpha)$ is the prior, and $\mathbf{y}_i^n$ is the vector of responses on n administered items. Any information regarding the attribute relationships could be summarized in $\pi_0(\alpha)$. For instance, when all attributes are independent, then $\pi_0(\alpha)$ is simply a vector of $\frac{1}{2^{K_t}}$'s. When there exist attribute hierarchies, then some combinations of α become impermissible, leaving corresponding elements in $\pi_0(\alpha)$ 0. Moreover, if the attributes are correlated, such as in the higher-order CDMs (de la Torre & Douglas, 2004; Wang, Zheng, & Chang, 2014), $\pi_0(\alpha)$ will take on unique patterns to reflect the correlation. From (3), it could be easily shown that

$$\pi_i^n(\alpha | \mathbf{y}_i^n) = \frac{L(y_{in} | \alpha) \pi_i^{n-1}(\alpha | \mathbf{y}_i^{n-1})}{\sum_{c=1}^{2^{K_t}} L(y_{in} | \alpha_c) \pi_i^{n-1}(\alpha | \mathbf{y}_i^{n-1})}. \quad (4)$$

Compared to Equation 3, Equation 4 circumvents the need of cycling through all n items to compute the likelihood, hence reducing the computation time, in particular when either n or K_t is large.

Second, as alluded to earlier, during an instructional unit, new skills are taught and evaluated periodically. Hence, the posterior density of $\alpha^{(t)}$ at the end of stage t , $\pi_i^{(t)}(\alpha^{(t)} | \mathbf{y}_i^{(t)})$, cannot serve as the prior of $\alpha^{(t+1)}$ at stage $(t+1)$ because $\alpha^{(t)}$ and $\alpha^{(t+1)}$ have different lengths. Even so, the responses from all preceding stages up to time (t) could still be gathered to form a more informative prior for person i at the beginning of the $(t+1)$ th interim CD-CAT. More specifically, the posterior density computed as follows will be used as the prior at stage $(t+1)$,

$$\pi_i^{(t+1)}(\alpha) = \frac{L(\mathbf{y}_i^{(1,2,\dots,t)} | \alpha) \pi_0(\alpha)}{\sum_{c=1}^{K_{t+1}} L(\mathbf{y}_i^n | \alpha_c) \pi_0(\alpha_c)}, \text{ where } \mathbf{y}_i^{(1,2,\dots,t)} \text{ denotes the stack of responses from all}$$

preceding stages, and $\pi_0(\alpha)$ is a uniform prior that reflects attribute hierarchies. That is, impermissible patterns that violate the hierarchical relationship have priors of 0, whereas all permissible patterns are equally likely. Of course the length of α becomes K_{t+1} to reflect all attributes measured till time $(t+1)$. $\pi_i^{(t+1)}(\alpha)$ will replace $\pi_0(\alpha)$ in Eq. (3) to start the $(t+1)$ th interim CD-CAT. In fact, the computation of $\pi_i^{(t+1)}(\alpha)$ could be done offline after the t^{th} CD-CAT finishes, hence it will not interrupt the flow of live CD-CAT if its computation takes some time.

One note to make is, even though the conditional correct response probability of an item depends only on the reduced vector (i.e., Eq. (1)) (and so is GDI), the posterior density of the full vector will be updated. In other words, if an item measures α_1 and α_2 whereas $K=3$, then only $\pi_i^n((\alpha_1, \alpha_2))$ counts in computing GDI, but the posterior mastery probability of α_3 will still be updated after administering this item. This point is important because for example, in the 2nd CD-CAT, although those attributes measured in the first test may not emerge again in the items given in the 2nd test, the mastery profile on these previously measured attributes may still update.

Therefore, it is recommended, in the interim CD-CAT context, to always work with full length α (i.e., K_t) rather than working with clusters of attributes in isolation.

Item Selection. Given the many desirable features of GDI, it could be used for item selection in interim CD-CAT. At stage t , the eligible items would be those that require one or multiple skills in $s^{(t)}$, some of which may also require skills taught prior to stage t .

When there are attribute hierarchies among the reduced vector α_{jc}^* , this information will be reflected in $\pi_0(\alpha_{jc}^*)$ that will be carried over to $\pi_i^n(\alpha_{jc}^*)$ in Equation 2.

Besides this simple modification, it would be interesting to delve deeper into the types of items that are preferable because this knowledge would guide future item bank design. When all attributes are independent, Xu et al. (2016) proposed a non-asymptotic theory-based approach to guide initial item selection in CD-CAT. In particular, they derived the minimum number of items required in order to identify the attribute pattern of a student as well as the specific types of initial items that are required to reach the optimal classification results, under both ideal and practical scenarios. Their proposal was later used by Chang, Chiu, and Tsai (2018) in their non-parametric CD-CAT design. While Xu et al. (2016)'s results are suitable for independent attribute structures, we extend the results to scenarios with attribute hierarchies. The two lemma presented below are for the ideal case of conjunctive CDM when students answer correctly all questions they are capable of and incorrectly otherwise. The conclusions will be empirically verified in non-ideal cases via simulation studies. On a side note, the Q-matrix does not necessarily have to conform to the attribute hierarchies, such as the Q-matrix from the Examination for the Certificate of Proficiency in English (ECPE) data (Templin & Bradshaw, 2014; Templin & Hoffman, 2013).

Lemma 1: Suppose a test intends to measure K attributes with the attribute hierarchies summarized in a K -by- K reachability matrix $\mathbf{R}$. To identify all possible attribute profiles, it is necessary and sufficient for a test to have K items with the Q-matrix satisfying the following rule: After row swapping, it is the transpose of the reachability matrix, with off-diagonal 1's replaced by "*" indicating that those entries could be either 1 or 0.

For instance, with $K=4$ and attributes exhibit a divergent structure shown in Figure 1, the Q-

matrix needs to take the following form, after row swapping: $\mathbf{Q} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ * & 1 & 0 & 0 \\ * & 0 & 1 & 0 \\ * & 0 & * & 1 \end{bmatrix}$. (5)

Figure 1. An illustration of the divergent structure

Proof. We first show the sufficiency of the Q-matrix. According to Chiu et al. (2009), when the Q-matrix is indeed an identity matrix, then it could differentiate all 2^K possible patterns had there be no attribute relationships. However, when there are attribute hierarchies, those unspecified elements will not affect separability of different permissible patterns. Using the Q-matrix in (5) as an example. If the ideal response pattern for item 2 is 1, then α_2 is 1 and so is α_1 , regardless of whether item 2 requires α_1 or not. Indeed, the inclusion of item $q = \mathbf{e}_2$ in Chiu et al. (2009) is to differentiate $\boldsymbol{\alpha} = \mathbf{0}_K$ from $\boldsymbol{\alpha} = \mathbf{e}_2$, but if α_2 has a pre-requisite, then $\mathbf{e}_2$ is no longer a permissible pattern. Hence, it could be verified that the Q-matrix specified in Lemma 1 will generate all different ideal response patterns for all permissible attribute profiles.

We then show the specified Q-matrix is necessary using method of contradiction. If $\mathbf{q}_k$, the k th row of the Q-matrix is missing, then $\alpha_n = \mathbf{e}_k + \mathbf{e}'_k$ and $\alpha_m = \mathbf{e}'_k$, where $\mathbf{e}'_k$ denotes the 1-by- K vector with the pre-requisites of α_k as 1, are not separable because they generate the same ideal response patterns. This completes the proof.

Lemma 2: (Extension of Theorem 1 in Xu et al., 2016). Consider a person with attribute profile $\alpha=0$. In the ideal case, to identify $\alpha=0$, it is necessary and sufficient for a test to have C items, where C denotes the total “root” attributes in all clusters of attributes. The corresponding Q-matrix of the C items, after row swapping, follows the following form: $\mathbf{Q} = (\mathbf{I}_C, \mathbf{0})_{C \times K}$, where the first C columns refer to the “root” attributes.

Moreover, consider a person with attribute profile $\alpha \neq \mathbf{0}$. Suppose there are k elements in α equal to 1. Then via a column swapping of the attribute labels, α is equivalent to $\mathbf{m}_k$. To identify $\alpha=\mathbf{m}_k$, it is necessary and sufficient for a test to satisfy the following conditions:

C1. There are d_1 items in the test such that the corresponding Q-matrix after row swapping takes the form

$$\mathbf{Q} = (\mathbf{0}_*, \mathbf{I}_{d_1}, \mathbf{0})_{d_1 \times K}$$

where d_1 denotes the number of “independent, root” attributes among the last $(K - k)$ attributes. $\mathbf{I}_{d_1}$ is a $d_1 \times d_1$ identity matrix, and $\mathbf{0}_*$ denotes an unspecified $d_1 \times k$ binary matrix. The remaining $d_1 \times (K - d_1 - k)$ elements are all 0’s. Here, the independent attributes mean they do not have a direct linkage. When a set of attributes are directly related, then $d_1 = 1$ for the “root” attribute.

C2. There is a set of item(s) that requires none of the last $(K - k)$ attributes but requires all of the “leaf” attributes among the first k attributes.

Again, use $K=4$ and the divergent structure in Figure 1 as an example. Since all attributes form a single cluster, when $\mathbf{a}=0$, only one item with $\mathbf{q}=(1,0,0,0)$ is needed so that any profile $\mathbf{a} \neq \mathbf{0}$ has different response patterns from that of $\mathbf{a}=0$. For $\mathbf{a}=(1,0,0,0)$, the “independent, root” attributes among $(\alpha_2, \alpha_3, \alpha_4)$ are α_2 and α_3 . Hence, the Q-matrix satisfying the two conditions in

Lemma 2 takes the form of $\mathbf{Q} = \begin{pmatrix} * & 1 & * & 0 \\ * & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$. If $\mathbf{a} = (1,0,1,0)$, then according to Lemma 2, we

need 3 items with the following Q-matrix to identify this profile: $\mathbf{Q} = \begin{pmatrix} * & 1 & * & 0 \\ * & 0 & * & 1 \\ * & 0 & 1 & 0 \end{pmatrix}$. The first

two rows satisfy condition C1. Since α_2 and α_4 are not directly linked, so $d_1 = 2$. The last row satisfies condition C2. For α_1 and α_3 , α_3 is the “leaf” attribute.

Similarly, if $\mathbf{a} = (1,1,0,0)$, then we only need 2 items with the following Q-matrix: $\mathbf{Q}$

$= \begin{pmatrix} * & * & 1 & 0 \\ * & 1 & 0 & 0 \end{pmatrix}$. Because α_3 and α_4 are directly related, $d_1 = 1$ and α_3 is the root attribute

between the two.

Proof. When all attributes are independent, then each attribute forms a separate cluster and each of them is considered a “root attribute”. For $\mathbf{a}=0$, then the Q-matrix in Lemma 2, $\mathbf{Q} = (I_C, \mathbf{0})_{C \times K}$, becomes the identity matrix, which is consistent with Theorem 1 in Xu et al. (2016) as well as Lemma 1 in Chiu et al. (2009).

For $\mathbf{a} \neq \mathbf{0}$, we first show that if conditions C1 and C2 are satisfied, then any permissible attribute profile $\mathbf{a} \neq \mathbf{m}_k$ has different response patterns from that of $\mathbf{m}_k$. Without loss of generality, suppose there is 1 item that satisfies condition C2. Then, for the target $\mathbf{a}=\mathbf{m}_k$, the ideal response pattern would be $(0, \dots, 0, 1)$, a set of d_1 0’s followed by a 1. Then for any permissible pattern $\mathbf{a}' \prec$

α , the ideal response pattern would be all 0's. The notation " $<$ " means any attributes mastered by α' are also mastered by α , but there is at least one attribute mastered by α that is not mastered by α' . For any permissible pattern $\alpha' > \alpha$, then regarding the ideal response pattern for α' , there is at least one 0 among the first d_1 0's that is 1 because α' must have mastered one more "root" attribute among the last $(K - k)$ attributes.

We then need to show that conditions C1 and C2 are necessary. We can again prove this by the method of contradiction. If C1 is not satisfied, then $\mathbf{m}_k$ and $\mathbf{m}_{k+1}$ have the same response vector for all items; we cannot distinguish them. Similarly, if C2 is not satisfied, then $\mathbf{m}_k$ and $\mathbf{m}_{k-1}$ have the same response vector for all items. This completes the proof.

For sequential item selection, Theorem 2 in Xu et al. (2016) cannot be easily extended due to the complication of the attribute hierarchies. That is, the first item does not have to be a single-attribute item any more. However, the general conclusion in Xu et al (2016) still holds: during item selection, when the answer to an item is wrong, implying that the corresponding attributes may not be mastered, then the non-mastered attributes should not be required by the next item. On the other hand, if the answer is right, then the mastered attribute could be "unspecified" for the next item.

Using Figure 1 as an example, one interesting take-away message from Lemma 2 is when α_1 is the pre-requisite attribute for all others, then except for people with $\alpha=0$ or $\alpha=(1,0, 0, 0)$, the sufficient Q-matrix for identifying true pattern for an individual really does not need item measuring α_1 . This fact is reflected by the '*' in the first column of the Q-matrix. This implies that items that only measure α_1 may not be as useful as items measure other, more advanced attributes.

When Learning Happens. Learning could happen between two interim CD-CATs, which means a student's α is no longer static. If this is the case, responses from prior stages may not be used in the current stage because those prior responses reflect a somewhat different α prior to learning. Instead, insights from learning models could play an important role in forming a better prior such that the ever changing α could still be recovered well within a handful of items. For instance, Chen, Culpepper, Wang, and Douglas (2017) proposed a first-order hidden Markov model for learning trajectories, and results from their model could be useful for interim CD-CAT. Similar models are also proposed by Madison and Bradshaw (2018).

To be specific, their model could produce either attribute level or pattern level first-order transition probabilities. That is, between any two time points, the attribute level transition probability spells out as: $\tau_{k,1|0} = P(\alpha_k^{(t+1)} = 1 | \alpha_k^{(t)} = 0)$, $\tau_{k,0|0} = P(\alpha_k^{(t+1)} = 0 | \alpha_k^{(t)} = 0)$, $\tau_{k,0|1} = P(\alpha_k^{(t+1)} = 0 | \alpha_k^{(t)} = 1)$, and $\tau_{k,1|1} = P(\alpha_k^{(t+1)} = 1 | \alpha_k^{(t)} = 1)$. Of course, $\tau_{k,1|0} + \tau_{k,0|0} = 1$ and $\tau_{k,1|1} + \tau_{k,0|1} = 1$. The transition probabilities do not have a subscript of time t on it because the current learning model assumes the transition probabilities to be time invariant (Li, Cohen, Bottge, & Templin, 2015; Chen et al., 2017). By knowing these attribute level transition probabilities, and assuming all attributes are learned independently, the pattern level transition probabilities could be obtained fairly easily. As an example, $\tau_{(1,1)|(0,0)} = \tau_{1,1|0} \times \tau_{2,1|0}$. On the other hand, both Chen et al. (2017) and Madison & Bradshaw (2018) discussed about estimating the pattern level transition probabilities directly, hence relaxing the assumption of independent learning.

Now with pattern level transition probabilities known from a learning model, we propose to update priors of α at the beginning of a new CD-CAT as follows. Say, at time t , the posterior

density of α for person i is written as $\pi_i^{(t)}(\alpha)$ for succinctness. Then, the individual prior that could be used for person i 's $(t + 1)^{th}$ CD-CAT is $\pi_{i,0}^{(t+1)}(\alpha) = \pi_i^{(t)}(\alpha) \times P(\alpha^{(t+1)}|\alpha^{(t)})$, where $P(\alpha^{(t+1)}|\alpha^{(t)})$ is the pattern transition matrix. This individual, informative prior synthesizes the information from a person's learning history as well as the learning model, and therefore, it is expected that the interim CD-CAT designed this way could produce accurate α estimates efficiently.

Simulation Studies

Two simulation studies were conducted to mimic the typical scenarios in learning context. In both cases, two time points and additive CDM (ACDM) were considered for illustration purpose, but the methods could be used for any number of time points and any CDMs. In the first scenario, new attributes are added at time II, and these new attributes may require the previously taught attributes as pre-requisites. In the second and more interesting scenario, not only new attributes are added at time II, students' true mastery profile on the attributes measured at time I also change due to learning. The simulation designs and results for these two simulation studies are presented in detail below.

Study I Design

In study I, the two manipulated factors are (1) the total number of attributes measured across two time points, $K=6$ or 10 , and (2) the relationship between the attributes tested at two time points, i.e., independent or hierarchical. Specifically, the attribute hierarchies are shown in Figure 2 for both two levels of K . When the attributes are independent, it implies that there are no direct linkages among any attributes. Fully crossing the two manipulated factors results in 4 simulated conditions.

Figure 2. Attribute hierarchies in the simulation study. The design in (a) results in 24 permissible patterns and the design in (b) results in 168 permissible patterns.

The item bank size was created in proportion to K . When $K=6$, there are 480 items in the item bank, with the Q-matrix designed as follows: the first 150 items only measure α_1 to α_3 , the second 150 items only measure α_4 to α_6 , and the remaining 180 items measure α_1 to α_6 . More specifically, the first 150 items' Q-matrix is the 15 copies of the first 10 rows of Table 1, whereas the second 150 items' Q-matrix is the 15 copies of the last 10 rows of Table 1. Regarding the remaining 180 items, one third of the items measure one attribute from α_1 to α_3 (i.e., randomly select) and one attribute from α_4 to α_6 , one third of the items measure two attributes from α_1 to α_3 and one attribute from α_4 to α_6 , and the last one third of the items measure one attribute from α_1 to α_3 and two attributes from α_4 to α_6 . As a result, the item selection at time I will be restrained to select items only from the first 150 items, whereas the item selection at time II will be from either the entire bank excluding the items used at time I (i.e., the same items will not be administered twice to the same examinee) or the remaining 330 items (i.e., the items that measure at least one of the attribute from α_4 to α_6). We name these two conditions as non-restricted and restricted conditions respectively. It is expected that the former scenario will yield higher precision of α estimates, whereas the latter scenario may reflect the real assessment design more closely because teachers usually focus on newly taught attributes (e.g., α_4 to α_6) in interim quizzes. Test length was set up to be proportional to K . When $K=6$, the

test lengths are 6 and 12 at time I and II respectively, whereas when $K=10$, the test lengths are 10 and 20.

As to the item parameters for ACDM with identity link, the intercepts were simulated from a uniform distribution $U(0.01, 0.2)$, and the main effects were simulated as $\delta_{jk} = \frac{0.65}{K_j} + \varepsilon_{jk}$ where $\varepsilon_{jk} \sim U(0.01, 0.05)$ is a random value so that the main effects are not equivalent. Setting the main effect this way ensures that the conditional correct response probabilities are between 0.2 and 0.85 (Xu et al., 2016).

Table 1. Sample Q-matrix for Study I ($K=6$)

	α_1	α_2	α_3	α_4	α_5	α_6
1	1	0	0	0	0	0
2	0	1	0	0	0	0
3	0	0	1	0	0	0
4	1	0	0	0	0	0
5	0	1	0	0	0	0
6	0	0	1	0	0	0
7	1	1	0	0	0	0
8	1	0	1	0	0	0
9	0	1	1	0	0	0
10	1	1	1	0	0	0
11	0	0	0	1	0	0
12	0	0	0	0	1	0
13	0	0	0	0	0	1
14	0	0	0	1	0	0
15	0	0	0	0	1	0
16	0	0	0	0	0	1
17	0	0	0	1	1	0
18	0	0	0	1	0	1
19	0	0	0	0	1	1
20	0	0	0	1	1	1

When $K=10$, there are 600 items in the item bank, with the Q-matrix designed similarly to the previous condition, as follows. The first 180 items only measure α_1 to α_5 , and there are 6 copies of a 30-row Q-matrix. The first 10 rows represent two copies of 5-by-5 identity matrix, followed by 10 rows of q-vector that has 2 entries of 1 (i.e., there are 10 combinations of 2 out of

5), and another 10 rows of q-vector that has 3 entries of 1. Similarly, the second 180 items only measure α_6 to α_{10} , with the Q-matrix constructed exactly as above. The remaining 240 items' q-vectors were constructed the same way as in $K=6$ scenario, with 80 items in each cluster. Again, all item parameters were simulated the same as before. Then at time I CD-CAT, item selection is restricted to the first 180 items, whereas item selection at time II, items are selected from either the remaining 420 items (restricted condition) or 590 items (non-restricted condition, excluding the previously administered 10 items).

The sample size in the four conditions were simulated to be proportional to the number of permissible latent classes, i.e., $N=100 \times \text{number of permissible conditions}$. The sample size was not hold equal across conditions because in CAT, every individual is treated independently, and therefore sample size really does not matter much as long as the sample size is large and the sample is representative. The total sample size is: 2400 ($K=6$, hierarchical), 6400 ($K=6$, independent), 16800 ($K=10$, hierarchical), and 102400 ($K=10$, independent).

The current method refers to treating time II CD-CAT as a separate test of time I CD-CAT, hence $\pi_0(\alpha)$ (respecting attribute hierarchies, if exist) will be used as the prior. In contrast, the new method refers to start time II CD-CAT using an individualized prior that takes into account each student's likelihood (i.e., learning history) from time I CD-CAT.

Study I Results

Table 2 and 3 present the pattern and attribute recovery rates under different scenarios when $K=6$ and 10 respectively. Both tables consistently show that the new method outperforms the current method by producing higher pattern and attribute recovery rates. When we restrict item selections in time II CD-CAT to items that measure at least one of the last three attributes (i.e., $\alpha_4 \sim \alpha_6$), the attribute and pattern recovery rates drop for both the new and current methods,

and this decrement is more dramatic for the current method. This is because when restricting item selection in time II, none of the items will measure one of the first three attribute (i.e., $\alpha_1 \sim \alpha_3$) alone, and hence the attribute level recovery of α_1 to α_3 is largely sacrificed. But if we only focus on the attribute recovery of α_4 to α_6 , whether or not restricting item selections to a subset of the item pool does not make much of a difference. Another interesting finding is, for the restricted item selection scenario using the current method, time II CD-CAT is treated independently of time I CD-CAT. In this case, even when the item selection focuses only on attributes $\alpha_4 \sim \alpha_6$, the persons' mastery on attributes $\alpha_1 \sim \alpha_3$ are still updated as reflected by higher than 50% attribute recovery rates for $\alpha_1 \sim \alpha_3$. This is consistent with our expectation. Lastly and unsurprisingly, when the attributes exhibit a hierarchical structure and the hierarchy is known, the pattern and attribute recovery rates are a lot higher than those from the independent attribute structure scenario, simply because the permissible space of α is reduced with known hierarchies.

Table 2. Pattern and attribute recovery rates under different scenarios when $K=6$

	Hierarchical					Independent				
	Time I J=6	Time II J=12				Time I J=6	Time II J=12			
		Restricted		Non-restricted			Restricted		Non-restricted	
Method		New	Current	New	Current		New	Current	New	Current
Pattern	0.725	0.853	0.595	0.894	0.759	0.721	0.730	0.328	0.764	0.591
α_1	0.904	0.975	0.890	0.986	0.954	0.894	0.924	0.696	0.951	0.908
α_2	0.903	0.951	0.833	0.979	0.949	0.901	0.931	0.708	0.956	0.905
α_3	0.894	0.961	0.845	0.986	0.951	0.897	0.929	0.724	0.960	0.911
α_4		0.978	0.972	0.973	0.947		0.967	0.968	0.954	0.938
α_5		0.985	0.974	0.974	0.960		0.968	0.962	0.960	0.918
α_6		0.980	0.963	0.978	0.947		0.969	0.972	0.957	0.914

Note. “Restricted” condition refers to the scenario where item selection in stage II CD-CAT is from the 330 items that measure at least one of the three attributes, i.e., $\alpha_4 \sim \alpha_6$; whereas “non-restricted” condition refers to the scenario where items in stage II CD-CAT can be selected from any eligible items in the bank. The eligible items are those that are not administered to the same person in stage I.

Table 3. Pattern and attribute recovery rates under different scenarios when $K=10$

	Hierarchical					Independent				
	Time I J=10	Time II J=20				Time I J=10	Time II J=20			
		Restricted		Non-restricted			Restricted		Non-restricted	
Method		New	Current	New	Current		New	Current	New	Current
Pattern	0.658	0.815	0.486	0.886	0.713	0.666	0.653	0.166	0.718	0.453
α_1	0.938	0.981	0.915	0.991	0.963	0.942	0.951	0.739	0.966	0.919
α_2	0.916	0.960	0.823	0.987	0.955	0.918	0.938	0.704	0.968	0.927
α_3	0.915	0.975	0.902	0.992	0.970	0.927	0.935	0.724	0.967	0.929
α_4	0.913	0.972	0.889	0.990	0.961	0.909	0.932	0.703	0.967	0.920
α_5	0.907	0.957	0.849	0.985	0.957	0.909	0.935	0.719	0.966	0.917
α_6		0.986	0.975	0.981	0.956		0.976	0.972	0.964	0.920
α_7		0.985	0.974	0.985	0.965		0.980	0.969	0.969	0.924
α_8		0.989	0.980	0.986	0.967		0.979	0.969	0.971	0.924
α_9		0.988	0.975	0.984	0.960		0.978	0.968	0.966	0.916
α_{10}		0.989	0.979	0.986	0.967		0.975	0.968	0.966	0.938

To further explore the types of items (i.e., items with certain q-vectors) that are more likely to be selected, we draw the heat map of item exposure in time II conditioning on the different true α patterns, as shown in Figure 3. We focus on time II because in time I, the attributes only display an independent structure. Only results from $K=6$ are shown because the results from $K=10$ are similar. When $K=6$, the item bank contains items with 41 different q-vectors. Interestingly, when we do not restrict item selections in stage II, the items that are more often selected are those that measure a single attribute, and among them, items measuring attributes α_4 to α_6 are the most often selected. Using the new method, the item exposure rate of items measuring only one of the attributes among α_1 to α_6 are 10.7%, 13.5%, 12.3%, 21.1%, 15.2%, and 21.2% respectively (relating to Figure 2a). The same pattern holds for the current method (i.e., Figure 2b). This is consistent with Lemma 1 and 2 discussed in the previous section. In contrast, when item selection in stage II is restricted, then items measuring only α_1 to α_3 are no longer eligible. In this case, items measuring one of the three last attributes continue to be the most popular items, followed by those items whose q-vectors conforming to Lemma 1 and 2.

(a) new method, non-restricted

(b) current method, non-restricted

(c) new method, restricted

Figure 3. Item exposure heat map per true α group in stage II CD-CAT ($K=6$)

Study II Design

The same four conditions were considered in study II. The item bank was also the same as in study I. In addition, to model learning transitions between time I and time II, we assumed the attribute level learning rate, $l_k \equiv \tau_{k,1|0}$, and attribute level forget rate, $f_k \equiv \tau_{k,0|1}$ are from two levels: 0.5/0.2 and 0.8/0.1 (Madison & Bradshaw, 2018, Table 9). Taking $K=6$ as an example. When attributes exhibit hierarchical structure, student mastery status of $\alpha_4 \sim \alpha_6$ depends on their mastery status of $\alpha_1 \sim \alpha_3$ at the beginning of CD-CAT at time II.

To form the $2^{K_1} \times 2^K$ transition matrix, where K_1 denotes the number of attributes measured at time 1, we can compute the pattern level transition probabilities as follows. For each element in the transition matrix, we have

$$P(\alpha^{(2)} | \alpha^{(1)}) \equiv P(\alpha_1^2, \alpha_2^2, \alpha_3^2, \alpha_4^2, \alpha_5^2, \alpha_6^2 | \alpha_1^1, \alpha_2^1, \alpha_3^1)$$

$$\begin{aligned}
&= P(\alpha_1^2, \alpha_2^2, \alpha_3^2 | \alpha_1^1, \alpha_2^1, \alpha_3^1) \times P(\alpha_4^2, \alpha_5^2, \alpha_6^2 | \alpha_1^2, \alpha_2^2, \alpha_3^2) \\
&= \prod_{k=1}^{K_1} \left[(1 - f_k)^{\alpha_k^1} l_k^{(1-\alpha_k^1)} \right]^{\alpha_k^2} \left[f_k^{\alpha_k^1} (1 - l_k)^{(1-\alpha_k^1)} \right]^{1-\alpha_k^2} \times P(\alpha_4^2 | \alpha_1^2) \times P(\alpha_5^2 | \alpha_1^2, \alpha_2^2) \times P(\alpha_6^2 | \alpha_3^2)
\end{aligned}$$

Assuming the attribute hierarchies follow the pattern in Figure 2, we have

$$P(\alpha_4^2 | \alpha_1^2) = \left[l_4^{\alpha_1^2} \times 0^{1-\alpha_1^2} \right]^{\alpha_4^2} \left[(1 - l_4)^{\alpha_1^2} \times 1^{1-\alpha_1^2} \right]^{1-\alpha_4^2} \text{ and } P(\alpha_6^2 | \alpha_3^2) \text{ can be computed in the}$$

similar fashion, whereas $P(\alpha_5^2 | \alpha_1^2, \alpha_2^2) = \left[l_5^{\alpha_1^2 \alpha_2^2} \times 0^{2-\alpha_1^2-\alpha_2^2} \right]^{\alpha_5^2} \left[(1 - l_5)^{\alpha_1^2 \alpha_2^2} \right]^{1-\alpha_5^2}$. If we know the pattern-level transition probabilities directly, such as in Madison and Bradshaw (2018), then we no longer need to assume that each attribute is learned and forgotten independently (e.g., Chen et al., 2017). In the event that there are no attribute hierarchies, we have

$$P(\alpha^{(2)} | \alpha^{(1)}) \equiv \prod_{k=1}^K \left[(1 - f_k)^{\alpha_k^1} l_k^{(1-\alpha_k^1)} \right]^{\alpha_k^2} \left[f_k^{\alpha_k^1} (1 - l_k)^{(1-\alpha_k^1)} \right]^{1-\alpha_k^2}.$$

Study II Results

Tables 4 and 5 presents the pattern and attribute recovery rates for time II CD-CAT from different conditions. Several patterns can be summarized from the table. First, the new method continue to outperform the current method consistently in all conditions. Second, when the learning rate is high and forget rate is low, the pattern and attribute recovery rates are a lot higher and the improvement of the new method over the current method is also larger. This is because when the learning rate is high and forget rate is low, the individual prior distribution in time II CD-CAT is more informative and concentrated, as reflected by its smaller entropy (see Figure 4 for instance). Although the current method does not take advantage of the individual priors, having higher learning rate leads to higher correlations among the attributes, which in turn, leads to more accurate results (see Wang, 2013). Third and consistent with the findings in study I, when attributes have a hierarchical relationship instead of an independent structure, the recovery

rates are uniformly higher. Fourth and again consistent with the findings in study I, restricting item selection in time II CD-CAT to a subset of the item pool lowers the pattern recovery rate mainly due to the poor recovery of the first 3 (or 5) attributes.

Figure 4. Histograms of the entropy of individual priors from low and high learning rate conditions (Study II, K=6)

Discussion

As emphasized in *Knowing What Students Know* (Pellegrino, Chudowsky, & Glaser, 2001), there is a need to move from assessment at a single point in time to assessment practices that guide “*additional teaching, supports, or interventions that will help students master challenging material*” (Fact Sheet: Testing Acting Plan). This step is consistent with the spirit of the learning-assessment cycle, calling for the need of longitudinal, dynamic assessments that assist teachers in understanding how knowledge and skill grow in sophistication (National Education Technology Plan, 2017; Wilson, 2009). In particular, research has shown that providing timely, informative feedback can greatly improve learning (Hanna, 1976; Kluger &

DeNisi, 1996). Interim assessments that are given frequently throughout the instructional period need to be short and highly efficient, which makes CAT promising (Kingsbury, Freeman, & Nesterak, 2014).

To embed CD-CAT in weekly instructions, an interim CD-CAT differs from the current available CD-CAT designs primarily because students' α changes due to learning, and new attributes are added periodically. This study is the first that discusses the various important design aspects of the interim CD-CAT, including the update of α and item selections in the presence of attribute hierarchies. In particular, when new attributes are added intermittently during interim CD-CATs, we propose an individual prior of α using a person's history responses in previous CATs. Simulation study I shows that using this individual prior greatly improves measurement precision of α . This type of design is useful in an intensive longitudinal setting where the mastery status of previously tested α does not change. In contrast, when the interim CD-CATs are spread out such that learning and forgetting could occur in between, we propose to update individual priors of α by taking into account the learning transitions from a learning model. When a student's true α keeps changing and when the length of α also keeps increasing, it is hard to recover α well with a short test. Simulation study II shows great promise of the new method. When the learning rate is high and forget rate is low, the pattern recovery rate is as high as 85.7% when $K=6$ (test length =12) and 83% when $K=10$ (test length =20). The attribute recovery rates are all above 96%.

Throughout the study, we use the general diagnostic index (GDI; Kaplan et al., 2015) as the item selection index because it is easy to compute and works well. One particular advantage of GDI is that its computation complexity only increases with K_j not K , hence even with large K , as long as K_j is small, the computation time is only a fraction of a second, making this index

practically attractive. In this study, we also extend Xu et al. (2016)'s conclusions to the scenarios when there are attribute hierarchies. As shown in Lemma 1 and 2 as well as Figure 2, items that measure a single attribute continue to be the most popular items and moreover, items that measure more advanced attributes are generally preferred. In the simulation studies, we compare two conditions for stage II item selection, namely, restricted vs. non-restricted selection.

Unsurprisingly, as shown in tables 1-4, when we restrict item selection in stage II to items that at least measure one of the last 3 (or 5) attributes, the recovery of the first 3 (or 5) attributes are sacrificed, so is the pattern recovery rate. The attribute recovery of the last 3 (or 5) attributes is not affected. This finding indicates that, if our focus in time II CD-CAT is to recover the entire α vector well, no restriction on item selection is preferred. On the other hand, if we are interested in recovering only the newly taught attributes (i.e., $\alpha_4 \sim \alpha_6$ or $\alpha_5 \sim \alpha_{10}$) well, then restricting item selection in time II is preferred.

This study provides a proof-of-concept illustration of embedding CD-CAT in an interim assessment setting. There are several new directions that are worth exploring in the future to further solidify this application. First, we compare the hierarchical versus independent structures in the simulation studies and find that the former scenario yields more precise α recovery. This improved precision is due to the known attribute hierarchies that are reflected in the structural 0's in the prior. Therefore, knowing the attribute relationship is the premise to the success, and future studies should be devoted to exploring and validating attribute hierarchies from data, such as the exploratory approaches proposed in Wang and Lu (2019), Lu and Wang (2019), or the confirmatory approaches proposed in Templin and Bradshaw (2014). Second, throughout the study, we assume K_t is the same for every individual. However, in a learning context, we could relax K_t as $K_{i,t+1}$, which implies that it can vary individually because a student who shows non-

mastery on some skills in a previous CD-CAT will have those non-mastered skills retested. For

this purpose, not only will the individual prior be updated as follows, $\pi_{i0}^{(t+1)}(\alpha_l) =$

$\frac{L(y_i^{(1,2,\dots,t)}|\alpha_l)\pi_0(\alpha_l)}{\sum_{l=1}^{K_{i,t+1}} L(y_i^n|\alpha_l)\pi_0(\alpha_l)}$, but also different subsets of the item pool will be used for different persons.

Third, the learning model considered in this study assumes that all attributes, aside from their hierarchical relationship, are learned or forgotten independently. Other learning models that allow for pattern-level transitions (e.g., Chen et al., 2017; Madison & Bradshaw, 2018) could be easily incorporated in the individual prior as well. One limitation of these current models, however, is that the number of attributes are the same across different measurement occasions. Future studies should develop learning models that can handle varying number of attributes across time.

Table 4. Pattern and attribute recovery rates from Time II under different scenarios when $K=6$

	Hierarchical								Independent							
Learning/Forget Rate	0.5/0.2		0.8/0.1		0.5/0.2		0.8/0.1		0.5/0.2		0.8/0.1		0.5/0.2		0.8/0.1	
	Time II (J=12), restricted				Time II, non-restricted				Time II (J=12), restricted				Time II, non-restricted			
	Current	New	Current	New	Current	New	Current	New	Current	New	Current	New	Current	New	Current	New
Pattern	0.556	0.611	0.722	0.775	0.768	0.775	0.802	0.857	0.328	0.418	0.360	0.611	0.607	0.653	0.568	0.763
α_1	0.874	0.900	0.951	0.958	0.945	0.955	0.981	0.981	0.676	0.752	0.714	0.868	0.913	0.926	0.922	0.953
α_2	0.818	0.838	0.907	0.912	0.946	0.954	0.966	0.978	0.729	0.778	0.731	0.862	0.914	0.926	0.914	0.952
α_3	0.838	0.861	0.920	0.915	0.945	0.947	0.980	0.970	0.728	0.776	0.730	0.866	0.911	0.933	0.913	0.957
α_4	0.971	0.968	0.965	0.976	0.959	0.956	0.945	0.960	0.969	0.971	0.964	0.972	0.920	0.927	0.908	0.952
α_5	0.976	0.976	0.957	0.982	0.962	0.960	0.937	0.964	0.958	0.961	0.968	0.983	0.930	0.934	0.915	0.956
α_6	0.969	0.970	0.962	0.974	0.957	0.953	0.941	0.963	0.968	0.970	0.970	0.973	0.936	0.939	0.900	0.961

Note. The time I α recovery is omitted here to save space because the results are very similar to those in Study I.

Table 5. Pattern and attribute recovery rates from Stage II under different scenarios when $K=10$

	Hierarchical								Independent							
Learning/Forget Rate	0.5/0.2		0.8/0.1		0.5/0.2		0.8/0.1		0.5/0.2		0.8/0.1		0.5/0.2		0.8/0.1	
	Time II (J=12), restricted				Time II, non-restricted				Time II (J=12), restricted				Time II, non-restricted			
	Current	New	Current	New	Current	New	Current	New	Current	New	Current	New	Current	New	Current	New
Pattern	0.452	0.516	0.654	0.733	0.710	0.746	0.740	0.830	0.152	0.244	0.188	0.481	0.439	0.479	0.471	0.696
α_1	0.867	0.894	0.962	0.968	0.960	0.971	0.980	0.982	0.694	0.761	0.750	0.886	0.915	0.932	0.926	0.960
α_2	0.842	0.863	0.897	0.913	0.956	0.962	0.973	0.977	0.713	0.783	0.741	0.881	0.926	0.934	0.925	0.956
α_3	0.894	0.908	0.960	0.965	0.959	0.968	0.982	0.984	0.703	0.776	0.757	0.886	0.921	0.932	0.930	0.962
α_4	0.862	0.885	0.948	0.957	0.960	0.969	0.980	0.984	0.688	0.767	0.694	0.878	0.913	0.929	0.934	0.963
α_5	0.834	0.855	0.913	0.925	0.953	0.959	0.973	0.976	0.703	0.772	0.703	0.874	0.924	0.931	0.931	0.959
α_6	0.978	0.978	0.974	0.984	0.960	0.964	0.956	0.969	0.968	0.968	0.974	0.979	0.920	0.923	0.937	0.971
α_7	0.982	0.982	0.963	0.981	0.971	0.970	0.948	0.973	0.969	0.969	0.975	0.978	0.917	0.923	0.916	0.961
α_8	0.981	0.981	0.975	0.982	0.971	0.969	0.956	0.972	0.971	0.970	0.970	0.979	0.926	0.927	0.941	0.975
α_9	0.980	0.980	0.971	0.980	0.960	0.962	0.954	0.975	0.976	0.977	0.966	0.978	0.921	0.921	0.915	0.957
α_{10}	0.981	0.982	0.972	0.982	0.968	0.971	0.950	0.974	0.969	0.971	0.969	0.977	0.921	0.924	0.915	0.969

References

- Chang, Y.-P., Chiu, C.-Y., & Tsai, R.-C. (2018). Nonparametric CAT for CD in Educational Settings With Small Samples. *Applied psychological measurement*, 0146621618813113.
- Chen, Y., Culpepper, S. A., Wang, S., & Douglas, J. (2018). A hidden Markov model for learning trajectories in cognitive diagnosis with application to spatial rotation skills. *Applied psychological measurement*, 42(1), 5-23.
- Cheng, Y. (2009). When cognitive diagnosis meets computerized adaptive testing: CD-CAT. *Psychometrika*, 74(4), 619.
- Cheng, Y. (2010). Improving cognitive diagnostic computerized adaptive testing by balancing attribute coverage: The modified maximum global discrimination index method. *Educational and Psychological Measurement*, 70(6), 902-913.
- Chiu, C.-Y., Douglas, J. A., & Li, X. (2009). Cluster analysis for cognitive diagnosis: Theory and applications. *Psychometrika*, 74(4), 633.
- De La Torre, J. (2011). The generalized DINA model framework. *Psychometrika*, 76(2), 179-199.
- De La Torre, J., & Douglas, J. A. (2004). Higher-order latent trait models for cognitive diagnosis. *Psychometrika*, 69(3), 333-353.
- Embretson, S. (1984). A general latent trait model for response processes. *Psychometrika*, 49(2), 175-186.
- Hanna, G. S. (1976). Effects of total and partial feedback in multiple-choice testing upon learning. *The Journal of Educational Research*, 69(5), 202-205.
- Henson, R. A., Templin, J. L., & Willse, J. T. (2009). Defining a family of cognitive diagnosis models using log-linear models with latent variables. *Psychometrika*, 74(2), 191.
- Huebner, A., & Wang, C. (2011). A note on comparing examinee classification methods for cognitive diagnosis models. *Educational and Psychological Measurement*, 71(2), 407-419.
- Kaplan, M., de la Torre, J., & Barrada, J. R. (2015). New item selection methods for cognitive diagnosis computerized adaptive testing. *Applied psychological measurement*, 39(3), 167-188.
- Kaya, Y., & Leite, W. L. (2017). Assessing change in latent skills across time with longitudinal cognitive diagnosis modeling: an evaluation of model performance. *Educational and Psychological Measurement*, 77(3), 369-388.
- Kingsbury, G. (2009). *Adaptive item calibration: A process for estimating item parameters within a computerized adaptive test*. Paper presented at the Proceedings of the 2009 GMAC Conference on Computerized Adaptive Testing.
- Kingsbury, G., Freeman, E., & Nesterak, M. (2014). The potential of adaptive assessment. *Educational Leadership*, 71(6). Retrieved from <http://www.ascd.org/publications/educational-leadership/mar14/vol71/num06/The-Potential-of-Adaptive-Assessment.aspx>
- Kluger, A. N., & DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological bulletin*, 119(2), 254-284.

- Köhn, H. F., & Chiu, C. Y. (2018). Attribute hierarchy models in cognitive diagnosis: identifiability of the latent attribute space and conditions for completeness of the Q-matrix. *Journal of Classification*, 1-25. <https://doi.org/10.1007/s00357-018-9278-6>
- Leighton, J., & Gierl, M. (2007). *Cognitive diagnostic assessment for education: Theory and applications*: Cambridge University Press.
- Li, F., Cohen, A., Bottge, B., & Templin, J. (2016). A latent transition analysis model for assessing change in cognitive skills. *Educational and Psychological Measurement*, 76(2), 181-204.
- Lu, J., & Wang, C. (2019). *A sequential exploratory approach to learning attribute hierarchies from data*. Paper presented at the 2019 International Meeting of Psychometric Society, Santiago, Chile.
- Madison, M. J., & Bradshaw, L. P. (2018). Assessing growth in a diagnostic classification model framework. *Psychometrika*, 83(4), 963-990.
- National Education Technology Plan (2017). Retrieved from <https://tech.ed.gov/netp/>
- Pellegrino, J.W., Chudowsky, N., Glaser, R. (Eds.) (2001). *Knowing what students know: the science and design of educational assessment*. Washington, DC: National Academy Press.
- Popham, W. J. (2008). Formative Assessment: Seven Stepping-Stones to Success. *Principal Leadership*, 9(1), 16-20.
- Quellmalz, E. S., & Pellegrino, J. W. (2009). Technology and testing. *Science*, 323(5910), 75-79.
- Shepard, L. A. (2009). Commentary: Evaluating the validity of formative and interim assessment. *Educational Measurement: Issues and Practice*, 28(3), 32-37.
- Templin, J., & Bradshaw, L. (2014). Hierarchical diagnostic classification models: A family of models for estimating and testing attribute hierarchies. *Psychometrika*, 79(2), 317-339.
- Templin, J., & Hoffman, L. (2013). Obtaining diagnostic classification model estimates using Mplus. *Educational Measurement: Issues and Practice*, 32(2), 37-50.
- Wang, C. (2013). Mutual information item selection method in cognitive diagnostic computerized adaptive testing with short test length. *Educational and Psychological Measurement*, 73(6), 1017-1035.
- Wang, C., & Lu, J. (2019). Learning attribute hierarchies from data: two exploratory approaches. PsyArXiv preprint. Retrived from <https://osf.io/5j2wu>
- Wang, C., Zheng, C., & Chang, H. H. (2014). An enhanced approach to combine item response theory with cognitive diagnosis in adaptive testing. *Journal of Educational Measurement*, 51(4), 358-380.
- Wilson, M. (2009). Measuring progressions: Assessment structures underlying a learning progression. *Journal of Research in Science Teaching*, 46(6), 716-730.
- Xu, G., Wang, C., & Shang, Z. (2016). On initial item selection in cognitive diagnostic computerized adaptive testing. *British Journal of Mathematical and Statistical Psychology*, 69(3), 291-315.