

Theory of self-learning Q -matrix

JINGCHEN LIU^{*}, GONGJUN XU^{**} and ZHILIANG YING[†]

Department of Statistics, Columbia University, 1255 Amsterdam Avenue, New York, NY 10027, USA.
E-mail: ^{*} jcliu@stat.columbia.edu; ^{**} gongjun@stat.columbia.edu; [†] zying@stat.columbia.edu

Cognitive assessment is a growing area in psychological and educational measurement, where tests are given to assess mastery/deficiency of attributes or skills. A key issue is the correct identification of attributes associated with items in a test. In this paper, we set up a mathematical framework under which theoretical properties may be discussed. We establish sufficient conditions to ensure that the attributes required by each item are learnable from the data.

Keywords: classification model; cognitive assessment; consistency; diagnostic; Q -matrix; self-learning

1. Introduction

Cognitive diagnosis has recently gained prominence in educational assessment, psychiatric evaluation, and many other disciplines. A key task is the correct specification of item-attribute relationships. A widely used mathematical formulation is the well known Q -matrix [27]. Under the setting of the Q -matrix, a typical modeling approach assumes a latent variable structure in which each subject possesses a vector of k attributes and responds to m items. The so-called Q -matrix is an $m \times k$ binary matrix establishing the relationship between responses and attributes by indicating the required attributes for each item. The entry in the i th row and j th column indicates if item i requires attribute j (see Example 2.3 for a demonstration of a Q -matrix). A short list of further developments of cognitive diagnosis models (CDMs) based on the Q -matrix includes the rule space method [28,29], the reparameterized unified/fusion model (RUM) [5,7,30], the conjunctive (noncompensatory) DINA and NIDA models [3,4,12,26,31], the compensatory DINO and NIDO models [31,32], the attribute hierarchy method [13], and clustering methods [1]; see also [11,23,33] for more approaches to cognitive diagnosis.

Statistical analysis with CDMs typically assumes a known Q -matrix provided by experts such as those who developed the questions [10,19,20,25]. Such a priori knowledge when correct is certainly very helpful for both model estimation and eventually identification of subjects' latent attributes. On the other hand, model fitting is usually sensitive to the choice of Q -matrix and its misspecification could seriously affect the goodness of fit. This is one of the main sources for lack of fit. Various diagnostic tools and testing procedures have been developed [2,8,9,14,21]. A comprehensive review of diagnostic classification models can be found in [22].

Despite the importance of the Q -matrix in cognitive diagnosis, its estimation problem is largely an unexplored area. Unlike typical inference problems, the inference for the Q -matrix is particularly challenging for the following reasons. First, in many cases, the Q -matrix is simply nonidentifiable. One typical situation is that multiple Q -matrices lead to an identical response distribution. Therefore, we only expect to identify the Q -matrix up to some equivalence relation

(Definition 2.2). In other words, two Q -matrices in the same equivalence class are not distinguishable based on data. Our first task is to define a meaningful and identifiable equivalence class. Second, the Q -matrix lives on a discrete space – the set of $m \times k$ matrices with binary entries. This discrete nature makes analysis particularly difficult because calculus tools are not applicable. In fact, most analyses are combinatorics based. Third, the model makes explicit distributional assumptions on the (unobserved) attributes, dictating the law of observed responses. The dependence of responses on attributes via Q -matrix is a highly nonlinear discrete function. The nonlinearity also adds to the difficulty of the analysis.

The primary purpose of this paper is to provide theoretical analyses on the learnability of the underlying Q -matrix. In particular, we obtain definitive answers to the identifiability of Q -matrix for one of the most commonly used models – the DINA model – by specifying a set of sufficient conditions under which the Q -matrix is identifiable up to an explicitly defined equivalence class. We also present the corresponding consistent estimators. We believe that the results (especially the intermediate results) and analysis strategies can be extended to other conjunctive models [12, 15, 18, 31, 32].

The rest of this paper is organized as follows. In Section 2, we present the basic inference result for Q -matrices in a conjunctive model with no slipping or guessing. In addition, we introduce all the necessary terminologies and technical conditions. In Section 3, we extend the results in Section 2 to the DINA model with known slipping and guessing parameters. In Section 4, we further generalize the results to the DINA model with unknown slipping parameters. Further discussion is provided in Section 5. Proofs are given in Section 6. Lastly, the proofs of two key propositions are given in the [Appendix](#).

2. Model specifications and basic results

We start the discussion with a simplified situation, under which the responses depend on the attribute profile deterministically (with no uncertainty). We describe our estimation procedure under this simple scenario. The results for the general cases are given in Sections 3 and 4.

2.1. Basic model specifications

The model specifications consist of the following concepts.

Attributes: subject's (unobserved) mastery of certain skills. Consider that there are k attributes. Let $\mathbf{A} = (A^1, \dots, A^k)^\top$ be the vector of attributes and $A^j \in \{0, 1\}$ be the indicator of the presence or absence of the j th attribute.

Responses: subject's binary responses to items. Consider that there are m items. Let $\mathbf{R} = (R^1, \dots, R^m)^\top$ be the vector of responses and $R^i \in \{0, 1\}$ be the response to the i th item.

Both \mathbf{A} and \mathbf{R} are subject specific. We assume that the integers m and k are known.

Q -matrix: the link between item responses and attributes. We define an $m \times k$ matrix $Q = (Q_{ij})_{m \times k}$. For each i and j , $Q_{ij} = 1$ when item i requires attribute j and 0 otherwise.

Furthermore, we define

$$\xi^i = \prod_{j=1}^k (A^j)^{Q_{ij}} = \mathbf{1}(A^j \geq Q_{ij} : j = 1, \dots, k), \tag{2.1}$$

which indicates whether a subject with attribute \mathbf{A} is capable of providing a positive response to item i . This model is conjunctive, meaning that it is necessary and sufficient to master all the necessary skills to be capable of solving one problem. Possessing additional attributes does not compensate for the absence of necessary attributes. In this section, we consider the simplest situation that there is no uncertainty in the response, that is,

$$R^i = \xi^i \tag{2.2}$$

for $i = 1, \dots, m$. Therefore, the responses are completely determined by the attributes. We assume that all items require at least one attribute. Equivalently, the Q -matrix does not have zero row vectors. Subjects who do not possess any attribute are not capable of responding positively to any item.

We use subscripts to indicate different subjects. For instance, $\mathbf{R}_r = (R_r^1, \dots, R_r^m)^\top$ is the response vector of subject r . Similarly, \mathbf{A}_r is the attribute vector of subject r . We observe $\mathbf{R}_1, \dots, \mathbf{R}_N$, where we use N to denote sample size. The attributes \mathbf{A}_r are not observed. Our objective is to make inference on the Q -matrix based on the observed responses.

2.2. Estimation of Q -matrix

We first introduce a few quantities for the presentation of an estimator.

T-matrix

In order to provide an estimator of Q , we first introduce one central quantity, the T -matrix, which connects the Q -matrix with the response and attribute distributions. Matrix $T(Q)$ has $2^k - 1$ columns each of which corresponds to one nonzero attribute vector, $\mathbf{A} \in \{0, 1\}^k \setminus \{(0, \dots, 0)\}$. Instead of labeling the columns of $T(Q)$ by ordinal numbers, we label them by binary vectors of length k . For instance, the \mathbf{A} th column of $T(Q)$ is the column that corresponds to attribute \mathbf{A} , for all $\mathbf{A} \neq (0, \dots, 0)$.

Let I_i be a generic notation for positive responses to item i . Let “ \wedge ” stand for “and” combination. For instance, $I_{i_1} \wedge I_{i_2}$ denotes positive responses to both items i_1 and i_2 . Each row of $T(Q)$ corresponds to one item or one “and” combination of items, for instance, I_{i_1} , $I_{i_1} \wedge I_{i_2}$ or $I_{i_1} \wedge I_{i_2} \wedge I_{i_3}, \dots$. If $T(Q)$ contains all the single items and all “and” combinations, $T(Q)$ contains $2^m - 1$ rows. We will later say that such a $T(Q)$ is *saturated* (Definition 2.1 in Section 2.4).

We now describe each row vector of $T(Q)$. We define that $B_Q(I_i)$ is a $2^k - 1$ dimensional row vector. Using the same labeling system as that of the columns of $T(Q)$, the \mathbf{A} th element of $B_Q(I_i)$ is defined as $\prod_{j=1}^k (A^j)^{Q_{ij}}$, which indicates if a subject with attribute \mathbf{A} is able to solve item i .

Using a similar notation, we define that

$$B_Q(I_{i_1} \wedge \dots \wedge I_{i_l}) = \Upsilon_{h=1}^l B_Q(I_{i_h}), \tag{2.3}$$

where the operator “ $\Upsilon_{h=1}^l$ ” is element-by-element multiplication from $B_Q(I_{i_1})$ to $B_Q(I_{i_l})$. For instance,

$$W = \Upsilon_{h=1}^l V_h$$

means that $W^j = \prod_{h=1}^l V_h^j$, where $W = (W^1, \dots, W^{2^k-1})$ and $V_h = (V_h^1, \dots, V_h^{2^k-1})$. Therefore, $B_Q(I_{i_1} \wedge \dots \wedge I_{i_l})$ is the vector indicating the attributes that are capable of responding positively to items i_1, \dots, i_l . The row in $T(Q)$ corresponding to $I_{i_1} \wedge \dots \wedge I_{i_l}$ is $B_Q(I_{i_1} \wedge \dots \wedge I_{i_l})$.

α -vector

We let α be a column vector the length of which equals to the number of rows of $T(Q)$. Each element of α corresponds to one row vector of $T(Q)$. The element in α corresponding to $I_{i_1} \wedge \dots \wedge I_{i_l}$ is defined as $N_{I_{i_1} \wedge \dots \wedge I_{i_l}} / N$, where $N_{I_{i_1} \wedge \dots \wedge I_{i_l}}$ denotes the number of people who have positive responses to items i_1, \dots, i_l , that is

$$N_{I_{i_1} \wedge \dots \wedge I_{i_l}} = \sum_{r=1}^N I(R_r^{ij} = 1 : j = 1, \dots, l).$$

For each $\mathbf{A} \in \{0, 1\}^k$, we let

$$\hat{p}_{\mathbf{A}} = \frac{1}{N} \sum_{r=1}^N I(\mathbf{A}_r = \mathbf{A}). \tag{2.4}$$

If (2.2) is strictly respected, then

$$T(Q)\hat{\mathbf{p}} = \alpha, \tag{2.5}$$

where $\hat{\mathbf{p}} = (\hat{p}_{\mathbf{A}} : \mathbf{A} \in \{0, 1\}^k \setminus \{(0, \dots, 0)\})$ is arranged in the same order as the columns of $T(Q)$. This is because each row of $T(Q)$ indicates the attribute profiles corresponding to subjects capable of responding positively to that set of item(s). Vector $\hat{\mathbf{p}}$ contains the proportions of subjects with each attribute profile. For each set of items, matrix multiplication sums up the proportions corresponding to each attribute profile capable of responding positively to that set of items, giving us the total proportion of subjects who respond positively to the corresponding items.

The estimator of the Q -matrix

For each $m \times k$ binary matrix Q' , we define

$$S(Q') = \inf_{\mathbf{p} \in [0, 1]^{2^k-1}} |T(Q')\mathbf{p} - \alpha|, \tag{2.6}$$

where $\mathbf{p} = (p_{\mathbf{A}} : \mathbf{A} \neq (0, \dots, 0))$. The above minimization is subject to the constraint that $\sum_{\mathbf{A} \neq (0, \dots, 0)} p_{\mathbf{A}} \in [0, 1]$. $|\cdot|$ is the Euclidean distance. An estimator of Q can be obtained by

minimizing $S(Q')$,

$$\hat{Q} = \arg \inf_{Q'} S(Q'), \tag{2.7}$$

where ‘‘arg inf’’ is the minimizer of the minimization problem over all $m \times k$ binary matrices. Note that the minimizers are not unique. We will later prove that the minimizers are in the same meaningful equivalence class. Because of (2.5), the true Q -matrix is always among the minimizers because $S(Q) = 0$.

2.3. Example

We illustrate the above construction by one simple example. We emphasize that this example is discussed to explain the estimation procedure for a concrete and simple example. The proposed estimator is certainly able to handle much larger Q -matrices. We consider the following 3×2 Q -matrix,

		Addition	Multiplication	
$Q =$	2 + 3	1	0	
	5 × 2	0	1	
	(2 + 3) × 2	1	1	(2.8)

There are two attributes and three items. We consider the contingency table of attributes,

	Multiplication	
Addition	\hat{p}_{00}	\hat{p}_{01}
	\hat{p}_{10}	\hat{p}_{11}

In the above table, \hat{p}_{00} is the proportional of people who do not master either addition or multiplication. Similarly, we define \hat{p}_{01} , \hat{p}_{10} and \hat{p}_{11} . $\{\hat{p}_{ij}; j = 0, 1\}$ is not observed.

Just for illustration, we construct a simple nonsaturated T -matrix. Suppose the relationship in (2.2) is strictly respected. Then, we should be able to establish the following identities:

$$N(\hat{p}_{10} + \hat{p}_{11}) = N_{I_1}, \quad N(\hat{p}_{01} + \hat{p}_{11}) = N_{I_2}, \quad N\hat{p}_{11} = N_{I_3}. \tag{2.9}$$

Therefore, if we let $\hat{\mathbf{p}} = (\hat{p}_{10}, \hat{p}_{01}, \hat{p}_{11})$, the above display imposes three linear constraints on the vector $\hat{\mathbf{p}}$. Together with the natural constraint that $\sum_{ij} \hat{p}_{ij} = 1$, $\hat{\mathbf{p}}$ solves linear equation,

$$T(Q)\hat{\mathbf{p}} = \alpha, \tag{2.10}$$

subject to the constraints that $\hat{\mathbf{p}} \in [0, 1]^3$ and $\hat{p}_{10} + \hat{p}_{01} + \hat{p}_{11} \in [0, 1]$, where

$$T(Q) = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \quad \alpha = \begin{pmatrix} N_{I_1}/N \\ N_{I_2}/N \\ N_{I_3}/N \end{pmatrix}. \tag{2.11}$$

Each column of $T(Q)$ corresponds to one attribute profile. The first column corresponds to $\mathbf{A} = (1, 0)$, the second column to $\mathbf{A} = (0, 1)$, and the third column to $\mathbf{A} = (1, 1)$. The first row corresponds to item 2 + 3, the second row to 5×2 and the last row to $(2 + 3) \times 2$. For this particular situation, $T(Q)$ has full rank and there exists one unique solution to (2.10). In fact, we would not expect the constrained solution to the linear equation in (2.10) to always exist unless (2.2) is strictly followed. This is the topic of the next section.

The identities in (2.9) only consider the marginal rate of each question. There are additional constraints if one considers “combinations” among items. For instance,

$$N \hat{p}_{11} = N_{I_1 \wedge I_2}.$$

People who are able to solve problem 3 must have both attributes and therefore are able to solve both problems 1 and 2. Again, if (2.2) is not strictly followed, this is not necessarily respected in the real data, though it is a logical conclusion. The DINA in the next section handles such a case. Upon considering I_1, I_2, I_3 and $I_1 \wedge I_2$, the new T -matrix is

$$T(Q) = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \quad \alpha = \begin{pmatrix} N_{I_1}/N \\ N_{I_2}/N \\ N_{I_3}/N \\ N_{I_1 \wedge I_2}/N \end{pmatrix}. \tag{2.12}$$

The last row is added corresponding to $I_1 \wedge I_2$. With (2.2) in force, we have

$$S(Q) = \inf_{\mathbf{p} \in [0,1]^3} |T(Q)\mathbf{p} - \alpha| = |T(Q)\hat{\mathbf{p}} - \alpha| = 0, \tag{2.13}$$

if Q is the true matrix.

2.4. Basic results

Before stating the main result, we provide a list of notations, which will be used in the discussions.

- Linear space spanned by vectors V_1, \dots, V_l :

$$\mathcal{L}(V_1, \dots, V_l) = \left\{ \sum_{j=1}^l a_j V_j : a_j \in \mathbb{R} \right\}.$$

- For a matrix M , $M_{1:l}$ denotes the submatrix containing the first l rows and all columns of M .
- Vector e_i denotes a column vector such that the i th element is one and the rest are zero. When there is no ambiguity, we omit the length index of e_i .
- Matrix \mathcal{I}_l denotes the $l \times l$ identity matrix.
- For a matrix M , $C(M)$ is the linear space generated by the column vectors of M . It is usually called the *column space* of M .

- C_M denotes the set of column vectors of M .
- R_M denotes the set of row vectors of M .
- Vector $\mathbf{0}$ denotes the zero vector, $(0, \dots, 0)$. When there is no ambiguity, we omit the index of length.
- Scalar $p_{\mathbf{A}}$ denotes the probability that a subject has attribute profile \mathbf{A} . For instance, p_{10} is the probability that a subject has attribute one but not attribute two.
- Define a $2^k - 1$ dimensional vector

$$\mathbf{p} = (p_{\mathbf{A}} : \mathbf{A} \in \{0, 1\}^k \setminus \{\mathbf{0}\}).$$

- Let c and g be two m dimensional vectors. We write $c > g$ if $c_i > g_i$ for all $1 \leq i \leq m$.
- We write $c \not\geq g$ if $c_i \neq g_i$ for all $i = 1, \dots, m$.
- Matrix Q denotes the true matrix and Q' denotes a generic $m \times k$ binary matrix.

The following definitions will be used in subsequent discussions.

Definition 2.1. We say that $T(Q)$ is saturated if all combinations of form $I_{i_1} \wedge \dots \wedge I_{i_l}$, for $l = 1, \dots, m$, are included in $T(Q)$.

Definition 2.2. We write $Q \sim Q'$ if and only if Q and Q' have identical column vectors, which could be arranged in different orders; otherwise, we write $Q \not\sim Q'$.

Remark 2.1. It is not hard to show that “ \sim ” is an equivalence relation. $Q \sim Q'$ if and only if they are identical after an appropriate permutation of the columns. Each column of Q is interpreted as an attribute. Permuting the columns of Q is equivalent to relabeling the attributes. For $Q \sim Q'$, we are not able to distinguish Q from Q' based on data.

Definition 2.3. A Q -matrix is said to be complete if $\{e_i : i = 1, \dots, k\} \subset R_Q$ (R_Q is the set of row vectors of Q); otherwise, we say that Q is incomplete.

A Q -matrix is complete if and only if for each attribute there exists an item only requiring that attribute. Completeness implies that $m \geq k$. We will show that completeness is among the sufficient conditions to identify Q .

Remark 2.2. One of the main objectives of cognitive assessment is to identify the subjects' attributes; see [22] for other applications. It has been established in [1] that the completeness of the Q -matrix is a sufficient and necessary condition for a set of items to consistently identify attributes if (2.2) is strictly followed. Thus, it is usually recommended to use a complete Q -matrix. For a precise formulation, see [1].

Listed below are assumptions which will be used in subsequent development.

- C1 Q is complete.
- C2 $T(Q)$ is saturated.

C3 $\mathbf{A}_1, \dots, \mathbf{A}_N$ are i.i.d. random vectors following distribution

$$P(\mathbf{A}_r = \mathbf{A}) = p_{\mathbf{A}}^*.$$

We further let $\mathbf{p}^* = (p_{\mathbf{A}}^* : \mathbf{A} \in \{0, 1\} \setminus \{\mathbf{0}\})$.

C4 $(p_{\mathbf{0}}^*, \mathbf{p}^*) > \mathbf{0}$.

C5 Each attribute has been required by at least two items.

With these preparations, we are ready to introduce the first theorem, the proof of which is given in Section 6.

Theorem 2.4. *Assume that conditions C1–C5 are in force. Suppose that for subject r the response corresponding to item i follows*

$$R_r^i = \xi_r^i = \prod_{j=1}^k (A_r^j)^{Q_{ij}}.$$

Let \hat{Q} , defined in (2.7), be a minimizer of $S(Q')$ among all $m \times k$ binary matrices, where $S(Q')$ is defined in (2.6). Then,

$$\lim_{N \rightarrow \infty} P(\hat{Q} \sim Q) = 1. \tag{2.14}$$

Further, let

$$\tilde{\mathbf{p}} = \arg \inf_{\mathbf{p}} |T(\hat{Q})\mathbf{p} - \alpha|^2. \tag{2.15}$$

With an appropriate rearrangement of the columns of \hat{Q} , for any $\varepsilon > 0$

$$\lim_{N \rightarrow \infty} P(|\tilde{\mathbf{p}} - \mathbf{p}^*| \leq \varepsilon) = 1.$$

Remark 2.3. If $Q_1 \sim Q_2$, the two matrices only differ by a column permutation and will be considered to be the “same”. Therefore, we expect to identify the equivalence class that Q belongs to. Also, note that $S(Q_1) = S(Q_2)$ if $Q_1 \sim Q_2$.

Remark 2.4. In order to obtain the consistency of \hat{Q} (subject to a column permutation), it is necessary to have \mathbf{p}^* not living on some sub-manifold. To see a counter example, suppose that $P(\mathbf{A}_r = (1, \dots, 1)^\top) = p_{1\dots 1}^* = 1$. Then, for all Q , $P(\mathbf{R}_r = (1, \dots, 1)^\top) = 1$, that is, all subjects are able to solve all problems. Therefore, the distribution of \mathbf{R} is independent of Q . In other words, the Q -matrix is not identifiable. More generally, if there exist A_r^i and A_r^j such that $P(A_r^i = A_r^j) = 1$, then the Q -matrix is not identifiable based on the data. This is because one cannot tell if an item requires attribute i alone, attribute j alone, or both; see [16,17] for similar cases for the multidimensional IRT models.

Remark 2.5. Note that the estimator of the attribute distribution, $\tilde{\mathbf{p}}$, in (2.15) depends on the order of columns of \hat{Q} . In order to achieve consistency, we will need to arrange the columns of \hat{Q} such that $\hat{Q} = Q$ whenever $\hat{Q} \sim Q$.

Remark 2.6. One practical issue associated with the proposed procedure is the computation. For a specific Q , the computation of $S(Q)$ only involves a constraint minimization of a quadratic function. However, if m or k is large, the computation overhead of searching the minimizer of $S(Q)$ over the space of $m \times k$ matrices could be substantial. One practical solution is to break the Q -matrix into smaller sub-matrices. For instance, one may divide the m items into l groups (possibly with nonempty overlap across different groups). Then apply the proposed estimator to each of the l group of items. This is equivalent to breaking a big m by k Q -matrix into several smaller matrices and estimating each of them separately. Lastly, combine the l estimated sub-matrices together to form a single estimate. The consistency results can be applied to each of the l sub-matrices and therefore the combined matrix is also a consistent estimator. A similar technique has been discussed in Chapter 8.6 of [29].

Remark 2.7. Conditions C1 and C2 are imposed to guarantee consistency. They may not be always necessary. Furthermore, constructing a saturated T -matrix is sometimes computationally not feasible, especially when the number of items is large. In practice, one may include the combinations of one item, two items, \dots , j items. The choice of j depends on the sample size and the computation resources. The condition C5 is required for technical purposes. Nonetheless, one can in fact construct counterexamples, that is, the Q -matrix is not identifiable up to the relationship “ \sim ”, if C5 is violated.

3. DINA model with known slipping and guessing parameters

3.1. Model specification

In this section, we extend the inference results in the previous section to the situation under which the responses do not deterministically depend on the attributes. In particular, we consider the DINA (Deterministic Input, Noisy Output “AND” gate) model [12]. We would like to introduce two parameters: the slipping parameter (s_i) and the guessing parameter (g_i). Here $1 - s_i$ (g_i) is the probability of a subject’s responding positively to item i given that s/he is capable of solving that problem. To simplify the notations, we denote $1 - s_i$ by c_i . An extension of (2.2) to include slipping and guessing specifies the response probabilities as

$$P(R^i = 1 | \xi^i) = c_i^{\xi^i} g_i^{1 - \xi^i}, \quad (3.1)$$

where ξ^i is the capability indicator defined in (2.1). In addition, conditional on $\{\xi^1, \dots, \xi^m\}$, $\{R^1, \dots, R^m\}$ are jointly independent.

In this context, the T -matrix needs to be modified accordingly. Throughout this section, we assume that both c_i ’s and g_i ’s are known. We discuss the case that c_i ’s are unknown in the next section.

We first consider the case that $g_i = 0$ for all $i = 1, \dots, m$. We introduce a diagonal matrix D_c . If the h th row of matrix $T_c(Q)$ corresponds to $I_{i_1} \wedge \dots \wedge I_{i_l}$, then the h th diagonal element of D_c is $c_{i_1} \times \dots \times c_{i_l}$. Then, we let

$$T_c(Q) = D_c T(Q), \tag{3.2}$$

where $T(Q)$ is the binary matrix defined previously. In other words, we multiply each row of $T(Q)$ by a common factor and obtain $T_c(Q)$. Note that in absence of slipping ($c_i = 1$ for each i) we have that $T_c(Q) = T(Q)$.

There is another equivalent way of constructing $T_c(Q)$. We define

$$B_{c,Q}(I_j) = c_j B_Q(I_j)$$

and

$$B_{c,Q}(I_{i_1} \wedge \dots \wedge I_{i_l}) = \Upsilon_{h=1}^l B_{c,Q}(I_{i_h}), \tag{3.3}$$

where “ Υ ” refers to element by element multiplication. Let the row vector in $T_c(Q)$ corresponding to $I_{i_1} \wedge \dots \wedge I_{i_l}$ be $B_{c,Q}(I_{i_1} \wedge \dots \wedge I_{i_l})$.

For instance, with $c = (c_1, c_2, c_3)$, the $T_c(Q)$ corresponding to the T -matrix in (2.12) would be

$$T_c(Q) = \begin{pmatrix} c_1 & 0 & c_1 \\ 0 & c_2 & c_2 \\ 0 & 0 & c_3 \\ 0 & 0 & c_1 c_2 \end{pmatrix}. \tag{3.4}$$

Lastly, we consider the situation that both the probability of making a mistake and the probability of guessing correctly could be strictly positive. By this, we mean that the probability that a subject responds positively to item i is c_i if s/he is capable of doing so; otherwise the probability is g_i . We create a corresponding $T_{c,g}(Q)$ by slightly modifying $T_c(Q)$. We define row vector

$$\mathbb{E} = (1, \dots, 1).$$

When there is no ambiguity, we omit the length index of \mathbb{E} . Now, let

$$B_{c,g,Q}(I_i) = g_i \mathbb{E} + (c_i - g_i) B_Q(I_i)$$

and

$$B_{c,g,Q}(I_{i_1} \wedge \dots \wedge I_{i_l}) = \Upsilon_{h=1}^l B_{c,g,Q}(I_{i_h}). \tag{3.5}$$

Let the row vector in $T_{c,g}(Q)$ corresponding to $I_{i_1} \wedge \dots \wedge I_{i_l}$ be $B_{c,g,Q}(I_{i_1} \wedge \dots \wedge I_{i_l})$. For instance, the matrix $T_{c,g}$ corresponding to the $T_c(Q)$ in (3.4) is

$$T_{c,g}(Q) = \begin{pmatrix} c_1 & g_1 & c_1 \\ g_2 & c_2 & c_2 \\ g_3 & g_3 & c_3 \\ c_1 g_2 & g_1 c_2 & c_1 c_2 \end{pmatrix}. \tag{3.6}$$

3.2. Estimation of the Q -matrix and consistency results

Having concluded our preparations, we are now ready to introduce our estimators for Q . Given c and g , we define

$$S_{c,g}(Q) = \inf_{\mathbf{p}' \in [0,1]^{2^k-1}} |T_{c,g}(Q)\mathbf{p}' + p'_0\mathbf{g} - \alpha|, \tag{3.7}$$

where $\mathbf{p}' = (p'_A : A \in \{0, 1\}^k \setminus \{\mathbf{0}\})$, $p'_0 = p'_{0\dots 0}$ and

$$\mathbf{g} = \begin{pmatrix} g_1 \\ \vdots \\ g_k \\ g_1 g_2 \\ \vdots \\ g_{k-1} g_k \\ g_1 g_2 g_3 \\ \vdots \end{pmatrix} \begin{matrix} I_1 \\ \vdots \\ I_k \\ I_1 \wedge I_2 \\ \vdots \\ I_{k-1} \wedge I_k \\ I_1 \wedge I_2 \wedge I_3 \\ \vdots \end{matrix} \tag{3.8}$$

The labels to the right of the vector indicate the corresponding row vectors in $T_{c,g}(Q)$. The minimization in (3.7) is subject to constraints that

$$p'_A \in [0, 1] \quad \text{and} \quad \sum_A p'_A = 1.$$

The vector \mathbf{g} contains the probabilities of providing positive responses to items simply by guessing. We propose an estimator of the Q -matrix through a minimization problem, that is,

$$\hat{Q}(c, g) = \arg \inf_{Q'} S_{c,g}(Q'). \tag{3.9}$$

We write c and g in the argument to emphasize that the estimator depends on c and g . The computation of the minimization in (3.7) consists of minimizing a quadratic function subject to finitely many linear constraints. Therefore, it can be done efficiently.

Theorem 3.1. *Suppose that c and g are known and that conditions C1–C5 are in force. For subject r , the responses are generated independently such that*

$$P(R_r^i = 1 | \xi_r^i) = c_i^{\xi_r^i} g_i^{1-\xi_r^i}, \tag{3.10}$$

where ξ_r^i is defined as in Theorem 2.4. Let $\hat{Q}(c, g)$ be defined as in (3.9). If $c_i \neq g_i$ for all i and $T_{c-g}(Q)\mathbf{p}^*$ does not have zero elements, then

$$\lim_{N \rightarrow \infty} P(\hat{Q}(c, g) \sim Q) = 1.$$

Furthermore, let

$$\tilde{\mathbf{p}}(c, g) = \arg \inf_{\mathbf{p}} |T_{c,g}(\hat{Q}(c, g))\mathbf{p} + p_0\mathbf{g} - \alpha|^2,$$

subject to constraint that $\sum_{\mathbf{A}} p_{\mathbf{A}} = 1$. Then, with an appropriate rearrangement of the columns of \hat{Q} , for any $\varepsilon > 0$,

$$\lim_{N \rightarrow \infty} P(|\tilde{\mathbf{p}}(c, g) - \mathbf{p}^*| \leq \varepsilon) = 1.$$

Remark 3.1. There are various metrics one can employ to measure the distance between the vectors $T_{c,g}(\hat{Q}(c, g))\mathbf{p} + p_0\mathbf{g}$ and α . In fact, any metric that generates the same topology as the Euclidian metric is sufficient to obtain the consistency results in the theorem. For instance, a principled choice of objective function would be the likelihood with \mathbf{p} profiled out. The reason we prefer the Euclidian metric (versus, for instance, the full likelihood) is that the evaluation of $S(Q)$ is easier than the evaluation based on other metrics. More specifically, the computation of current $S(Q)$ consists of quadratic programming types of well oiled optimization techniques.

4. Extension to the situation with unknown slipping probabilities

In this section, we further extend our results to the situation where the slipping probabilities are unknown and the guessing probabilities are known. In the context of standard exams, the guessing probabilities can typically be set to zero for open problems. For instance, the chance of guessing the correct answer to “ $(3 + 2) \times 2 = ?$ ” is very small. On the other hand, for multiple choice problems, the guessing probabilities cannot be ignored. In that case, g_i can be considered as $1/n$ when there are n choices; see Remark 4.2 for more description.

4.1. Estimator of c

We provide two estimators of c given Q and g . One is applicable to all Q -matrices, but computationally intensive. The other is computationally easy, but requires certain structures of Q . Then, we combine them into a single estimator.

A general estimator

We first provide an estimator of c that is applicable to all Q -matrices. Considering that the estimator of Q minimizes the objective function $S_{c,g}(Q)$, we propose the following estimator of c :

$$\tilde{c}(Q, g) = \arg \inf_{c \in [0,1]^m} S_{c,g}(Q). \tag{4.1}$$

A moment estimator

The computation of $\tilde{c}(Q, g)$ is typically intensive. When the Q -matrix has a certain structure, we are able to estimate c consistently based on estimating equations.

For a particular item i , suppose that there exist items i_1, \dots, i_l (different from i) such that

$$B_Q(I_i \wedge I_{i_1} \wedge \dots \wedge I_{i_l}) = B_Q(I_{i_1} \wedge \dots \wedge I_{i_l}), \tag{4.2}$$

that is, the attributes required by item i are a subset of the attributes required by i_1, \dots, i_l .

Let $c - g = (c_1 - g_1, \dots, c_m - g_m)$ and

$$\tilde{T}_{c,g}(Q) = \begin{pmatrix} \mathbf{g} & T_{c,g}(Q) \\ \mathbf{1} & \mathbb{E} \end{pmatrix}.$$

We borrow a result which will be given in the proof of Proposition 6.6 (Section 6.1) to say that there exists a matrix D (only depending on g) such that

$$D\tilde{T}_{c,g}(Q) = (\mathbf{0}, T_{c-g}(Q)).$$

Let a_g and a_{*g} be the row vectors in D corresponding to $I_{i_1} \wedge \dots \wedge I_{i_l}$ and $I_i \wedge I_{i_1} \wedge \dots \wedge I_{i_l}$ (in $T_{c-g}(Q)$).

Then,

$$\begin{aligned} \frac{a_{*g}^\top(\alpha)}{a_g^\top(\alpha)} &= \frac{a_{*g}^\top \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix}}{a_g^\top \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix}} + o_p(1) \\ &= \frac{B_{c-g,Q}(I_i \wedge I_{i_1} \wedge \dots \wedge I_{i_l}) \mathbf{p}^*}{B_{c-g,Q}(I_{i_1} \wedge \dots \wedge I_{i_l}) \mathbf{p}^*} + o_p(1) \xrightarrow{p} (c_i - g_i), \end{aligned} \tag{4.3}$$

where the vectors a_g and a_{*g} only depend on g .

Therefore, the corresponding estimator of c_i would be

$$\bar{c}_i(Q, g) = g_i + \frac{a_{*g}^\top(\alpha)}{a_g^\top(\alpha)}. \tag{4.4}$$

Note that the computation of $\bar{c}_i(Q, g)$ only consists of affine transformations and therefore is very fast.

Proposition 4.1. *Suppose conditions C3, (3.10) and (4.2) are true. Then $\bar{c}_i \rightarrow c_i$ in probability as $N \rightarrow \infty$.*

Proof. By the law of large numbers,

$$a_{*g}^\top \begin{pmatrix} \alpha \\ \mathbf{1} \end{pmatrix} - a_{*g}^\top \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} \rightarrow 0, \quad a_g^\top \begin{pmatrix} \alpha \\ \mathbf{1} \end{pmatrix} - a_g^\top \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} \rightarrow 0,$$

in probability as $N \rightarrow \infty$. By the construction of a_{*g} and a_g , we have

$$a_{*g}^\top \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} = B_{c-g,Q}(I_i \wedge I_{i_1} \wedge \dots \wedge I_{i_l}) \mathbf{p}^*,$$

$$a_g^\top \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} = B_{c-g,Q}(I_{i_1} \wedge \dots \wedge I_{i_l}) \mathbf{p}^*.$$

Thanks to (4.2), we have

$$\frac{a_{*g}^\top \begin{pmatrix} \alpha \\ 1 \end{pmatrix}}{a_g^\top \begin{pmatrix} \alpha \\ 1 \end{pmatrix}} \rightarrow c_i - g_i. \quad \square$$

Combined estimator

Lastly, we combine \bar{c}_i and \tilde{c}_i . For each Q , we write $c = (c^*, c^{**})$. For each c_i in the sub-vector c^* , (4.2) holds. Let $\bar{c}^*(Q, g)$ be defined in (4.4) (element by element). For c^{**} , we let $\tilde{c}^{**}(Q, g) = \arg \inf_{c^{**}} S_{(\bar{c}^*(Q,g), c^{**}),g}(Q)$. Finally, let $\hat{c}(Q, g) = (\bar{c}^*(Q, g), \tilde{c}^{**}(Q, g))$. Furthermore, each element of $\hat{c}(Q, g)$ greater than one is set to be one and each element less than zero is set to be zero. Equivalently, we impose the constraint that $\hat{c}(Q, g) \in [0, 1]^m$.

4.2. Consistency result

Theorem 4.2. *Suppose that g is known and the conditions in Theorem 3.1 hold. Let*

$$\hat{Q}_{\hat{c}}(g) = \arg \inf_{Q'} S_{\hat{c}(Q',g),g}(Q'), \quad \tilde{\mathbf{p}}_{\hat{c}}(g) = \arg \inf_{\mathbf{p}} |T_{\hat{c}(\hat{Q},g),g}(\hat{Q}_{\hat{c}}(g)) \mathbf{p} + p_0 \mathbf{g} - \alpha|.$$

The second optimization is subject to constraint that $\sum_{\mathbf{A}} p_{\mathbf{A}} = 1$. Then,

$$\lim_{N \rightarrow \infty} P(\hat{Q}_{\hat{c}}(g) \sim Q) = 1.$$

Furthermore, if the estimator $\tilde{c}(Q, g)$, defined in (4.1), is consistent, then by appropriately rearranging the columns of $\hat{Q}_{\hat{c}}(g)$, for any $\varepsilon > 0$,

$$\lim_{N \rightarrow \infty} P(|\tilde{\mathbf{p}}_{\hat{c}}(g) - \mathbf{p}^*| \leq \varepsilon) = 1.$$

Remark 4.1. The consistency of $\hat{Q}_{\hat{c}}(g)$ does not rely on the consistency of $\tilde{c}(Q, g)$, which is mainly because of the central intermediate result in Proposition 6.6. The consistency of $\tilde{c}(Q, g)$ is a necessary condition for the consistency of $\tilde{\mathbf{p}}_{\hat{c}}(g)$.

For most usual situations, (\mathbf{p}^*, c) is estimable based on the data given a correctly specified Q -matrix. Nonetheless, there are some rare occasions in which nonidentifiability does exist. We provide one example, explained at the intuitive level, to illustrate that it is not always possible to consistently estimate c and \mathbf{p}^* . This example is simply to justify that the existence of the consistent estimator for c in the above theorem is not an empty assumption. Consider a complete

matrix $Q = \mathcal{I}_k$. The degrees of freedom of a k -way binary table is $2^k - 1$. On the other hand, the dimension of parameters (\mathbf{p}^*, c) is $2^k - 1 + k$. Therefore, \mathbf{p}^* and c cannot be consistently identified without additional information. This problem is typically tackled by introducing additional parametric assumptions such as \mathbf{p}^* satisfying certain functional form or in the Bayesian setting (weakly) informative prior distributions [6]. Given that the emphasis of this paper is the inference of Q -matrix, we do not further investigate the identifiability of (\mathbf{p}^*, c) . Nonetheless, estimation for (\mathbf{p}^*, c) is definitely an important issue.

Remark 4.2. Assuming that the guessing probability g_i being known is somewhat strong. For complicated situations, such as for multiple choice problems the incorrect choices do not look “equally incorrect”, the guessing probability is typically not $1/n$. In Theorem 4.2, we make this assumption mostly for technical reasons.

One can certainly provide the same treatment to the unknown guessing probabilities just as to the slipping probabilities by plugging in a consistent estimator of g_i or profiling it out (like \tilde{c}). However, the rigorous establishment of the consistency results is certainly much more difficult and additional technical conditions may be needed. We leave the analysis of the problem with unknown guessing probability to the future study.

5. Discussion

This paper provides basic theoretical results of the estimation of Q -matrix, a key element in modern cognitive diagnosis. Under the conjunctive model assumption, sufficient conditions are developed for the Q -matrix to be identifiable up to an equivalence relation and the corresponding consistent estimators are constructed. The equivalence relation defines a natural partition of the space of Q -matrices and may be viewed as the finest “resolution” that is possibly distinguishable based on the data, unless there is additional information about the specific meaning of each attribute. Our results provide the first steps for statistical inference about Q -matrices by explicitly specifying the conditions under which two Q -matrices lead to different response distributions. We believe that these results, especially the intermediate results in Section 6, can also be applied to general conjunctive models.

There are several directions along which further exploration may be pursued. First, some conditions may be modified to reflect practical circumstance. For instance, if the population is not fully diversified, meaning that certain attribute profiles may never exist, then condition C4 cannot be expected to hold. To ensure identifiability, we will need to impose certain structures on the Q -matrix. In the addition-multiplication example of Section 2.3, if individuals capable of multiplication are also capable of addition, then we may need to impose the natural constraint that every item that requires multiplication should also require addition, which also implies that the Q -matrix is never complete.

Second, when an a priori “expert’s” knowledge of the Q -matrix is available, we may wish to incorporate such information into the estimation. This could be in the form of an additive penalty function attached to the objective function S . Such information, if correct, not only improves estimation accuracy but also reduces the computational complexity – one can just perform a minimization of $S(Q)$ in a neighborhood around the expert’s Q -matrix.

Third, throughout this paper we assume that the number of attributes (dimension) is known. In practice, it would be desirable to develop a data driven way to estimate the dimension, not only to deal with the situation of unknown dimension, but also to check if the assumed dimension is correct. One possible way to tackle the problem is to introduce a penalty function similar to that of BIC [24] which would give a consistent estimator of the Q -matrix even if the dimension is unknown.

Fourth, one issue of both theoretical and practical importance is the inference of the parameters additional to the Q -matrix, such as the slipping ($s = 1 - c$), guessing (g) parameters and the attribute distribution \mathbf{p}^* . In the current paper, given that the main interesting parameter is the Q -matrix, the estimations of \mathbf{p}^* and c are treated as by-product of the main results. On the other hand, given a known Q , the identifiability and estimation of these parameters are important topics. In the previous discussion, we provided a few examples for potential identifiability issues. Further careful investigation is definitely of great importance and challenges.

Fifth, the rate of convergence of the estimator \hat{Q} is not only of theoretical importance. From a practical point of view, it is crucial to study the rate of convergence as the scale of the problem becomes large in terms of the number of attributes and number of items.

Lastly, the optimization of $S(Q)$ over the space of $m \times k$ binary matrices is a nontrivial problem. It consists of evaluating the function S $2^{m \times k}$ times. This is a substantial computational load if m and k are reasonably large. As mentioned previously, this computation might be reduced by additional information about the Q -matrix or splitting the Q -matrix into small sub-matrices. Nevertheless, it would be highly desirable to explore the structures of the Q -matrix and the function S so as to compute \hat{Q} more efficiently.

6. Proofs of the theorems

6.1. Several propositions and lemmas

To make the discussion smooth, we postpone several long proofs to the [Appendix](#).

Proposition 6.1. *Suppose that Q is complete and matrix $T(Q)$ is saturated. Then, we are able to arrange the columns and rows of Q and $T(Q)$ such that $T(Q)_{1:(2^k-1)}$ has full rank and $T(Q)$ has full column rank.*

Proof. Provided that Q is complete, without loss of generality we assume that the i th row vector of Q is e_i^\top for $i = 1, \dots, k$, that is, item i only requires attribute i for each $i = 1, \dots, k$. Let the first $2^k - 1$ rows of $T(Q)$ be associated with $\{I_1, \dots, I_k\}$. In particular, we let the first k rows correspond to I_1, \dots, I_k and the first k columns of $T(Q)$ correspond to \mathbf{A} 's that only have one attribute. We further arrange the next C_2^k rows of $T(Q)$ to correspond to combinations of two items, $I_i \wedge I_j$, $i \neq j$. The next C_2^k columns of $T(Q)$ correspond to \mathbf{A} 's that only have two positive attributes. Similarly, we arrange $T(Q)$ for combinations of three, four, and up to k items. Therefore, the first $2^k - 1$ rows of $T(Q)$ admit a block upper triangle form. In addition, we are able to further arrange the columns within each block such that the diagonal matrices are

identities, so that $T(Q)$ has form

$$\begin{matrix} I_1, I_2, \dots \\ I_1 \wedge I_2, I_1 \wedge I_3, \dots \\ I_1 \wedge I_2 \wedge I_3, \dots \\ \vdots \end{matrix} \begin{pmatrix} \mathcal{I}_k & * & * & * & \dots \\ 0 & \mathcal{I}_{C_2^k} & * & * & \\ 0 & 0 & \mathcal{I}_{C_3^k} & * & \\ \vdots & \vdots & \vdots & & \end{pmatrix}. \tag{6.1}$$

Note that $T(Q)$ has $2^k - 1$ columns and $T(Q)_{1:(2^k-1)}$ obviously has full rank, therefore $T(Q)$ has full column rank. □

From now on, we assume that $Q_{1:k} = \mathcal{I}_k$ and the first $2^k - 1$ rows of $T(Q)$ are arranged in the order as in (6.1).

Proposition 6.2. *Suppose that Q is complete, $T(Q)$ is saturated, and $c \not\cong \mathbf{0}$. Then, $T_c(Q)$ and $T_c(Q)_{1:(2^k-1)}$ have full column rank.*

Proof. By Proposition 6.1, (3.2) and the fact that D_c is a diagonal matrix of full rank as long as $c \not\cong \mathbf{0}$,

$$T_c(Q) = D_c T(Q),$$

is of full column rank. □

The following two propositions, which compare the column spaces of $T_c(Q)$ and $T_c(Q')$, are central to the proof of all the theorems. Their proofs are delayed to the [Appendix](#).

The first proposition discusses the case where $Q'_{1:k}$ is complete. We can always rearrange the columns of Q' so that $Q_{1:k} = Q'_{1:k}$. In addition, according to the proof of Proposition 6.1, the last column vector of $T_c(Q)$ corresponds to attribute $\mathbf{A} = (1, \dots, 1)^\top$. Therefore, this column vector is all of nonzero entries.

Proposition 6.3. *Assume that Q is a complete matrix and $T(Q)$ is saturated. Without loss of generality, let $Q_{1:k} = \mathcal{I}_k$. Assume that the first k rows of Q' form a complete matrix. Further, assume that $Q_{1:k} = Q'_{1:k} = \mathcal{I}_k$. If $Q' \neq Q$ and $c \not\cong \mathbf{0}$, under the conditions in Theorem 4.2, $T_c(Q)\mathbf{p}^*$ is not in the column space $C(T_{c'}(Q'))$ for all $c' \in \mathbb{R}^m$.*

The next proposition discusses the case where $Q'_{1:k}$ is incomplete.

Proposition 6.4. *Assume that Q is a complete matrix and $T(Q)$ is saturated. Without loss of generality, let $Q_{1:k} = \mathcal{I}_k$. If $c \not\cong \mathbf{0}$ and $Q'_{1:k}$ is incomplete, under the conditions in Theorem 4.2, $T_c(Q)\mathbf{p}^*$ is not in the column space $C(T_{c'}(Q'))$ for all $c' \in \mathbb{R}^m$.*

The next result is a direct corollary of these two propositions. It follows by setting $c_i = 1$ and $g_i = 0$ for all $i = 1, \dots, m$.

Corollary 6.5. *If $Q \approx Q'$, under the conditions of Theorem 4.2, $T_c(Q)\mathbf{p}^*$ is not in the column space $C(T_{c'}(Q'))$ for all $c' \in [0, 1]^m$.*

To obtain a similar proposition for the cases where the g_i 's are nonzero, we will need to expand the $T_{c,g}(Q)$ as follows. As previously defined, let

$$\tilde{T}_{c,g}(Q) = \begin{pmatrix} \mathbf{g} & T_{c,g}(Q) \\ \mathbf{1} & \mathbb{E} \end{pmatrix}. \tag{6.2}$$

The last row of $\tilde{T}_{c,g}(Q)$ consists entirely of ones. Vector \mathbf{g} is defined as in (3.8).

Proposition 6.6. *Suppose that Q is a complete matrix, $Q' \approx Q$, T is saturated and $c \not\approx g$. Let $\mathbf{p}_0^* = (p_0^*, (\mathbf{p}^*)^\top)^\top$. Under the conditions of Theorem 4.2, $\tilde{T}_{c,g}(Q)\mathbf{p}_0^*$ is not in the column space $C(\tilde{T}_{c',g}(Q'))$ for all $c' \in [0, 1]^m$. In addition, $\tilde{T}_{c,g}(Q)$ is of full column rank.*

To prove Proposition 6.6, we will need the following lemma.

Lemma 6.7. *Consider two matrices T_1 and T_2 of the same dimension. If $T_1\mathbf{p} \in C(T_2)$, then for any matrix D of appropriate dimension for multiplication, we have*

$$DT_1\mathbf{p} \in C(DT_2).$$

Conversely, if for some D , $DT_1\mathbf{p}$ does not belong to $C(DT_2)$, then $T_1\mathbf{p}$ does not belong to $C(T_2)$.

Proof. Note that DT_i is just a linear row transform of T_i for $i = 1, 2$. The conclusion is immediate by basic linear algebra. □

Proof of Proposition 6.6. Thanks to Lemma 6.7, we only need to find a matrix D such that $D\tilde{T}_{c,g}(Q)\mathbf{p}_0^*$ does not belong to the column space of $D\tilde{T}_{c',g}(Q')$ for all $c' \in [0, 1]^m$.

We define

$$\begin{aligned} c - g &= (c_1 - g_1, \dots, c_m - g_m), \\ c' - g &= (c'_1 - g_1, \dots, c'_m - g_m). \end{aligned}$$

We claim that there exists a matrix D such that

$$D\tilde{T}_{c,g}(Q) = (0, T_{c-g}(Q))$$

and

$$D\tilde{T}_{c',g}(Q') = (0, T_{c'-g}(Q')),$$

where the choice of D does not depend on c or c' . In the rest of the proof, we construct such a D -matrix for $\tilde{T}_{c,g}(Q)$. The verification for $\tilde{T}_{c',g}(Q')$ is completely analogous. Note that each row in $D\tilde{T}_{c,g}(Q)$ is just a linear combination of rows of $\tilde{T}_{c,g}(Q)$. Therefore, it suffices to show that every row vector of the form

$$(0, B_{c-g,Q}(I_{i_1} \wedge \cdots \wedge I_{i_l}))$$

can be written as a linear combination of the row vectors of $\tilde{T}_{c,g}(Q)$. We prove this by induction. First note that for each $1 \leq i \leq m$,

$$(0, B_{c-g,Q}(I_i)) = (c_i - g_i)(0, B_Q(I_i)) = (g_i, B_{c,g,Q}(I_i)) - g_i \mathbb{E}. \tag{6.3}$$

Suppose that all rows of the form

$$(0, B_{c-g,Q}(I_{i_1} \wedge \cdots \wedge I_{i_l}))$$

for all $1 \leq l \leq j$ can be written as linear combinations of the row vectors of $\tilde{T}_{c,g}(Q)$ with coefficients only depending on g_1, \dots, g_m . Thanks to (6.3), the case of $j = 1$ holds. Suppose the statement holds for some general j . We consider the case of $j + 1$. By definition,

$$\begin{aligned} (g_{i_1} \cdots g_{i_{j+1}}, B_{c,g,Q}(I_{i_1} \wedge \cdots \wedge I_{i_{j+1}})) &= \Upsilon_{h=1}^{j+1}(g_{i_h}, B_{c,g,Q}(I_{i_h})) \\ &= \Upsilon_{h=1}^{j+1}(g_{i_h} \mathbb{E} + (0, B_{c-g,Q}(I_{i_h}))). \end{aligned} \tag{6.4}$$

Let “ $*$ ” denote element-by-element multiplication. For every generic vector V' of appropriate length,

$$\mathbb{E} * V' = V'.$$

We expand the right-hand side of (6.4). The last term would be

$$(0, B_{c-g,Q}(I_{i_1} \wedge \cdots \wedge I_{i_{j+1}})) = \Upsilon_{h=1}^{j+1}(0, B_{c-g,Q}(I_{i_h})).$$

From the induction assumption and definition (3.3), the other terms on both sides of (6.4) belong to the row space of $\tilde{T}_{c,g}(Q)$. Therefore, $(0, B_{c-g,Q}(I_{i_1} \wedge \cdots \wedge I_{i_{j+1}}))$ is also in the row space of $\tilde{T}_{c,g}(Q)$. In addition, all the corresponding coefficients only consist of g_i . Therefore, one can construct a $(2^m - 1) \times 2^m$ matrix D such that

$$D\tilde{T}_{c,g}(Q) = (0, T_{c-g}(Q)).$$

Because D is free of c and Q , we have

$$D\tilde{T}_{c',g}(Q') = (0, T_{c'-g}(Q')).$$

In addition, thanks to Propositions 6.3 and 6.4, $D\tilde{T}_{c,g}(Q)\mathbf{p}_0^* = T_{c-g}(Q)\mathbf{p}^*$ is not in the column space $C(T_{c'-g}(Q')) = C(D\tilde{T}_{c',g}(Q'))$ for all $c' \in [0, 1]^m$. Therefore, by Lemma 6.7, $\tilde{T}_{c,g}(Q)\mathbf{p}_0^*$ is not in the column space $C(\tilde{T}_{c',g}(Q'))$ for all $c' \in [0, 1]^m$.

In addition,

$$\begin{pmatrix} D \\ e_{2^m}^\top \end{pmatrix} \tilde{T}_{c,g}(Q)$$

is of full column rank, where $e_{2^m}^\top$ is a 2^m dimension row vector with last element being one and rest being zero. Therefore, $\tilde{T}_{c,g}(Q)$ is also of full column rank. \square

6.2. Proof of the theorems

Using the results of the previous propositions and lemmas, we now proceed to prove our theorems.

Proof of Theorem 2.4. Consider $Q' \approx Q$ and $T(\cdot)$ saturated. Recall that $\hat{\mathbf{p}}$ is the vector containing \hat{p}_A 's with $\mathbf{A} \neq \mathbf{0}$, where

$$\hat{p}_A = \frac{1}{N} \sum_{r=1}^N \mathbf{1}(\mathbf{A}_r = \mathbf{A}).$$

For any $\mathbf{p}^* > \mathbf{0}$, since $\hat{\mathbf{p}} \rightarrow \mathbf{p}^*$ almost surely, according to Corollary 6.5, $\alpha = T(Q)\hat{\mathbf{p}}$ by (2.5), and $T(Q)\mathbf{p}^* \notin C(T(Q'))$, there exists $\delta > 0$ such that,

$$\lim_{N \rightarrow \infty} P\left(\inf_{\mathbf{p} \in [0, 1]^{2^k-1}} |T(Q')\mathbf{p} - \alpha| > \delta\right) = 1$$

and

$$P\left(\inf_{\mathbf{p} \in [0, 1]^{2^k-1}} |T(Q)\mathbf{p} - \alpha| = 0\right) = 1.$$

Given that there are finitely many $m \times k$ binary matrices, $P(\hat{Q} \sim Q) \rightarrow 1$ as $N \rightarrow \infty$. In fact, we can arrange the columns of \hat{Q} such that $P(\hat{Q} = Q) \rightarrow 1$ as $N \rightarrow \infty$.

Note that $\hat{\mathbf{p}}$ satisfies the identity

$$T(Q)\hat{\mathbf{p}} = \alpha.$$

In addition, since $T(Q)$ is of full rank (Proposition 6.1), the solution to the above linear equation is unique. Therefore, the solution to the optimization problem $\inf_{\mathbf{p}} |T(Q)\mathbf{p} - \alpha|$ is unique and is $\hat{\mathbf{p}}$. Notice that when $\hat{Q} = Q$, $\tilde{\mathbf{p}} = \arg \inf_{\mathbf{p}} |T(\hat{Q})\mathbf{p} - \alpha| = \hat{\mathbf{p}}$. Therefore,

$$\lim_{N \rightarrow \infty} P(\tilde{\mathbf{p}} = \hat{\mathbf{p}}) = 1.$$

Together with the consistency of $\hat{\mathbf{p}}$, the conclusion of the theorem follows immediately. \square

Proof of Theorem 3.1. Note that for all Q'

$$T_{c,g}(Q')\mathbf{p} + p_0\mathbf{g} - \alpha = (\mathbf{g}, T_{c,g}(Q')) \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \alpha.$$

By the law of large numbers,

$$|T_{c,g}(Q)\mathbf{p}^* + p_0^*\mathbf{g} - \alpha| = \left| (\mathbf{g}, T_{c,g}(Q)) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} - \alpha \right| \rightarrow 0$$

almost surely as $N \rightarrow \infty$. Therefore,

$$S_{c,g}(Q) \rightarrow 0$$

almost surely as $N \rightarrow \infty$.

For any $Q' \approx Q$, note that

$$\begin{pmatrix} \alpha \\ 1 \end{pmatrix} \rightarrow \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix}.$$

According to Proposition 6.6 and the fact that $\mathbf{p}^* > \mathbf{0}$, there exists $\delta(c') > 0$ such that $\delta(c')$ is continuous in c' and

$$\inf_{\mathbf{p}, p_0} \left| \tilde{T}_{c',g}(Q') \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} \right| > \delta(c').$$

By elementary calculus,

$$\delta \triangleq \inf_{c' \in [0,1]^m} \delta(c') > 0$$

and

$$\inf_{c', \mathbf{p}, p_0} \left| \tilde{T}_{c',g}(Q') \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} \right| > \delta.$$

Therefore,

$$P \left(\inf_{c', \mathbf{p}, p_0} \left| \tilde{T}_{c',g}(Q') \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \begin{pmatrix} \alpha \\ 1 \end{pmatrix} \right| > \delta/2 \right) \rightarrow 1,$$

as $N \rightarrow \infty$. For the same δ , we have

$$P \left(\inf_{c', \mathbf{p}, p_0} \left| (\mathbf{g}, T_{c',g}(Q')) \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \alpha \right| > \delta/2 \right) = P \left(\inf_{c'} S_{c',g}(Q') > \delta/2 \right) \rightarrow 1.$$

The above minimization on the left of the equation is subject to the constraint that

$$\sum_{\mathbf{A} \in \{0,1\}^k} p_{\mathbf{A}} = 1.$$

Together with the fact that there are only finitely many $m \times k$ binary matrices, we have

$$P(\hat{Q}(c, g) \sim Q) = 1.$$

We arrange the columns of $\hat{Q}(c, g)$ so that $P(\hat{Q}(c, g) = Q) \rightarrow 1$ as $N \rightarrow \infty$.

Now we proceed to the proof of consistency for $\tilde{\mathbf{p}}(c, g)$. Note that

$$\begin{aligned} \left| \tilde{T}_{c,g}(\hat{Q}(c, g)) \begin{pmatrix} \tilde{p}_0(c, g) \\ \tilde{\mathbf{p}}(c, g) \end{pmatrix} - \begin{pmatrix} \alpha \\ 1 \end{pmatrix} \right| &\xrightarrow{P} 0, \\ \left| \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0^* \\ \mathbf{p}^* \end{pmatrix} - \begin{pmatrix} \alpha \\ 1 \end{pmatrix} \right| &\xrightarrow{P} 0. \end{aligned}$$

Since $\tilde{T}_{c,g}(Q)$ is a full column rank matrix and $P(\hat{Q}(c, g) = Q) \rightarrow 1$, $\tilde{\mathbf{p}}(c, g) \rightarrow \mathbf{p}^*$ in probability. □

Proof of Theorem 4.2. Assuming g is known, note that

$$\inf_{p_0, \mathbf{p}} \left| \tilde{T}_{c,g}(Q) \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \begin{pmatrix} \alpha \\ 1 \end{pmatrix} \right|$$

is a continuous function of c . According to the results of Proposition 4.1, the definition in (4.1), and the definition of \hat{c} in Section 4.1, we obtain that

$$\inf_{p_0, \mathbf{p}} \left| \tilde{T}_{\hat{c}(Q,g),g}(Q) \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} - \begin{pmatrix} \alpha \\ 1 \end{pmatrix} \right| \rightarrow 0,$$

in probability as $N \rightarrow \infty$. In addition, thanks to Proposition 6.6 and with a similar argument as in the proof of Theorem 3.1, $\hat{Q}_{\hat{c}}(g)$ is a consistent estimator.

Furthermore, if $\tilde{c}(Q, g)$ is a consistent estimator, then $\hat{c}(Q, g)$ is also consistent. Then, the consistency of $\tilde{\mathbf{p}}_{\hat{c}}(g)$ follows from the facts that $\hat{Q}_{\hat{c}}(g)$ is consistent and $\tilde{T}_{\hat{c},g}(Q)$ is of full column rank. □

Appendix: Technical proofs

Proof of Proposition 6.3. Note that $Q_{1:k} = Q'_{1:k} = \mathcal{I}_k$. Let $T(\cdot)$ be arranged as in (6.1). Then, $T(Q)_{1:(2^k-1)} = T(Q')_{1:(2^k-1)}$. Given that $Q \neq Q'$, we have $T(Q) \neq T(Q')$. We assume that $T(Q)_{li} \neq T(Q')_{li}$, where $T(Q)_{li}$ is the entry in the l th row and i th column. Since $T(Q)_{1:(2^k-1)} = T(Q')_{1:(2^k-1)}$, it is necessary that $l \geq 2^k$.

Suppose that the l th row of the $T(Q')$ corresponds to an item that requires attributes $i_1, \dots, i_{l'}$. Then, we consider $1 \leq h \leq 2^k - 1$, such that the h th row of $T(Q')$ is $B_{Q'}(I_{i_1} \wedge \dots \wedge I_{i_{l'}})$. Then, the h th row vector and the l th row vector of $T(Q')$ are identical.

Since $T(Q)_{1:(2^k-1)} = T(Q')_{1:(2^k-1)}$, we have $T(Q)_{hj} = T(Q')_{hj} = T(Q')_{lj}$ for $j = 1, \dots, 2^k - 1$. If $T(Q)_{li} = 0$ and $T(Q')_{li} = 1$, the matrices $T(Q)$ and $T(Q')$ look like

$$T(Q') = \begin{matrix} & \text{column } i \\ & \downarrow \\ \text{row } h \rightarrow & \begin{pmatrix} \mathcal{I} & * & \dots & * & \dots \\ \vdots & \vdots & & \dots & \dots \\ \vdots & \vdots & \mathcal{I} & \dots & \dots \\ \vdots & \vdots & \vdots & & \\ \text{row } l \rightarrow & * & 1 & * \\ * & * & * \end{pmatrix} \end{matrix}$$

and

$$T(Q) = \begin{matrix} & \text{column } i \\ & \downarrow \\ \text{row } h \rightarrow & \begin{pmatrix} \mathcal{I} & * & \dots & * & \dots \\ \vdots & \vdots & & \dots & \dots \\ \vdots & \vdots & \mathcal{I} & \dots & \dots \\ \vdots & \vdots & \vdots & & \\ \text{row } l \rightarrow & * & 0 & * \\ * & * & * \end{pmatrix}. \end{matrix}$$

Case 1. Either the h th or l th row vector of $T_{c'}(Q')$ is a zero vector. The conclusion is immediate because all the entries of $T_c(Q)\mathbf{p}^*$ are nonzero.

Case 2. The h th and l th row vectors of $T_{c'}(Q')$ are nonzero vectors. Suppose that the l th row corresponds to an item. There are three different situations: according to the true Q -matrix (a) the item in row l requires strictly more attributes than row h , (b) the item in row l requires strictly fewer attributes than row h , (c) otherwise. We consider these three situations, respectively.

- (a) Under the true Q -matrix, there are two types of sub-populations in consideration: people who are able to answer item(s) in row h (p_1) only and people who are able to answer items in both row h and row l (p_2). Then, the sub-matrix of $T_c(Q)$ and $T_{c'}(Q)$ are like

	$T_c(Q)$			$T_{c'}(Q')$	
	p_1	p_2		p_1	p_2
row h	c_h	c_h	row h	c'_h	c'_h
row l	0	c_l	row l	c'_l	c'_l

We now claim that c_l and c'_l must be equal (otherwise the conclusion hold) for the following reason.

Consider the following two rows of $T(Q)$: row A corresponding to the combination that contains all the items; row B corresponding to the row that contains all the items except for the one in row l .

Rows A and B are in fact identical in $T(Q)$. This is because all the attributes are used at least twice (condition C5). Then, the attributes in row l are also required by some other item(s) and rows A and B require the same combination of items. Thus, the corresponding entries of all the column vectors of $T_c(Q)$ are different by a factor of c_l .

For $T(Q')$, rows A and B are also identical. This is because row h and row l have identical attribute requirements. Then, thus, the corresponding entries of all the column vectors of $T_{c'}(Q')$ are different by a factor of c'_l . Thus, c'_l and c_l must be identical otherwise $T_c(Q)\mathbf{p}^*$ is not in the column space of $T_{c'}(Q)$.

Similarly, we obtain that $c_h = c'_h$. Given that $c_h = c'_h$ and $c_l = c'_l$, we now consider row h and row l . Notice that all the column vectors in $T_{c'}(Q')$ have their entries in row h and row l different by a factor of c_h/c_l . On the other hand, the h and l th entries of $T_c(Q)\mathbf{p}^*$ are NOT different by a factor of c_h/c_l as long as the proportion of p_1 is positive. Thereby, we conclude this case.

- (b) Consider the following two types of sub-populations: people who are able to answer item(s) in row l (p_1) only and people who are able to answer items in both row h and row l (p_2). Similar to the analysis of (a), the sub-matrices look like:

$$\begin{array}{cc|cc} & & T_c(Q) & T_{c'}(Q') \\ & & p_1 & p_2 & p_1 & p_2 \\ \text{row } h & & 0 & c_h & \text{row } h & 0 & c'_h \\ \text{row } l & & c_l & c_l & \text{row } l & 0 & c'_l \end{array}$$

With exactly the same argument as in (a), we conclude that $c_j = c'_j$, $c_h = c'_h$, and further $T_c(Q)\mathbf{p}^*$ is not in the column space of $T_{c'}(Q')$.

- (c) Consider the following three types of sub-populations: people who are able to answer item(s) in row l only (p_1), people who are able to answer item(s) in row h only (p_2), and people who are able to answer items in both row h and row l (p_3). The sub-matrices look like:

$$\begin{array}{cc|ccc|ccc} & & T_c(Q) & & T_{c'}(Q') & & & \\ & & p_1 & p_2 & p_3 & p_1 & p_2 & p_3 \\ \text{row } h & & 0 & c_h & c_h & \text{row } h & 0 & c'_h & c'_h \\ \text{row } l & & c_l & 0 & c_l & \text{row } l & 0 & c'_l & c'_l \\ \text{row } l \wedge h & & 0 & 0 & c_h c_l & \text{row } l \wedge h & 0 & c'_h c'_l & c'_h c'_l \end{array}$$

With the same argument, we have that $c_l = c'_l$ and $c_h = c'_h$. On considering row h and row $l \wedge h$, we conclude that $T_c(Q)\mathbf{p}^*$ is not in the column space of $T_{c'}(Q')$. \square

Proof of Proposition 6.4. $T(\cdot)$ is arranged as in (6.1). Consider Q' such that $Q'_{1:k}$ is incomplete. We discuss the following situations.

1. There are two row vectors, say the h th and l th row vectors ($1 \leq i, j \leq k$), in $Q'_{1:k}$ that are identical. Equivalently, two items require exactly the same attributes according to Q' .

With exactly the same argument as in the previous proof, under condition C5, we have that $c_h = c'_h$ and $c_l = c'_l$. We now consider the rows corresponding to l and $l \wedge h$. Note that the elements corresponding to row l and row $l \wedge h$ for all the vectors in the column space of $T_{c'}(Q')$ are different by a factor of c_h . However, the corresponding elements in the vector $T_c(Q)\mathbf{p}^*$ are NOT different by a factor of c_h as long as the population is fully diversified.

2. No two row vectors in $Q'_{1:k}$ are identical. Then, among the first k rows of Q' there is at least one row vector containing two or more nonzero entries. That is, there exists $1 \leq i \leq k$ such that

$$\sum_{j=1}^k Q'_{ij} > 1.$$

This is because if each of the first k items requires only one attribute and $Q'_{1:k}$ is not complete, there are at least two items that require the same attribute. Then, there are two identical row vectors in $Q'_{1:k}$ and it belongs to the first situation. We define

$$a_i = \sum_{j=1}^k Q'_{ij},$$

the number of attributes required by item i according to Q' .

Without loss of generality, assume $a_i > 1$ for $i = 1, \dots, n$ and $a_i = 1$ for $i = n + 1, \dots, k$. Equivalently, among the first k items, only the first n items require more than one attribute while the $(n + 1)$ th through the k th items require only one attribute each, all of which are distinct. Without loss of generality, we assume $Q'_{ii} = 1$ for $i = n + 1, \dots, k$ and $Q_{ij} = 0$ for $i = n + 1, \dots, k$ and $i \neq j$.

- (a) $n = 1$. Since $a_1 > 1$, there exists an $l > 1$ such that $Q'_{1l} = 1$. We now consider rows 1 and l . With the same argument as before (i.e., the attribute required by row l is also required by item 1 in Q'), we have that $c_l = c'_l$ (be careful that we cannot claim that $c_1 = c'_1$). We now consider the rows 1 and $1 \wedge l$. Note that in $T_{c'}(Q')$ these two rows are different by a factor of c_l ; while the corresponding entries in $T_c(Q)\mathbf{p}^*$ are NOT different by a factor of c_l . Thereby, we conclude the result in this situation.
- (b) $n > 1$ and there exists $j > n$ and $i \leq n$ such that $Q'_{ij} = 1$. The argument is identical to that in (2a).
- (c) $n > 1$ and for each $j > n$ and $i \leq n$, $Q'_{ij} = 0$. Let the i^* th row in $T(Q')$ correspond to $I_1 \wedge \dots \wedge I_n$. Let the i_h^* th row in $T(Q')$ correspond to $I_1 \wedge \dots \wedge I_{h-1} \wedge I_{h+1} \wedge \dots \wedge I_n$ for $h = 1, \dots, n$.

We claim that there exists an h such that the i^* th row and the i_h^* th row are identical in $T(Q')$, that is

$$B_{Q'}(I_1 \wedge \dots \wedge I_{h-1} \wedge I_{h+1} \wedge \dots \wedge I_n) = B_{Q'}(I_1 \wedge \dots \wedge I_n). \tag{A.1}$$

If the above claim is true, then the attributes required by item h have been required by some other items. Then, we conclude that c_h and c'_h must be identical. In addition, rows in $T_{c'}(Q')$ corresponding to $I_1 \wedge \dots \wedge I_{h-1} \wedge I_{h+1} \wedge \dots \wedge I_n$ and $I_1 \wedge \dots \wedge I_n$ are

different by a factor of c_h . On the other hand, the corresponding entries in $T_c(Q)\mathbf{p}^*$ are NOT different by a factor of c_h . Then, we are able to conclude the results for all the cases.

In what follows, we prove the claim in (A.1) by contradiction. Suppose that there does not exist such an h . This is equivalent to saying that for each $j \leq n$ there exists an α_j such that $Q'_{j\alpha_j} = 1$ and $Q'_{i\alpha_j} = 0$ for all $1 \leq i \leq n$ and $i \neq j$. Equivalently, for each $j \leq n$, item j requires at least one attribute that is not required by other first n items. Consider

$$C_i = \{j : \text{there exists } i \leq i' \leq n \text{ such that } Q'_{i',j} = 1\}.$$

Let $\#(\cdot)$ denote the cardinality of a set. Since for each $i \leq n$ and $j > n$, $Q'_{ij} = 0$, we have that $\#(C_1) \leq n$. Note that $Q'_{1\alpha_1} = 1$ and $Q'_{i\alpha_1} = 0$ for all $2 \leq i \leq n$. Therefore, $\alpha_1 \in C_1$ and $\alpha_1 \notin C_2$. Therefore, $\#(C_2) \leq n - 1$. By a similar argument and induction, we have that $a_n = \#(C_n) \leq 1$. This contradicts the fact that $a_n > 1$. Therefore, there exists an h such that (A.1) is true. As for $T(Q)$, we have that

$$B_Q(I_1 \wedge \cdots \wedge I_{h-1} \wedge I_{h+1} \wedge \cdots \wedge I_n) \neq B_Q(I_1 \wedge \cdots \wedge I_n).$$

Summarizing the cases in 1, (2a), (2b) and (2c), we conclude the proof. \square

Acknowledgements

This research was supported in part by Grants NSF CMMI-1069064, SES-1123698, Institute of Education Sciences R305D100017 and NIH R37 GM047845.

References

- [1] Chiu, C.Y., Douglas, J.A. and Li, X. (2009). Cluster analysis for cognitive diagnosis: Theory and applications. *Psychometrika* **74** 633–665. [MR2565331](#)
- [2] de la Torre, J. (2008). An empirically-based method of Q -matrix validation for the DINA model: Development and applications. *Journal of Educational Measurement* **45** 343–362.
- [3] de la Torre, J. (2008). The generalized DINA model. In *International Meeting of the Psychometric Society*, Durham, NH.
- [4] de la Torre, J. and Douglas, J.A. (2004). Higher-order latent trait models for cognitive diagnosis. *Psychometrika* **69** 333–353. [MR2272454](#)
- [5] DiBello, L.V., Stout, W.F. and Roussos, L. (1995). Unified cognitive psychometric assessment likelihood-based classification techniques. In *Cognitively Diagnostic Assessment* 361–390. Hillsdale, NJ: Erlbaum.
- [6] Gelman, A., Jakulin, A., Pittau, M.G. and Su, Y.S. (2008). A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat.* **2** 1360–1383. [MR2655663](#)
- [7] Hartz, S.M. (2002). A Bayesian framework for the unified model for assessing cognitive abilities: Blending theory with practicality. Ph.D. dissertation, Univ. Illinois, Urbana–Champaign.
- [8] Henson, R. and Douglas, J. (2005). Test construction for cognitive diagnosis. *Appl. Psychol. Meas.* **29** 262–277. [MR2142701](#)

- [9] Henson, R., Roussos, L., Douglas, J. and He, X. (2008). Cognitive diagnostic attribute-level discrimination indices. *Appl. Psychol. Meas.* **32** 275–288. [MR2432210](#)
- [10] Henson, R.A. and Templin, J.L. (2005). Hierarchical log-linear modeling of the skill joint distribution. Technical report, External Diagnostic Research Group.
- [11] Junker, B.W. (1999). Some statistical models and computational methods that may be useful for cognitively-relevant assessment. Technical report. Available at <http://www.stat.cmu.edu/~brian/nrc/cfa/documents/final.pdf>.
- [12] Junker, B.W. and Sijtsma, K. (2001). Cognitive assessment models with few assumptions, and connections with nonparametric item response theory. *Appl. Psychol. Meas.* **25** 258–272. [MR1842982](#)
- [13] Leighton, J.P., Gierl, M.J. and Hunka, S.M. (2004). The attribute hierarchy model for cognitive assessment: A variation on Tatsuoka's rule-space approach. *Journal of Educational Measurement* **41** 205–237.
- [14] Liu, Y., Douglas, J.A. and Henson, R. (2007). Testing person fit in cognitive diagnosis. In *The Annual Meeting of the National Council on Measurement in Education (NCME)*, Chicago, IL.
- [15] Maris, E. (1999). Estimating multiple classification latent class models. *Psychometrika* **64** 187–212. [MR1700708](#)
- [16] Reckase, M.D. (1990). Unidimensional data from multidimensional tests and multidimensional data from unidimensional tests. In *The Annual Meeting of the American Educational Research Association*, Boston, MA, April 16–20.
- [17] Reckase, M.D. (2009). *Multidimensional Item Response Theory*. New York: Springer.
- [18] Roussos, L., DiBello, L.V., Stout, W., Hartz, S., Henson, R. and Templin, J. (2007). The fusion model skills diagnosis system. In *Cognitively Diagnostic Assessment for Education: Theory and Practice* (J. P. Leighton and M. J. Gierl, eds.) 275–318. New York: Cambridge Univ. Press.
- [19] Roussos, L.A., Templin, J.L. and Henson, R.A. (2007). Skills diagnosis using IRT-based latent class models. *Journal of Educational Measurement* **44** 293–311.
- [20] Rupp, A.A. (2002). Feature selection for choosing and assembling measurement models: A building-block-based organization. *Psychometrika* **2** 311–360.
- [21] Rupp, A.A. and Templin, J.L. (2008). Effects of Q -matrix misspecification on parameter estimates and misclassification rates in the DINA model. *Educational and Psychological Measurement* **68** 78–98.
- [22] Rupp, A.A. and Templin, J.L. (2008). Unique characteristics of diagnostic classification models: A comprehensive review of the current state-of-the-art. *Measurement: Interdisciplinary Research and Perspective* **6** 219–262.
- [23] Rupp, A.A., Templin, J.L. and Henson, R.A. (2010). *Diagnostic Measurement: Theory, Methods, and Applications*. New York: Guilford Press.
- [24] Schwarz, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6** 461–464. [MR0468014](#)
- [25] Stout, W. (2007). Skills diagnosis using IRT-based continuous latent trait models. *Journal of Educational Measurement* **44** 313–324.
- [26] Tatsuoka, C. (2002). Data analytic methods for latent partially ordered classification models. *J. Roy. Statist. Soc. Ser. C* **51** 337–350. [MR1920801](#)
- [27] Tatsuoka, K.K. (1983). Rule space: An approach for dealing with misconceptions based on item response theory. *Journal of Educational Measurement* **20** 345–354.
- [28] Tatsuoka, K.K. (1985). A probabilistic model for diagnosing misconceptions in the pattern classification approach. *Journal of Educational Statistics* **12** 55–73.
- [29] Tatsuoka, K.K. (2009). *Cognitive Assessment: An Introduction to the Rule Space Method*. Florence, KY: Routledge.
- [30] Templin, J., He, X., Roussos, L. and Stout, W. (2003). The pseudo-item method: A simple technique for analysis of polytomous data with the fusion model. Technical report, External Diagnostic Research Group.

- [31] Templin, J.L. (2006). CDM: Cognitive diagnosis modeling with Mplus. Computer software. Available at <http://jtemplin.coe.uga.edu/research/>.
- [32] Templin, J.L. and Henson, R.A. (2006). Measurement of psychological disorders using cognitive diagnosis models. *Psychol. Methods* **11** 287–305.
- [33] von Davier, M. (2005). A general diagnosis model applied to language testing data. Research report, Educational Testing Service.

Received March 2011 and revised November 2011