Semi-Automated Rasch Analysis with Differential Item Functioning

Feri Wijayanto^{a,b}, Ioan Gabriel Bucur^a, Karlien Mul^c, Perry Groot^a, Baziel G.M. van Engelen^c, Tom Heskes^a

^a*Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands*

^b*Department of Informatics, Universitas Islam Indonesia, Sleman, Indonesia*

^c*Department of Neurology, Donders Institute for Brain, Cognition, and Behaviour, Radboud University Medical Center, Nijmegen, The Netherlands*

Abstract

Rasch analysis is a procedure to develop and validate instruments that aim to measure persons' traits. However, manual Rasch analysis is a complex and time-consuming task, even more so when the possibility of differential item functioning (DIF) is taken into consideration. Furthermore, manual Rasch analysis by construction relies on a modeller's subjective choices. As an alternative approach, we introduce a semi-automated procedure that is based on the optimization of a new criterion, called in-plus-out-of-questionnaire log likelihood with differential item functioning (IPOQ-LL-DIF) which extends our previous criterion. We illustrate our procedure on artificially generated data as well as on several real-world datasets containing potential DIF items. On these real-world datasets, our procedure found instruments with similar clinimetric properties as those suggested by experts through manual analyses.

Keywords: Semi-automated Rasch analysis, Rasch model, Generalized partial credit model, Penalized JMLE, GPCMLasso, GPCM-DIF, Differential item functioning, DIF detection

1. Introduction

In measurement theory, personal aspects may contain latent constructs or traits which can not be approached directly, such as “intelligence” and “quality of life”. In an effort to measure these latent constructs, many scales have been developed from uniquely designed questionnaires. Rasch analysis is one of the scientific methods to transform the original survey into a linear-weighted, clinimetrically sound scale. Using inherent criteria, e.g., goodness-of-fit, unidimensionality, and local dependency [10, 12, 36], manual Rasch analysis follows a step-by-step procedure, repeatedly fitting the observed responses to the Rasch model. The worst item(s) are generally removed, after which the remaining items are reevaluated, until a clinimetrically optimal itemset has been obtained.

Rasch analysis becomes even more complex when the original survey contains items that can function differently due to the respondents' backgrounds (e.g., age, gender, and nationality). This phenomenon is known as differential item functioning

(DIF). DIF occurs if respondents from a particular group tend to score higher or lower on a particular item compared to other group(s), despite having otherwise similar characteristics. This type of items are often found in clinical observations, for example in running, old people tend to have more trouble compared to young people. Erroneous ignorance of such biases leads to a biased instrument [5, 16]. DIF assessment has become one of the standard ingredients of Rasch analysis and has been implemented in various ways, e.g., [13, 15, 20, 26, 29, 32].

In current practice, step-by-step procedures are carried out manually by the experts, which can be relatively time-consuming even with the support from the available software packages, such as [3, 6, 21, 37]. Decisions on how to prioritize the various evaluation criteria and which items to include partly rely on human judgments blended with clinical expertise, and different experts may obtain different but equally suitable instruments. These procedures will become even more complex when the DIF items have to be resolved iteratively [2, 11].

The objective of this research is to incorporate the DIF assessment procedure while automating the Rasch analysis. Doing so, we extend our previous method, which automates the Rasch analysis using the in-plus-out-of-questionnaire log likelihood (IPOQ-LL) criterion [35]. Our previous method naturally incorporates standard Rasch criteria, e.g., item goodness-of-fit, unidimensionality, and local dependency [35]. Accordingly, we will show that our new procedure in addition naturally incorporates the standard DIF assessment in Rasch analysis. Our novel procedure makes use of a generalization of the IPOQ-LL criterion, which we will refer to as the in-plus-out-of-questionnaire log likelihood with DIF (IPOQ-LL-DIF).

The rest of this article is structured as follows. Section 2 describes the central model in our implementation, the GPCMLasso model [26], its transformation to other models, and the idea to solve its estimation problem using the $L1$ (lasso) penalty together with the coordinate descent. Section 3 discusses the main part of our proposed method, the in-plus-out-of questionnaire log likelihood with DIF (IPOQ-LL-DIF), which extends the previous method and argues for the method in comparison with the typical assessment of DIF items in standard Rasch analysis. Section 4 reports our experimental results on an artificial and three real-world datasets. Section 5 discusses general aspects of our procedure and the results it obtained, and concludes our research.

2. Preliminary

2.1. Generalized Partial Credit Model with DIF

Differential item functioning (DIF) occurs as a result of an inconsistency between estimated abilities and true abilities for given groups. If the inconsistency uniformly affects all subjects in the group, then it is known as a uniform DIF, otherwise, it is a non-uniform DIF [11, 18, 33].

In this work, we focus on uniform DIF and adopt the GPCMLasso model, introduced in [26], which extends the generalized partial credit model (GPCM) [23] after parameterizing the DIF effects. Rooted to the GPCM, the GPCMLasso has the ability to model responses that are coded into two or more ordered categories. We write $x_{ni} \in \{0, 1, \dots, m_i\}$ for the observed response of subject n on item i , where

item i consists of $m_i + 1$ ordered categories. We have $m_i = 1$ for dichotomous test items and $m_i > 1$ for polytomous items.

The GPCMlasso model contains the same type of parameters as the GPCM: θ_n for the *ability* of subject n , β_{ij} , with $j = 1, \dots, m_i$ for the *difficulties* or *thresholds* of item i , and α_i for the *discrimination parameter* of item i . Additionally, to model the difference in difficulty on item i between the members and non-members of focal group f , the *DIF parameters* δ_{if} are introduced. Furthermore, κ_{nf} , with $f = 1, \dots, m_f$ and where m_f represents the number of potential DIF-inducing covariates, is a binary matrix that maps subject n into group f with $\kappa_{nf} = 1$ if respondent n is a member of group f and $\kappa_{nf} = 0$ otherwise.

Given these definitions, the probability of subject n gives response x on item i reads

$$P(X_{ni} = x | \theta, \beta, \alpha, \delta) = \frac{\exp \left[\alpha_i \sum_{j=1}^x (\theta_n - (\beta_{ij} - \sum_{f=1}^{m_f} \delta_{if} \kappa_{nf})) \right]}{1 + \sum_{k=1}^{m_i} \exp \left[\alpha_i \sum_{j=1}^k (\theta_n - (\beta_{ij} - \sum_{f=1}^{m_f} \delta_{if} \kappa_{nf})) \right]}, \quad (1)$$

for $x > 0$, and

$$P(X_{ni} = 0 | \theta, \beta, \alpha, \delta) = \frac{1}{1 + \sum_{k=1}^{m_i} \exp \left[\alpha_i \sum_{j=1}^k (\theta_n - (\beta_{ij} - \sum_{f=1}^{m_f} \delta_{if} \kappa_{nf})) \right]}. \quad (2)$$

From now on, we will refer to this as the generalized partial credit model with differential item functioning, GPCM-DIF. Setting $\alpha_i = 1$ for $i = 1, \dots, \mathbb{P}$ in the GPCM-DIF model gives what we will refer to as the partial credit model with DIF (PCM-DIF). Using the PCM-DIF to estimate the respondents' traits is comparable to the use of the partial credit model (PCM) on items after the DIF has been resolved. With $\delta_{if} = 0$ for $i = 1, \dots, \mathbb{P}$ and $f = 1, \dots, m_f$ we obtain the GPCM. By then also fixing $\alpha_i = 1$ for $i = 1, \dots, \mathbb{P}$, we get the PCM. In the case of binary responses, with $m_i = 1$ for $i = 1, \dots, \mathbb{P}$, the GPCM transforms to the 2-parameter logistic (2PL) model and the PCM transforms to the original Rasch model [19, 22, 24].

2.2. Coordinate Descent

Given observed responses x_{ni} , the log likelihood of all model parameters for a given set of items $\mathcal{S} \subset \{1, \dots, \mathbb{P}\}$ reads

$$L_{\mathcal{S}}(\theta, \beta, \alpha, \delta) = \sum_{i \in \mathcal{S}} \sum_{n=1}^N \log P(X = x_{ni} | \theta, \beta, \alpha, \delta), \quad (3)$$

with $P(X = x_{ni} | \theta, \beta, \alpha, \delta)$ from (1) and (2). This log likelihood measures how well the parameters predict the subjects' observed responses on the items from set $\mathcal{S}$.

We turn the log likelihood into a penalized log likelihood by adding penalty terms. As in [35], we add Tikhonov regularization for the abilities θ , to regularize

these towards zero, as well as for $\ln \alpha$, to drive the discrimination parameters towards one. Inspired by [26], we further add a Lasso ($L1$) penalty for the DIF parameters δ , so that irrelevant DIF parameters are optimized to zero:

$$F_{\mathcal{S}}(\theta, \beta, \alpha, \delta) = L_{\mathcal{S}}(\theta, \beta, \alpha, \delta) - \lambda_{\theta} \sum_{n=1}^N \theta_n^2 - \lambda_{\alpha} \sum_{i \in \mathcal{S}} (\ln \alpha_i)^2 - \lambda_{\delta} \sum_{i \in \mathcal{S}} \sum_{f=1}^{m_f} |\delta_{if}|. \quad (4)$$

To optimize (4), we propose to apply two-level coordinate descent [9]. At the top level, we treat the GPCM parameters θ , α , and β as one coordinate, and the DIF parameters δ as another. Given fixed DIF parameters, the GPCM parameters are optimized using penalized joint maximum likelihood estimation (PJMLE) (see [35] for details). Given fixed GPCM parameters, we optimize the DIF parameters through coordinate descent at the second level, treating each δ_{if} for $i = 1, \dots, \mathbb{P}$ and $f = 1, \dots, m_f$ as a unique coordinate. Details of the coordinate descent algorithm applied to (4) are provided in Appendix B.1.

3. The Proposed Method

3.1. In-plus-out-of-questionnaire Log Likelihood with DIF

In instrument design, we are given an initial set of $\mathbb{P}$ items that, based on responses on a survey including all these items, we would like to reduce to a smaller set of items that make up the final questionnaire. We will refer to the set of included items as the included itemset, denoted $\mathcal{S}_{\text{in}}$, and to its complement as the excluded itemset, denoted $\mathcal{S}_{\text{out}} = \{1, \dots, \mathbb{P}\} \setminus \mathcal{S}_{\text{in}}$. In [35], we introduced a novel criterion called in-plus-out-of-questionnaire log likelihood (IPOQ-LL) for evaluating the quality of any split into $\mathcal{S}_{\text{in}}$ and $\mathcal{S}_{\text{out}}$ given the observed responses on the original survey. Following the same rationale, we here extend this criterion to also incorporate the possibility of item(s) with differential functioning.

For a given final questionnaire, only the items in the included itemset $\mathcal{S}_{\text{in}}$ can be used to estimate the subjects' abilities θ . We propose to obtain these abilities, and at the same time the discrimination parameters, thresholds, and DIF parameters corresponding to the included items, by maximizing the penalized log likelihood in (4):

$$\{\hat{\theta}_{\mathcal{S}_{\text{in}}}, \hat{\beta}_{\mathcal{S}_{\text{in}}}, \hat{\alpha}_{\mathcal{S}_{\text{in}}}, \hat{\delta}_{\mathcal{S}_{\text{in}}}\} = \underset{\{\theta, \beta, \alpha, \delta\}}{\operatorname{argmax}} L_{\mathcal{S}_{\text{in}}}(\theta, \beta, \alpha, \delta) - \lambda_{\theta} \sum_{n=1}^N \theta_n^2 - \lambda_{\text{in}} \sum_{i \in \mathcal{S}_{\text{in}}} (\ln \alpha_i)^2 - \lambda_{\delta} \sum_{i \in \mathcal{S}_{\text{in}}} \sum_{f=1}^{m_f} |\delta_{if}|. \quad (5)$$

We refer to the log likelihood of these fitted parameters on the included itemset as the in-questionnaire log likelihood with DIF:

$$\text{IQ-LL-DIF}(\mathcal{S}_{\text{in}}) = L_{\mathcal{S}_{\text{in}}}(\hat{\theta}_{\mathcal{S}_{\text{in}}}, \hat{\beta}_{\mathcal{S}_{\text{in}}}, \hat{\alpha}_{\mathcal{S}_{\text{in}}}, \hat{\delta}_{\mathcal{S}_{\text{in}}}). \quad (6)$$

This IQ-LL-DIF resembles standard test statistics in Rasch analysis (e.g., outfit and infit to the resolved DIF items) [30].

Next, although we may not need the excluded items to arrive at a clinimetrically relevant scale, we would like the abilities estimated on $\mathcal{S}_{\text{in}}$ to properly represent the

observed responses on $\mathcal{S}_{\text{out}}$ as well, if only because the original survey was designed to also include these items. We therefore fix the abilities $\hat{\theta}_{\mathcal{S}_{\text{in}}}$ and optimize the penalized log likelihood given the responses on the excluded items w.r.t. the thresholds, the discrimination parameters, and the DIF parameters:

$$\{\hat{\beta}_{\mathcal{S}_{\text{out}}}, \hat{\alpha}_{\mathcal{S}_{\text{out}}}, \hat{\delta}_{\mathcal{S}_{\text{out}}}\} = \underset{\{\beta, \alpha\}}{\operatorname{argmax}} L_{\mathcal{S}_{\text{out}}}(\hat{\theta}_{\mathcal{S}_{\text{in}}}, \beta, \alpha) - \lambda_{\text{out}} \sum_{i \in \mathcal{S}_{\text{out}}} (\ln \alpha_i)^2 - \lambda_{\delta} \sum_{i \in \mathcal{S}_{\text{out}}} \sum_{f=1}^{m_f} |\delta_{if}|. \quad (7)$$

We refer to

$$\text{OQ-LL-DIF}(\mathcal{S}_{\text{out}}) = L_{\mathcal{S}_{\text{out}}}(\hat{\theta}_{\mathcal{S}_{\text{in}}}, \hat{\beta}_{\mathcal{S}_{\text{out}}}, \hat{\alpha}_{\mathcal{S}_{\text{out}}}, \hat{\delta}_{\mathcal{S}_{\text{out}}}). \quad (8)$$

as the out-of-questionnaire log likelihood with DIF. Our new criterion, the in-plus-out-of-questionnaire log likelihood with DIF (IPOQ-LL-DIF), is the total of both log likelihoods:

$$\text{IPOQ-LL-DIF}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{out}}) = \text{IQ-LL-DIF}(\mathcal{S}_{\text{in}}) + \text{OQ-LL-DIF}(\mathcal{S}_{\text{out}}).$$

Algorithm 1 outlines the procedure for computing the in-plus-out-of-questionnaire log likelihood with DIF given a subdivision of all items into the included itemset $\mathcal{S}_{\text{in}}$ and excluded itemset $\mathcal{S}_{\text{out}}$.

Algorithm 1 Pseudocode for computing the in-plus-out-of-questionnaire log likelihood with DIF for a particular included itemset $\mathcal{S}_{\text{in}}$ and excluded itemset $\mathcal{S}_{\text{out}}$.

- 1: Fit θ , β , α , and δ parameters through coordinate descent on the included itemset $\mathcal{S}_{\text{in}}$ as in Eq. (5).
 - 2: Compute the in-questionnaire log likelihood with DIF of the parameters obtained in step 1 on the included itemset $\mathcal{S}_{\text{in}}$ as in Eq. (6).
 - 3: Keep the θ parameters fixed to those obtained in step 1 and fit β , α , and δ parameters through coordinate descent on the excluded items $\mathcal{S}_{\text{out}}$ as in Eq. (7).
 - 4: Compute the out-of-questionnaire log likelihood with DIF of the β , α , and δ parameters obtained in step 3 and the θ parameters obtained in step 1 on the excluded itemset $\mathcal{S}_{\text{out}}$ as in Eq. (8).
 - 5: Add the out-of-questionnaire log likelihood with DIF from step 4 to the in-questionnaire log likelihood with DIF from step 2.
-

In our earlier work [35], we noticed that the outcome of our fitting procedure without the additional DIF parameters is relatively insensitive to the setting of the regularization parameters, as long as the regularization parameter λ_{in} of the β and α parameters for the included itemset is an order of magnitude larger than the regularization parameter λ_{out} for the excluded itemset. In this paper, we therefore stick to the same settings: $\lambda_{\theta} = 0.05$, $\lambda_{\text{in}} = 50$, and $\lambda_{\text{out}} = 1$.

Whether or not non-zero DIF parameters are obtained, does depend on the precise setting of the regularization parameter λ_{δ} : the larger λ_{δ} , the fewer non-zero DIF parameters will remain. Unless specified otherwise, in this paper we set $\lambda_{\delta} = 10$. With this setting, our procedure yields more or less the same DIF items in

the three real-world datasets compared to those obtained with a manual analysis. An arguably more principled, but computationally much more intensive approach would be a cross-validation procedure for finding the optimal value of λ_δ , as described in Appendix B.2.

3.2. Comparison with Other Approaches for DIF Assessment

There are two main approaches for handling DIF items.

Blending in with the Rasch analysis. In many practices of Rasch analysis, DIF detection is infused as an additional step in the estimation procedure [25, 33, 34]. Resolved DIF items are treated as any other items: if they fit the Rasch model well they are kept, otherwise they are removed. In accordance with our previous method [35], our new method has a tendency to put predictive split-items in the included itemset: these items help to obtain a better estimate of the subjects’ ability not only in the included itemset, but also in the excluded itemset.

Treating the DIF items separately. Andrich and Hagquist [2] distinguish between ‘real’ and ‘artificial’ DIF items. A real DIF item is stable, independent of the inclusion or exclusion of other potential DIF items. An artificial DIF item, on the other hand, only becomes a DIF item by virtue of the presence of other (real) DIF items. Therefore, Andrich and Hagquist [2] suggest to resolve the DIF items iteratively, starting with the largest effect, in an attempt to neutralize the effect of artificial DIF items. Our procedure also applies a thorough strategy to identify and resolve all potential DIF items. However, instead of doing this sequentially, we simultaneously estimate the DIF effects for all items that are still included. The Lasso ($L1$) penalty helps to distinguish between DIF and non-DIF items by nullifying the insignificant DIF effects.

3.3. Itemset Selection

In this paper, we introduce a single criterion, IPOQ-LL-DIF, to measure the quality of a final instrument by considering the differential functioning over items. With this criterion, we can in principle apply any optimization procedure to determine which items to keep in the included itemset $\mathcal{S}_{\text{in}}$ and which items to put in the excluded itemset $\mathcal{S}_{\text{out}}$. In our experiments, we consider the same optimization procedure as in our previous work (see [35] for details), i.e., stepwise selection. Stepwise selection alternates between backward elimination, which starts from the full set of items, and forward selection, which starts from the empty set. Starting from a full itemset, backward elimination will eliminate the item that corresponds to the highest IPOQ-LL-DIF. Forward selection gives the search procedure the ability to recover items later in the process.

4. Experimental Study

To evaluate our new method, we experiment on an artificial dataset and on three publicly available real-world datasets.

4.1. Application to Artificial Data

This simulation aims to show that our semi-automated algorithm aligns with the standard Rasch analysis procedure for dealing with DIF, i.e., identifies, resolves, and removes split items that are relatively hard to predict. In this experiment, we consider an artificial dataset that consists of responses to 14 items from 490 subjects. The dataset is composed of two inhomogeneous subsets with 6 items (12 items in total) and two DIF items. To simulate the DIF effect, the subjects are split into two different groups of 245. Responses are generated independently from the generalized partial credit model for the polytomous case with $m_i = 5$ ordered categories.

The response in all inhomogeneous subsets are generated with the same person ability and item difficulty scores, $\theta_n = (0.02(n - 1) - 3)$ for $n = 1, \dots, 245$ and $\beta_{ij} = ((i - 1) - (1.3 + 0.8(j - 1)))$ for $j = 1, \dots, 4$ and $i = 1, \dots, 6$, respectively. However, the discrimination parameters are varied among subsets, $\alpha = \{0.2, 2\}$ for the first and second subsets, respectively. This should make responses on items in the first subset (item₁₋₆) relatively hard to predict, and relatively easy to predict in the second subset (item₇₋₁₂). To simulate the effect of DIF, responses of both DIF items are generated with different difficulty parameters for the two subgroups. We set the thresholds to $\beta = \{-3.7, -2.9, -2.1, -1.3\}$ for the first and $\beta = \{1.3, 2.1, 2.9, 3.7\}$ for the second subgroup of the first DIF item (item₁₃). As for the second DIF item (item₁₄), we set $\beta = \{-5.2, -4.4, -3.6, -2.8\}$ for the first subgroup, and $\beta = \{2.8, 3.6, 4.4, 5.2\}$ for the second. Furthermore, we choose item₁₃ to be predictive and item₁₄ to be hard-to-predict for both subgroups, by setting $\alpha = \{0.8\}$ and $\alpha = \{0.1\}$, respectively.

For a given dataset containing DIF item(s), the PCM-DIF (see Section 2.1) can be applied to estimate the DIF parameters. Figure 1 shows that the PCM-DIF can identify the DIF items (item₁₃ and item₁₄) when the value of λ_δ is not too high. However, for a high value of λ_δ , these DIF effects disappear and the PCM-DIF leads to the same estimated parameters as a standard PCM without DIF. The PCM-DIF correctly estimates the DIF parameters (δ) of all non-DIF items to equal zero for any value of λ_δ .

Figure 1: The estimated DIF parameters ($\hat{\delta}$) along $\log(\lambda_\delta)$ for DIF and non-DIF items. The dashed grey line represents the value of λ_δ that is used in the estimation.

Infit is one of the item fit statistics that is commonly used to judge the goodness-of-fit of items to the Rasch model and to the PCM. In Figure 2, we show that this statistic relates to the discriminative power of the items, represented by the discrimination parameter $\hat{\alpha}$. A hard-to-predict item with low discriminative power normally has a high Infit, which indicates misfit. On the contrary, an easy-to-predict item with high discriminative power normally has a low infit. Further, we also show that the misfitting item₁₃ in Figure 2a (estimated using the PCM) improves its Infit after considering the DIF effect (estimated using PCM-DIF). The PCM-DIF clearly models the responses of item₁₃ better than the PCM. As expected, applying the PCM-DIF does not improve the Infit of item₁₄, the hard-to-predict DIF item. For the non-DIF items, the PCM-DIF estimates are indistinguishable from the PCM estimates.

Figure 2: The estimated discrimination parameters ($\hat{\alpha}$) against the Infit statistics of (a) the PCM and (b) the PCM-DIF. Discrimination parameters (α) are estimated using (a) the GPCM and (b) the GPCM-DIF. Horizontal line represents the expected Infit value and vertical line represents the expected discrimination value.

When DIF is present in particular items, standard Rasch analysis tends to detect and resolve these items. This step is then followed by removing misfit items, including those that are hard to predict even after splitting the DIF items. As shown in Figures 3b, our semi-automated algorithm does the same for reasons explained in Section 3.2: the IPOQ-LL-DIF score favors the DIF item that has a good fit after being split and puts the one that has a low discrimination parameter in the excluded set. The maximum of the IPOQ-LL-DIF score as a function of the number of items in the included set in this simulation is obtained when seven items are still included, including the resolved item₁₃.

As a comparison, we also apply our previous criterion, the IPOQ-LL, to this dataset. Figure 3a shows that the IPOQ-LL detects item₁₉ as a hard-to-predict item since it cannot estimate the DIF effect. Consequently, when a potential DIF effect is ignored item₁₃ will be put to the excluded itemset. As for item₁₄, being designed as a hard-to-predict DIF item, both IPOQ-LL and IPOQ-LL-DIF agree to put it in the excluded itemset. This fact is also supported by Figure 2 which shows that the Infit statistic of item₁₄ does not improve even after the DIF effect has been identified.

Figure 3: The highest IPOQ-LL and IPOQ-LL-DIF scores as a function of the number of included items $|S_{in}|$ for the artificial dataset. (a) IPOQ-LL computed with GPCM. (b) IPOQ-LL-DIF computed with GPCM-DIF. The numbers on the plot shows in which order the items are removed.

4.2. Application to Real-World Datasets

To validate our method on real-world data, we searched for datasets that satisfy the following criteria:

- The original dataset (survey with responses) is publicly available.
- A manual Rasch analysis has been applied to develop an instrument.
- According to the manual Rasch analysis, the initial survey contains differential item functioning.
- None of the authors of the current paper have been involved in the development of the instrument.
- The corresponding publication is not more than 5 years old.

We found three such datasets: the Osteopathy Clinical Teaching Questionnaire dataset [33], the Interdisciplinary Education Perception Scale dataset [34], and the Multiple Sclerosis Quality of Life Scale dataset [25]. To these three datasets, we applied our semi-automated procedure with the new criterion, the IPOQ-LL-DIF. For comparison, we also use the IPOQ-LL criterion.

4.2.1. The Osteopathy Clinical Teaching Questionnaire (OCTQ) Dataset

The Osteopathy Clinical Teaching Questionnaire (OCTQ) is an instrument that was developed to assess the quality of the clinical educators [33]. The original survey contains 30 items with five-point Likert scale and 3 global questions that have been answered by 399 participants. Vaughan [33] performed a manual Rasch analysis and ended up with 12 items, $S_{in} = \{2, 5, 7, 9, 10, 12, 15, 16, 18, 20, 23, 30\}$, as the final instrument. We will refer to this set of 12 items as the OCTQ manual instrument.

In the original survey, Vaughan [33] identified some items with disordered thresholds, four items with DIF (item₁₄, item₁₉, item₂₇, and item₂₈), and 122 misfitting persons. After resolving the few items with disordered thresholds in the original survey (item₁, item₉, item₂₇, and item₃₀), we applied the semi-automated procedure

for both criteria. Running the whole stepwise procedure leads to the result shown in Figure 4. Both criteria agree that the maximum of the IPOQ-LL-DIF occurs when the same 26 items are still included. The vertical lines give the location of maximum scores, $|\mathcal{S}_{\text{in}}| = 26$ and $|\mathcal{S}_{\text{in}}| = 12$. The horizontal lines give the location of the corresponding scores and show the score differences among instruments. Figure 4b zooms in on the search result near $|\mathcal{S}_{\text{in}}| = 12$, the number of items in the OCTQ manual instrument.

Figure 4: The highest IPOQ-LL-DIF scores as a function of the number of included items $|\mathcal{S}_{\text{in}}|$ when running the semi-automated procedure using the IPOQ-LL-DIF and IPOQ-LL criteria on the OCTQ dataset. Graph (b) zooms in on the number of included items close to those of the manual instrument.

For a fair and easy comparison with the manual instrument, we zoom in on the semi-automated instruments that are based on the same number of included items. We will refer to these as the IPOQ-LL-DIF ($\mathcal{S}_{\text{in}} = \{2, 3, 5, 7, 10, 12, 14, 16, 18, 22, 26, 30\}$) and IPOQ-LL ($\mathcal{S}_{\text{in}} = \{2, 3, 5, 7, 10, 12, 14, 16, 17, 22, 26, 30\}$) instruments, respectively. The semi-automated instruments only differ in one item: 17 versus 18. The overlap between the IPOQ-LL-DIF and the manual instrument is 8 items, which can be considered large: the probability of having an overlap of 8 or more items just by chance is smaller than 0.05. The overlap between the IPOQ-LL instrument and the manual instrument is 7 items.

In the initial analysis, Vaughan [33] suspected some items to be DIF items, i.e., item₂₇, and item₂₈ for institution, item₁₄ for institution and educator gender, and item₁₉ for student gender. As part of the Rasch analysis, Vaughan [33] chose to remove all DIF items to ensure that the final version of OCTQ would be applicable to a range of teaching institutions and free from gender influence. Employing the IPOQ-LL-DIF criterion, our semi-automated procedure retains three DIF items, i.e., item₁₄, item₂₂, and item₂₆.

To further illustrate the clinimetric quality of the three (i.e., manual, IPOQ-LL-DIF, and IPOQ-LL) instruments, we consider standard Rasch statistics such as goodness-of-fit, local independency, reliability, and unidimensionality. For comparison we also compute these statistics for 10000 randomly drawn 12-item instruments. The statistics of the three instruments are all well within the acceptable range and, in this case, are better than most of the random 12-item instruments in

local independency (see Appendix A.1 for details).

Figure 5 compares all instruments using our own IPOQ-LL-DIF criterion. By definition, the IPOQ-LL-DIF instrument is very well optimized for this criterion and slightly higher than the IPOQ-LL instrument. However, the manual instrument also does well and better than most of the randomly drawn 12-item instruments, which shows that as an extension of the IPOQ-LL criterion, the IPOQ-LL-DIF intrinsically captures many of the properties that a typical Rasch analysis cares about, including the presence of the DIF.

Figure 5: In-plus-out-questionnaire log likelihood with DIF (IPOQ-LL-DIF) values for the IPOQ-LL-DIF instrument (green dotted-dashed line), the IPOQ-LL instrument (brown dashed line), the manual instrument from [33] (red dotted line), and random 12-item instruments (histogram) on the OCTQ dataset.

Considering the standard Rasch statistics, which are averages over all items and all subjects, we conclude that the manual and the semi-automated instruments are clinimetrically all very similar. We then also expect that the abilities estimated for individual subjects based on the manual and the IPOQ-LL-DIF instruments will be very much alike. Figure 6 plots these estimated ability parameters for the two instruments against each other. Indeed, the estimated ability parameters for the two instruments are highly correlated ($\rho = 0.975$).

4.2.2. The Interdisciplinary Education Perception Scale (IEPS) Dataset

The interdisciplinary education perception scale is an instrument to evaluate students' professional perception in a particular program [34]. The complete survey consists of 18 items of 6-points Likert scale that are answered by 319 participants. Adopting the work of Leitch [17], Vaughan [34] excluded item₁₂ and item₁₇ and applied a manual Rasch analysis to 16 items as the initial survey. During the analysis, Vaughan [34] identified 51 misfitting persons, resolved four items with disordered thresholds (i.e., item₁₀, item₁₃, item₁₅, and item₁₆), removed 8 items, and ended up with 8 items as the final instrument. We will refer to this set of 8 items, $\mathcal{S}_{\text{in}} = \{1, 2, 7, 10, 13, 14, 15, 16\}$, as the IEPS manual instrument.

After resolving the four items with disordered thresholds, we applied our semi-automated procedure to the remaining 16-items IEPS responses. Running the whole

Figure 6: Estimated abilities for individual subjects based on the IPOQ-LL-DIF against those based on the manual instrument. The average root mean squared standard error for the estimates on both axes is visualized through the errorbars at the top left.

stepwise procedure using the IPOQ-LL-DIF and IPOQ-LL criteria leads to the graph shown in Figure 7. Both criteria agree that the maximum of the IPOQ-LL-DIF occurs when the same 12 items are still included. Using the same setup, the vertical lines give the location of the maximum scores, $|\mathcal{S}_{\text{in}}| = 12$ and $|\mathcal{S}_{\text{in}}| = 8$. The horizontal lines give the location of the corresponding scores and show the score differences among instruments.

Figure 7: The highest IPOQ-LL-DIF scores obtained for each number of included items $|\mathcal{S}_{\text{in}}|$ when running the semi-automated procedure using both the IPOQ-LL-DIF and IPOQ-LL criteria on the IEPS dataset. (a) Of all available $|\mathcal{S}_{\text{in}}|$. (b) Zoomed in version from $|\mathcal{S}_{\text{in}}| = 6$ to $|\mathcal{S}_{\text{in}}| = 10$. The horizontal grey lines display the differences of the IPOQ-LL-DIF scores obtained by the three instruments, i.e., the manual, IPOQ-LL, and IPOQ-LL-DIF instruments.

For a fair and easy comparison with the manual instrument, we again zoom in on the semi-automated instruments that are based on the same number of included items as the manual instrument. We will refer to these as the IPOQ-LL-DIF and IPOQ-LL instruments, respectively, which happen to contain the exact same items ($\mathcal{S}_{\text{in}} = \{1, 2, 4, 5, 7, 13, 15, 16\}$). Figure 7b zooms in on the search result near $|\mathcal{S}_{\text{in}}| = 8$, the number of items in the IEPS manual instrument. The overlap between the

IPOQ-LL-DIF and the manual instrument is 6 items. The probability of having an overlap of 6 or more items just by chance is 0.07.

Vaughan [34] also reported the presence of three DIF items, i.e., item₆ for year level, item₁₁ for gender, and item₁₈ for university. Vaughan [34] chose to remove all DIF items in order to produce a questionnaire that is free of demographic influence. Our semi-automated procedure also led to the removal of these three potential DIF items, but for a different reason: they did not survive the selection procedure when optimizing the IPOQ-LL-DIF and IPOQ-LL.

The figures in Appendix A.2 show that the three (i.e., manual, IPOQ-LL-DIF, and IPOQ-LL) instruments are clearly better than most of the 10000 randomly drawn 8-item instruments, especially on person separation reliability (PSR), local dependency, and unidimensionality. Furthermore, the statistics for these three instruments are clinimetrically very similar and all well within the acceptable range. Figure 8 shows that the manual instrument obtains an IPOQ-LL-DIF score that is only slightly smaller than the one for the semi-automated instruments. The estimated abilities of the manual and the semi-automated instruments indeed turn out to be very similar ($\rho = 0.94$) (see Figure 9).

Figure 8: In-plus-out-questionnaire log likelihood with DIF (IPOQ-LL-DIF) values for the IPOQ-LL-DIF instrument (green dotted-dashed line), the IPOQ-LL instrument (brown dashed line), the manual instrument from [34] (red dotted line), and random 8-item instruments (histogram) on the IEPS dataset.

4.2.3. The Multiple Sclerosis Quality of Life (MSQOL) Dataset

The multiple sclerosis quality of life questionnaire is an individual’s or a group’s perceived physical and mental health over time for people with multiple sclerosis [25]. The initial MSQOL survey consists of 54 items with different numbers of categories that was answered by 473 patients. The items are grouped into 12 multi-item and two single-item subscales. Rosato et al. [25] applied separate manual Rasch analyses to 11 subscales, each originally containing at least three items. For two of these subscales, no items survived the analysis. We will refer to the remaining sets of nine subscales as the MSQOL manual instruments. They are listed in Table A.3.

For two subscales (“Bodily Pain” and “Sexual Function”), the manual Rasch analysis kept all items. Running our semi-automated stepwise procedure on these

Figure 9: Estimated abilities for individual subjects based on the IPOQ-LL-DIF against those based on the manual instrument. The average root mean squared standard error for the estimates on both axes is visualized through the errorbars at the top left.

same subscales leads to the results shown in Figure 10. It can be seen that the semi-automated procedure agrees to keep all items from both subscales: the maximum IPOQ-LL-DIF scores are obtained with all items still in the included set.

Next, we applied our semi-automated procedure on all 11 subscales, constraining the semi-automated instruments to end up with the same number of items as the corresponding manual instruments. For optimization with the IPOQ-LL-DIF and IPOQ-LL criteria we arrived at the exact same included itemsets. We will refer to these as the MSQOL semi-automated instruments, also listed in Table A.3. As can be seen, the semi-automated and manual instruments have 21 out of 27 items overlapping, which can be considered a lot: the probability of having an overlap of 21 or more items just by chance is smaller than 0.01.

Figure 10: The highest IPOQ-LL-DIF scores obtained for each number of included items $|S_{in}|$ when running the semi-automated procedure using either the IPOQ-LL-DIF or IPOQ-LL criterion to the MSQOL dataset on (a) subscale “Bodily pain” and (b) subscale “Sexual Function”.

Table A.4 compares the psychometric quality of the manual and the semi-automated MSQOL instruments for all subscales. It can be seen that the standard Rasch statistics for both instruments are more or less the same.

Rosato et al. [25] also reported the presence of eight DIF items, i.e., item_1 , item_8 , item_{10} , item_{23} , item_{32} , item_{36} , item_{40} , and item_{51} and decided to remove these by hand. Our semi-automated procedure, however, retains two of these items, i.e., item_8 and item_{40} , albeit with corresponding DIF parameters $\hat{\delta}$ set to zero, i.e., without treating these as DIF items.

5. Discussion and Conclusion

In this work, we have successfully enhanced our semi-automated procedure to deal with DIF items. We extend our previous criterion, the so-called in-plus-out-of-questionnaire log likelihood (IPOQ-LL) to a new criterion named in-plus-out-of-questionnaire log likelihood with DIF (IPOQ-LL-DIF). The new criterion is based on the same ideas as the IPOQ-LL [35]: a good final instrument should reliably estimate people’s abilities. Although this ability estimate is fitted on the responses of the items in the final instrument, it should still represent the items that are left out.

The effectiveness of our extended procedure to yield clinimetrically similar results as the standard Rasch analyses relies on four essential ingredients. Two are passed down from the previous procedure, while the other two are new. The inherited ingredients are the flexible discrimination parameters and stronger regularization for this discrimination parameter on the included itemset compared to the excluded itemset [35]. The new ingredients are the DIF parameters together with a lasso penalty that distinguishes between the DIF and non-DIF items. We have shown that the new procedure also naturally incorporates essential aspects of Rasch analysis, where it tends to favor DIF items with a good Infit value and to exclude the ones without.

In our simulations, we have shown that DIF item(s) can indeed help obtain a better parameter estimate, in accordance with [2]. In a validation with real-world datasets our procedure yields similar instruments to the manual analyses, clinimetrically and physically when we constrain the number of included items to be the same as the number of included items in the manual instrument. With reasonable settings for the regularization parameters, our procedure tends to be somewhat more conservative in that it typically prefers to keep more items than the manual instrument.

In our experiments on real-world data, the IPOQ-LL-DIF and the IPOQ-LL criterion lead to very similar, often even the same instruments. Even though our procedure does detect and include DIF items, properly modeling this DIF effect has a relatively small effect on the selection of other items. However, this does depend on the setting of the regularization parameter λ_δ . In our main experiments, we chose $\lambda_\delta = 10$, which is relatively small, to arrive at more or less the same DIF items as in the manual analysis. An alternative approach would be to optimize λ_δ in a cross-validation procedure (see Appendix B.2). Applying this cross-validation procedure to the real-world data, we obtain a much larger λ_δ . With this stronger penalty, our procedure no longer finds any DIF items (see Figure B.1).

To summarize, in our real-world experiments the manual, IPOQ-LL-DIF, and IPOQ-LL instruments are largely comparable and have clinimetrically similar qual-

ities. Compared to randomly generated instruments, they all score well on standard Rasch statistics (see Figures A.1 through A.4).

In this paper, we have assumed that the DIF groups are specified beforehand. For binary types of information (e.g., gender), these groups are naturally defined. For continuous type information (e.g., age), our procedure can be easily extended with a recursive partitioning method based on the IPOQ-LL-DIF to find the optimal groups, along the lines of some recent methods for detecting DIF [15, 28, 31].

We are careful to frame our procedure as semi-automated rather than fully automated. We are aware that either using the manual or the (semi-)automated method, there are still several aspects to consider, e.g., clinical expertise in developing questions, the pre- and post-analysis, etc. With a decent original survey and clinical expertise, our procedure could potentially develop a valid, reliable, and clinimetrically robust instrument in a less time-consuming and more objective manner. Consequently, it challenges the current practice of Rasch analyses and questions the necessity of cumbersome manual procedures.

References

- [1] Akaike, H., 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19, 716–723. doi:10.1109/TAC.1974.1100705.
- [2] Andrich, D., Hagquist, C., 2015. Real and artificial differential item functioning in polytomous items. *Educational and Psychological Measurement* 75, 185–207. doi:10.1177/0013164414534258.
- [3] Bollmann, S., Berger, M., Tutz, G., 2018. Item-focused trees for the detection of differential item functioning in partial credit models. *Educational and Psychological Measurement* 78, 781–804. doi:10.1177/0013164417722179.
- [4] Bond, T.G., Fox, C.M., 2015. Applying the Rasch Model: Fundamental Measurement in the Human Sciences. doi:10.1017/CB09781107415324.004.
- [5] Borsboom, D., 2006. When does measurement invariance matter? *Medical Care* 44. doi:10.1097/01.mlr.0000245143.08679.cc.
- [6] Choi, S.W., Gibbons, L.E., Crane, P.K., 2011. lordif: An R package for detecting differential item functioning using iterative hybrid ordinal logistic regression/item response theory and monte carlo simulations. *Journal of statistical software* 39, 1–30. doi:10.18637/jss.v039.i08.
- [7] Christensen, K.B., Makransky, G., Horton, M., 2017. Critical values for Yen’s Q3: Identification of local dependence in the Rasch model using residual correlations. *Applied Psychological Measurement* 41, 178–194. doi:10.1177/0146621616677520.
- [8] Duncan, P.W., Bode, R.K., Lai, S.M., Perera, S., 2003. Rasch analysis of a new stroke-specific outcome scale: The stroke impact scale. *Archives of Physical Medicine and Rehabilitation* 84, 950–963.

- [9] Friedman, J., Hastie, T., Höfling, H., Tibshirani, R., 2007. Pathwise coordinate optimization. *The Annals of Applied Statistics* 1, 302–332. doi:10.1214/07-aos131, arXiv:0708.1485.
- [10] Gustafsson, J.e., 1980. Testing and obtaining fit of data to the Rasch model. *British Journal of Mathematical and Statistical Psychology* 33, 205–233. doi:10.1111/j.2044-8317.1980.tb00609.x.
- [11] Hagquist, C., Andrich, D., 2017. Recent advances in analysis of differential item functioning in health research using the Rasch model. *Health and Quality of Life Outcomes* 15, 1–8. doi:10.1186/s12955-017-0755-0.
- [12] Hermans, M.C., Faber, C.G., De Baets, M.H., de Die-Smulders, C.E.M., Merkies, I.S.J., 2010. Rasch-built myotonic dystrophy type 1 activity and participation scale (DM1-Activ). *Neuromuscular Disorders* 20, 310–318. doi:10.1016/j.nmd.2010.03.010.
- [13] Holland, P.W., Thayer, D.T., 1986. Differential item functioning and the Mantel-Haenszel procedure. *ETS Research Report Series* 1986, i–24. doi:10.1002/j.2330-8516.1986.tb00186.x.
- [14] Hu, L.T., Bentler, P.M., 1999. Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling* 6, 1–55. doi:10.1080/10705519909540118.
- [15] Komboz, B., Strobl, C., Zeileis, A., 2018. Tree-based global model tests for polytomous Rasch models. *Educational and Psychological Measurement* 78, 128–166. doi:10.1177/0013164416664394.
- [16] Kopf, J., Zeileis, A., Strobl, C., 2015. A framework for anchor methods and an iterative forward approach for DIF detection. *Applied Psychological Measurement* 39, 83–103. doi:10.1177/0146621614544195.
- [17] Leitch, J., 2014. Exploring psychometric properties of the interdisciplinary education perception scale in health graduate students. *Journal of Interprofessional Care* 28, 52–57. doi:10.3109/13561820.2013.820691.
- [18] Lönnfjörð, V., Hagquist, C., 2018. The psychometric properties of the Swedish version of the General Self-Efficacy Scale: A Rasch analysis based on adolescent data. *Current Psychology* 37, 703–715. doi:10.1007/s12144-016-9551-y.
- [19] Lord, F.M., Novick, M.R., 1968. *Statistical Theories of Mental Test Scores*. Addison-Wesley.
- [20] Magis, D., Facon, B., 2013. Item purification does not always improve DIF detection. *Educational and Psychological Measurement* 73, 293–311. doi:10.1177/0013164412451903.
- [21] Magis, D., Facon, B., 2014. deltaPlotR : An R package for differential item functioning analysis with Angoff’s delta plot . *Journal of Statistical Software* 59. doi:10.18637/jss.v059.c01.

- [22] Masters, G.N., 1982. A Rasch model for partial credit scoring. *Psychometrika* 47, 149–174. doi:10.1007/BF02296272.
- [23] Muraki, E., 1992. A generalized partial credit model: Application of an EM algorithm. *Applied Psychological Measurement* 16. doi:10.1177/014662169201600206.
- [24] Rasch, G.G., 1980. *Probabilistic Models for Some Intelligence and Attainment Tests*. University of Chicago Press.
- [25] Rosato, R., Testa, S., Bertolotto, A., Confalonieri, P., Patti, F., Lugaresi, A., Grasso, M.G., Toscano, A., Giordano, A., Solari, A., 2016. Development of a short version of MSQOL- 54 using factor analysis and item response theory. *PLoS ONE* 11, 1–15. doi:10.1371/journal.pone.0153466.
- [26] Schauberg, G., Mair, P., 2019. A regularization approach for the detection of differential item functioning in generalized partial credit models. *Behavior Research Methods* , 1–16doi:10.3758/s13428-019-01224-2.
- [27] Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- [28] Strobl, C., Kopf, J., Zeileis, A., 2015. Rasch Trees: A New Method for Detecting Differential Item Functioning in the Rasch Model. *Psychometrika* 80, 289–316. doi:10.1007/s11336-013-9388-3.
- [29] Swaminathan, H., Rogers, H.J., 1990. Detecting differential item functioning using logistic regression procedures. *Journal of Educational Measurement* 27, 361–370.
- [30] Tennant, A., Penta, M., Tesio, L., Grimby, G., Thonnard, J.L., Slade, A., Lawton, G., Simone, A., Carter, J., Lundgren-nilsson, Å., Tripolski, M., Ring, H., Biering-sørensen, F., 2004. Assessing and adjusting for cross-cultural validity of impairment and activity limitation scales through differential item functioning within the framework of the Rasch model. *Medical Care* 42, 37–48. doi:10.1097/01.mlr.0000103529.63132.77.
- [31] Tutz, G., Berger, M., 2016. Item-focussed Trees for the Identification of Items in Differential Item Functioning. *Psychometrika* 81, 727–750. doi:10.1007/s11336-015-9488-3.
- [32] Tutz, G., Schauberg, G., 2015. A penalty approach to differential item functioning in Rasch models. *Psychometrika* 80, 21–43. doi:10.1007/s11336-013-9377-6.
- [33] Vaughan, B., 2018. Exploring the measurement properties of the osteopathy clinical teaching questionnaire using Rasch analysis. *Chiropractic and Manual Therapies* 26, 1–16.

- [34] Vaughan, B., 2019. Measurement properties of the Interdisciplinary Education Perception Scale in an Australian allied health student cohort. *Health Professions Education* 5, 275–280. doi:10.1016/j.hpe.2018.07.005.
- [35] Wijayanto, F., Mul, K., Groot, P., van Engelen, B.G., Heskes, T., 2021. Semi-automated Rasch analysis using in-plus-out-of-questionnaire log likelihood. *British Journal of Mathematical and Statistical Psychology* 74, 313–339. doi:10.1111/bmsp.12218.
- [36] Woudstra, A.J., Meppelink, C.S., Pander Maat, H., Oosterhaven, J., Fransen, M.P., Dima, A.L., 2019. Validation of the short assessment of health literacy (SAHL-D) and short-form development: Rasch analysis. *BMC Medical Research Methodology* 19, 122. doi:10.1186/s12874-019-0762-4.
- [37] www.rasch.org, 2014. Rasch Measurement Analysis Software Directory. URL: <https://www.rasch.org/software.htm>.

Appendix A

This section describes the results of the standard Rasch statistics applied to the OCTQ, IEPS, and MSQOL datasets. The comparison considers manual, IPOQ-LL-DIF, IPOQ-LL, and randomly drawn instruments. We examine goodness-of-fit, local dependency, unidimensionality, and reliability. For goodness-of-fit and local dependency, we take the mean over all items. For Outfit and Infit, acceptable values are 0.6 - 1.4 for for all items in a Likert type survey [4]. As for residual correlations, commonly used threshold is lower than 0.2 or 0.3. However, higher thresholds (e.g., 0.5 and 0.7) are also frequently used [7].

Confirmatory factor analysis (CFA) is utilized to investigate the unidimensionality and the person separation reliability (PSR) measures the reliability of the estimated parameters. The combination of comparative fit index (CFI), Tucker-Lewis index (TLI), root mean square of approximation (RMSEA), and standardized root mean square residual (SRMR) is often used to describe the unidimensionality with particular thresholds, e.g., $CFI > 0.95$, $TLI > 0.95$, $RMSEA < 0.06$, and $SRMR < 0.08$ [14]. As for reliability, $PSR \geq 0.7$ is considered acceptable, $PSR \geq 0.8$ means good reliability, and $PSR \geq 0.9$ is excellent [8].

A.1. OCTQ Dataset Analysis Results

Table A.1: The results of Rasch standard statistics on the OCTQ dataset for the three instruments, i.e., the manual, IPOQ-LL-DIF, and IPOQ-LL instruments.

No.	Statistics	Manual items	IPOQ-LL-DIF items	IPOQ-LL items
1	Outfit	0.61 - 0.97 (mean = 0.732)	0.53 - 1.01 (mean = 0.726)	0.53 - 1.03 (mean = 0.721)
2	Infit	0.75 - 1.13 (mean = 0.930)	0.73 - 1.18 (mean = 0.921)	0.72 - 1.21 (mean = 0.918)
3	Residual correlation	mean = -0.08	mean = -0.08	mean = -0.08
4	PSR	0.818	0.826	0.829
5	Unidimensionality	CFI = 0.995, TLI = 0.993, RMSEA = 0.084, SRMR = 0.051	CFI = 0.991, TLI = 0.989, RMSEA = 0.122, SRMR = 0.075	CFI = 0.990, TLI = 0.988, RMSEA = 0.128, SRMR = 0.078

Figure A.1: Statistics of the original 12-item instrument from [33] obtained through manual Rasch analysis (red dotted line), the optimal 12-item instrument according to the IPOQ-LL-DIF criterion (green dotted-dashed line), IPOQ-LL criterion (brown dashed line), and 10000 random 12-item instruments (histogram) on the OCTQ dataset. (a) Mean of Outfit MnSq. (b) Mean of Infit MnSq. (c) Mean of residual correlation. (d) Person separation reliability (PSR).

Figure A.2: Unidimensionality test indices for the original 12-item instrument from [33] obtained through confirmatory factor analysis (red dotted line), the optimal 12-item instrument according to the IPOQ-LL-DIF criterion (green dotted-dashed line), IPOQ-LL criterion (brown dashed line), and 10000 random 12-item instruments (histogram) on the osteopathy clinical teaching questionnaire dataset. (a) Comparative fit index (CFI). (b) Tucker-Lewis index (TLI). (c) Root mean-square error of approximation (RMSEA). (d) Standardized root mean squared residual (SRMR).

A.2. IEPS Dataset Analysis Results

Table A.2: The result of Rasch standard statistics on the IEPS dataset for the three instruments, i.e., the manual, IPOQ-LL-DIF, and IPOQ-LL instruments.

No.	Statistics	Manual items	IPOQ-LL-DIF items	IPOQ-LL items
1	Outfit	0.82 - 1.06 (mean = 0.931)	0.83 - 1.21 (mean = 0.94)	0.83 - 1.21 (mean = 0.94)
2	Infit	0.86 - 1.03 (mean = 0.946)	0.85 - 1.19 (mean = 0.954)	0.85 - 1.19 (mean = 0.954)
3	Residual correlation	mean = -0.139	mean = -0.137	mean = -0.137
4	PSR	0.83	0.81	0.81
5	Unidimensionality	CFI = 0.97, TLI = 0.96, RMSEA = 0.11, SRMR = 0.07	CFI = 0.96, TLI = 0.94, RMSEA = 0.12, SRMR = 0.08	CFI = 0.96, TLI = 0.94, RMSEA = 0.12, SRMR = 0.08

Figure A.3: Statistics of the original 12-item instrument from [34] obtained through manual Rasch analysis (red dotted line), the optimal 8-item instrument according to the IPOQ-LL-DIF criterion (green dotted-dashed line), IPOQ-LL criterion (brown dashed line), and 10000 random 8-item instruments (histogram) on the interdisciplinary education perception scale dataset. (a) Mean of Outfit MnSq. (b) Mean of Infit MnSq. (c) Mean of residual correlation. (d) Person separation reliability (PSR).

Figure A.4: Unidimensionality test indices for the original 8-item instrument from [34] obtained through confirmatory factor analysis (red dotted line), the optimal 8-item instrument according to the IPOQ-LL-DIF criterion (green dotted-dashed line), IPOQ-LL criterion (brown dashed line), and 10000 random 8-item instruments (histogram) on the interdisciplinary education perception scale dataset. (a) Comparative fit index (CFI). (b) Tucker-Lewis index (TLI). (c) Root mean-square error of approximation (RMSEA). (d) Standardized root mean squared residual (SRMR).

A.3. MSQOL Dataset Analysis Results

This section describes the analysis results applied to the 11 subscales of the MSQOL dataset. Table A.3 shows the final instruments obtained by the manual Rasch analysis and the semi-automated procedure for each subscale. For all subscales, the IPOQ-LL instruments contain the same items as the IPOQ-LL-DIF instruments. Table A.4 reports the results of the commonly used standard Rasch statistics for subscales of which the manual and the semi-automated instruments show some differences in item preferences.

Table A.3: Items included in the two instruments, i.e., the MSQOL manual, and IPOQ-LL-DIF, for 11 subscales. The numbers of the items in the IPOQ-LL-DIF instruments are constrained to the number of the items in the corresponding manual instruments.

No.	Subscale	Manual items	IPOQ-LL-DIF items
1	Physical Function	4,5,6,7,9,11	3,5,7,8,9,11
2	Role Limitations — Physical	-	-
3	Role Limitations — Emotional	-	-
4	Bodily Pain	21,22,52	21,22,52
5	Emotional Wellbeing	25,26,30	25,26,28
6	Energy	27,29,31	27,29,31
7	Health Perceptions	35	34
8	Social Function	33	20
9	Cognitive Function	42,43,44	42,43,44
10	Health Distress	38,39,41	38,39,40
11	Sexual Function	46,47,48,49	46,47,48,49

Table A.4: Standard Rasch statistics for the manual and IPOQ-LL-DIF MSQOL instruments. We only report those subscales for which the two instruments differ, i.e., “Physical Function”, “Emotional Wellbeing”, “Health Perceptions”, “Social Function”, and “Health Distress”.

No.	Statistics	Manual items	IPOQ-LL-DIF items
Physical Function			
1	Outfit	0.32 - 0.46 mean = 0.4	0.30 - 0.65 mean = 0.44
2	Infit	0.54 - 0.69 mean = 0.62	0.55 - 0.71 mean = 0.64
3	PSR	0.78	0.80

No.	Statistics	Manual items	IPOQ-LL-DIF items
4	Residual Correlation	no $r > 0.3$ mean = -0.11	no $r > 0.3$ mean = -0.12
5	CFA	CFI = 1 TLI = 1 RMSEA = 0.82 SRMR = 0.02	CFI = 1 TLI = 1 RMSEA = 0.3 SRMR = 0.014
6	IPOQ-LL-DIF	-2010.68	-1966.41
Emotional Wellbeing			
1	Outfit	0.65 - 0.72 mean = 0.67	0.50 - 0.81 mean = 0.61
2	Infit	0.67 - 0.73 mean = 0.69	0.51 - 0.85 mean = 0.63
3	PSR	0.79	0.82
4	Residual Correlation	no $\rho > 0.3$ mean = -0.45	no $\rho > 0.3$ mean = -0.44
5	CFA	CFI = 1 TLI = 1 RMSEA = 0 SRMR = 0	CFI = 1 TLI = 1 RMSEA = 0 SRMR = 0
6	IPOQ-LL-DIF	-2177.78	-2123.44
Health Perceptions			
1	Outfit	0.16	0.15
2	Infit	0.12	0.13
3	IPOQ-LL-DIF	-2641.10	-2559.88
Social Function			
1	Outfit	0.14	0.14
2	Infit	0.08	0.09
4	IPOQ-LL-DIF	-1291.04	-1281.129
Health Distress			
1	Outfit	0.41 - 0.49 mean = 0.46	0.45 - 0.57 mean = 0.49
2	Infit	0.45 - 0.53 mean = 0.50	0.47 - 0.61 mean = 0.53
3	PSR	0.83	0.84
4	Residual Correlation	no $\rho > 0.3$ mean = -0.36	no $\rho > 0.3$ mean = -0.38
5	CFA	CFI = 1 TLI = 1 RMSEA = 0 SRMR = 0	CFI = 1 TLI = 1 RMSEA = 0 SRMR = 0
6	IPOQ-LL-DIF	-1460.23	-1456.97

Appendix B

B.1. Coordinate Descent

Due to the addition of an $L1$ (lasso) penalty in the current model, joint maximum likelihood estimation (JMLE) requires a more involved procedure. We propose to apply a two-level coordinate descent as described in Algorithm B.1. In the outer loop, we consider two sets of parameters. Our first coordinate contains all GPCM parameters (i.e., θ , β , and α), our second coordinate the DIF parameters (δ). Fixing the DIF parameters, we can optimize for the GPCM parameters. Given fixed GPCM parameters, we then have an inner loop in which we consider the DIF parameters (δ) as separate coordinates and solve them sequentially.

Algorithm B.1 Pseudocode for two-level coordinate descent.

```

1: procedure CoordinateDescent
2:    $iteration \leftarrow 0$ 
3:    $L_{(t)} \leftarrow L_{\mathcal{D}}(\theta_{init}, \beta_{init}, \alpha_{init}, \delta_{init})$ 
4:    $L_{(t-1)} \leftarrow L_{(t)} + (2 \times \eta)$ 
5:   while ( $|L_{(t)} - L_{(t-1)}| > \eta$ ) OR ( $iteration < max.iter$ ) do ▷ Outer loop
6:      $L_{(t-1)} \leftarrow L_{(t)}$ 
7:      $[\hat{\theta}, \hat{\beta}, \hat{\alpha}] \leftarrow \text{JMLE}(\mathcal{D})$ 
8:     for  $j := 1$  to  $n_{\delta}$  do ▷ Inner loop
9:        $[\hat{\delta}_j] \leftarrow \text{MLE}(\mathcal{D}, \hat{\theta}, \hat{\beta}, \hat{\alpha}, \hat{\delta}_{[\text{other than } j]})$ 
10:    end for
11:     $L_{(t)} \leftarrow L_{\mathcal{D}}(\hat{\theta}, \hat{\beta}, \hat{\alpha}, \hat{\delta})$ 
12:     $iteration \leftarrow iteration + 1$ 
13:  end while
14: end procedure

```

B.2. Parameter Setting

To tune the λ_{δ} to its optimal value, some methods are available. These methods include model selection criteria, e.g., Akaike information criterion (AIC) [1] and Bayesian information criterion (BIC) [27], as well as k -fold cross-validation [26]. In this work, we implement the k -fold cross-validation procedure described in Algorithm B.2. The items and subjects are split into several groups whose combinations constitute small blocks of the dataset. During cross-validation, we use one block as a test set and the other as a training set for every iteration.

We apply Algorithm B.2 for cross-validation and the results are shown in Figure B.1. In Figure B.1a, we see that for the artificial dataset it is better to give no penalty to the DIF parameters (δ). In contrast, Figure B.1b, B.1c, and B.1d on the real-world datasets suggest giving a large penalty to the DIF parameters (δ), which makes the DIF model equivalent to a non-DIF model. This suggests that the responses on these real-world datasets are not significantly different between the different DIF groups.

Algorithm B.2 Pseudocode for k-fold cross-validation.

```

1: procedure CrossValidation
2:    $\mathcal{I}_{\text{test}} \leftarrow \text{split } \{1, \dots, P\} \text{ to } n_I \text{ groups of items}$ 
3:    $\mathcal{R}_{\text{test}} \leftarrow \text{split } \{1, \dots, V\} \text{ to } n_R \text{ groups of respondents}$ 
4:    $K \leftarrow n_I \times n_R$ 
5:   for every  $\lambda_\delta$  values do
6:     for  $k := 1$  to  $K$  do
7:        $\mathcal{I}_{\text{train}} \leftarrow \{1, \dots, P\} \setminus \mathcal{I}_{\text{test}}[k]$ 
8:        $\mathcal{R}_{\text{train}} \leftarrow \{1, \dots, V\} \setminus \mathcal{R}_{\text{test}}[k]$ 
9:        $\mathcal{D}_{\text{train}} \leftarrow \mathcal{D}[\mathcal{R}_{\text{train}}, \mathcal{I}_{\text{train}}]$ 
10:       $\mathcal{D}_{\text{test\_item}} \leftarrow \mathcal{D}[\mathcal{R}_{\text{train}}, \mathcal{I}_{\text{test}}[k]]$ 
11:       $\mathcal{D}_{\text{test\_theta}} \leftarrow \mathcal{D}[\mathcal{R}_{\text{test}}[k], \mathcal{I}_{\text{train}}]$ 
12:       $\mathcal{D}_{\text{test}} \leftarrow \mathcal{D}[\mathcal{R}_{\text{test}}[k], \mathcal{I}_{\text{test}}[k]]$ 
13:       $[\hat{\theta}_{\text{train}}, \hat{\beta}_{\text{train}}, \hat{\alpha}_{\text{train}}, \hat{\delta}_{\text{train}}] \leftarrow \text{CoordinateDescent}(\mathcal{D}_{\text{train}}, \lambda_\delta)$ 
14:       $[\hat{\beta}_{\text{test}}, \hat{\alpha}_{\text{test}}, \hat{\delta}_{\text{test}}] \leftarrow \text{CoordinateDescent}(\mathcal{D}_{\text{test\_theta}}, \lambda_\delta, \hat{\theta}_{\text{train}})$ 
15:       $[\hat{\theta}_{\text{test}}] \leftarrow \text{CoordinateDescent}(\mathcal{D}_{\text{test\_item}}, \lambda_\delta, \hat{\beta}_{\text{train}}, \hat{\alpha}_{\text{train}}, \hat{\delta}_{\text{train}})$ 
16:       $L_{cv_k} \leftarrow L_{\mathcal{D}_{\text{test}}}(\hat{\theta}_{\text{test}}, \hat{\beta}_{\text{test}}, \hat{\alpha}_{\text{test}}, \hat{\delta}_{\text{test}})$ 
17:    end for
18:     $L_{cv_{\lambda_\delta}} \leftarrow \text{sum}(L_{cv_k})$ 
19:  end for
20: end procedure

```

Figure B.1: Results of cross-validation of four datasets using different values of λ_δ . (a) The artificial dataset. (b) The OCTQ dataset. (c) The IEPS dataset. (d) The first dimension of the MSQOL dataset.