Comparison of Recent Acceleration Techniques for the EM algorithm in One- and Two-Parameter Logistic IRT models

Marie Beisemann^{*†}, Ortrud Wartlick^{*} and Philipp Doeblner

TU Dortmund University

Abstract

The Expectation-Maximisation (EM) algorithm is an important numerical method for maximum likelihood estimation in incomplete-data problems. However, convergence of the EM algorithm can be slow, and for this reason, many EM acceleration techniques have been proposed. After a review of acceleration techniques in a unified notation with illustrations, three recently proposed EM acceleration techniques are compared in detail: quasi-Newton methods (QN) (Zhou et al., 2011), “squared” iterative methods (SQUAREM) (Varadhan & Roland, 2004; Roland et al., 2007; Varadhan & Roland, 2008) and parabolic EM (PEM) (Berlinet & Roland, 2009). These acceleration techniques are applied to marginal maximum likelihood estimation with the EM algorithm in one- and two-parameter logistic item response theory (IRT) models for binary data, and their performance is compared. QN and SQUAREM methods accelerate convergence of EM algorithm for the two parameter logistic model significantly in high-dimensional data problems. Compared to standard EM, all three methods reduce the number of iterations, but increase the number of total marginal log-likelihood evaluations per iteration. Efficient approximations of the marginal log-likelihood are hence an important part of implementations.

^{*}These authors contributed equally to this work.

[†]Corresponding author. Contact information: TU Dortmund University, Department of Statistics, Vogelpothsweg 87, 44227 Dortmund, Germany; e-mail: beisemann@statistik.tu-dortmund.de

1 Introduction

Item response theory (IRT) models have long become a staple in psychometric research with applications in the investigation of test properties in smaller to medium samples (e.g., in personality, intelligence, or creativity research, Lee et al., 2014; Chermahini et al., 2012), but also in large-scale assessments (Berezner & Adams, 2017; Von Davier & Sinharay, 2014). Especially popular and topic of the present work are logistic IRT models for binary data, in particular the one parameter (1PL; also: Rasch) and two parameter (2PL) model. While conditional maximum likelihood is a feasible for the 1PL model, marginal maximum likelihood (MML) estimation for IRT models is the predominant estimation approach, in particular for the 2PL model. MML estimation is usually performed with the help of the expectation-maximization (EM) algorithm (Dempster et al., 1977; Bock and Aitkin, 1981; for a comprehensive introduction see McLachlan and Krishnan, 2008, for implementation details, see Baker and Kim, 2004). The EM algorithm is well-suited to carry out ML estimation for situations in which one can consider parts of the "complete data" to be unobserved. In IRT models, the unobserved data is identical to latent ability. The EM algorithm maximizes the expected complete-data log-likelihood using an iterative two-step approach (Dempster et al., 1977; Hastie et al., 2001) of estimating the current posterior distribution of latent variables given the data and current or initial parameter estimates by determining the expected complete-data log-likelihood function (expectation (E) step), and then maximizing this expected complete-data log-likelihood function to obtain parameter estimates (maximization (M) step), until convergence. However, one problem with the EM algorithm is that convergence can be slow (McLachlan & Krishnan, 2008). This can pose a serious challenge to applied researchers employing the EM algorithm for IRT models in rich data settings, such as large-scale assessments. Also, simulation is a routine strategy to evaluate statistical decisions in the IRT context, so speeding up the EM algorithm speeds up methodology development.

Research into the acceleration of the EM algorithm has yielded several promising techniques, such as new quasi-Newton methods (QN) (Zhou et al., 2011), squared iterative methods (SQUAREM) (Varadhan & Roland, 2004; Roland et al., 2007;

Varadhan & Roland, 2008), and parabolic EM (Berlinet & Roland, 2009). However, these recent methods have only been presented in isolation and the way they work has not been compared in detail. Further, they have not been previously presented and compared in the psychometric context of IRT models. In fact, to our knowledge, only one psychometric paper (Bartolucci et al., 1998) from over twenty years ago has been dedicated to a comparison of EM accelerators in the context of IRT models. As several new methods have been developed since then and as rich data settings are becoming increasingly more common, this work aims to provide a reasonably self-contained introduction to general purpose EM accelerators, and a systematic comparison of recent EM accelerators in a psychometric setting.

2 Accelerating the EM algorithm

2.1 The EM algorithm for logistic IRT models

The 2PL model models N participants' dichotomous responses (with 1 for a correct and 0 for an incorrect response) to M items on a test, where the response X_{ij} of person i to an item j depends on the person's ability Z_i and some item parameters θ_j (Embretson, 2000). Assuming X_{ij} follows a Bernoulli distribution, the 2PL model assumes the probability of a correct response of person i with ability z_i to item j is

$$P(X_{ij} = 1|Z_i = z_i) = \frac{1}{1 + \exp(-a_j(z_i - d_j))}, \quad i = 1, \dots, N, j = 1, \dots, M, \quad (2.1)$$

with item discrimination a_j (describing how well item j discriminates between persons of high and low ability) and item difficulty d_j (Embretson, 2000). Equivalent but numerically advantageous is the slope-intercept parameterization (Chalmers, 2012; Baker & Kim, 2004), for which we replace the expression inside the exponential function with the expression $-(\alpha_j z_i + \delta_j)$, with the following relationships between the parameters: $\alpha_j = a_j$ and $\delta_j = -a_j d_j$. The 1PL model is nested within the 2PL model as a special case; it is obtained when $\alpha_j = a_j = 1$ for all items $j = 1, \dots, M$. Let $Q_j(z_i) := P(X_{ij} = 0|Z_i = z_i) = 1 - P(X_{ij} = 1|Z_i = z_i)$ for $i = 1, \dots, N$, $j = 1, \dots, M$, denote the probability of observing the response vector

$\mathbf{x}_i = (x_{i1}, \dots, x_{iM})$, i.e., the responses of person i to all M items, if the person has ability z_i , is given by (Baker & Kim, 2004)

$$P(\mathbf{x}_i|Z_i = z_i) = \prod_{j=1}^M P_j(z_i)^{x_{ij}} Q_j(z_i)^{1-x_{ij}}, \quad \text{for } i = 1, \dots, N. \quad (2.2)$$

Assuming that latent ability is normally distributed around 0 with variance σ^2 , $Z \sim N(0, \sigma^2)$, and $\varphi(z|\sigma^2)$ denoting the corresponding probability density function, the observed data or marginal likelihood (i.e., the responses of all N participants to all M items) is given by (Baker & Kim, 2004)

$$L_m(\theta; x) = \prod_{i=1}^N P(\mathbf{x}_i) = \prod_{i=1}^N \left(\int P(\mathbf{x}_i|Z_i = z_i, \theta) \varphi(z_i|\sigma^2) dz_i \right). \quad (2.3)$$

Here θ is the item parameter vector which characterizes the distribution of X and is to be estimated. In slope-intercept notation, it is given by $\theta = (\alpha_1, \dots, \alpha_M, \delta_1, \dots, \delta_M)^T$. Note that σ^2 , the variance of person ability, is generally also unknown, i.e., it may also have to be estimated (1PL case) or needs to be fixed, usually to 1 in the 2PL model, for identification purposes (Baker & Kim, 2004).

The integral in Equation 2.3 proves difficult to evaluate and thus is usually approximated numerically. A common approach is Gauss-Hermite (GH) quadrature (Baker & Kim, 2004), see Chalmers (2012) for other methods. GH quadrature approximates the continuous integral by a discrete sum over K discrete ability levels $\zeta_k, k = 1, \dots, K$, each weighted with a weight $\tilde{\varphi}_k$, that takes into account the distance between discrete ability levels and the height of the density function φ in the neighbourhood of ζ_k (Baker & Kim, 2004). Rewriting Equation 2.3 in quadrature notation yields (Baker & Kim, 2004)

$$L_m(\theta; x) = \prod_{i=1}^N \left(\sum_{k=1}^K (P(\mathbf{x}_i|\zeta_k, \theta) \tilde{\varphi}_k) \right). \quad (2.4)$$

By considering response patterns and how often they occur for each ability level across participants rather than considering participants' responses individually, we obtain the complete-data likelihood

$$L_c(\theta; n, r) = \prod_{j=1}^M \left[\left(\prod_{k=1}^K \binom{n_{jk}}{r_{jk}} P_j(\zeta_k)^{r_{jk}} Q_j(\zeta_k)^{n_{jk}-r_{jk}} \right) \left(\left[\frac{n_j!}{n_{j1}! \cdots n_{jK}!} \right] \prod_{k=1}^K \varphi_k^{n_{jk}} \right) \right], \quad (2.5)$$

where n_{jk} denotes the total number of persons of ability level ζ_k attempting item j and r_{jk} the total number of correct responses to item j for persons of ability level ζ_k . Further $n_j = (n_{j1}, \dots, n_{jK})$ represents the number of persons of each of the K ability levels attempting item j and $r_j = (r_{j1}, \dots, r_{jK})$ the number of correct responses to item j by the persons at each of the K ability levels. The tuple (n_j, r_j) can be said to contain “complete” information about an item j . The values of the n_j and r_j vectors are unobserved. For this reason, the posterior expected complete-data log-likelihood, conditional on the observed response data and current parameter estimates, is used instead (Baker & Kim, 2004, Section 6.4.1), where we replace r_{jk} and n_{jk} with their posterior expected values (Baker & Kim, 2004, Chapter 6):

$$\mathbb{E}_{n,r|X,\theta,\sigma^2}(n_{jk}) = \sum_{i=1}^N P(\tilde{Z}_i = \zeta_k | \mathbf{x}_i, \theta, \sigma^2) := \bar{n}_k, \quad \text{and} \quad (2.6)$$

$$\mathbb{E}_{n,r|X,\theta,\sigma^2}(r_{jk}) = \sum_{i=1}^N x_{ij} P(\tilde{Z}_i = \zeta_k | \mathbf{x}_i, \theta, \sigma^2) := \bar{r}_{jk}, \quad (2.7)$$

for $k = 1, \dots, K$ and $j = 1, \dots, M$. $P(\tilde{Z}_i = \zeta_k | \mathbf{x}_i, \theta, \sigma^2)$ is the posterior probability of person i having ability ζ_k , given that person’s observed response vector, $\mathbf{x}_i$, the item parameters θ and the variance of ability in the population, σ^2 . The posterior membership probabilities $P(\tilde{Z}_i = \zeta_k | \mathbf{x}_i, \theta, \sigma^2)$ for $i = 1, \dots, N$ and $k = 1, \dots, K$ are unknown and must be estimated during each E-step of the EM algorithm, given the response data $\mathbf{x}$, as well as the $\hat{\varphi}_{k,t}$ and current item parameter estimates $\hat{\theta}_t = (\hat{\alpha}_{1,t}, \dots, \hat{\alpha}_{M,t}, \hat{\delta}_{1,t}, \dots, \hat{\delta}_{M,t})^T$. The estimated posterior membership probabilities can then be used to compute estimates of $\bar{n}_k$ and $\bar{r}_{jk}$. Each E-step is then followed by an M-step, in which the posterior expected complete-data log-likelihood, that is the logarithm of term in Equation 2.5 with the estimates of posterior expected values $\bar{n}_k$ and $\bar{r}_{jk}$ replacing the respective n_{jk} and r_{jk} terms, is maximized with respect to the vector of item parameters $\hat{\theta}_t$. This is done iteratively by finding the root of the

derivatives of the posterior expected complete-data log-likelihood independently for each item, i.e., solving (Baker & Kim, 2004)

$$(\alpha_j) \quad \sum_{k=1}^K \zeta_k \left[\widehat{r}_{jk,t} - \widehat{n}_{k,t} P_j(\zeta_k) \right] W_{jk} = 0 \quad (2.8)$$

$$(\delta_j) \quad \sum_{k=1}^K \left[\widehat{r}_{jk,t} - \widehat{n}_{k,t} P_j(\zeta_k) \right] W_{jk} = 0, \quad (2.9)$$

for $j = 1, \dots, M$. Here $W_{jk} = \frac{P_j^*(\zeta_k) Q_j^*(\zeta_k)}{P_j(\zeta_k) Q_j(\zeta_k)}$, with $P_j^*(\zeta_k) = P_j(\zeta_k | c_j = 0)$ and $Q_j^*(\zeta_k) = 1 - P_j^*(\zeta_k)$ denoting the probabilities to get a correct and incorrect answer, respectively. Please note that unless the ability variance σ^2 is fixed, it also needs to be estimated during each M-step. E- and M-steps are then carried out in alternating fashion until convergence. It is worth noting that the convergence of the EM algorithm is not guaranteed if the IRT model employed is not an exponential family model (Baker & Kim, 2004, Section 6.4.1). Only the 1PL (or Rasch) model is a member of the exponential family. However, empirical work suggests that the EM algorithm generally also converges for other IRT models (Mislevy & Bock, 1982).

2.2 EM accelerators

Varadhan and Roland (2008) divide existing techniques into two broad classes: The first class, which they title monotone acceleration techniques, are modifications of specific EM algorithms. These methods include data augmentation (Meng & Van Dyk, 1997), parameter expansion EM (PX-EM) (Liu et al., 1998), Expectation/Conditional Maximisation (ECM) (Meng & Rubin, 1993) and Expectation/Conditional Maximisation Either (ECME) (Liu & Rubin, 1994). These methods have been shown to substantially speed up convergence rates for the EM algorithm (McLachlan & Krishnan, 2008). However, they have to be construed and implemented for each application of the EM algorithm, which means they are not feasible to use for applied researchers unless they have been specifically implemented for the model they are using. As this work aims to make recommendations for applied researchers, these methods will not be discussed here. Please see McLachlan and Krishnan (2008) for an introduction and an overview. The second class, referred to as

non-monotone by Varadhan and Roland (2008), includes methods which are based on finding the root of a function and are universally applicable to applications of the EM algorithm. This class includes quasi-Newton methods (Lange, 1995; Jamshidian & Jennrich, 1997), conjugate gradient methods (Jamshidian & Jennrich, 1993) and multivariate Aitken’s procedure (Louis, 1982). Only the quasi-Newton methods (QN1 and QN2 in the current paper’s terminology) were studied in a psychometric study on the topic (Bartolucci et al., 1998). In recent years, new acceleration methods of this class have been proposed and have shown promising performance: faster quasi-Newton methods (QN) (Varadhan & Roland, 2008; Zhou et al., 2011), squared iterative methods (SQUAREM) (Varadhan & Roland, 2004; Roland et al., 2007; Varadhan & Roland, 2008), and parabolic EM (PEM) (Berlinet & Roland, 2009). In the following, we will present these methods in unified notation for easy comparability and compare their performances in a numerical simulation.

2.2.1 EM accelerators and fixed-point mappings

The recently proposed EM accelerators discussed and compared in this work take advantage of the fact that the EM algorithm can be considered a fixed-point mapping F from the parameter space Ω onto itself (Berlinet & Roland, 2009),

$$\begin{aligned} F : \Omega \subset \mathbb{R}^p &\mapsto \Omega \\ \theta &\mapsto F(\theta). \end{aligned} \tag{2.10}$$

By using the mapping iteratively, as the EM algorithm does, a sequence $\theta_0, \theta_1, \dots, \theta_t, \dots$ is generated from a starting point θ_0 (Berlinet & Roland, 2009), where $\theta_{t+1} = F(\theta_t)$, $t \geq 0$. In other words, the mapping F corresponds to one evaluation of the E- and M-step of the EM algorithm. Under certain conditions (Wu, 1983), the sequence $\{\theta_i\}$ satisfies $L_c(\theta_{t+1}; x, z) \geq L_c(\theta_t; x, z)$, $t \geq 0$, and thus converges to θ^* , which corresponds to either a saddle point or a local maximum of L_c (and L_m) (McLachlan & Krishnan, 2008; Dempster et al., 1977; Wu, 1983). If F is differentiable, θ^* is a fixed point of F , i.e. $F(\theta^*) = \theta^*$. Linearisation of F in a small neighbourhood of θ^*

yields (Berlinet & Roland, 2009)

$$F(\theta_t) = F(\theta^*) + J|_{\theta^*} (\theta_t - \theta^*) + o(\|\theta_t - \theta^*\|), \quad \text{i.e.,} \quad (2.11)$$

$$\theta_{t+1} - \theta^* \approx J|_{\theta^*} (\theta_t - \theta^*), \quad \text{or}$$

$$e_{t+1} \approx J|_{\theta^*} e_t \quad \text{where} \quad e_s = \theta_s - \theta^*, \quad \text{for } s = t, t+1, \quad \text{and} \quad (2.12)$$

where $J|_{\theta^*}$ is the Jacobian matrix of F evaluated at θ^* . Usually, J cannot be expressed explicitly. The vector $e_{t+1} = \theta_{t+1} - \theta^*$ is the error after the $(t+1)$ th EM update. Equation 2.12 stresses that θ_{t+1} converges approximately linearly to θ^* . The rate of convergence depends on the eigenvalues of the rate matrix $J|_{\theta^*}$, with the latter having been shown to measure the fraction of missing information (Dempster et al., 1977).

As explained above, the EM algorithm effectively creates a sequence $(\theta_t)_t$ that converges to θ^* . To accelerate convergence, another sequence, $(\tilde{\theta}_t)_t$, has to be generated from the same starting point, but with a faster rate of convergence to θ^* , so that $\lim_{t \rightarrow \infty} (\tilde{\theta}_t - \theta^*)/(\theta_t - \theta^*) = 0$. Faster convergence also implies that fewer evaluations of F are required until a convergence criterion, such as, $\|\theta_{t+1} - \theta_t\| < \epsilon$ or $|\log L(\theta_{t+1}) - \log L(\theta_t)| < \epsilon$ is met (Berlinet & Roland, 2009). Here ϵ denotes a small, fixed, positive, real number. Note one could use either the observed or the complete-data log-likelihood because both are maximized by the same parameter estimates. The former is often easier to evaluate and thus often used to probe for convergence. In the following, for simplicity, the symbol θ is used throughout instead of $\tilde{\theta}$ to indicate a member of the accelerated parameter estimation sequence.

2.2.2 Steffenson-type methods

Before giving a detailed description of the accelerators studied in the present work, we are going to provide a brief introduction to Quasi-Newton (QN), in particular Steffenson-type methods (STEM), which form the basis for some of the more recently proposed accelerator and thus will be crucial in the understanding of them. For a more detailed description of these classes of methods please consult McLachlan and Krishnan (2008) or Varadhan and Roland (2008). Newton's method for locating

the fixed point of F is based on finding the root of a linear approximation of the residual function $g(\theta) = \theta - F(\theta)$ around θ_t , i.e. $0 = g(\theta^*) \approx g(\theta_t) + (I - J|_{\theta_t})(\theta^* - \theta_t)$. Iteratively, this can be expressed as $\theta_{t+1} = \theta_t - (I - J|_{\theta_t})^{-1}g(\theta_t)$, $t \geq 0$ (Varadhan & Roland, 2008; Zhou et al., 2011), which we are going to refer to as a Newton update. If $J|_{\theta_t}$ is unknown, it has to be approximated by a low-rank matrix, so that θ_{t+1} can be easily computed (leading to quasi-Newton methods). To find a good approximation for $J|_{\theta_t}$, secant constraints can be used (Varadhan & Roland, 2008; Zhou et al., 2011): If two successive iterations of the EM algorithm, $F(F(\theta_t))$, are performed close to the fixed point based on the current parameter estimate θ_t , Equation 2.11 yields $J|_{\theta^*} u \approx v$, with $u = F(\theta_t) - \theta_t$ and $v = F(F(\theta_t)) - F(\theta_t)$ (Varadhan & Roland, 2008).

Steffenson-type (STEM) methods make use of a single secant constraint to obtain an approximation $M \approx J|_{\theta_t}$, in the form of $M = [1 - (1/\alpha_t)]I$ (Varadhan & Roland, 2008). With this approximation for $J|_{\theta_t}$, the Newton update simplifies to a quasi-Newton update of $\theta_{t+1} = \theta_t + \alpha_t u =: STEM(\theta_t)$, $t \geq 0$. The parameter α_t is called the steplength (Varadhan & Roland, 2008). The STEM update therefore consists of a step of length α_t taken from θ_t in the direction of u (Fig. 1). Because the direction is fixed, a single STEM update cannot account for curvature in trajectories to the fixed point (Fig 1). Instead, the direction is continually corrected in subsequent updates.

The steplength α_t must be chosen in such a way that the secant constraint $Mu = v$ is satisfied, resulting in $u = -\alpha_t(v - u)$ as the constraint for α_t . Varadhan and Roland (2008) provide three suggestions for the steplength, all satisfying the imposed secant constraint:

$$\text{S1: } \alpha_t = \underset{\alpha_t \in \mathbb{R}}{\operatorname{argmin}} ||u + \alpha_t(v - u)||^2 = -\frac{u^T(v - u)}{(v - u)^T(v - u)} \quad (2.13)$$

$$\text{S2: } \alpha_t = \underset{\alpha_t \in \mathbb{R}}{\operatorname{argmin}} \frac{||u + \alpha_t(v - u)||^2}{\alpha_t^2} = -\frac{u^T u}{u^T(v - u)} \quad (2.14)$$

$$\text{S3: } \alpha_t = \underset{\alpha_t \in \mathbb{R}^-}{\operatorname{argmin}} \frac{||u + \alpha_t(v - u)||^2}{\alpha_t} = \frac{||u||}{||v - u||}. \quad (2.15)$$

The STEM method with steplength S1 is also equivalent to first order Reduced Rank Extrapolation (RRE1), a polynomial extrapolation method (Roland et al.,

Figure 1: Illustration of the STEM method in two dimensions, with $\theta = (a, b)^T$. Blue: a STEM update with θ_t as a starting point. Grey: successive EM updates from θ_t to θ^* . Black arrows: the vectors $u = F(\theta_t) - \theta_t$ and $v = F(F(\theta_t)) - F(\theta_t)$.

2007). Similarly, STEM with S2 is equivalent to first order Minimal Polynomial Extrapolation (MPE1) (Roland et al., 2007). In general, for a fixed θ_t and $u^T(v-u) < 0$, $\alpha_t^{S2} \geq \alpha_t^{S3} \geq \alpha_t^{S1} > 0$ (Varadhan & Roland, 2008). Note, however, that for acceleration to occur, α_t must be larger than 1 (Varadhan & Roland, 2008); for $\alpha_t = 1$, the quasi-Newton STEM update becomes a simple EM update, i.e., there is no acceleration.

A STEM update with one of the above steplengths does not necessarily increase the likelihood, i.e., STEM is not necessarily globally convergent. To obtain global convergence, the steplength may have to be adapted in such a way that the likelihood $L(\alpha_t) := L(\theta_{t+1}(\alpha_t))$ increases in every iteration. The strategy proposed by Varadhan and Roland (2008) for globally convergent STEM is as follows: If $\alpha_t < 1$, α_t is set to 1. If $\alpha_t > 1$ but $L(\alpha_t) < L(0)$, α_t is decreased towards 1 until $L(\alpha_t) > L(0)$. If $\alpha_t > 1$ and $L(\alpha_t) \geq L(0)$, α_t is accepted as steplength.

2.2.3 Advanced quasi-Newton methods (QN)

Zhou et al. (2011) propose a more advanced quasi-Newton method, which, instead of a single secant constraint, uses q secant constraints, yielding $J|_{\theta^*} U \approx V$, with $U = (u_1, \dots, u_q)$ and $V = (v_1, \dots, v_q)$, where $u_i = F(\theta_{t+1-i}) - \theta_{t+1-i}$ and $v_i = F(F(\theta_{t+1-i})) - F(\theta_{t+1-i})$, for $i = 1, \dots, q$. For $q = 1$ this simplifies to $J|_{\theta^*} u \approx v$ (compare above). Note that U and V are both $p \times q$ -matrices, where p is the dimension of θ . For the matrices to have full rank, q must not be bigger than p (Zhou et al., 2011).

An approximation $M \approx J|_{\theta_t}$ is determined by minimising the Frobenius norm $\|M\|_F^2$ under the secant constraints $MU = V$ (Zhou et al., 2011). This yields $M = V(U^T U)^{-1} U^T$, with the inverse $(I - M)^{-1} = I - V(U^T V - U^T U)^{-1} U^T$, and the quasi-Newton update

$$\theta_{t+1} = F(\theta_t) - V(U^T V - U^T U)^{-1} U^T (F(\theta_t) - \theta_t) =: QN(\theta_t), \quad t \geq 0. \quad (2.16)$$

For $q = 1$, this simplifies to $\theta_{t+1} = F(\theta_t) + \alpha_t v$, $t \geq 0$, with $\alpha_t = -(u^T u)/(u^T (v - u))$. In other words, the case $q = 1$ is similar to a STEM update with an additional EM update: One EM update is performed and then a step of length α_t is taken in the direction of v (Fig. 2). The steplength used is S2 (2.14).

Zhou et al. (2011) show that QN methods can be very efficient: With these methods, the fixed point is reached in 10-100 times fewer iterations than with the standard (unaccelerated) EM algorithm. Furthermore, these methods are suitable for high-dimensional problems, because of minimal storage requirements. In contrast to earlier quasi-Newton methods (Lange, 1995; Jamshidian & Jennrich, 1997), STEM and QN methods do not require the storage or manipulation of the observed information matrix or the Hessian of F . Combined with globalisation strategies, they are globally convergent (Zhou et al., 2011).

Figure 2: Illustration of the QN ($q = 1$) method in two dimensions, with $\theta = (a, b)^T$. Blue: a QN ($q = 1$) update with θ_t as a starting point. Grey: successive EM updates from θ_t to θ^* . Black arrows: the vectors $u = F(\theta_t) - \theta_t$ and $v = F(F(\theta_t)) - F(\theta_t)$.

2.2.4 Squared iterative methods (SQUAREM)

In a small neighbourhood of θ^* , STEM has the following error equation (Varadhan & Roland, 2008):

$$STEM(\theta_t) = \theta^* + \left. \frac{\partial STEM(\theta_t)}{\partial \theta} \right|_{\theta^*} (\theta_t - \theta^*) + o(\|\theta_t - \theta^*\|)$$

$$e_{t+1} \approx (I + \alpha_t (J|_{\theta^*} - I)) e_t, \quad (2.17)$$

where $e_{t+1} = \theta_{t+1} - \theta^*$ is the error after the $(t + 1)$ th STEM update. Equation (2.17) shows that the rate matrix of STEM is $(I + \alpha_t (J - I))$. The squared iterative method proposed by Varadhan and Roland (2008) is derived from STEM by squaring this rate matrix and demanding that the error of the new method satisfies $e_{t+1} = (I + \alpha_t (J|_{\theta^*} - I))^2 e_t$. From $0 = g(\theta^*) \approx g(\theta_t) + (I - J|_{\theta_t})(\theta^* - \theta_t)$ follows that $u = -g(\theta_t) = (J|_{\theta^*} - I) e_t$ and $v - u = -g(F(\theta_t)) + g(\theta_t) = (J|_{\theta^*} - I)^2 e_t$. This

results in the SQUAREM update (Varadhan & Roland, 2008):

$$\theta_{t+1} = \theta_t + 2\alpha_t u + \alpha_t^2(v - u) =: \text{SQUAREM}(\theta_t), \quad t \geq 0. \quad (2.18)$$

Figure 3 shows a SQUAREM update, consisting of two steps: one of length $2\alpha_t$ in the direction of u and one of length α_t^2 in the direction of $v - u$. The second step can accommodate a change in direction between u and v , which allows SQUAREM to account for curvature in the trajectory to the fixed point. The sum of the steps can also be written as $\alpha_t(2 - \alpha_t)u + \alpha_t^2 v$, i.e., as a weighted sum of u and v . The weights are positive if $0 < \alpha_t < 2$.

Figure 3: Illustration of the SQUAREM ($k = 1$) method in two dimensions, with $\theta = (a, b)^T$. Blue: a SQUAREM ($k = 1$) update with θ_t as a starting point. Grey: successive EM updates from θ_t to θ^* . Black arrows: the vectors $u = F(\theta_t) - \theta_t$ and $v = F(F(\theta_t)) - F(\theta_t)$.

Together with the three steplengths proposed for α_t in equations (2.13-2.15) and in conjunction with globalisation strategies, Expression 2.18 yields three SQUAREM methods: gSqS1, gSqS2 and gSqS3, where S1-S3 indicate which steplength is used (Varadhan & Roland, 2008). Because of the squaring of the rate matrix, SQUAREM usually converges faster than STEM or QN with $q = 1$. Of the three global SQUAREM

methods, both gSqS1 and gSqS3 perform better than gSqS2, and gSqS3 is consistently faster than gSqS1 (Varadhan & Roland, 2008). However, QN methods with $q > 1$ can be faster than SQUAREM (Zhou et al., 2011).

The SQUAREM method described thusfar is of the first order, $k = 1$. The order parameter, k , is related (but not equal) to the number of successive EM updates used during the accelerated update. Higher order SQUAREM methods ($k > 1$) can easily be derived from first order SQUAREM methods (Roland et al., 2007). Let F^i denote the iteration of F i times, with $F^0(\theta_t) = \theta_t$, $F^1(\theta_t) = F(\theta_t)$, $F^2(\theta_t) = F(F(\theta_t))$ and so on. Then the SQUAREM update (2.18) of order $k = 1$ can be rewritten as (Roland et al., 2007)

$$\begin{aligned}\theta_{t+1} &= \theta_t + 2\alpha_t u + \alpha_t^2(v - u) \\ &= (1 - \alpha_t)^2 \theta_t + 2\alpha_t(1 - \alpha_t)F^1(\theta_t) + \alpha_t^2 F^2(\theta_t) \\ &= (\gamma_{0,t})^2 F^0(\theta_t) + 2\gamma_{0,t}\gamma_{1,t}F^1(\theta_t) + (\gamma_{1,t})^2 F^2(\theta_t) \\ &= \sum_{i=0}^1 \sum_{j=0}^1 \gamma_{i,t}\gamma_{j,t}F^{i+j}(\theta_t), \quad t \geq 0,\end{aligned}\tag{2.19}$$

with $\gamma_{0,t} = (1 - \alpha_t)$ and $\gamma_{1,t} = \alpha_t$. Equation 2.19 can be generalised as follows to give rise to higher order SQUAREM methods ($k > 1$) (Roland et al., 2007):

$$\theta_{t+1} = \sum_{i=0}^k \sum_{j=0}^k \gamma_{i,t}\gamma_{j,t}F^{i+j}(\theta_t) := \text{SQUAREM}_k(\theta_t), \quad t \geq 0.\tag{2.20}$$

The coefficients $\gamma_{i,t}$ and $\gamma_{j,t}$ have to fulfill (Roland et al., 2007):

$$\sum_{i=0}^k \gamma_{i,t} = 1\tag{2.21}$$

$$\sum_{i=0}^k \beta_{i,j,t}\gamma_{i,t} = 0, \quad \text{for } j = 0, \dots, k-1.\tag{2.22}$$

See also Roland et al. (2007) for an alternative derivation and how higher order SQUAREM methods relate to other acceleration methods, a discussion of which would be beyond the scope of this work. Like QN methods, SQUAREM methods

are suitable for high-dimensional problems (Roland et al., 2007). Combined with globalisation strategies, they are globally convergent (Varadhan & Roland, 2008). Furthermore, first order SQUAREM ($k = 1$) does not have additional storage costs compared to STEM or QN ($q = 1$). This means first order SQUAREM is highly competitive. However, higher order SQUAREM methods ($k > 1$) do require $2k$ evaluations of F , which incurs additional cost, so k should be chosen as the smallest number for which reasonable acceleration can be achieved (Varadhan & Roland, 2008). Note that first-order SQUAREM ($k = 1$) can already accelerate the EM algorithm four- to ten-fold, depending on the problem studied (Varadhan & Roland, 2008).

2.2.5 Parabolic EM (PEM)

Unlike the QN and SQUAREM methods, parabolic EM is derived from purely geometric considerations. It is based on a Bezier parabola, a parametric curve controlled by three points, here θ_{t-2} , θ_{t-1} and θ_t , and a parameter, s (Berlinet & Roland, 2009). The Bezier parabola is given by $B(s) = (1-s)^2\theta_{t-2} + 2s(1-s)\theta_{t-1} + s^2\theta_t$, originally defined for $0 \leq s \leq 1$ (Berlinet & Roland, 2009). For PEM, we choose $s \geq 1$, so that values beyond θ_t can be explored (Berlinet & Roland, 2009). It passes through θ_{t-2} and through θ_t and is tangent at these points to the line that passes through θ_{t-2} and θ_{t-1} and the line that passes through θ_{t-1} and θ_t (also see Fig. 4). For ease of exposition, we assume the Bezier parabola is a subspace of the parameter space, omitting a discussion of parameter constraints. The parameter subspace defined by the parabola is explored by means of either arithmetic or geometric search (Berlinet & Roland, 2009). In either case, s is incremented from $s = 1$ until the log-likelihood, $L(B(s))$, no longer increases. At this point, the search is stopped and the PEM update is performed (Berlinet & Roland, 2009):

$$\theta_{t+1} = B(\hat{s}) := PEM(\theta_{t-2}, \theta_{t-1}, \theta_t), \quad t \geq 2, \quad (2.23)$$

where $\hat{s} = \operatorname{argmax}_{s \geq 1} L(B(s))$.

Each explored value of s , prompts a log-likelihood evaluation, so log-likelihood evaluations are ideally cheap. Note that, if θ_{t-2} , θ_{t-1} and θ_t lie on a straight line, PEM

will simply explore that line (Berlinet & Roland, 2009). This does not cause problems for the algorithm. A few EM updates are usually performed before starting PEM (Berlinet & Roland, 2009), as PEM tends to be more efficient closer to the fixed point. In contrast to QN and SQUAREM methods, the parabola allows PEM to explore a larger parameter subspace for the best possible update (compare Figure 4).

The exploration of the parameter space by arithmetic or geometric search means that PEM has one or two tuning parameters: the mesh size h and the ratio a can be used to control the incrementation of s (Berlinet & Roland, 2009), with $s = 1 + ih$ for the arithmetic and $s = 1 + a^i h$ for the geometric search, where i is a dummy variable. Berlinet and Roland (2009) show that large values for h or a can destabilize the algorithm and recommend the geometric search with $h = 0.1$ and $a = 1.5$.

Figure 4: Illustration of the PEM method (2.18) in two dimensions, with $\theta = (a, b)^T$. Blue: a PEM update based on a Bezier parabola with $\theta_{t-2}, \theta_{t-1}$ and θ_t as starting points. Grey: successive EM updates from θ_{t-2} to θ^* .

The Bezier parabola function can be re-written as $B(s) = \theta_{t-2} + 2s\tilde{u} + s^2(\tilde{v} - \tilde{u})$,

with $\tilde{u} = \theta_{t-1} - \theta_{t-2}$ and $\tilde{v} = \theta_t - \theta_{t-1}$. This is very similar to the SQUAREM update (2.18) with $s = \alpha_t$. In fact, Berlinet and Roland (2009) show that first order SQUAREM is a special case of PEM, using a single point on the Bezier parabola, $B(\alpha_t)$, as the next update. Unlike SQUAREM, PEM does not use a predefined point, but instead explores the neighbourhood of $s \geq 1$ until a local maximum of $L(B(s))$ has been found. For this reason, PEM is more stable than SQUAREM and can be faster, although this has only been measured in CPU time, not in terms of number of F and L evaluations (Berlinet & Roland, 2009).

Like QN and SQUAREM methods, PEM has low storage requirements and is therefore suitable for high-dimensional problems. PEM may perform a much larger number of log-likelihood evaluations during the exploration step. However, since the evaluation of L is often costly less than an evaluation of F , this may not be a disadvantage. Varadhan and Roland (2008) show that PEM can accelerate the EM algorithm four- to ten-fold, depending on the problem studied.

2.2.6 Expected results of the numerical comparison

Based on the detailed descriptions and considerations above, we expect SQUAREM ($k = 1$) to perform better than QN ($q = 1$) because of the squared rate matrix. PEM, in turn, should perform better than SQUAREM ($k = 1$) because instead of using a fixed update it explores the parameter space for the best possible update during each iteration. However, it is important to keep in mind that the runtime of all three acceleration techniques is determined by a trade-off between a reduced number of iterations to the fixed point and an increased number of F and L evaluations per iteration. In particular, higher-order QN ($q > 1$) and SQUAREM ($k > 1$) methods may suffer from a large number of F evaluations, while PEM may suffer from a large number of log-likelihood evaluations.

3 Simulation study design

3.1 Simulation conditions

We varied the IRT model used between a 1PL model with estimated ability variance σ^2 , a 1PL model with σ^2 fixed to 1, and a 2PL model (σ^2 fixed to 1 for identification purposes). Further, we varied the number of items ($M = 9$ vs. $M = 100$), focusing on more extreme ends of the spectrum to get a good idea of the behavior of the accelerators. Thus, we compared the accelerators, i.e., QN with $q \in \{1, 2, 3, 4\}$, SQUAREM with $k \in \{1, 2, 3, 4\}$, and PEM, in six different scenarios (see Table 1).

Condition	Model	Number of Items
1	1PL	9
2	1PL (fixed var.)	9
3	2PL	9
4	1PL	100
5	1PL (fixed var.)	100
6	2PL	100

Table 1: The six simulation conditions. var. = variance

3.2 Generation of item responses

To generate the item parameters of the 2 PL model in each trial, we used a bivariate truncated normal distribution to sample the discrimination (a_j) and difficulty (d_j) parameters, and then computed the intercept parameters d_j as $\delta_j = -a_j d_j$ and slopes as $\alpha_j = a_j$ for each item $j = 1, \dots, M$. This approach was inspired by but slightly adapted from Paek and Cai, 2014. As the correlation between difficulties and discriminations we chose 0.35 (but note this will only be approximately the underlying correlation as we implemented the truncation via rejection sampling, see below), based on the application of a 2PL model to math achievement test data Stone, 1992; Nader et al., 2011. For the means and variances for the bivariate normal distribution, we used means of 0 and 1 for difficulties and discriminations, respectively, and variances of 1, again similar to Paek and Cai, 2014. We also adopted their lower bounds of -2 and 0.5 for difficulties and discriminations, respectively, and upper

bounds of 2 for both. We implemented the truncation by means of rejection sampling, that is, for each item $j = 1, \dots, M$, a pair of discrimination and difficulty was drawn and assigned as item parameters if both values were within the respective bounds, otherwise rejected and redrawn. For the 1PL model, item discrimination was fixed at 1 and difficulty drawn from a univariate truncated normal distribution with the same parameters and bounds as for the 2PL model (note that the respective intercept parameters are then given by $\delta_j = -b_j$). Person abilities were sampled randomly from a standard normal distribution. Using the slope-intercept parameterization of Equation 2.1 to calculate probabilities, response data $\mathbf{x}$ for the N participants was generated from a binomial distribution in each trial. The R code used for the simulation is available as supplementary material.

3.3 Simulation procedure

In each condition and trial, we generated response data for $N = 10,000$ participants as described above. We obtained the starting values for the EM algorithm from the `mirt::mirt` function in each trial, to mimic a realistic application of the EM algorithm. We then used a combination of the `mirt` and `turboEM` packages Bobb and Varadhan, 2018; Chalmers, 2012, as described below, to fit the respective model of the condition with the EM algorithm in conjunction with the different accelerators. For the conditions in which the latent factor variance was estimated for the 1PL model, we used a parameter value constraint of $\sigma^2 > 0$ in alignment with the `mirt` package. The R code used for the simulation is available as supplementary material.

3.4 Implementation of models and accelerators

The EM accelerators, i.e., QN with $q \in \{1, 2, 3, 4\}$, SQUAREM with $k \in \{1, 2, 3, 4\}$, and PEM, were all implemented in our simulation study via the R package `turboEM` Bobb and Varadhan, 2018 which implements these methods in conjunction with globalisation strategies in order to ensure global convergence if necessary (for an illustration, see Algorithms A1, A2, and A3 in Appendix A). To run the EM algorithm, we used a combination of the `mirt`- and `turboEM`-packages Bobb and Varadhan, 2018; Chalmers, 2012. The `mirt::mirt` function implements the EM-algorithm

for IRT models Chalmers, 2012. A single (unaccelerated, i.e., standard) EM-step version of the `mirt::mirt` function for one EM update, i.e., the fixed point function F , as well as the objective function (i.e., the observed log-likelihood for the current parameter estimates), were extracted from the `mirt`-package. These functions are then used as inputs for the `turboEM::turboem` function, which is in turn used to run the EM algorithm with the various acceleration techniques as well as the standard EM algorithm (used as a benchmark in our numeric comparison). Note that this results in standard EM being implemented differently to how it is implemented in `mirt`, which is likely to be less efficient than the complete `mirt` implementation but yields a perfectly comparable implementation of standard EM and all versions of accelerated EM. All acceleration methods were run with default tuning parameters Bobb and Varadhan, 2018. The convergence criterion used is based on the parameter only, i.e., $||\hat{\theta}_{t+1} - \hat{\theta}_t|| < 10^{-7}$, in order to avoid additional evaluations of the log-likelihood. This is fairly strict, allowing the acceleration techniques the opportunity to show their full power. Maximum number of iterations were 1500 and maximum running time 60 seconds. Illustrative pseudo-code for (some of) the algorithms is provided in Appendix A (Algorithms A1, A2, and A3). The R code used for the simulation is available as supplementary material.

3.5 Examination of trajectories: Properties of F

Trajectories of F are sequences from initial starting values θ_0 to the fixed point θ^* , as given by the standard EM algorithm. In order to be able to describe these for our different simulation conditions to better understand the latter, we used the following properties of F . The relative steplength is defined as the distance between successive EM updates, relative to the total distance between the initial parameter estimates $\hat{\theta}_0$ and the final maximum likelihood estimates θ^* :

$$r(t) = \frac{||F(\hat{\theta}_t) - \hat{\theta}_t||}{\sum_{i=0}^{N_0^*} ||F(\hat{\theta}_i) - \hat{\theta}_i||}, \quad t \geq 0, \quad (3.24)$$

where N_0^* is the total number of iterations required to reach θ^* from the given $\hat{\theta}_0$. As $\hat{\theta}_t$ approaches θ^* , $||F(\hat{\theta}_t) - \hat{\theta}_t||$ approaches zero and $\sum_{t \geq 0} r(t)$ approaches 1. The relative steplength $r(t)$ therefore can be used to show how fast θ_t approaches the fixed

point. The normalisation to the total length of the trajectory, $\sum_{i=0}^{N_i} \|F(\hat{\theta}_i) - \hat{\theta}_i\|$, ensures that the quantity is comparable for different starting points and models. We further define the curvature of F as the angle between successive iterations,

$$\cos \psi(t) = \frac{u_t^T v_t}{\|u_t\| \|v_t\|}, \quad t \geq 0, \quad (3.25)$$

where $u_t := F(\hat{\theta}_t) - \hat{\theta}_t$ and $v_t := F(F(\hat{\theta}_t)) - F(\hat{\theta}_t) = u_{t+1}$. The angle ψ is zero if the vectors u_t and v_t point in the same direction, i.e., if the fixed point is approached in a straight line. If the angle is different from zero, the fixed point is approached on a curve.

4 Results

4.1 Examination of trajectories

As item parameters can be optimized independently for each item during the M-step, we only visualize the results regarding the trajectories for item 1 (Figures 5 and 6). Figures 5 and 6 show the cumulative relative steplength (left columns) and the angle between successive EM updates (right columns) in all six conditions for all 100 trials. Note that for each trial, the underlying true parameter value(s) for item 1 were sampled randomly and therefore different in every trial. Across the left-hand columns in Figures 5 and 6, we can see that most trajectories approach the proximity of the fixed point in relatively few steps, with some, especially for the 1PL models, starting off very close to it, attesting to the choice of starting values provided by `mirt`. While others do take a few more steps to reach the proximity of the fixed point, all trajectories across all conditions show that before half the steps of the complete trajectory have been taken, the EM algorithm is already very close to the fixed point and only makes very small adjustment in the parameter estimates with each further EM cycle. The right-hand columns in Figures 5 and 6 show that instabilities occur in some trajectories: the angle between successive iterations appears to switch from zero to π and back. An angle of π indicates that after a step in one direction, a step is taken in the opposite direction. It is unclear to

the authors why these instabilities occur; however, they do not impede the ability of the trajectories to find the fixed point. While these instabilities occur across all trajectory steps for conditions with only nine items (Figure 5), they only occur for earlier steps in conditions with 100 items in conjunction with the 1PL model(s) and only occur during later steps for a few trajectories in conditions with 100 items in conjunction with the 2PL model. If the angle is different from zero, the fixed point is approached on a curve. Such a pattern occurs most notably for the 2 PL model, both in conjunction with nine and with 100 items (see C' and F' in Figures 5 and 6). In these cases, acceleration methods such as SQUAREM and PEM may be more suitable because they explicitly allow for changes in direction.

Figure 5: Properties of F for conditions 1–3 (i.e., with nine items). For illustrative purposes, only the parameter spaces of item 1 are shown here. Note that the colours are transparent so that darker colours indicate that more trajectories pass through the same spot. For each trajectory, the cumulative relative steplength (A for condition 1 (1PL), B for condition 2 (1PL with fixed variance), C for condition 3 (2PL)) and the angle between successive EM updates (A' for condition 1, B' for condition 2, C' for condition 3) are plotted as a function of the iteration index, t .

Figure 6: Properties of F for conditions 4–6 (i.e., with 100 items). For illustrative purposes, only the parameter spaces of item 1 are shown here. Note that the colours are transparent so that darker colours indicate that more trajectories pass through the same spot. For each trajectory, the cumulative relative steplength (D for condition 4 (1PL), E for condition 5 (1PL with fixed variance), F for condition 6 (2PL)) and the angle between successive EM updates (D' for condition 4, E' for condition 5, F' for condition 6) are plotted as a function of the iteration index, t .

4.2 Comparison of accelerators

Tables 2 and 3 show the average number of iterations from the starting value to the fixed point, the average number of F evaluations, the average number of log-likelihood evaluations, the fraction of converged trials, and the average CPU time in ms for all nine acceleration methods as well as the standard EM algorithm as a benchmark in all six conditions. The first three conditions, i.e., all models with nine

items, are shown in Table 2, and the other three conditions, i.e., all models with 100 items, are shown in Table 3.

Notably, the first-order SQUAREM accelerator failed to converge in any of the trials in condition 1 and 4, that is, for the 1PL model where we estimate the variance. This may be due to the parameter constraints used, as we did not observe such complete failure of convergence in prior trial runs without such restrictions. However, without parameter constraints, higher-order SQUAREM methods showed more flawed convergence rates and comparability to the estimation with the `mirt` would be impaired. Note that when fixing the variance to 1 in the 1PL model (conditions 2 and 4), this no longer was an issue.

Performance of the accelerators varied considerably between conditions, with some drastic speed-ups and even some increases in computing time. For conditions with nine items (conditions 1–3, see 2), we only observe speed-ups in CPU time across all accelerators (except PEM which even shows an increase in computing time) for condition 2, that is, for the 1PL model with the fixed latent variance. In this condition, higher order QN methods (especially $q = 4$) as well as first-order SQUAREM and SQUAREM with $k = 4$ show the best results in terms of CPU time decreases, with QN ($q = 4$) taking first place and halving the CPU time of standard EM. For condition 3 (2PL with nine items), we see consistently bad performance of all accelerators in terms of CPU time, with all of the increasing computing time. For condition 1 (1PL with nine items), results are mixed with regard to CPU time. Only SQUAREM ($k = 4$), SQUAREM ($k = 3$), and QN ($q = 4$) show a (modest) decrease in CPU time in the pre-decimal digits. The other accelerators either just about match the CPU time of standard EM or even show increases. In terms of average number of iterations from the starting value to the fixed items, all accelerators show notable reductions. The best performances with regard to the average number of iterations for nine items (i.e., across conditions 1–3) are exhibited by the higher-order SQUAREM methods in condition 2, reducing the number of iterations from 37 to only 3. PEM and higher-order QN also show competitive performances here. In terms of the number of F evaluations, the QN methods as well as first-order SQUAREM (but only for conditions 2–3) show the best results across conditions 1–3. Across conditions 1–3, PEM is characterized by a very high number of log-

likelihood evaluations, which may explain its performance in terms of CPU time.

For the conditions with 100 items (conditions 4–6, see 3), only condition 4, that is, the 1PL model, showed similar unfortunate patterns of increased CPU time compared to standard EM as we have seen in the first three conditions. Conditions 5 and 6 (i.e., the 1PL model with fixed latent variance and the 2PL model), however, showed very good results in terms of speed-up of CPU times for all accelerators. Out of all of them, first-order QN and PEM showed the least computing time reductions, however they still exhibited notable speed-ups. For the 1PL model with fixed latent variance (condition 5), higher-order SQUAREM and especially higher-order QN showed the best results in terms of computing time. For the 2PL model (condition 6), higher-order QN methods still performed very well, but were beat by lower-order SQUAREM methods, in particular, first-order SQUAREM. In terms of the average number of iterations and the number of F evaluations, the pattern of results observed for 100 items is similar to the one outlined above for the conditions with nine items. PEM (and for conditions 5 and 6, also first-order QN) still suffers from a very high number of log-likelihood evaluations.

5 Discussion

In the present study, we numerically compared three recently proposed variants of model-independent EM accelerators, “squared” iterative methods (SQUAREM; $k \in \{1, 2, 3, 4\}$; Varadhan and Roland, 2004; Roland et al., 2007; Varadhan and Roland, 2008), advanced quasi-Newton (QN; $q \in \{1, 2, 3, 4\}$; Zhou et al., 2011) methods, and parabolic EM (PEM; Berlinet and Roland, 2009), evaluated in large-sample psychometric settings with either few or a large number of items (nine vs. one hundred) for the 1PL and 2PL model. The underlying idea of all three acceleration methods studied is to view the EM algorithm as a fixed-point mapping, which, applied iteratively, generates a sequence of parameter estimates from the starting value to – upon convergence – the fixed point. Acceleration is achieved by all three

methods discussed here by creating another sequence of parameter estimates from the same starting point converging to the same fixed point, but with a faster rate of convergence. While all accelerators decreased the number of iterations as well as (for the most part) the number of EM cycles required to reach the fixed point many-fold, not all accelerators were able to achieve speed-ups in CPU time (as compared to standard EM) across all conditions. The most promising decreases in CPU time were observed for settings with 100 items and for models which did not estimate the latent ability variance (but instead had it fixed at 1). Higher-order QN and SQUAREM (depending on the model either higher- or lower-order) methods performed the best in terms of reduction of CPU time. PEM did neither perform as well as expected nor as well as the other methods.

5.1 Properties of trajectories for standard EM

Studying the trajectories of the standard EM algorithm in our simulation study, we observed that, using starting values provided by `mirt`, the majority of trajectories found themselves in proximity of the fixed point right from the starting point. Even those which started further away reached proximity of the fixed point in only relatively few iterations. Thus, quite the number of EM steps actually are rather small. This yields the opportunity for the accelerators to replace a number of small EM steps by one larger EM step. How successful this has the potential to be for each acceleration method also depends on the shape of the trajectory. To this end, we also examined the angles between successive EM updates, in the course of which we observed some instabilities in the trajectories with angles indicating that the direction of the trajectory was changed into the opposite from one step to the next. This pattern occurred for several trajectories, more so for settings with nine compared to 100 items, and more so for the 2PL compared to the 1PL model. While such instabilities occurred at any point throughout the trajectories for only nine items, they only occurred early on for the 1PL model when 100 items were included and occurred predominantly early and only occasionally later on for the 2PL model in conjunction with 100 items. Note that with the EM algorithm, the log-likelihood is increased with each EM step. Thus, situations in which backwards steps are necessary within the trajectories likely only arise when the algorithm overshoots the

optimum but still lands at an estimate which increases the log-likelihood, so is closer to the fixed point, just approaching it from the other side compared to the previous estimate. However, since all trajectories reach the fixed point eventually, these instabilities are unlikely to affect the acceleration results. [NOTE:] Das war meine Ueberlegung als Erklaerung fuer die Overshoots. Ich bin mir nicht sicher, ob das ganz richtig ist, was meinst du? Apart from the backwards steps, the examination of the angles of successive EM steps revealed that while the fixed points are mostly approached in straight lines for the 1PL model (again, except for when backwards steps have to be taken), the angles indicate some curvilinearity in the trajectories for the 2PL model. This would suggest that in scenarios with the 2PL models, accelerators which explicitly take the non-linear shape of the trajectories into account, that is, SQUAREM and PEM, should show superior performance.

5.2 Comparison of EM accelerators

Contrary to what the observations made on the trajectories for standard EM led us to believe, we found that PEM did not provide competitive if any speed-ups of computing time in any of our simulation scenarios. In settings with nine items, we either observed roughly unaltered (1PL) or even decelerated (2PL) computing times as compared to standard EM. In settings with one hundred items, we found that PEM did achieve acceleration in computing times (except for when the latent variance was also estimated in the 1PL along with the item parameters – then we again observed deceleration), but these could not compete with the accelerations achieved by some of the other methods. This is unexpected, because for one, the Bezier parabola allows PEM to search the parameter space for the best possible update in each iteration. Additionally, especially the non-linear trajectories in conditions with the 2PL model should provide a setting in which PEM can shine, in particular compared to methods such as QN, which cannot account for the non-linearity in the trajectories. However, the number of evaluations of the fixed point function and of the log-likelihood is markedly increased, which incurs additional cost and increases the CPU time. It should be noted that these results are not consistent with previous studies, which applied PEM to the estimation of different models (a linear mixed model, a poisson mixture model, and a multivariate Student distribution), where the

CPU time with PEM was reliably three to four times smaller than with standard EM (Berlinet & Roland, 2009). Even when PEM did accelerate the EM algorithm in our simulations, the speed-up did not even manage to halve the computing time. This may certainly be explained by the large number of log-likelihood evaluations which PEM required. This number was consistently the greatest for PEM out of all accelerators and up to almost 10 times larger than that of the other methods. One possible explanation for the results is that we followed the generic recommendations of (Berlinet & Roland, 2009) for the choice of tuning parameters in our PEM implementation. While these performed well for models studied by Berlinet and Roland, 2009, parameter values more optimally suited to IRT models may exist. This is still a good recommendation to derive from our simulations: Applied researchers might wish to tune parameters for PEM specifically to their application or choose another acceleration method rather than rely on the recommended parameters by Berlinet and Roland, 2009.

For QN and SQUAREM, we found performance depended (even more strongly) on the simulation setting. The performance in the settings with the 1PL for which the latent ability variance was estimated was generally rather poor across all acceleration methods. For nine items, only higher-order SQUAREM and QN methods provided acceleration in terms of CPU time (but note that first-order QN even provided deceleration and first-order SQUAREM failed to converge in any of the trials in this condition, see also Section 5.3). For 100 items, all methods showed decelerations. We attribute this impaired performance to the additional estimation of the latent ability variance, as we observed considerably better performances for conditions with the 1PL model for which the latent variance was fixed to 1. For nine items, we here saw the best results for higher-order QN ($q = 4$) and first-order SQUAREM with CPU times cut in half compared to standard EM. For 100 items, higher-order QN ($q = 4$) still came out on top for the 1PL model with the fixed latent variance, but this time with CPU time reduced to a fifth of what it was for standard EM. These magnitudes of the reduction in run time are in line with results for other models (Varadhan & Roland, 2008; Zhou et al., 2011). Generally, the higher-order QN and SQUAREM methods performed better than the lower-order

variants for the 1PL model with the fixed latent variance. For the 2PL model, all methods decelerated the EM algorithm when only nine items were used. The performance of the accelerators was considerably better for the 2PL model in conjunction with 100 items. Here, we observed the best performance for first-order SQUAREM, which reduced the CPU time to nearly a third of the CPU time for standard EM. While we still observed the pattern that higher-order QN methods performed better than lower-order QN methods in this condition, we observed the reversed pattern for SQUAREM.

Better performances of SQUAREM compared to QN, as we observed for the 2PL model with 100 items, are in line with our expectations, as SQUAREM is characterized by a “squared” rate matrix compared to QN. However, this pattern did not emerge consistently across conditions with (higher-order) QN methods performing better than SQUAREM methods for the 1PL model (with a fixed latent ability variance). We observed a general tendency across conditions of SQUAREM requiring more fixed-point function but less log-likelihood evaluations and QN showing the reversed pattern. This may well provide an explanation for the different rank orders of methods for the 1PL and 2PL model: When the log-likelihood is more costly to evaluate, i.e., in our case for the 2PL model, SQUAREM benefits more noticeably from the fewer evaluations of the log-likelihood, ranking in front of QN. This is much less of an advantage when log-likelihood evaluation is cheaper, as it is for the 1PL model.

5.3 Limitations and avenues for future research

One important limitation of the present study constitutes the exclusive usage of default tuning parameters as provided by (Bobb & Varadhan, 2018). Our decision was motivated by our desire to closely mimic real applications of the EM algorithm and its accelerators in psychometric settings, where applied researchers likely would also rely on generically recommended tuning parameters. However, this may have also kept the accelerators, most notably PEM, from unfolding their full acceleration potential. It would therefore be interesting for future research to investigate which tuning parameter constellations are optimal for IRT models. Based on our

results which found differences in rank orders of the accelerators between the 1PL and 2PL model, we would also encourage future research to derive specific PEM tuning parameter recommendations for different IRT models. In the same vein, another interesting avenue for future research presents itself in the shape of exploring the possibility of self-tuning accelerators. This may be especially advantageous in the light of the observed differing performance for different IRT models and would provide general as well as user-friendly solutions to applied researchers.

Another limitation is constituted by the failed convergence of first-order SQUAREM across all trials for the 1PL model when the latent ability variance is estimated. In the respective conditions, we used a constraint on the parameter space for the latent ability variance ($\sigma^2 \geq 0$), as we (1) observed some trials with failed convergence for higher-order SQUAREM without the constraint, and (2) wanted to mimic the implementation in `mirt`. While we did therewith obtain perfect convergence rates for higher-order SQUAREM, we were then left with no convergence for first-order SQUAREM. As there is no reason evident to us why first-order SQUAREM should not be working in this condition in conjunction with a constraint on the parameter space, this is likely an error in our implementation.

Further, the implementation of the accelerators in this simulation study via `turboEM` and an extracted EM cycle from `mirt` was not optimized for speed. In fact, it is not very efficient because single EM-steps have to be passed from `mirt` to `turboEM` in every iteration. Thus, our results are not necessarily comparable to e.g. the runtime of standard EM as implemented and runtime optimized in `mirt`. This was mostly due to ensure comparability between all acceleration methods, as they are not all available in `mirt`. This limitation leads us down another avenue for future research: Exploring how to increase efficiency of the implementation of the different accelerators for IRT models. A problem to this end is the increased number of costly evaluations of the log-likelihood. One major reason why the log-likelihood is so costly to evaluate is the integral over the latent ability which requires some type of numerical approximation, e.g., via Gauss-Hermite (GH) quadrature. A worthwhile endeavor for future research might be to use GH quadrature with increasing numbers

of quadrature points along the trajectory, starting out with cruder approximations in the beginning. This may help reduce the time cost of log-likelihood evaluations, at least for the earlier trajectory steps. To this end, further research is needed to find an optimal balance between the imprecision of approximation and speedup of the EM algorithm.

In terms of the examination of the trajectories of the EM algorithm, our results are limited to the trajectories for standard EM. It might be interesting for future research to examine the trajectories for the different accelerators in order to better understand their behavior. For example, it would be very interesting to see if accelerated trajectories are straighter and make larger steps per iteration than standard EM trajectories. Finally, our simulation study used 100 trials per simulation condition. This constituted a good trade-off between precision of average results and computational feasibility for us, however, as holds for any simulation study, results would be more reliable if the number of trials were larger. Lastly, as Kriegel et al., 2017 stresses, any simulation studies comparing runtime of algorithms should be taken with (maybe a little more than) a grain of salt. We have already discussed that our implementation is likely not the most efficient, and naturally any CPU times also depend on the CPU used. Thus, we would like to stress that attention should be paid more to the ranking of the accelerators rather the absolute CPU times. However, even these might be different if all accelerators were implemented differently as well as run with the optimal choices for tuning parameters (Kriegel et al., 2017).

5.4 Conclusion

To summarize, in what is to our knowledge the first comparison of recently proposed acceleration methods for the EM algorithm in an IRT context, we have found good performance of higher-order QN and also higher-order SQUAREM for the 1PL model, and of first-order SQUAREM and also higher-order QN for the 2PL model. Surprisingly, PEM performed neither as well as expected nor was able to compete with the other acceleration methods. Generally, acceleration was more successful when the latent ability variance was fixed rather than estimated. Speed-ups in com-

puting time were more substantial in settings with one hundred than with nine items.

A Appendix A

Algorithm A1 shows pseudocode for the QN method, based on a simplified representation of the QN implementation within the `turboEM::accelerate` and `turboEM::bodyQuasiNewton` methods Bobb and Varadhan, 2018. The algorithm starts with initial parameters θ_0 . First, the q columns of matrix U are created from a sequence of EM updates (F ; Alg. A1, lines 1-11). Then matrix V is created based on U and a further EM update (Alg. A1, lines 12-16). Finally, QN updates are performed until convergence (Alg. A1, lines 17-27). Algorithm A1 only shows a skeleton of the QN method. The implementation of the QN update in `turboEM::bodyQuasiNewton` is a little more complex. For instance, parameter constraints can be supplied by the user to keep the parameter updates in `turboEM::bodyQuasiNewton` within certain bounds Bobb and Varadhan, 2018. Perhaps more importantly, there is an additional check that the negative log-likelihood is decreased (i.e. the log-likelihood is increased) by the QN update θ_{QN} compared to the last EM update θ_2 . If $L(\theta_{QN}) < L(\theta_2)$, the value of θ_{QN} is set to θ_2 instead Bobb and Varadhan, 2018. This modification is used to ensure global convergence.

Algorithm 1: Pseudocode for QN

```
/* Prepare parameters for QN method based on
   input of initial parameter estimates  $\theta_0$ .
   */

1 Create empty matrix  $U$  (with no columns);
2 for  $i \leftarrow 1$  to  $q$  do
3    $\theta_1 = F(\theta_0)$ ;
4    $u = \theta_1 - \theta_0$ ;
5   Add new column  $u$  to the right hand side of matrix
      $U$ ;
6    $\theta_0 = \theta_1$ ;
7 end
8  $\theta_1 = F(\theta_0)$ ;
9  $u = \theta_1 - \theta_0$ ;
10 Remove the leftmost column of matrix  $U$ ;
11 Add new column  $u$  to the right hand side of matrix  $U$ ;
12  $\theta_2 = F(\theta_1)$ ;
13  $v = \theta_2 - \theta_1$ ;
14 Create matrix  $V = U$ ;
15 Remove the leftmost column of matrix  $V$ ;
16 Add new column  $v$  to the right hand side of matrix  $V$ ;
   /* Perform QN updates until convergence,
      using inputs  $\theta_0, \theta_1, \theta_2, U$  and  $V$ .      */
17 while !Converged do
18   Compute QN update:
      $\theta_{QN} = \theta_2 - V(U^T V - U^T U)^{-1} U^T (\theta_2 - \theta_1)$  ;
   /* Update  $\theta_0, \theta_1, \theta_2, U$  and  $V$  with new
      values for next iteration.      */
19    $\theta_0 = \theta_{QN}$ ;
20    $\theta_1 = F(\theta_0)$ ;
21    $\theta_2 = F(\theta_1)$ ;
22    $u = \theta_1 - \theta_0$ ;
23    $v = \theta_2 - \theta_1$ ;
24   Remove the leftmost columns of matrices  $U$  and  $V$ ;
25   Add new columns  $u$  and  $v$  to the right hand sides
     of matrices  $U$  and  $V$ ;
26   Check for convergence; if converged set value of
     Converged to true;
27 end
```

Algorithm A2 shows pseudocode for the first order SQUAREM method, based on Varadhan and Roland, 2008. It is consistent with the SQUAREM implementation within the `turboEM::accelerate` and `turboEM::bodySquarem1` methods Bobb and Varadhan, 2018. The algorithm operates on an input of initial parameter estimates θ_0 . Based on these parameters, the algorithm is run until convergence. First, two successive EM updates are performed (Alg. A2, lines 2-3) to calculate u and v (Alg. A2, lines 4-5). Based on u and v , the steplength α is computed (Alg. A2, line 6). For global convergence, the steplength is modified if necessary (Alg. A2, line 7). Then the SQUAREM update is computed (Alg. A2, line 8). Finally, an EM update of the SQUAREM update provides a new parameter estimate for the next iteration Varadhan and Roland, 2008; Bobb and Varadhan, 2018. The algorithm can be run in three versions, “1”, “2” or “3”, depending on which steplength is to be used. By default, the `turboEM::bodySquarem1` method uses steplength S3, as recommended by Varadhan and Roland, 2008. The implementation of Algorithm A2 in the `turboEM::bodySquarem1` method includes additional checks on the SQUAREM update Bobb and Varadhan, 2018. For instance, if the user has specified constraints on the parameter space, the update must fall within the constraints. If the SQUAREM update does not increase the log-likelihood, the last EM update is used instead. Finally, upper and lower bounds are set on the steplength and modified dynamically to increase stability and convergence Bobb and Varadhan, 2018. Pseudo-code for higher-order SQUAREM is not shown here.

Algorithm 2: Pseudocode for SQUAREM ($k = 1$)

```
/* Algorithm is based on initial parameter estimates  $\theta_0$ .          */
1 while !Converged do
2    $\theta_1 = F(\theta_0);$ 
3    $\theta_2 = F(\theta_1);$ 
4    $u = \theta_1 - \theta_0;$ 
5    $v = \theta_2 - \theta_1;$ 
6   Compute steplength  $\alpha$ : (2.13),(2.14) or (2.15);
7   For global convergence, modify  $\alpha$  (section ??);
8   Compute SQUAREM update:  $\theta_{SQ} = \theta_0 + 2\alpha u + \alpha^2(v - u)$  (2.18);
   /* Update  $\theta_0$  with new value for next iteration.          */
9    $\theta_0 = F(\theta_{SQ});$ 
10  Check for convergence; if converged set value of Converged to true;
11 end
```

Algorithm A3 shows pseudocode for the PEM acceleration method, consistent with the implementation of PEM in the `turboEM::bodyParaEM` method Berlinet and Roland, 2009; Bobb and Varadhan, 2018. First, an initial PEM update is computed for an initial value of the search parameter s (Alg. A3, lines 2-4). If the log-likelihood is not increased by this PEM update, then the PEM update is disregarded and regular EM updates are performed instead (Alg. A3, lines 5-9). If the log-likelihood does increase, further PEM updates are performed for increasing values of s , until the log-likelihood no longer increases (Alg. A3, lines 11-18). The last PEM update for which the log-likelihood increased is retained and two regular EM updates are performed on this PEM update to create the next member of the parameter estimation sequence (Alg. A3, lines 19-22). (According to Berlinet and Roland, 2009, these two EM updates performed on the final PEM update stabilize the algorithm.) Finally, the procedure is repeated until convergence.

Algorithm 3: Pseudocode for PEM

```
/* Algorithm requires three points as input:  $\theta_0$ ,  
    $\theta_1$  and  $\theta_2$  as well as the log-likelihood value  
    $L_2$  at the latest point  $\theta_2$ . */  
1 while !Converged do  
    /* Initialise  $s$  for geometric or arithmetic  
       search. */  
2    Set  $i = 1$  and  $s = 1 + ih$  for arithmetic search or  $i = 0$   
       and  $s = 1 + a^i h$  for geometric search;  
    /* Compute initial  $\theta_{PEM}$  update and its  
       log-likelihood. */  
3     $\theta_{PEM} = (1 - s)^2 \theta_0 + 2s(1 - s)\theta_1 + s^2 \theta_2$  (??);  
4     $L_{PEM} = L(\theta_{PEM})$ ;  
5    if  $L_{PEM} \leq L_2$  then  
        /* Since the  $\theta_{PEM}$  update does not increase  
            $L$ , update  $\theta_0$ ,  $\theta_1$ ,  $\theta_2$  and  $L_2$  with regular  
           EM for the next iteration. */  
6         $\theta_0 = \theta_2$ ;  
7         $\theta_1 = F(\theta_0)$ ;  
8         $\theta_2 = F(\theta_1)$ ;  
9         $L_2 = L(\theta_2)$ ;  
10    else  
        /* While  $L_{PEM} > L_2$ , perform  $\theta_{PEM}$  updates  
           for increasing  $s$ . */  
11        while  $L_{PEM} > L_2$  do  
12             $\theta_{OLD} = \theta_{PEM}$ ;  
13             $L_2 = L_{PEM}$ ;  
14             $i = i + 1$ ;  
15             $s = 1 + ih$  (arithmetic search) or  $s = 1 + a^i h$   
                (geometric search);  
16             $\theta_{PEM} = (1 - s)^2 \theta_0 + 2s(1 - s)\theta_1 + s^2 \theta_2$  (??);  
17             $L_{PEM} = L(\theta_{PEM})$ ;  
18        end  
        /* Update  $\theta_0$ ,  $\theta_1$ ,  $\theta_2$  and  $L_2$  with values for  
           next iteration. */  
19         $\theta_0 = \theta_1$ ;  
20         $\theta_1 = \theta_2$ ;  
21         $\theta_2 = F(F(\theta_{OLD}))$ ;  
22         $L_2 = L(\theta_2)$ ;  
23    end  
24    Check for convergence; if converged set value of  
        Converged to true;  
25 end
```

References

- Baker, F. B., & Kim, S.-H. (2004). *Item response theory: Parameter estimation techniques*. New York, Marcel Dekker, Inc.
- Bartolucci, F., Forcina, A., & Stanghellini, E. (1998). A comparison of recent em accelerators within item response theory, In *Compstat*, Springer.
- Berezner, A., & Adams, R. J. (2017). Implementation of large-scale education assessments. In P. Lietz, J. C. Cresswell, K. F. Rust, & R. J. Adams (Eds.). Wiley Online Library.
- Berlinet, A., & Roland, C. (2009). Parabolic acceleration of the em algorithm. *Stat Comput.*, 19(1), 35–47.
- Bobb, J. F., & Varadhan, R. (2018). *Turboem: A suite of convergence acceleration schemes for em, mm and other fixed-point algorithms* [R package version 2018.1]. R package version 2018.1. <https://CRAN.R-project.org/package=turboEM>
- Bock, R. D., & Aitkin, M. (1981). Marginal maximum likelihood estimation of item parameters: Application of an em algorithm. *Psychometrika*, 46(4), 443–459.
- Chalmers, R. P. (2012). mirt: A multidimensional item response theory package for the R environment. *Journal of Statistical Software*, 48(6), 1–29.
- Chermahini, S. A., Hickendorff, M., & Hommel, B. (2012). Development and validity of a dutch version of the remote associates task: An item-response theory approach. *Thinking Skills and Creativity*, 7(3), 177–186.
- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1), 1–38.
- Embretson, S. E. (2000). *Item response theory for psychologists* (1st ed.). Berlin, Heidelberg, Psychology Press.
- Hastie, T., Tibshirani, R., & Friedman, J. (2001). The elements of statistical learning. New York, Springer.
- Jamshidian, M., & Jennrich, R. I. (1993). Conjugate gradient acceleration of the em algorithm. *Journal of the American Statistical Association*, 88(421), 221–228.

- Jamshidian, M., & Jennrich, R. I. (1997). Acceleration of the em algorithm by using quasi-newton methods. *Journal of the Royal Statistical Society. Series B (Methodological)*, 59(3), 569–587.
- Kriegel, H.-P., Schubert, E., & Zimek, A. (2017). The (black) art of runtime evaluation: Are we comparing algorithms or implementations? *Knowledge and Information Systems*, 52(2), 341–378.
- Lange, K. (1995). A quasi-newton acceleration of the em algorithm. *Statistica Sinica*, 5(1), 1–18.
- Lee, C. S., Huggins, A. C., & Therriault, D. J. (2014). A measure of creativity or intelligence? examining internal and external structure validity evidence of the remote associates test. *Psychology of Aesthetics, Creativity, and the Arts*, 8(4), 446–460.
- Liu, C., & Rubin, D. B. (1994). The ecme algorithm: A simple extension of em and ecm with faster monotone convergence. *Biometrika*, 81(4), 633–648.
- Liu, C., Rubin, D. B., & Wu, Y. N. (1998). Parameter expansion to accelerate EM: The PX-EM algorithm. *Biometrika*, 85(4), 755–770.
- Louis, T. A. (1982). Finding the observed information matrix when using the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 44(2), 226–233.
- McLachlan, G. J., & Krishnan, T. (2008). *The em algorithm and extensions* (2nd ed.). New York, Wiley-Interscience.
- Meng, X.-L., & Rubin, D. B. (1993). Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika*, 80(2), 267–278.
- Meng, X.-L., & Van Dyk, D. (1997). The em algorithm—an old folk-song sung to a fast new tune. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(3), 511–567.
- Mislevy, R. J., & Bock, R. D. (1982). Implementation of the em algorithm in the estimation of item parameters: The bilog computer program.
- Nader, I. W., Tran, U. S., & Formann, A. K. (2011). Sensitivity to initial values in full non-parametric maximum-likelihood estimation of the two-parameter logistic model. *British Journal of Mathematical and Statistical Psychology*, 64(2), 320–336.

- Paek, I., & Cai, L. (2014). A comparison of item parameter standard error estimation procedures for unidimensional and multidimensional item response theory modeling. *Educational and Psychological Measurement*, *74*(1), 58–76.
- Roland, C., Varadhan, R., & Frangakis, C. (2007). Squared polynomial extrapolation methods with cycling: An application to the positron emission tomography problem. *Numer. Algor.*, *44*, 159–172.
- Stone, C. A. (1992). Recovery of marginal maximum likelihood estimates in the two-parameter logistic response model: An evaluation of multilog. *Applied Psychological Measurement*, *16*(1), 1–16.
- Varadhan, R., & Roland, C. (2004). Squared extrapolation methods (squarem): A new class of simple and efficient numerical schemes for accelerating the convergence of the em algorithm. *Johns Hopkins University, Department of Biostatistics Working Papers*, (Paper 63).
- Varadhan, R., & Roland, C. (2008). Simple and globally convergent methods for accelerating the convergence of any em algorithm. *Scand J Stat.*, *35*(2), 335–353.
- Von Davier, M., & Sinharay, S. (2014). Handbook of international large-scale assessment: Background, technical issues, and methods of data analysis. In M. Rutkowski, M. von Davier, & D. Rutkowski (Eds.). CRC Press Boca Raton, FL.
- Wu, C. F. J. (1983). On the convergence properties of the em algorithm. *The Annals of Statistics*, *11*(1), 95–103.
- Zhou, H., Alexander, D., & Lange, K. (2011). A quasi-newton acceleration for high-dimensional optimization algorithms. *Stat Comput.*, *21*(2), 261–273.

Condition 1: 1PL, 9 items					
Method	N_0^*	Fevals	Levals	Conv.	CPU (ms)
EM	29	29	1	1.00	12.68
QN (q = 1)	14	17	15	1.00	18.62
QN (q = 2)	14	18	14	1.00	12.67
QN (q = 3)	13	18	13	1.00	12.92
QN (q = 4)	13	19	13	1.00	10.46
SQUAREM (k = 1)	—	—	—	0.00	—
SQUAREM (k = 2)	8	38	7	1.00	12.15
SQUAREM (k = 3)	6	37	5	1.00	10.91
SQUAREM (k = 4)	5	38	4	1.00	9.93
PEM	10	30	22	1.00	12.53
Condition 2: 1PL (fixed var.), 9 items					
Method	N_0^*	Fevals	Levals	Conv.	CPU (ms)
EM	37	37	1	1.00	8.52
QN (q = 1)	8	11	17	1.00	8.41
QN (q = 2)	6	10	11	1.00	5.67
QN (q = 3)	5	10	10	1.00	5.47
QN (q = 4)	4	10	8	1.00	4.10
SQUAREM (k = 1)	6	11	7	1.00	4.51
SQUAREM (k = 2)	4	22	5	1.00	6.09
SQUAREM (k = 3)	3	25	4	1.00	6.18
SQUAREM (k = 4)	3	27	4	1.00	5.52
PEM	5	20	30	1.00	9.67
Condition 3: 2PL, 9 items					
Method	N_0^*	Fevals	Levals	Conv.	CPU (ms)
EM	60	60	1	1.00	14.11
QN (q = 1)	30	33	61	1.00	23.39
QN (q = 2)	30	34	60	1.00	22.21
QN (q = 3)	29	34	57	1.00	20.27
QN (q = 4)	28	34	55	1.00	24.28
SQUAREM (k = 1)	25	32	26	1.00	15.67
SQUAREM (k = 2)	19	98	20	1.00	21.42
SQUAREM (k = 3)	13	90	14	1.00	21.07
SQUAREM (k = 4)	11	99	12	1.00	21.92
PEM	20	50	73	1.00	21.11

Table 2: EM acceleration results for conditions 1, 2, and 3 (all with nine items). The numbers shown represent a rounded average over 100 simulation trials. No results are shown for SQUAREM (k = 1) in condition 1 as none of the trials converged. N_0^* : Average number of iterations from the starting value to the fixed point; Fevals: Average number of F evaluations; Levals: Average number of log-likelihood evaluations; Conv.: Fraction of converging runs; CPU: Average CPU time in milliseconds.

Condition 4: 1PL, 100 items					
Method	N_0^*	Fevals	Levals	Conv.	CPU (ms)
EM	82	82	1	1.00	764.23
QN (q = 1)	41	44	41	1.00	1117.71
QN (q = 2)	41	45	41	1.00	1114.32
QN (q = 3)	40	45	40	1.00	1118.99
QN (q = 4)	40	46	40	1.00	1112.11
SQUAREM (k = 1)	—	—	—	0.00	—
SQUAREM (k = 2)	21	104	21	1.00	1139.75
SQUAREM (k = 3)	15	99	14	1.00	1032.70
SQUAREM (k = 4)	11	98	11	1.00	998.70
PEM	37	83	74	1.00	1403.95
Condition 5: 1PL (fixed var.), 100 items					
Method	N_0^*	Fevals	Levals	Conv.	CPU (ms)
EM	241	241	1	1.00	2183.43
QN (q = 1)	41	44	82	1.00	1490.32
QN (q = 2)	21	25	41	1.00	773.61
QN (q = 3)	13	18	27	1.00	510.93
QN (q = 4)	10	16	21	1.00	413.30
SQUAREM (k = 1)	27	40	28	1.00	842.71
SQUAREM (k = 2)	11	57	12	1.00	645.99
SQUAREM (k = 3)	9	64	10	1.00	666.74
SQUAREM (k = 4)	6	57	7	1.00	597.33
PEM	33	77	103	1.00	1569.65
Condition 6: 2PL, 100 items					
Method	N_0^*	Fevals	Levals	Conv.	CPU (ms)
EM	451	451	1	1.00	3947.93
QN (q = 1)	78	81	156	1.00	2713.00
QN (q = 2)	47	51	93	1.00	1664.42
QN (q = 3)	47	52	94	1.00	1671.58
QN (q = 4)	44	50	87	1.00	1566.44
SQUAREM (k = 1)	42	70	43	1.00	1350.81
SQUAREM (k = 2)	29	146	30	1.00	1513.24
SQUAREM (k = 3)	24	168	25	1.00	1687.02
SQUAREM (k = 4)	21	192	22	1.00	1862.17
PEM	50	111	207	1.00	2716.41

Table 3: EM acceleration results for conditions 4, 5, and 6 (all with 100 items). The numbers shown represent a rounded average over 100 simulation trials. No results are shown for SQUAREM (k = 1) in condition 4 as none of the trials converged. N_0^* : Average number of iterations from the starting value to the fixed point; Fevals: Average number of F evaluations; Levals: Average number of log-likelihood evaluations; Conv.: Fraction of converging runs; CPU: Average CPU time in milliseconds.