A Deep Learning Algorithm for High-Dimensional Exploratory Item Factor Analysis*

Christopher J. Urban[†] and Daniel J. Bauer

*L.L. Thurstone Psychometric Laboratory in the Department of Psychology and Neuroscience
University of North Carolina at Chapel Hill*

Abstract

Marginal maximum likelihood (MML) estimation is the preferred approach to fitting item response theory models in psychometrics due to the MML estimator’s consistency, normality, and efficiency as the sample size tends to infinity. However, state-of-the-art MML estimation procedures such as the Metropolis-Hastings Robbins-Monro (MH-RM) algorithm as well as approximate MML estimation procedures such as variational inference (VI) are computationally time-consuming when the sample size and the number of latent factors are very large. In this work, we investigate a deep learning-based VI algorithm for exploratory item factor analysis (IFA) that is computationally fast even in large data sets with many latent factors. The proposed approach applies a deep artificial neural network model called an importance-weighted autoencoder (IWAE) for exploratory IFA. The IWAE approximates the MML estimator using an importance sampling technique wherein increasing the number of importance-weighted (IW) samples drawn during fitting improves the approximation, typically at the cost of decreased computational efficiency. We provide a real data application that recovers results aligning with psychological theory across random starts. Via simulation studies, we show that the IWAE yields more accurate estimates as either the sample size or the number of IW samples increases (although factor correlation and intercepts estimates exhibit some bias) and obtains similar results to MH-RM in less time. Our simulations also suggest that the proposed approach performs similarly to and is potentially faster than constrained joint maximum likelihood estimation, a fast procedure that is consistent when the sample size and the number of items simultaneously tend to infinity.

Key words: Deep learning, artificial neural network, variational inference, variational autoencoder, importance sampling, importance weighted autoencoder, item response theory, categorical factor analysis, latent variable modeling

1 Introduction

Psychology and education researchers often collect large-scale test data with many respondents and many items in order to measure unobserved latent constructs such as personality traits or cognitive abilities. When test items are dichotomous (e.g., “Yes” or “No”) or polytomous (e.g., “Always”, “Frequently”, “Occasionally”, or “Never”), item factor analysis (IFA) is a principled alternative to linear factor analysis for summarizing the items using a smaller number of continuous latent factors. Exploratory IFA (Bock et al., 1988) in particular is an indispensable tool for uncovering the latent structure underlying a test by estimating the

*This is a preprint of an article published in Psychometrika. The final authenticated version is available online at: <https://doi.org/doi:10.1007/s11336-021-09748-3>

[†]Correspondence to cjurban@live.unc.edu.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1650116.

associations between items and latent factors (i.e., the factor loadings) in a data-driven manner. See Bolt (2005) or Wirth and Edwards (2007) for overviews of exploratory IFA.

Exploratory IFA parameters are most often estimated using Bock and Aitkin’s (1981) marginal maximum likelihood (MML) estimator, which enjoys consistency, normality, and efficiency as the sample size approaches infinity. The MML approach estimates the item parameters by maximizing the marginal likelihood of the observed item responses, which is obtained by integrating out the latent factors. Problematically, the computational complexity of evaluating this integral is exponential in the dimension of the latent space, making direct evaluation of the marginal likelihood computationally burdensome in the high-dimensional setting. A variety of methods for approximating the integrals have been proposed, including adaptive Gaussian quadrature (Rabe-Hesketh et al., 2005; Schilling & Bock, 2005), Laplace approximation (e.g., Huber et al., 2004), Monte Carlo integration (e.g., Meng & Schilling, 1996; Song & Lee, 2005), Markov Chain Monte Carlo (e.g., Béguin & Glas, 2001; Edwards, 2010), and stochastic approximation (SA; e.g., Cai, 2010a; Cai, 2010b; Zhang, Chen, & Liu, 2020). The Metropolis-Hastings Robbins-Monro (MH-RM; Cai, 2010a; 2010b) SA algorithm has been particularly widely used in psychology and in education due to its computational efficiency, and the recent stochastic expectation-maximization (stEM; Zhang, Chen, & Liu, 2020) algorithm performs comparably to MH-RM. However, even these state-of-the-art SA algorithms are computationally intensive when the sample size and the number of latent factors are very large (e.g., more than 10 latent factors and more than 10000 respondents).

Other marginal likelihood-based parameter estimation methods for exploratory IFA avoid approximating high-dimensional integrals and are therefore more computationally efficient. Limited-information approaches such as the bivariate composite likelihood estimator (Jöreskog & Moustaki, 2001) and the weighted least squares estimator (Muthén, 1978; 1984) yield fast, consistent, and asymptotically normally distributed estimates but are not asymptotically efficient. Approaches based on variational inference (VI; Jordan et al., 1998; Wainwright & Jordan, 2008) perform approximate MML estimation by optimizing a lower bound on the marginal likelihood rather than the marginal likelihood itself. More specifically, VI posits a family of approximate latent variable (LV) posterior distributions, then finds the member of this family that is closest to the true LV posterior distribution by optimizing the aforementioned lower bound; the variational estimator is equivalent to the MML estimator when the approximate and true LV posterior distributions are precisely equal. Since a separate set of approximate LV posterior distribution parameters is estimated for each data point, VI’s computational complexity depends on the sample size and on the complexity of the approximating family. Variational methods for IFA have demonstrated competitive performance with SA algorithms such as MH-RM for small sample sizes (Cho, 2020; Hui et al., 2017). Additionally, concurrent work by Cho (2020) has established consistency of the variational estimator for the multidimensional two-parameter logistic (M2PL) model in the double asymptotic setting where both the sample size and the number of items simultaneously tend to infinity. However, the variational estimator’s theoretical properties have not yet been established for other IFA models or in the classical asymptotic setting where only the sample size tends to infinity.

The MML estimator’s computational inefficiency arises from treating the latent factors as random effects that must be integrated out of the marginal likelihood. An alternative class of computationally efficient estimators treats the latent factors as fixed parameters, thereby avoiding the need for specifying a prior distribution on the latent factors and for evaluating high-dimensional integrals. However, these estimators pay a price for their computational efficiency: namely, they are only consistent in the double asymptotic setting described above. The constrained joint maximum likelihood estimator (CJMLE; Chen, Li, & Zhang, 2019) is the state-of-the-art estimator in this class; it is faster than MML-based approaches and is efficient in the double asymptotic setting. Zhang, Chen, and Li’s (2020) estimator based on singular value decomposition (SVD) is faster than CJMLE and does not suffer from convergence issues, although it is not (double) asymptotically efficient.

It is clear that an estimation procedure combining the asymptotic properties of the MML estimator with the computational efficiency of CJMLE is lacking from the IFA literature. In this work, we investigate a VI-based procedure that offers a step toward achieving these properties. This procedure employs techniques from two active deep learning (DL) research areas: amortized variational inference (AVI; Gershman & Goodman,

2014) and importance weighted variational inference (IWVI; Burda et al., 2016; Domke & Sheldon, 2018). AVI modifies traditional VI by using a powerful function approximator called an inference model to predict the parameters of the LV posterior for each data point rather than estimating these parameters directly. AVI is faster than traditional VI for large data sets, although it can be less flexible in practice (Cremer et al., 2018). IWVI decreases the gap between the variational lower bound and the true marginal likelihood by drawing multiple importance-weighted (IW) samples from the approximate LV posterior during fitting, thereby trading computational efficiency for a better lower bound. When the number of IW samples equals infinity, IWVI is theoretically equivalent to MML estimation and thus inherits the asymptotic properties of the MML estimator (Burda et al., 2016).

The proposed algorithm is based on the importance-weighted autoencoder (IWAE; Burda et al., 2016), an algorithm for amortized IWVI whose inference model is a deep artificial neural network (ANN). The IWAE is in turn an extension of the variational autoencoder (VAE; Kingma & Welling, 2014; Rezende et al., 2014), a foundational AVI algorithm that also employs an ANN inference model. Our work extends that of Curi et al. (2019), who used a VAE to estimate item parameters in a confirmatory M2PL model. Our work is also related to concurrent work by Wu et al. (2020), who applied a VAE for confirmatory M2PL item parameter estimation in the Bayesian setting. Our major contributions are as follows: (1) We introduce the IWAE to the IFA literature and describe how it may be applied for exploratory analysis of polytomous item response data in the frequentist setting, and (2) we conduct simulation studies to investigate the finite sample behavior the IWAE and to compare the IWAE to MH-RM and CJMLE.

Our paper is organized as follows. Section 2 provides a brief overview of ANNs. Section 3 introduces the problem of fitting IFA models with polytomous responses. Section 4 describes AVI and IWVI for IFA. The full algorithm is proposed in Section 5 and computational details are discussed. Section 6 includes an empirical example and simulation studies. Extensions of the method are described in Section 7 and discussions are given in Section 8.

2 A Brief Overview of Artificial Neural Networks

Deep learning (DL) models are machine learning models that map a set of predictor variables through a sequence of transformations called layers to predict a set of outcome variables. Much of DL’s success in recent years can be attributed to a family of nonlinear statistical models called artificial neural networks (ANNs; LeCun et al., 2015). ANNs are essential building blocks for the algorithm described in this work.

2.1 Feedforward Neural Networks

Feedforward neural networks (FNNs) are a simple class of ANNs. In practice, they are used as powerful function approximators because they can approximate any Borel measurable function between finite dimensional spaces to any desired degree of accuracy (Cybenko, 1989). Consider a data set $\{\mathbf{y}_i, \mathbf{x}_i\}_{i=1}^N$ where $\mathbf{x}_i$ is the i^{th} observed $J \times 1$ vector of predictor variables and $\mathbf{y}_i$ is the i^{th} observed $P \times 1$ vector of outcome variables. Note that here we define $\mathbf{x}_i$ as a vector of observed variables in line with typical treatments of FNNs, although we will redefine it as a vector of LVs in Section 3. FNNs map the predictor variables through a sequence of L transformations to predict the outcome variables as follows:

$$\mathbf{h}_i^{(l)} = f^{(l)}(\mathbf{W}^{(l)}\mathbf{h}_i^{(l-1)} + \mathbf{b}^{(l)}), \quad l = 1, \dots, L, \quad (1)$$

where $\mathbf{h}_i^{(0)} = \mathbf{x}_i$, $\mathbf{h}_i^{(L)} = \mathbf{y}_i - \boldsymbol{\varepsilon}_i$ where $\boldsymbol{\varepsilon}_i$ is the i^{th} $J \times 1$ vector of errors, $\mathbf{h}_i^{(l)}$ is a $P_l \times 1$ vector of LVs for layers $l = 2, \dots, L-1$, $\mathbf{W}^{(l)}$ is a $P_l \times P_{l-1}$ matrix of regression weights for layer l , $\mathbf{b}^{(l)}$ is a $P_l \times 1$ vector of intercepts for layer l , and $f^{(l)}$ is an almost everywhere differentiable activation function for layer l . $\mathbf{x}_i$ is called the input layer, $\mathbf{h}_i^{(1)}, \dots, \mathbf{h}_i^{(L-1)}$ are called hidden layers, and $\mathbf{y}_i$ is called the output layer. Figure 1 shows an FNN schematic diagram.

Notice that FNNs are recursive generalized linear models where each activation function $f^{(l)}$ is an inverse link function linking a linear combination of the variables at layer $l-1$ to the mean of the variables at layer l .

Figure 1: Schematic representation of a feedforward neural network with a single hidden layer. The input layer is a 3×1 vector, the hidden layer is a 4×1 vector, and the output layer is a 2×1 vector. Case subscripts i are omitted to avoid clutter.

In this work, we set the hidden layer activation functions $f^{(1)}, \dots, f^{(L-1)}$ to the exponential linear unit (ELU) function

$$f(z) = \begin{cases} z, & \text{if } z \geq 0 \\ \gamma(\exp(z) - 1), & \text{if } z < 0 \end{cases}, \quad z \in \mathbb{R}, \quad (2)$$

where $\gamma \in \mathbb{R}$ is a hyperparameter (we set $\gamma = 1$ in this work) and f is applied to vectors element-wise. FNNs with ELU hidden layer activation functions are easy to fit and perform well in practice (Clevert et al., 2016). We set the final activation function $f^{(L)}$ to the identity function:

$$f(z) = z, \quad z \in \mathbb{R}, \quad (3)$$

which is applied to vectors element-wise and corresponds to a linear regression of the layer $L - 1$ LVs on the outcomes.

2.2 Fitting FNNs Using AMSGrad

FNNs are typically fitted using stochastic gradient (SG) methods, a class of algorithms that iteratively update model parameters using stochastic estimates of the gradient of the objective function. Readers are referred to Bottou et al. (2018) for an overview of SG methods. In this work, we use the AMSGrad SG algorithm (Reddi et al., 2018), a method that adapts the magnitudes of its parameter updates using exponential moving averages of past stochastic gradient estimates. This approach allows AMSGrad to dynamically utilize information from the observed data to update each parameter a different amount at each iteration (Duchi et al., 2011; McMahan & Streeter, 2010). AMSGrad has theoretical convergence guarantees and performs well in practice with little tuning. In contrast, approaches based on the Robbins-Monro SA algorithm require the user to pre-specify a sequence of parameter update magnitudes that are fixed across parameters at each fitting iteration. These pre-specified update schemes typically require fine-tuning to the observed data and are often unstable on implementation (Nemirovski et al., 2009; Spall, 2003).

Let ξ_t be a $d \times 1$ vector of parameter values at fitting iteration t , $t = 0, \dots, T$, and let $\mathcal{J} : \mathbb{R}^d \mapsto \mathbb{R}$ be a possibly non-convex objective function that decomposes as a sum over observations:

$$\mathcal{J}(\xi_t) = \frac{1}{N} \sum_{i=1}^N \mathcal{J}_i(\xi_t), \quad (4)$$

where $\mathcal{J}_i$ is a per-observation objective function. Let $\{\mathbf{y}_i, \mathbf{x}_i\}_{i=1}^M$ where $M < N$ be a subsample of observations called a mini-batch. Then at iteration t , an unbiased estimator of the gradient of the objective function for the full data set is

$$\mathbf{g}_t = \frac{1}{M} \nabla_{\xi_t} \sum_{i=1}^M \mathcal{J}_i(\xi_t), \quad (5)$$

where ∇_{ξ_t} returns a $d \times 1$ vector of first-order partial derivatives w.r.t. ξ_t . AMSGrad proposes iterative parameter updates as follows:

$$\xi_{t+1} = \xi_t - \eta \frac{\mathbf{m}_t}{\sqrt{\hat{\mathbf{v}}_t}}, \quad (6)$$

where

$$\begin{aligned} \mathbf{m}_t &= \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \mathbf{g}_t; \\ \mathbf{v}_t &= \beta_2 \mathbf{v}_{t-1} + (1 - \beta_2) \mathbf{g}_t^2; \\ \hat{\mathbf{v}}_t &= \max(\hat{\mathbf{v}}_{t-1}, \mathbf{v}_t); \end{aligned} \quad (7)$$

$\mathbf{m}_0 = \mathbf{0}$; $\mathbf{v}_0 = \mathbf{0}$; $\hat{\mathbf{v}}_0 = \mathbf{0}$; $\mathbf{m}_t$ and $\mathbf{v}_t$ are $d \times 1$ vectors containing exponential moving averages of the gradient and the squared gradient at iteration t , respectively; $\beta_1 \in [0, 1]$ and $\beta_2 \in [0, 1]$ are forgetting factors for the gradient and the squared gradient, respectively; $\eta > 0$ is a step size called the learning rate; and square, square root, division, and maximum operations are applied to vectors element-wise. When mini-batches are sampled uniformly at random with replacement and the learning rate is sufficiently small, AMSGrad is guaranteed to converge to a local stationary point for smooth, non-convex objective functions. Zhou et al. (2018) as well as Chen, Liu, Sun, and Hong (2019) provide conditions required for first-order convergence, while Staib et al. (2019) discuss second-order convergence. Importantly, computation time per iteration does not increase with the sample size, allowing for convergence even with very large-scale data (Bottou et al., 2018).

When ξ_t are FNN parameters, an algorithm called backpropagation (BP) is used to efficiently compute the gradient estimator in equation 5. BP is an application of the chain rule of calculus and is a special case of reverse mode automatic differentiation (Linnainmaa, 1970). Goodfellow et al. (2016) provide a detailed discussion of BP.

3 The Problem of Fitting High-Dimensional Item Factor Analysis Models

3.1 The Graded Item Response Model

Samejima's (1969) graded response model (GRM) is a widespread model for polytomous item responses. We introduce notation for the GRM following Cai (2010a). Suppose there are $i = 1, \dots, N$ distinct respondents and $j = 1, \dots, J$ items. Let $y_{i,j} \in \{0, 1, \dots, C_j - 1\}$ denote the response for respondent i to item j in C_j graded (i.e., ordinal) categories. Note that when $C_j = 2$ for all j , the GRM reduces to the M2PL (McKinley & Reckase, 1983).

Suppose we have P LVs; let $\mathbf{x}_i$ denote the $P \times 1$ vector of factor scores (i.e., LV values) for respondent i . Let β_j denote the $P \times 1$ vector of loadings, let $\alpha_j = (\alpha_{j1}, \dots, \alpha_{j,C_j-1})^\top$ denote the $(C_j - 1) \times 1$ vector of strictly ordered category intercepts, and let $\theta_j = (\alpha_j^\top, \beta_j^\top)^\top$ be a vector collecting all parameters for item j . The GRM defines a set of boundary response probabilities conditional on the item parameters θ_j and the factor scores $\mathbf{x}_i$:

$$\Pr(y_{i,j} \geq k \mid \theta_j, \mathbf{x}_i) = \frac{1}{1 + \exp[-D(\alpha_{j,k} + \beta_j^\top \mathbf{x}_i)]}, \quad k \in \{1, \dots, C_j - 1\}, \quad (8)$$

where $\Pr(y_{i,j} \geq 0 \mid \theta_j, \mathbf{x}_i) = 1$, $\Pr(y_{i,j} \geq C_j \mid \theta_j, \mathbf{x}_i) = 0$, and D is a scaling constant (typically 1.702) used to help the logistic metric better approximate the normal ogive metric (Reckase, 2009). The conditional

probability for a particular response $y_{i,j} = k$, $k \in \{0, \dots, C_j - 1\}$ is

$$\pi_{i,j,k} = P(y_{i,j} = k \mid \boldsymbol{\theta}_j, \mathbf{x}_i) = \Pr(y_{i,j} \geq k \mid \boldsymbol{\theta}_j, \mathbf{x}_i) - \Pr(y_{i,j} \geq k+1 \mid \boldsymbol{\theta}_j, \mathbf{x}_i). \quad (9)$$

3.2 Observed Data Likelihood

It follows from equation 9 that the conditional distribution of $y_{i,j}$ is multinomial with C_j cells, trial size 1, and cell probabilities $\pi_{i,j,k}$:

$$p_{\boldsymbol{\theta}_j}(y_{i,j} \mid \mathbf{x}_i) = \prod_{k=0}^{C_j-1} \pi_{i,j,k}^{\mathbb{1}_k(y_{i,j})}, \quad (10)$$

where we define the indicator function

$$\mathbb{1}_k(y) = \begin{cases} 1, & \text{if } y = k \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

for $k \in \{0, \dots, C_j - 1\}$. Let $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,n})^\top$ be respondent i 's response pattern. By the usual conditional independence assumption, the conditional distribution of $\mathbf{y}_i$ is

$$p_{\boldsymbol{\theta}}(\mathbf{y}_i \mid \mathbf{x}_i) = \prod_{j=1}^J p_{\boldsymbol{\theta}_j}(y_{i,j} \mid \mathbf{x}_i), \quad (12)$$

where $\boldsymbol{\theta}$ is a vector collecting the estimable parameters for all J items.

Assume the prior distribution of factor scores $\mathbf{x}_i$ is standard multivariate normal with density function $\mathcal{N}(\mathbf{x}_i)$. Then the marginal distribution of $\mathbf{y}_i$ is given by

$$p_{\boldsymbol{\theta}}(\mathbf{y}_i) = \int \prod_{j=1}^J p_{\boldsymbol{\theta}_j}(y_{i,j} \mid \mathbf{x}) \mathcal{N}(\mathbf{x}) d\mathbf{x}, \quad (13)$$

where the above integral is over $\mathbb{R}^p$. Let $\mathbf{Y}$ be an $N \times J$ matrix of independent response patterns whose i^{th} row is $\mathbf{y}_i^\top$. The observed data likelihood is

$$\mathcal{L}(\boldsymbol{\theta} \mid \mathbf{Y}) = \prod_{i=1}^N \left[\int \prod_{j=1}^J p_{\boldsymbol{\theta}_j}(y_{i,j} \mid \mathbf{x}) \mathcal{N}(\mathbf{x}) d\mathbf{x} \right]. \quad (14)$$

Maximizing $\mathcal{L}(\boldsymbol{\theta} \mid \mathbf{Y})$ directly is difficult because we must approximate the N integrals over $\mathbb{R}^p$ numerically. In this work, we avoid this difficulty by deriving an analytical lower bound on $\log \mathcal{L}(\boldsymbol{\theta} \mid \mathbf{Y})$ using VI. We then maximize this lower bound using a DL algorithm.

4 Variational Methods for Item Factor Analysis

Variational inference (VI) is an approach to approximate maximum likelihood estimation for LV models that has recently gained traction in the machine learning community (Blei et al., 2017; Zhang, Butepage, Kjellstrom, & Mandt, 2019). VI has been applied for confirmatory IFA in both the frequentist setting (Cho, 2020; Curi et al., 2019) and the Bayesian setting (Chen, Filho, Prudêncio, Diethe, & Flach, 2019; Natesan et al., 2016; Wu et al., 2020) as well as for exploratory IFA in the frequentist setting (Cho, 2020; Hui et al., 2017). In this section, we describe variational methods for IFA as well as an importance sampling technique for helping the variational estimator better approximate the MML estimator.

4.1 Variational Inference

We now describe VI in the context of a general LV model, then apply VI to IFA in the following sections. The main idea behind VI is to treat LV inference as an optimization problem. More formally, let $\mathbf{y} \in \mathcal{Y}$ and $\mathbf{x} \in \mathcal{X}$ denote observed and LVs, respectively, where $\mathcal{Y}$ and $\mathcal{X}$ are sample spaces. VI introduces a family $\mathcal{Q}$ of approximate densities over LVs and aims to find the member $q_{\psi^*(\mathbf{y})}(\mathbf{x}) \in \mathcal{Q}$ that minimizes the Kullback-Leibler (KL) divergence¹ from itself to the true LV posterior:

$$q_{\psi^*(\mathbf{y})}(\mathbf{x}) = \underset{q_{\psi(\mathbf{y})}(\mathbf{x}) \in \mathcal{Q}}{\operatorname{argmin}} D_{\text{KL}} [q_{\psi(\mathbf{y})}(\mathbf{x}) \| p(\mathbf{x} | \mathbf{y})], \quad (15)$$

where $\psi(\mathbf{y})$ is a vector of variational parameters. Note that $\psi(\mathbf{y})$ depends on $\mathbf{y}$, indicating that a different vector of variational parameters is estimated for each observation. For models with continuous LVs, an often tractable choice for the approximate posterior is the isotropic normal density (Kingma & Welling, 2014):

$$q_{\psi(\mathbf{y})}(\mathbf{x}) = \mathcal{N}(\mathbf{x} | \mu(\mathbf{y}), \sigma^2(\mathbf{y}) \mathbf{I}_P), \quad (16)$$

where $\mu(\mathbf{y})$ is a $P \times 1$ vector of means, $\sigma^2(\mathbf{y})$ is a $P \times 1$ vector of variances, and $\mathbf{I}_P$ is a $P \times P$ identity matrix. Minimizing the KL divergence from the isotropic normal approximate posterior to the true LV posterior produces the “best” isotropic normal approximation to the true LV posterior. In practice, however, tractable approximate posteriors such as the isotropic normal density are rarely flexible enough to perfectly approximate the true LV posterior and thereby minimize the KL divergence to zero. The importance sampling technique described later in this section improves the accuracy of VI by implicitly increasing the flexibility of the approximate posterior.

4.2 Evidence Lower Bound

The log-likelihood of the observed data under the GRM can be written as a sum over the marginal likelihood of each observation:

$$\ell(\theta | \mathbf{Y}) = \sum_{i=1}^N \log p_{\theta}(\mathbf{y}_i), \quad (17)$$

where $\ell(\theta | \mathbf{Y}) = \log \mathcal{L}(\theta | \mathbf{Y})$. Let the approximate LV posterior be the isotropic normal density as in equation 16. We can re-write a single summand in equation 17 as

$$\log p_{\theta}(\mathbf{y}_i) = D_{\text{KL}} [q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i) \| p_{\theta}(\mathbf{x}_i | \mathbf{y}_i)] + \mathbb{E}_{q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i)} [\log p_{\theta}(\mathbf{x}_i, \mathbf{y}_i) - \log q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i)]. \quad (18)$$

The first term on the r.h.s. of equation 18 is the KL divergence from the approximate to the true LV posterior (i.e., it is the term we wish to minimize from equation 15). Since this term is non-negative, the second term on the r.h.s. of equation 18 is a lower bound on the marginal likelihood of a single observation. This term is called the evidence lower bound (ELBO) and can be re-written as

$$\log p_{\theta}(\mathbf{y}_i) \geq \mathbb{E}_{q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i)} [\log p_{\theta}(\mathbf{x}_i, \mathbf{y}_i) - \log q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i)] \quad (19)$$

$$= \mathbb{E}_{q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i)} [\log p_{\theta}(\mathbf{y}_i | \mathbf{x}_i)] - D_{\text{KL}} [q_{\psi(\mathbf{y}_i)}(\mathbf{x}_i) \| p_{\theta}(\mathbf{x}_i)] \quad (20)$$

$$= \text{ELBO}_i. \quad (21)$$

The first term in the ELBO on line 20 is an expected conditional log-likelihood that encourages $q^*(\mathbf{x}_i | \mathbf{y}_i)$ to place mass on LVs that explain the observed data well, while the second term encourages densities that are close to the LV prior $p_{\theta}(\mathbf{x}_i)$. Maximizing the ELBO over all observations w.r.t. the item parameters θ and the variational parameters $\psi(\mathbf{y}_i)$ both approximately maximizes the observed data log-likelihood and minimizes the KL divergence from the approximate to the true LV posterior.

¹ For distributions q and p , the KL divergence is defined as $D_{\text{KL}} [q \| p] = \mathbb{E}_q [\log q] - \mathbb{E}_q [\log p]$. It can be shown that $D_{\text{KL}} [q \| p] \geq 0$ with equality if and only if $p = q$ almost everywhere w.r.t. q .

4.3 Amortized Variational Inference

Traditional VI fits a different approximate LV posterior for each observation, which quickly becomes computationally infeasible for large data sets. It is also not straightforward to apply models fitted using VI to previously unseen observations (e.g., to perform LV inference for or to compute the log-likelihood of the unseen observations). Amortized variational inference (AVI) is a computationally efficient alternative to VI that uses a powerful function approximator called an inference model to parameterize the approximate posterior. By sharing the parameters of the inference model across observations, AVI estimates a constant number of parameters regardless of the sample size, whereas VI estimates a number of parameters that at best grows linearly as a function of the sample size. Further, models fitted using AVI can easily be applied to previously unseen observations by simply feeding the observations to the inference model.

The variational autoencoder (VAE; Kingma & Welling, 2014; Rezende et al., 2014) is an AVI algorithm whose inference model is an ANN. We can use a VAE for IFA by parameterizing the approximate LV posterior as follows:

$$\begin{aligned} (\mu_i^\top, \log \sigma_i^\top)^\top &= \text{FNN}_\psi(\mathbf{y}_i), \\ q_\psi(\mathbf{x}_i | \mathbf{y}_i) &= \mathcal{N}(\mathbf{x}_i | \mu_i, \sigma_i^2 \mathbf{I}_P), \end{aligned} \quad (22)$$

where μ_i is a $P \times 1$ predicted vector of means, $\log \sigma_i$ is a $P \times 1$ predicted vector of log-standard deviations, and FNN_ψ is an L -layer FNN parameterized by ψ . Rather than estimating a set of variational parameters $\psi(\mathbf{y}_i)$ for each observation, the FNN parameters ψ are now shared across observations. That is, rather than maximizing equation 20 over observations, we now maximize

$$\text{ELBO} = \mathbb{E}_{q_\psi(\mathbf{x}|\mathbf{y})} [\log p_\theta(\mathbf{y} | \mathbf{x})] - D_{\text{KL}} [q_\psi(\mathbf{x} | \mathbf{y}) \| p_\theta(\mathbf{x})] \quad (23)$$

over observations. Note that we now drop the case index i since the FNN parameters ψ are shared across $\{\mathbf{y}_i, \mathbf{x}_i\}_{i=1}^N$. In theory, the VAE is equivalent to VI when the FNN is sufficiently flexible (e.g., when the FNN has one infinitely large hidden layer). In practice, the FNN has finite capacity and may prevent the VAE from performing as well as VI. This performance difference is called the amortization gap and may be reduced by increasing the flexibility of the approximate LV posterior (Cremer et al., 2018).

4.4 Fitting the Amortized Model

Fitting the VAE for IFA can be accomplished with AMSGrad and BP after obtaining an unbiased estimator for the gradient of the ELBO w.r.t. the model parameters θ and ψ . An unbiased estimator for the gradient of the ELBO w.r.t. the item parameters θ is

$$\nabla_\theta \text{ELBO} = \nabla_\theta \mathbb{E}_{q_\psi(\mathbf{x}|\mathbf{y})} [\log p_\theta(\mathbf{x}, \mathbf{y}) - \log q_\psi(\mathbf{x} | \mathbf{y})] \quad (24)$$

$$= \mathbb{E}_{q_\psi(\mathbf{x}|\mathbf{y})} [\nabla_\theta \log p_\theta(\mathbf{x}, \mathbf{y})] \quad (25)$$

$$\approx \frac{1}{S} \sum_{s=1}^S \nabla_\theta \log p_\theta(\mathbf{y}, \mathbf{x}_s), \quad (26)$$

where line 26 approximates the expectations in line 25 with a size S Monte Carlo sample of factor scores from the approximate LV posterior.² Obtaining an unbiased estimator for the gradient of the ELBO w.r.t. the FNN parameters ψ is more challenging because, in general,

$$\nabla_\psi \text{ELBO} = \nabla_\psi \mathbb{E}_{q_\psi(\mathbf{x}|\mathbf{y})} [\log p_\theta(\mathbf{x}, \mathbf{y}) - \log q_\psi(\mathbf{x} | \mathbf{y})] \quad (27)$$

$$\neq \mathbb{E}_{q_\psi(\mathbf{x}|\mathbf{y})} [\nabla_\psi \log p_\theta(\mathbf{x}, \mathbf{y}) - \nabla_\psi \log q_\psi(\mathbf{x} | \mathbf{y})], \quad (28)$$

²We move the gradient inside the expectation in line 25 using the fact that $q_\psi(\mathbf{x} | \mathbf{y})$, $\log q_\psi(\mathbf{x} | \mathbf{y})$, and $\log p_\theta(\mathbf{x}, \mathbf{y})$ satisfy certain regularity conditions. For details, see Lehmann and Casella (1998).

since the expectations are taken w.r.t. $q_\psi(\mathbf{x} | \mathbf{y})$, which is a function of ψ . To overcome this problem, we reparameterize $\mathbf{x}$ as follows:

$$\begin{aligned}\boldsymbol{\varepsilon} &\sim \mathcal{N}(\boldsymbol{\varepsilon}), \\ \mathbf{x} &= \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\varepsilon},\end{aligned}\tag{29}$$

where $\boldsymbol{\varepsilon}$ is a $P \times 1$ sample from a standard multivariate normal density, $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the outputs of the FNN inference model given in equations 22, and $\odot$ denotes element-wise multiplication. This reparameterization “externalizes” the randomness in $\mathbf{x}$ by writing $\mathbf{x}$ as a deterministic function of ψ . We can now obtain an unbiased estimator for the gradient of the ELBO w.r.t. ψ as follows:

$$\nabla_\psi \text{ELBO} = \nabla_\psi \mathbb{E}_{\mathcal{N}(\boldsymbol{\varepsilon})} [\log p_\theta(\mathbf{x}, \mathbf{y}) - \log q_\psi(\mathbf{x} | \mathbf{y})]\tag{30}$$

$$= \mathbb{E}_{\mathcal{N}(\boldsymbol{\varepsilon})} [\nabla_\psi \log p_\theta(\mathbf{x}, \mathbf{y}) - \nabla_\psi \log q_\psi(\mathbf{x} | \mathbf{y})]\tag{31}$$

$$\approx \frac{1}{S} \sum_{s=1}^S [\nabla_\psi \log p_\theta(\mathbf{x}_s, \mathbf{y}) - \nabla_\psi \log q_\psi(\mathbf{x}_s | \mathbf{y})],\tag{32}$$

where the expectations are now taken w.r.t. $\mathcal{N}(\boldsymbol{\varepsilon})$ and line 32 is a Monte Carlo approximation to the expectation in line 31. Figure 2 illustrates how computation proceeds in a VAE for IFA. We note that the KL divergence term shown has a closed form that is efficient to compute (Kingma & Welling, 2014):

$$D_{\text{KL}} [\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \boldsymbol{\sigma}^2 \mathbf{I}_P) \| \mathcal{N}(\mathbf{x})] = \frac{1}{2} \sum_{p=1}^P [\mu_p^2 + \sigma_p^2 - 1 - \log \sigma_p^2].\tag{33}$$

4.5 Importance-Weighted Variational Inference

Importance-weighted variational inference (IWVI; Burda et al., 2016; Domke & Sheldon, 2018) is a VI strategy that can approximate the true log-likelihood arbitrarily well. Amortized IWVI for IFA maximizes a lower bound called the importance-weighted ELBO (IW-ELBO):

$$\log p_\theta(\mathbf{y}) \geq \text{IW-ELBO}\tag{34}$$

$$= \mathbb{E}_{\mathbf{x}_{1:R}} \left[\log \frac{1}{R} \sum_{r=1}^R w_r \right],\tag{35}$$

where $\mathbf{x}_{1:R} \sim \prod_{r=1}^R q_\psi(\mathbf{x}_r | \mathbf{y})$, $w_r = p_\theta(\mathbf{x}_r, \mathbf{y}) / q_\psi(\mathbf{x}_r | \mathbf{y})$ are unnormalized importance weights for the joint distribution of latent and observed variables, and R is the number of importance-weighted (IW) samples. When $R = 1$, the IW-ELBO reduces to the ELBO. As $R \rightarrow \infty$, the IW-ELBO converges monotonically to the marginal log-likelihood (Burda et al., 2016). IWVI also implicitly defines a flexible approximate LV posterior $q_\psi^{\text{IW}}(\mathbf{x} | \mathbf{y})$ that converges to the true LV posterior pointwise as $R \rightarrow \infty$ (Cremer et al., 2017). These facts imply that IWVI is equivalent to MML estimation when the number of importance samples R equals infinity, in which case IWVI inherits the MML estimator’s asymptotic properties. When the inference model is an FNN, the associated IWVI algorithm is called the importance-weighted autoencoder (IWAE; Burda et al., 2016).

Optimizing the IW-ELBO permits trading computational efficiency for a better approximation to the MML estimator by increasing R . As with the ELBO, we can obtain an unbiased estimator for the gradient of the IW-ELBO w.r.t. $\xi = (\boldsymbol{\theta}^\top, \boldsymbol{\psi}^\top)^\top$ via the reparameterization trick:

$$\nabla_\xi \mathbb{E}_{\mathbf{x}_{1:R}} \left[\log \frac{1}{R} \sum_{r=1}^R w_r \right] = \mathbb{E}_{\boldsymbol{\varepsilon}_{1:R}} \left[\sum_{r=1}^R \tilde{w}_r \nabla_\xi \log w_r \right]\tag{36}$$

$$\approx \frac{1}{S} \sum_{s=1}^S \left[\sum_{r=1}^R \tilde{w}_{r,s} \nabla_\xi \log w_{r,s} \right],\tag{37}$$

Figure 2: Schematic diagram of a variational autoencoder for item factor analysis with $J = 6$ items, $C_j = 2$ categories per item, $P = 2$ factors, $S = 1$ Monte Carlo sample from the approximate latent variable posterior, and an inference model consisting of a feedforward neural network with a single hidden layer. The reparameterization trick is not illustrated for simplicity. LV = latent variable.

where $\varepsilon_{1:R} \sim \prod_{r=1}^R \mathcal{N}(\varepsilon_r)$ and $\tilde{w}_r = w_r / \sum_{r'=1}^R w_{r'}$ are normalized importance weights. Problematically, however, increasing R degrades performance of the above gradient estimator for the inference model parameters ψ (but not for the item parameters θ). Specifically, Rainforth et al. (2018) theoretically and empirically show that as R increases, the signal-to-noise ratio (SNR) of the inference model gradient estimator tends to zero so that the estimator becomes completely random. We resolve this issue using Tucker et al.’s (2019) doubly reparameterized gradient (DReG) estimator:

$$\nabla_{\psi} \mathbb{E}_{\mathbf{x}_{1:R}} \left[\log \frac{1}{R} \sum_{r=1}^R w_r \right] = \mathbb{E}_{\varepsilon_{1:R}} \left[\sum_{r=1}^R \tilde{w}_r^2 \frac{\partial \log w_r}{\partial \mathbf{x}_r} \frac{\partial \mathbf{x}_r}{\partial \psi} \right] \quad (38)$$

$$\approx \frac{1}{S} \sum_{s=1}^S \left[\sum_{r=1}^R \tilde{w}_{r,s}^2 \frac{\partial \log w_{r,s}}{\partial \mathbf{x}_{r,s}} \frac{\partial \mathbf{x}_{r,s}}{\partial \psi} \right]. \quad (39)$$

The DReG estimator is unbiased, has increasing SNR as $R \rightarrow \infty$, and empirically demonstrates lower variance than alternative estimators. In practice, the IW-ELBO θ -gradient and DReG estimators can be successfully approximated using a single Monte Carlo sample (e.g., Burda et al., 2016; Tucker et al., 2019), so we set $S = 1$ for all numerical examples in this work.

5 Implementation Details

5.1 Starting Values

The proposed algorithm is detailed in Algorithm 1. We now discuss choosing the algorithm starting values $\xi_0 = (\theta_0^\top, \psi_0^\top)^\top$.

Algorithm 1 *Deep Learning Algorithm for Exploratory Item Factor Analysis*

1. **Initialization** Input item responses $\mathbf{Y}$; dimension of latent space P ; mini-batch size M ; IW samples R ; MC samples S ; optimization hyperparameters η , β_1 , and β_2 ; and starting values $\xi_0 = (\theta_0^\top, \psi_0^\top)^\top$
2. At fitting iteration t , $t = 0, \dots, T$:
 - (a) **Computation** Randomly sample a mini-batch $\{\mathbf{y}_i\}_{i=1}^M$; compute objective function value for respondent i , $i = 1, \dots, M$:

$$(\mu_i^\top, \log \sigma_i^\top)^\top = \text{FNN}_{\psi_t}(\mathbf{y}_i)$$
 For IW sample r and MC sample s , $r = 1, \dots, R$, $s = 1, \dots, S$:

$$\begin{aligned} \varepsilon_{i,r,s} &\sim \mathcal{N}(\varepsilon_{i,r,s}) \\ \mathbf{x}_{i,r,s} &= \mu_i + \sigma_i \odot \varepsilon_{i,r,s} \\ \widetilde{\mathcal{L}}_1 &= \sum_{j=1}^J \sum_{k=0}^{C_j-1} \mathbb{1}_k(y_{i,j}) \log \pi_{i,j,k} &:= \log p_{\theta_t}(\mathbf{y}_i | \mathbf{x}_{i,r,s}) \\ \widetilde{\mathcal{L}}_2 &= \frac{1}{2} \sum_{p=1}^P [\mu_{i,p}^2 + \sigma_{i,p}^2 - 1 - \log \sigma_{i,p}^2] &:= D_{\text{KL}}[\mathcal{N}(\mathbf{x}_{i,r,s} | \mu_i, \sigma_i^2 \mathbf{I}_P) \| \mathcal{N}(\mathbf{x}_{i,r,s})] \\ w_{i,r,s} &= \exp[\widetilde{\mathcal{L}}_1 - \widetilde{\mathcal{L}}_2] \\ \text{IW-ELBO}_i &\approx \frac{1}{S} \sum_{s=1}^S \left[\log \frac{1}{R} \sum_{r=1}^R w_{i,r,s} \right] \end{aligned}$$
 - (b) **Optimization** Update model parameters using AMSGrad:

$$\begin{aligned} \mathbf{g}_t &= \frac{1}{M} \nabla_{\xi_t} \sum_{i=1}^M \text{IW-ELBO}_i \\ \mathbf{m}_t &= \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \mathbf{g}_t \\ \mathbf{v}_t &= \beta_2 \mathbf{v}_{t-1} + (1 - \beta_2) \mathbf{g}_t^2 \\ \hat{\mathbf{v}}_t &= \max(\hat{\mathbf{v}}_{t-1}, \mathbf{v}_t) \\ \xi_{t+1} &= \xi_t - \eta \frac{\mathbf{m}_t}{\sqrt{\hat{\mathbf{v}}_t}} \end{aligned}$$
3. **Output** Return $\hat{\xi} = \xi_T$

IW = importance-weighted, MC = Monte Carlo.

The inference model starting values ψ_0 include a $P_l \times P_{l-1}$ regression weight matrix $\mathbf{W}_0^{(l)}$ and a $P_l \times 1$ intercept vector $\mathbf{b}_0^{(l)}$ at FNN layers $l = 1, \dots, L$. We initialize these parameters using a variant of Kaiming initialization (He et al., 2015), which has demonstrated good performance when applied to ANNs with asymmetric activation functions (e.g., the ELU function). Let $\mathcal{U}(a, b)$ denote a uniform density with lower bound a and upper bound b . We randomly sample starting values as

$$w_{p_l, p_{l-1}, 0}^{(l)}, b_{p_l, 0}^{(l)} \sim \mathcal{U}\left(-\frac{1}{\sqrt{P_{l-1}}}, \frac{1}{\sqrt{P_{l-1}}}\right) \quad (40)$$

for $p_l = 1, \dots, P_l$, $p_{l-1} = 1, \dots, P_{l-1}$, $l = 1, \dots, L$. This initialization strategy often prevents the FNN hidden layer values from growing too large or too small at the start of fitting while accounting for the asymmetry of the ELU activation function around zero.

The starting values θ_0 include the $P \times 1$ factor loadings vector $\beta_{j,0}$ as well as the $(C_j - 1) \times 1$ intercept vector $\alpha_{j,0}$ for items $j = 1, \dots, J$. We initialize the factor loadings using Xavier initialization (Glorot &

Bengio, 2010), which performs well when applied to ANNs with symmetric activation functions:

$$\beta_{j,p,0}^{(l)} \sim \mathcal{U}\left(-\sqrt{\frac{6}{J+P}}, \sqrt{\frac{6}{J+P}}\right), \quad (41)$$

where $j = 1, \dots, J$ and $p = 1, \dots, P$. This approach stabilizes fitting in a manner similar to Kaiming initialization while accounting for the symmetry of the inverse logistic link function (i.e., equation 8) around zero. For $j = 1, \dots, J$, we initialize the elements of $\alpha_{j,0}$ to an increasing sequence such that the cumulative density of logistic distribution between consecutive elements is the same (Christensen, 2019).

5.2 Stabilizing Fitting and Checking Convergence

We use a KL annealing strategy to avoid entrapment in local optima at the start of fitting (Bowman et al., 2016; Sønderby et al., 2016). KL annealing multiplies the KL divergence term by t/τ for the first τ fitting iterations where $t = 0, \dots, \tau - 1$. We conduct KL annealing for $\tau = 1000$ fitting iterations for all models.

Once KL annealing is completed, we determine convergence similarly to Cremer et al. (2018). At each fitting iteration, we store the IW-ELBO computed for the associated mini-batch. After every 100 fitting iterations, we compute the average of the previous 100 mini-batch IW-ELBOs and compare this average to the previous best achieved average. If the best achieved average IW-ELBO does not improve after 100 such comparisons, fitting is terminated.

It is sometimes necessary to assess whether different optimization runs have converged to equivalent stationary points. We conduct these checks using the estimated loadings matrices. We compare loadings matrices across runs by first rotating the factor solution using the Geomin oblique rotation method (Yates, 1988). Next, we invert factors if the sum of their loadings is negative (Asparouhov & Muthén, 2009). We then select a reference matrix and find the column permutation of each comparison matrix that minimizes the element-wise mean squared error (MSE). Finally, we compute Tucker’s congruence coefficient between the permuted matrices (Lorenzo-Seva & ten Berge, 2006). Solutions with congruence coefficients larger than 0.98 are deemed equivalent (MacCallum et al., 1999). We note that to compare factor correlation solutions, the same inversion and permutation procedure is applied to both columns and rows of the estimated factor correlation matrices.

5.3 Tuning Hyperparameters

Inference model hyperparameters include the number and size of the FNN hidden layers. After some experimentation, we found that performance was relatively insensitive to these values. We therefore use a single hidden layer for all models and set the hidden layer size to a value close to the mean of the input layer size and twice the latent dimension P . This choice is based on the observation that “the optimal size of the hidden layer is usually between the size of the input and size of the output layers” (Heaton, 2008).

Optimization hyperparameters include the forgetting factors for the gradient and squared gradient, β_1 and β_2 ; the learning rate η ; and the mini-batch size M . We set $\beta_1 = 0.9$ and $\beta_2 = 0.999$, which are default values typically recommended in practice (Reddi et al., 2018). We set $\eta = 0.005$ for most models. For some models with many factors and many items, the IW-ELBO diverged, so we set $\eta = 0.0025$. This approach is based on the observations that $\eta \leq 0.005$ typically performs well for adaptive SG methods and η should be reduced if the objective function diverges (Bengio, 2012). Keskar et al. (2017) note that mini-batch sizes $M \geq 32$ perform well in many applications, and Bengio (2012) notes that M mostly impacts time to convergence rather than model performance. We therefore set $M = 32$ as a default value for all analyses.³

Setting the number of IW samples R typically does not require extensive tuning but does require some consideration. Empirically, we found that increasing R increases the accuracy of θ estimates. However, even for small R (e.g., $R = 5$), amortized IWVI typically yields comparable θ estimates to state-of-the-art MML estimation procedures in less time. We also found that computational efficiency increases for small R then

³ M is typically set to a power of 2 to reduce fitting times by facilitating GPU (or CPU) memory allocation (Goodfellow et al., 2016).

decreases as R grows large (e.g., $R = 25$). These results suggest that R may be chosen according to whether (1) computational efficiency or (2) highly accurate θ estimates are desired: If (1), choose some small R for which the algorithm converges quickly; if (2), choose the largest R for which the algorithm converges in a reasonable amount of time.

The main hyperparameter that requires tuning is the latent dimension P . We tried tuning P using a pseudo-likelihood Bayesian information criterion (pseudo-BIC; Erosheva et al., 2007) as well as using a more computationally intensive 5-fold cross-validation (CV) approach (details available upon request) but found that both approaches performed poorly as N increased. We therefore use a more subjective scree plot approach based on the Monte Carlo CV method described by Hui et al. (2017). To construct each scree plot, we first create a holdout set by randomly sampling some percentage of the item responses without replacement. Let Ω denote the index set for the item responses in the holdout set and let Ω' denote the indices of the item responses excluding the holdout set. For a fixed P , we fit the model using the item responses indexed by Ω' . We denote the fitted parameters so obtained as $\hat{\xi} = (\hat{\theta}^\top, \hat{\psi}^\top)^\top$. Treating the IW-ELBO with $R = 5000$ IW samples as a close approximation to the true log-likelihood (Cremer et al., 2018), we predict the approximate log-likelihood for the holdout set as

$$\tilde{\ell}(P) = \sum_{i \in \Omega} \left[\log \frac{1}{5000} \sum_{r=1}^{5000} \frac{p_{\hat{\theta}}(\mathbf{x}_{i,r}, \mathbf{y}_i)}{q_{\hat{\psi}}(\mathbf{x}_{i,r} | \mathbf{y}_i)} \right], \quad (42)$$

which corresponds to equation 35 with $S = 1$ and $\xi = \hat{\xi}$. After performing the above procedure for several successive values of P , the scree plot is constructed by plotting $-\tilde{\ell}(P)$ against increasing P . The latent dimension coinciding with an “elbow” in the plot may be selected. We note that this approach differs from traditional scree plots in that we plot predicted approximate log-likelihoods rather than eigenvalues, although both approaches are interpreted similarly (i.e., look for the “elbow”). We empirically evaluate this approach in the following section.

6 Numerical Examples

Models were programmed with the machine learning library PyTorch (Version 1.1.6; Paszke et al., 2019) and were fitted on a laptop computer with a 2.8 GHz Intel Core i7 CPU and 16 GB of RAM. Although GPU computing is directly supported in PyTorch and often speeds up fitting, we opted for CPU computing to enable fairer comparisons with other methods and to assess performance using hardware more typically available for psychology and education research. All code is available as online supplemental material.

6.1 Application to a Big-Five Personality Questionnaire

We first demonstrate amortized IWVI via an empirical example intended to: (1) showcase IWVI’s computational efficiency when analyzing large-scale item response data and (2) obtain reasonable population values for conducting simulation studies. Specifically, we analyze 1015342 responses to Goldberg’s (1992) 50 Big-Five Factor Marker (FFM) items from the International Personality Item Pool (IPIP; Goldberg et al., 2006) downloaded from the Open-Source Psychometrics Project (<https://openpsychometrics.org/>). The IPIP-FFM items were designed to assess respondents’ levels of five personality factors: *Conscientiousness*, *openness*, *emotional stability*, *agreeableness*, and *extraversion*. Empirical Big-Five studies often yield substantial factor inter-correlations (e.g., Biesanz & West, 2004), so we permitted correlated factors by applying the Geomin oblique rotation method to all fitted loadings matrices. Each of the five factors included 10 five-category items anchored by “Disagree” (1), “Neutral” (3), and “Agree” (5). Item responses were recoded as necessary so that the highest numerical value of the response scale indicated a high level of the corresponding factor. After pre-processing the data (details available upon request), our final sample size was $N = 515\,708$ responses.

Computation was carried out following the procedures described in Section 5. A scree plot of $-\tilde{\ell}(P)$ computed on a holdout set of 2.5% of observations for $P \in \{1, \dots, 10\}$ (Figure 3) demonstrated an “elbow” at

Figure 3: Scree plot of predicted approximate negative log-likelihood as a function of the number of latent factors. The “elbow” at 5 factors is marked with a dotted line.

Table 1: Factor correlations for IPIP-FFM data set.

Factor	Factor				
	1	2	3	4	5
1. Extraversion	1.00				
2. Emotional Stability	-.20	1.00			
3. Agreeableness	.15	-.01	1.00		
4. Conscientiousness	.08	-.22	.11	1.00	
5. Openness	.12	-.05	.10	-.01	1.00

$P = 5$, suggesting that 5 latent factors accounted for most of the correlation between item responses. We set the inference model hidden layer size to 130 (i.e., the mean of the input layer size and $2P$) and the learning rate to $\eta = 0.005$. We set the number of IW samples $R = 5$ to demonstrate the importance-weighting approach. We fitted the full data set 100 times to assess the replicability of the parameter estimates across random seeds. Only equivalent factor solutions were compared.

We report results from the fitted model that attained the highest IW-ELBO. Figure 4 contains a heatmap of the Geomin-rotated factor loadings estimates, which fit with the expected five factor structure. Factor correlations in Table 1 also fit with the typical finding that emotional stability is negatively correlated with the other factors. Notably, fitting was fast: Mean fitting time across random seeds was 170 seconds ($SD = 47$ seconds). Further, parameter estimates were fairly stable: Across random seeds, mean loadings root-mean-square error (RMSE) was 0.018 ($SD = 0.006$), mean intercepts RMSE was 0.042 ($SD = 0.018$), and mean factor correlation RMSE was 0.028 ($SD = 0.009$).

Figure 4: Heat map of factor loadings for IPIP-FFM items. EXT = extraversion, EST = emotional stability, AGR = agreeableness, CON = conscientiousness, OPN = openness.

6.2 Simulation Studies

6.2.1 Evaluation of the Importance-Weighting Procedure

In this study, we investigate amortized IWVI’s performance in terms of parameter recovery and computational efficiency as the number of IW samples R increases (i.e., as the approximation to the marginal likelihood improves). We consider $R = 1, 5$, and 25 . The first setting uses the ELBO objective, while the latter settings use the IW-ELBO objective. Data generating loadings, intercepts, and factor correlations are rounded estimates from the IPIP-FFM example in section 6.1. We set $P = 5$, $J = 50$, and $C_j = 5$ for $j = 1, \dots, J$. Each factor loads on ten items with cross loadings set to zero to produce a perfect simple structure. To investigate IWVI’s performance as the sample size increases, we conduct 100 replications of simulation for $N = 500, 1000, 2000$, and 10000 . This leads to 12 different simulation settings for all possible combinations of R and N . We also assessed the model selection performance of the scree plot approach by plotting $-\tilde{\ell}(P)$ computed on a holdout set of 20% of observations for $P = 2, \dots, 8$ at each replication. All inference model and optimization hyperparameters from section 6.1 were reused for these analyses.

To assess parameter recovery, we computed bias for each parameter as the mean deviation of the estimated parameter from the data generating parameter across replications:

$$\text{bias}(\hat{\xi}, \xi) = \frac{1}{100} \sum_{a=1}^{100} [\hat{\xi}^{(a)} - \xi], \quad (43)$$

where $\hat{\xi}^{(a)}$ is the estimated parameter at replication a and ξ is the data generating parameter. Figure 5 uses boxplots to summarize the parameter biases separately for the factor loadings, factor correlations, and intercepts. All estimates become more accurate as N increases. Intercept and factor correlation estimates become more accurate with increasing R but exhibit slight bias across N settings. Loadings estimates are nearly unbiased and either become more accurate or obtain comparable accuracy as R increases. We also computed MSE for each parameter (i.e., by squaring the summands in equation 43). Results are summarized using boxplots in Figure 6. For each R setting, parameter MSE quickly decreases toward zero with increasing N . Increasing R tends to decrease MSE for each N setting, with factor correlation estimates demonstrating particularly large improvements as R increases.

Figure 7 contains line plots of fitting times for each simulation setting across replications. Increasing R from 1 to 5 leads to a decrease in median fitting time (around 80 seconds to around 65 seconds), while increasing R from 5 to 25 leads to a moderate increase ($R = 25$ takes around 90 to 120 seconds). Median fitting times for $R = 1$ and 5 were essentially constant as N increased. The median fitting time for $R = 25$ increased around 20 seconds from $N = 1000$ to 2000 , although absolute fitting times for this R setting were never unreasonably large. These results highlight the scalability of AMSGrad to very large data sets.

To assess factor score estimation accuracy at each replication, we first obtained expected *a posteriori* (EAP) factor score estimates for all models. For $R = 1$, we obtained EAPs by computing the approximate LV posterior mean μ for each respondent. For $R = 5$ and 25 , we obtained EAPs for each respondent by averaging $S = 50$ Monte Carlo samples drawn from $q_{\psi}^{\text{IW}}(\mathbf{x} | \mathbf{y})$ using sampling-importance-resampling (for details, see Cremer et al., 2017). After rotating the scores and applying the inversion and column permutation procedure used to compare loadings solutions, we computed the correlation between the true and estimated scores for each latent factor. Estimates were accurate: For fixed R , correlations ranged from 0.88 to 0.95 and tended to increase with increasing N . Correlations also increased slightly with increasing R and fixed N . The scree plot approach to tuning the latent dimension P appeared to perform well across simulation settings. Figure 8 presents scree plots for simulation settings where $N = 10000$, which possess sharp “elbows” at $P = 5$. Median $-\tilde{\ell}(P)$ values decreased slightly with increasing R , indicating that importance sampling helped models obtain slightly better fit to previously unseen data. Plots for other N settings were nearly identical and are not shown.

6.2.2 Comparison to MH-RM

In this study, we compare amortized IWVI to the MML estimator implemented via the MH-RM algorithm. We note that the stEM algorithm is somewhat faster than MH-RM (Zhang, Chen, & Liu, 2020) and may

Figure 5: Parameter bias for amortized importance-weighted variational inference (IWVI) computed based on 100 replications of simulation. Three settings for the number of importance-weighted (IW) samples are compared.

Figure 6: Parameter mean squared error (MSE) for amortized IWVI computed based on 100 replications of simulation.

Figure 7: Fitting times for amortized IWVI across 100 replications of simulation. For all line plots in this work, points indicate medians while error bars indicate 25% and 75% quantiles.

Figure 8: Approximate log-likelihood scree plots for amortized IWVI constructed for simulation settings with $N = 10\,000$. The “elbows” at 5 factors are marked with a horizontal dotted line.

therefore be a suitable alternative comparison method. However, given that MH-RM is relatively widely used and that stEM has only been implemented for the M2PL, we choose MH-RM for these analyses. MH-RM is implemented via the R package *mirt* (Version 1.32.1; Chalmers, 2012). Comparing the computational efficiency of the proposed approach and MH-RM is therefore fair in the sense that both *mirt* and Pytorch core functions are written in C++ and comparisons are conducted on the same computer.

We compare these methods in the high-dimensional setting where $P = 10$, $J = 100$, and $C_j = 5$ for $j = 1, \dots, J$. Data generating parameters are again rounded estimates from the IPIP-FFM example. We set the parameters for items 51-100 equal to the parameters for items 1-50. We construct the factor correlation matrix as a 10×10 block diagonal matrix with main-diagonal blocks equal to the rounded IPIP-FFM estimates and zeros elsewhere. Results of the previous simulation suggest that choosing a small $R > 1$ increases both estimation accuracy and computational efficiency, so we set $R = 5$ for these analyses. All other hyperparameters are set as in previous sections. MH-RM hyperparameters were set to the *mirt* package defaults. We conduct 100 replications of simulation for $N = 1000, 2000, 5000$, and 10000 .

Results are shown in Figures 9 and 10. Both methods obtain comparable loadings and intercepts estimates for $N \leq 5000$. MH-RM's loadings and intercepts estimates are slightly more accurate when $N = 10000$, although both methods are very accurate in this setting. Amortized IWVI produces more accurate factor correlation estimates across N settings. Additionally, IWVI is much faster than MH-RM: MH-RM's median fitting time increases from 8 minutes for $N = 1000$ to 21 minutes for $N = 10000$, whereas IWVI's median fitting time stays around 3 minutes regardless of N .

6.2.3 Comparison to CJMLE

We also compare amortized IWVI to CJMLE in the setting where N and J increase simultaneously. Chen, Li, and Zhang (2019) empirically show that the MML estimator implemented via MH-RM performs poorly when both N and J increase and that CJMLE attains much faster convergence via an alternating minimization algorithm. CJMLE is implemented in the R package *mirtjml* (Version 1.4; Zhang, Chen, & Li, 2019) and has core functions written in C++. Although CJMLE computation may be parallelized, we compare methods using a single core to ensure fairness.

We again set $P = 10$ and consider $(N, J) = (2000, 100), (10000, 200), (50000, 300)$, and $(100000, 400)$. CJMLE is only implemented for the M2PL, so we set $C_j = 2$ for $j = 1, \dots, J$. Data generating item parameters are again set by repeating the IPIP-FFM item parameters. For example, when $J = 400$, we repeat the parameters for items 1-50 seven times to get the parameters for items 51-400. Since each item needs only a single intercept, we randomly select an intercept from the four fitted IPIP-FFM intercepts for each item. The factor correlation matrix from section 6.2.2 is reused. We set $\eta = 0.005$ for settings with $J \leq 200$ but set $\eta = 0.0025$ when $J \geq 300$ because these larger models sometimes diverged otherwise. Other hyperparameters are set similarly to those in section 6.2.2.

Results are presented in Figures 11 and 12. Unlike the MML estimator, amortized IWVI performs better as N and J increase and attains comparable accuracy to CJMLE. CJMLE estimates loadings and intercepts inaccurately in the smallest (N, J) setting but is more accurate than IWVI in the highest setting (although both methods are accurate when $N \geq 50000$). Factor correlation estimates are not reported because CJMLE treats the LVs as fixed effects. Importantly, IWVI is always comparably fast or faster than CJMLE. IWVI's median fitting time is 63 seconds when $(N, J) = (2000, 100)$ and increases to just over 6 minutes when $(N, J) = (100000, 400)$, whereas CJMLE's median fitting time increases from 73 seconds to over 43 minutes in the same settings. When $(N, J) = (50000, 300)$, CJMLE sometimes took around 1500 seconds to converge rather than around 600 seconds (i.e., the median fitting time). This is possibly due to CJMLE converging to different local optima of the joint likelihood function. We note that CJMLE may achieve a significant speedup using parallel computing, although IWVI may achieve a similar speedup using a GPU.

MSE for Amortized IWVI vs. MH-RM

Figure 9: MSE for amortized importance-weighted variational inference (IWVI) and the marginal maximum likelihood estimator computed based on 100 replications of simulation. MH-RM = Metropolis-Hastings Robbins-Monro.

Figure 10: Fitting times for amortized IWVI and the marginal maximum likelihood estimator computed based on 100 replications of simulation.

Figure 11: MSE for amortized IWVI and the constrained joint maximum likelihood estimator (CJMLE) computed based on 100 replications of simulation.

Figure 12: Fitting times for amortized IWVI and CJMLE computed based on 100 replications of simulation.

7 Extensions

We now briefly consider a variety of interesting ways in which application of amortized IWVI could be expanded.

7.1 Confirmatory Item Factor Analysis

Confirmatory IFA is useful when sufficient prior theory exists to posit a specific factor structure for the items. More precisely, the measurement design may be indicated by a pre-specified $J \times P$ matrix $\mathbf{Q}$ with entries $q_{j,p} \in \{0, 1\}$ such that $q_{j,p} = 1$ if item j measures factor p (i.e., $\beta_{j,p}$ is freely estimated) and $q_{j,p} = 0$ otherwise (i.e., $\beta_{j,p}$ is set to zero). Anderson and Rubin (1957) provide sufficient conditions on $\mathbf{Q}$ to ensure the model is identified. The deep learning algorithm discussed here may be used to conduct confirmatory IFA by ensuring that factor loadings are either freely estimated or set to zero as specified in $\mathbf{Q}$ (Curi et al., 2019). In the confirmatory setting, it is also sometimes of interest to impose constraints on the factor covariance matrix Σ . Letting $\Sigma = \mathbf{L}\mathbf{L}^\top$ where $\mathbf{L}$ is a lower triangular matrix, we can estimate $\mathbf{L}$ using a hyperspherical parameterization (Pinheiro & Bates, 1996; Rapisarda et al., 2007), which enables unconstrained estimation of a variety of structured correlation matrices and has similar computational efficiency to estimating $\mathbf{L}$ directly (e.g., Ghosh et al., 2020; Tsay & Pourahmadi, 2017).

7.2 Regularized Exploratory Item Factor Analysis

Regularization has been proposed as a viable alternative to factor rotation for both exploratory linear factor analysis (e.g., Choi et al., 2010; Hirose & Konishi, 2012; Hirose & Yamamoto, 2014) and exploratory IFA (Hui et al., 2018; Sun et al., 2016). Many regularization approaches automatically rotate the factors to produce a sparse loadings structure. The regularized, amortized, importance-weighted variational estimator

is obtained by solving the optimization problem

$$\xi^* = \arg \max_{\xi} \left[\sum_{i=1}^N \text{IW-ELBO}_i - \mathcal{P}(\mathbf{B}) \right], \quad (44)$$

where $\mathbf{B}$ is a $J \times P$ factor loadings matrix whose j^{th} row is β_j and $\mathcal{P}$ is a penalty function that is potentially non-smooth and non-convex. This optimization problem may be solved using a proximal version of AMS-Grad based on the ProxGen procedure developed by Yun et al. (2020), which is guaranteed to converge to a local stationary point when mild conditions are satisfied.

7.3 Flexible Latent Density Estimation

Recent work by Monroe (2014) aims to relax the assumption that the LVs are multivariate normally distributed. An alternative approach developed recently in the deep learning literature is based on the concept of normalizing flows (NFs; Tabak & Turner, 2012; Tabak & Vanden-Eijnden, 2010). NFs apply a sequence of invertible mappings parameterized by ANNs that aims to transform a simple base density into an arbitrarily complicated density. Since the mappings are invertible, the transformed density can be explicitly evaluated via the change of variables formula. NFs scale well to high-dimensional spaces and may be used to increase the flexibility of AVI by building complicated latent prior or posterior distributions (e.g., Huang et al., 2018; Kingma et al., 2016; Rezende & Mohamed, 2015).

7.4 Nonlinear Factor Analysis

The full IWAE may be viewed as a model for nonlinear factor analysis (Yalcin & Amemiya, 2001) of the form

$$\mathbf{y}_i = g(\mathbf{x}_i) + \varepsilon_i, \quad (45)$$

for $i = 1, \dots, N$ where g is an arbitrary nonlinear function and ε_i is the i^{th} $J \times 1$ vector of errors. In the IWAE, g is approximated using an ANN. This corresponds to approximating the inverse link function between observed and latent variables while keeping the latent density fixed (Wu et al., 2020). This approach is typically less interpretable than approximating the latent density and fixing g , which provides equivalent model fit (e.g., Woods & Thissen, 2006).

8 Discussion

In this paper, we explored the suitability of an amortized importance-weighted variational inference algorithm for exploratory IFA. Numerical studies highlighted several benefits of the proposed approach. Analysis of a large-scale Big-Five personality factors data set yielded fast results that aligned with existing psychological theory across random starts. Our simulations suggested that, unlike other estimators, the amortized importance-weighted variational estimator performs comparably to state-of-the-art estimators in both the setting where the number of observations increases and in the setting where the number of items and the number of observations simultaneously increase. Amortized IWVI converges faster than existing approaches when optimized using the adaptive stochastic gradient algorithm AMSGrad, particularly with large-scale data. Factor score estimates were accurate and improved with increasing sample size. The sampling-based initialization procedures appeared to mitigate problems associated with convergence to local optima and performed comparably to the more computationally intensive initialization procedures used by MH-RM and CJMLE.

Two practical considerations not discussed here are standard errors (SEs) and missing data. Hui et al. (2017) note that for the former, approximate SEs may be obtained by evaluating the observed information matrix at the estimates $\hat{\theta}$ obtained by maximizing the IW-ELBO. Since this matrix has a block diagonal structure, it may be block-wise inverted to produce the covariance matrix from which SEs can be calculated. We note that SEs will likely be quite small for the large-scale applications considered here. Mattei and

Frellsen (2019) discuss a simple approach to handling missing-at-random data in amortized IWVI that can be easily applied to the models considered here.

The proposed approach has several limitations. The main practical difficulty we encountered was tuning the number of latent factors P . Although we tried tuning P using objective methods such as computing a pseudo-BIC and conducting 5-fold CV, these methods typically failed for large sample sizes, possibly due to log-likelihood approximation error exceeding sampling error. We therefore used subjective log-likelihood scree plots to tune P . Research is needed to develop objective criteria for selecting the latent dimension in large samples. In the meantime, the more subjective scree plot approach used here as well as approaches such as parallel analysis and retaining theoretically meaningful factors may serve as practical substitutes.

As noted by Hui et al. (2017), substantial theoretical work remains to be done to show that variational approximations produce consistent, asymptotically normal estimators and to obtain their rates of convergence. The importance sampling approach explored here provides a theoretical link between VI and MML estimation. Our simulations showed that obtaining a better approximation to the true marginal likelihood using importance sampling increases parameter estimation accuracy, although intercepts and factor correlation estimates exhibited some bias. Future theoretical work may support our empirical results by establishing amortized IWVI's asymptotic properties.

Notwithstanding these limitations, the present research suggests that amortized IWVI is a feasible and promising approach to high-dimensional exploratory IFA for psychological and educational measurement, permitting quick, accurate exploration of large-scale data sets. Additionally, amortized IWVI has many other compelling benefits that are worthy of further exploration. The rapidly developing DL literature includes a huge number of extensions that could enhance modeling and estimation in a wide variety of contexts. We view amortized IWVI as part of a progression that started with the linear models of classical test theory, transitioned to the partially nonlinear models of IRT, and is now advancing to utilize the fully nonlinear models available in machine learning. We hope our work will aid this progression by helping to spur a fruitful dialogue between the fields of machine learning and psychometrics.

References

- Anderson, T. W., & Rubin, H. (1957). Statistical inference in factor analysis. In J. Neyman (Ed.), *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability* (pp. 111-150). University of California Press.
- Asparouhov, T., & Muthén, B. (2009). Exploratory structural equation modeling. *Structural Equation Modeling: A Multidisciplinary Journal*, 16 (3), 397-438.
- Béguin, A. A., & Glas, C. A. W. (2001). MCMC estimation and some model-fit analysis of multidimensional IRT models. *Psychometrika*, 66 (4), 541-562.
- Bengio, Y. (2012). Practical recommendations for gradient-based training of deep architectures. In G. Montavon, G. Orr, & K.-R. Müller (Eds.), *Neural Networks: Tricks of the Trade* (pp. 437-478). Springer-Verlag.
- Biesanz, J. C., & West, S. G. (2004). Towards understanding assessments of the Big Five: Multitrait-multimethod analyses of convergent and discriminant validity across measurement occasion and type of observer. *Journal of Personality*, 72 (4), 845-876.
- Blei, D. M., Kucukelbir, A., & McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112 (518), 859-877.
- Bock, R. D., & Aitkin, M. (1981). Marginal maximum likelihood estimation of item parameters: Application of an EM algorithm. *Psychometrika*, 46 (4), 443-459.
- Bock, R. D., Gibbons, R., & Muraki, E. (1988). Full-information item factor analysis. *Applied Psychological Measurement*, 12 (3), 261-280.
- Bolt, D. M. (2005). Limited- and full-information estimation of item response theory models. In A. Maydeau-Olivares & J. J. McArdle (Eds.), *Contemporary psychometrics* (Chap. 2, pp. 27-72). Lawrence Erlbaum Associates, Inc.
- Bottou, L., Curtis, F. E., & Nocedal, J. (2018). Optimization methods for large-scale machine learning. *SIAM Review*, 60 (2), 223-311.
- Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., & Bengio, S. (2016). Generating sentences from a continuous space. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning* (pp. 10-21). Association for Computational Linguistics. Retrieved from <https://arxiv.org/pdf/1511.06349.pdf>.
- Burda, Y., Grosse, R. & Salakhutdinov, R. (2016). Importance weighted autoencoders. In *4th International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1509.00519.pdf>.
- Cai, L. (2010a). High-dimensional exploratory item factor analysis by a Metropolis-Hastings Robbins-Monro algorithm. *Psychometrika*, 75 (1), 33-57.
- Cai, L. (2010b). Metropolis-Hastings Robbins-Monro algorithm for confirmatory item factor analysis. *Journal of Educational and Behavioral Statistics*, 35 (3), 307-335.
- Chalmers, R. P. (2012). Mirt: A multidimensional item response theory package for the R environment. *Journal of Statistical Software*, 48 (6), 1-29.
- Chen, Y., Filho, T. S., Prudêncio, R. B. C., Diethe, T., & Flach, P. (2019). β^3 -IRT : A new item response model and its applications. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics* (pp. 1013-1021). Retrieved from <http://proceedings.mlr.press/v89/chen19b/chen19b.pdf>.
- Chen, Y., Li, X., & Zhang, S. (2019). Joint maximum likelihood estimation for high-dimensional exploratory item factor analysis. *Psychometrika*, 84 (1), 124-146.
- Chen, X., Liu, S., Sun, R., & Hong, M. (2019). On the convergence of a class of ADAM-type algorithms for non-convex optimization. In *7th International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1808.02941.pdf>.
- Cho, A. E. (2020). Gaussian variational estimation for multidimensional item response theory. [Doctoral dissertation, University of Michigan]. Deep Blue Data. Retrieved from https://deepblue.lib.umich.edu/bitstream/handle/2027.42/162939/aprilcho_1.pdf?sequence=1&isAllowed=y.

- Choi, J., Oehlert, G., & Zou, H. (2010). A penalized maximum likelihood approach to sparse factor analysis. *Statistics and Its Interface*, 3 (4), 429-436.
- Christensen, R. H. B. (2019). Cumulative link models for ordinal regression with the R package ordinal. Retrieved from https://cran.r-project.org/web/packages/ordinal/vignettes/clm_article.pdf.
- Clevert, D. A., Unterthiner, T., & Hochreiter, S. (2016). Fast and accurate deep network learning by exponential linear units (ELUs). In *4th International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1511.07289.pdf>.
- Cremer, C., Li, X., & Duvenaud, D. (2018). Inference suboptimality in variational autoencoders. In *Proceedings of the 35th International Conference on Machine Learning* (pp. 1078-1086). JMLR, Inc. and Microtome Publishing. Retrieved from <http://proceedings.mlr.press/v80/cremer18a/cremer18a.pdf>.
- Cremer, C., Morris, Q., & Duvenaud, D. (2017). Reinterpreting importance-weighted autoencoders. In *5th International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1704.02916.pdf>.
- Curi, M., Converse, G. A., Hajewski, J., & Oliveira, S. (2019). Interpretable variational autoencoders for cognitive models. *2019 International Joint Conference on Neural Networks*. <https://doi.org/10.1109/IJCNN.2019.8852333>
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2 (1), 303-314.
- Domke, J., & Sheldon, D. (2018). Importance weighting and variational inference. In *Advances in Neural Information Processing Systems 31* (pp. 4470-4479). Curran Associates, Inc. Retrieved from <https://papers.nips.cc/paper/2018/file/25db67c5657914454081c6a18e93d6dd-Paper.pdf>.
- Duchi, J. C., Hazan, E., Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12 (1), 2121-2159.
- Edwards, M. (2010). A Markov chain Monte Carlo approach to confirmatory item factor analysis. *Psychometrika*, 75 (3), 474-497.
- Erosheva, E. A., Fienberg, S. E., & Joutard, C. (2007). Describing disability through individual-level mixture models for multivariate binary data. *The Annals of Applied Statistics*, 1 (2), 502-537.
- Gershman, S., & Goodman, N. (2014). Amortized inference in probabilistic reasoning. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, (Vol. 1, pp. 517-522). Retrieved from https://escholarship.org/content/qt34j1h7k5/qt34j1h7k5_noSplash_8e5b24dd056d61b53b1170a1861e49d1.pdf?t=op9xkp.
- Ghosh, R. P., Mallick, B., & Pourahmadi, M. (2020). Bayesian estimation of correlation matrices of longitudinal data. *Bayesian Analysis*, 1-20. <https://doi.org/10.1214/20-ba1237>
- Glorot, X., & Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. *Journal of Machine Learning Research*, 9 (1), 249-256.
- Goldberg, L. R. (1992). The development of markers for the Big-Five factor structure. *Psychological Assessment*, 4 (1), 26-42.
- Goldberg, L. R., Johnson, J. A., Eber, H. W., Hogan, R., Ashton, M. C., Cloninger, C. R., & Gough, H. G. (2006). The international personality item pool and the future of public-domain personality measures. *Journal of Research in Personality*, 40 (1), 84-96.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *2015 IEEE International Conference on Computer Vision* (pp. 1026-1034). <https://doi.org/10.1109/ICCV.2015.123>
- Heaton, J. (2008). *Introduction to Neural Networks for Java* (2nd ed.). Heaton Research, Inc.
- Hirose, K., & Konishi, S. (2012). Variable selection via the weighted group lasso for factor analysis models. *The Canadian Journal of Statistics*, 40 (2), 345-361.
- Huang, C. W., Krueger, D., Lacoste, A., & Courville, A. (2018). Neural autoregressive flows. In *Proceedings of the 35th International Conference on Machine Learning* (pp. 2078-2087). Retrieved from <http://proceedings.mlr.press/v80/huang18d/huang18d.pdf>.

- Huber, P., Ronchetti, E., & Victoria-Feser, M.-P. (2004). Estimation of generalized linear latent variable models. *Journal of the Royal Statistical Society – Series B*, 66 (4), 893-908.
- Hui, F. K. C., Tanaka, E., & Warton, D. I. (2018). Order selection and sparsity in latent variable models via the ordered factor LASSO. *Biometrics*, 74 (4), 1311-1319.
- Hui, F. K. C., Warton, D. I., Ormerod, J. T., Haapaniemi, V., & Taskinen, S. (2017). Variational approximations for generalized linear latent variable models. *Journal of Computational and Graphical Statistics*, 26 (1), 35-43.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., & Saul, L. K. (1998). *Learning in Graphical Models*, 37 (1), 183-233.
- Jöreskog, K. G., & Moustaki, I. (2001). Factor analysis of ordinal variables: A comparison of three approaches. *Multivariate Behavioral Research*, 36 (3), 347-387.
- Keskar, N. S., Mudigere, D., Nocedal, J., Smelyanskiy, M., & Tang, P. T. P. (2017). On large-batch training for deep learning: Generalization gap and sharp minima. In *5th International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1609.04836.pdf>.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., & Welling, M. (2016). Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems 31* (pp. 4743-4751). Curran Associates, Inc. Retrieved from <https://papers.nips.cc/paper/2016/file/ddeebdeefdb7e7e7a697e1c3e3d8ef54-Paper.pdf>.
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. In *2nd International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1312.6114.pdf>.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature Methods*, 521 (1), 436-444. <https://doi.org/10.1038/nmeth.3707>
- Lehmann, E. L., & Casella, G. (1998). *Theory of point estimation*. Springer-Verlag.
- Linnainmaa, S. (1970). The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors. [Unpublished master's thesis (in Finnish)]. University of Helsinki.
- Lorenzo-Seva, U., & ten Berge, J. M. (2006). Tucker's congruence coefficient as a meaningful index of factor similarity. *Methodology: European Journal of Research Methods for The Behavioral and Social Sciences*, 2 (2), 57-64.
- MacCallum, R. C., Widaman, K. F., Zhang, S., & Hong, S. (1999). Sample size in factor analysis. *Psychological Methods*, 4 (1), 84-99.
- Mattei, P.-A., & Frellsen, J. (2019). MIWAE: Deep generative modelling and imputation of incomplete data. In *Proceedings of the 36th International Conference on Machine Learning*, (pp. 4413-4423). Retrieved from <http://proceedings.mlr.press/v97/mattei19a/mattei19a.pdf>.
- McKinley, R., & Reckase, M. (1983). *An extension of the two-parameter logistic model to the multidimensional latent space* (Research Report ONR83-2). The American College Testing Program.
- McMahan, H. B., & Streeter, M. (2010). Adaptive bound optimization for online convex optimization. In A. T. Kalai & M. Mohr (Eds.), *The 23rd Conference on Learning Theory* (pp. 244-256). Retrieved from <http://www.learningtheory.org/colt2010/conference-website/papers/COLT2010proceedings.pdf>.
- Meng, X.-L., & Schilling, S. (1996). Fitting full-information item factor models and an empirical investigation of bridge sampling. *Journal of the American Statistical Association*, 91 (435), 1254-1267.
- Monroe, S. L. (2014). Multidimensional item factor analysis with semi-nonparametric latent densities. [Unpublished doctoral dissertation]. University of California.
- Muthén, B. (1978). Contributions to factor analysis of dichotomous variables. *Psychometrika*, 43 (4), 551-560.
- Muthén, B. (1984). A general structural equation model with dichotomous, ordered categorical, and continuous latent variable indicators. *Psychometrika*, 49 (1), 115-132.
- Natesan, P., Nandakumar, R., Minka, T., & Rubright, J. D. (2016). Bayesian prior choice in IRT estimation using MCMC and variational Bayes. *Frontiers in Psychology*, 7 (1). <https://doi.org/10.3389/fpsyg.2016.01422>

- Nemirovski, A., Juditsky, A., Lan, G. & Shapiro, A. (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19 (4), 1574-1609.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Demaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., ...Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32* (pp. 8024–8035). Curran Associates, Inc. Retrieved from <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Pinheiro, J. C., & Bates, D. M. (1996). Unconstrained parametrizations for variance-covariance matrices. *Statistics and Computing*, 6 (3), 289-296.
- Rabe-Hesketh, S., Skrondal, A., & Pickles, A. (2005). Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics*, 128 (2), 301-323.
- Rainforth, T., Kosiorek, A. R., Le, T. A., Maddison, C. J., Igl, M., Wood, F., & Teh, Y. W. (2018). Tighter variational bounds are not necessarily better. In *Proceedings of the 35th International Conference on Machine Learning* (Vol. 80, pp. 4277-4285). Retrieved from <http://proceedings.mlr.press/v80/rainforth18b/rainforth18b.pdf>.
- Rapisarda, F., Brigo, D., & Mercurio, F. (2007). Parameterizing correlations: A geometric interpretation. *IMA Journal of Management Mathematics*, 18 (1), 55–73. <https://doi.org/10.1093/imaman/dpl010>
- Reckase, M. D. (2009). *Multidimensional item response theory*. Springer-Verlag.
- Reddi, S. J., Kale, S., & Kumar, S. (2018). On the convergence of ADAM and beyond. In *6th International Conference on Learning Representations. ICLR*. Retrieved from <https://arxiv.org/pdf/1904.09237.pdf>.
- Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning* (pp.1278-1286). Retrieved from <http://proceedings.mlr.press/v32/rezende14.pdf>.
- Rezende, D. J., & Mohamed, S. (2015). Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning* (pp. 530-1538). Retrieved from <http://proceedings.mlr.press/v37/rezende15.pdf>.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, 22 (3), 400-407.
- Samejima, F. (1969). Estimation of latent ability using a response pattern of graded scores. *Psychometrika*, 35 (1), 139.
- Schilling, R., & Bock, D. (2005). High-dimensional maximum marginal likelihood item factor analysis by adaptive quadrature. *Psychometrika*, 70 (3), 533-555.
- Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., & Winther, O. (2016). Ladder variational autoencoders. In *Advances in Neural Information Processing Systems* (pp. 3745-3753). Curran Associates, Inc. Retrieved from <https://papers.nips.cc/paper/2016/file/6ae07dcb33ec3b7c814df797cbda0f87-Paper.pdf>.
- Song, X., & Lee, S. (2005). A multivariate probit latent variable model for analyzing dichotomous responses. *Statistica Sinica*, 15 (3), 45-64.
- Spall, J. C. (2003). Introduction to stochastic search and optimization: estimation, simulation, and control. John Wiley & Sons, Inc.
- Staib, M., Reddi, S., Kale, S., Kumar, S., & Sra, S. (2019). Escaping saddle points with adaptive gradient methods. In *Proceedings of the 36th International Conference on Machine Learning* (pp. 5956-5965). Retrieved from <http://proceedings.mlr.press/v97/staib19a/staib19a.pdf>.
- Sun, J., Chen, Y., Liu, J., Ying, Z., & Xin, T. Latent variable selection for multidimensional item response theory models via L1 regularization. *Psychometrika*, 81 (4), 921-939.
- Tabak, E. G., & Turner, C. V. (2012). A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66 (2), 145-164.
- Tabak, E. G., & Vanden-Eijnden, E. (2010). Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8 (1), 217-233.

- Tsay, R. S., & Pourahmadi, M. (2017). Modelling structured correlation matrices. *Biometrika*, 104 (1), 237–242. <https://doi.org/10.1093/biomet/asw061>
- Tucker, G., Lawson, D., Gu, S., & Maddison, C. J. (2019). Doubly reparameterized gradient estimators for Monte Carlo objectives. In *7th International Conference on Learning Representations*. ICLR. Retrieved from <https://arxiv.org/pdf/1810.04152.pdf>.
- Wainwright, M. J., & Jordan, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1 (1-2), 1-305.
- Wirth, R. J., & Edwards, M. C. (2007). Item factor analysis: Current approaches and future directions. *Psychological Methods*, 12 (1), 58–79.
- Woods, C. M., & Thissen, D. (2006). Item response theory with estimation of the latent population distribution using spline-based densities. *Psychometrika*, 71 (2), 281-301.
- Wu, M., Davis, R. L., Domingue, B. W., Piech, C., & Goodman, N. (2020). Variational item response theory: Fast, accurate, and expressive. In A. N. Rafferty, J. Whitehill, C. Romero, & V. Cavalli-Sforza (Eds.), *Proceedings of the 13th International Conference on Educational Data Mining 2020* (pp. 257-268). Retrieved from <https://educationaldatamining.org/files/conferences/EDM2020/EDM2020Proceedings.pdf>.
- Yalcin, I., & Amemiya, Y. (2001). Nonlinear factor analysis as a statistical method. *Statistical Science*, 16 (3), 275-294.
- Yates, A. (1988). *Multivariate exploratory data analysis: A perspective on exploratory factor analysis*. State University of New York Press.
- Yun, J., Lozano, A. C., & Yang, E. (2020). A general family of stochastic proximal gradient methods for deep learning. *arXiv preprint*. Retrieved from <https://arxiv.org/pdf/2007.07484.pdf>.
- Zhang, C., Butepage, J., Kjellstrom, H., & Mandt, S. (2019). Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41 (8), 2008–2026.
- Zhang, S., Chen, Y., & Li, X. (2019). mirtjml [Computer software]. Retrieved from <https://cran.r-project.org/web/packages/mirtjml/index.html>.
- Zhang, H., Chen, Y., & Li, X. (2020). A note on exploratory item factor analysis by singular value decomposition. *Psychometrika*, 1-15.
- Zhang, S., Chen, Y., & Liu, Y. (2020). An improved stochastic EM algorithm for large-scale full-information item factor analysis. *British Journal of Mathematical and Statistical Psychology*, 73 (1), 44-71.
- Zhou, D., Tang, Y., Yang, Z., Cao, Y., & Gu, Q. (2018). On the convergence of adaptive gradient methods for nonconvex optimization. *arXiv preprint*. Retrieved from <https://arxiv.org/pdf/1808.05671.pdf>.