

Sparse Factor Analysis for Learning and Content Analytics

Andrew S. Lan*

MR.LAN@SPARFA.COM

Andrew E. Waters

WATERS@SPARFA.COM

Dept. Electrical and Computer Engineering

Rice University

Houston, TX 77005, USA

Christoph Studer

STUDER@SPARFA.COM

School of Electrical and Computer Engineering

Cornell University

Ithaca, NY 14853, USA

Richard G. Baraniuk

RICHB@SPARFA.COM

Dept. Electrical and Computer Engineering

Rice University

Houston, TX 77005, USA

Editor: Francis Bach

Abstract

We develop a new model and algorithms for machine learning-based *learning analytics*, which estimate a learner's knowledge of the concepts underlying a domain, and *content analytics*, which estimate the relationships among a collection of questions and those concepts. Our model represents the probability that a learner provides the correct response to a question in terms of three factors: their understanding of a set of underlying concepts, the concepts involved in each question, and each question's intrinsic difficulty. We estimate these factors given the graded responses to a collection of questions. The underlying estimation problem is ill-posed in general, especially when only a subset of the questions are answered. The key observation that enables a well-posed solution is the fact that typical educational domains of interest involve only a small number of key concepts. Leveraging this observation, we develop both a bi-convex maximum-likelihood-based solution and a Bayesian solution to the resulting *SPARse Factor Analysis* (SPARFA) problem. We also incorporate user-defined tags on questions to facilitate the interpretability of the estimated factors. Experiments with synthetic and real-world data demonstrate the efficacy of our approach. Finally, we make a connection between SPARFA and noisy, binary-valued (1-bit) dictionary learning that is of independent interest.

Keywords: factor analysis, sparse probit regression, sparse logistic regression, Bayesian latent factor analysis, personalized learning

1. Introduction

Textbooks, lectures, and homework assignments were the answer to the main educational challenges of the 19th century, but they are the main bottleneck of the 21st century. Today's textbooks are static, linearly organized, time-consuming to develop, soon out-of-date,

*. First authorship determined by coin toss

and expensive. Lectures remain a primarily passive experience of copying down what an instructor says and writes on a board (or projects on a screen). Homework assignments that are not graded for weeks provide poor feedback to learners (e.g., students) on their learning progress. Even more importantly, today’s courses provide only a “one-size-fits-all” learning experience that does not cater to the background, interests, and goals of individual learners.

1.1 The Promise of Personalized Learning

We envision a world where access to high-quality, personally tailored educational experiences is affordable to all of the world’s learners. The key to reaching this goal is to integrate textbooks, lectures, and homework assignments into a *personalized learning system* (PLS) that closes the learning feedback loop by (i) continuously monitoring and analyzing learner interactions with learning resources in order to assess their learning progress and (ii) providing timely remediation, enrichment, or practice based on that analysis. See Linden and Glas (2000), Murray et al. (2004), Stamper et al. (2007), Rafferty et al. (2011), Li et al. (2011), and Knewton (2012) for various visions and examples.

Some progress has been made over the past few decades on personalized learning; see, for example, the sizable literature on *intelligent tutoring systems* discussed in Psofka et al. (1988). To date, the lion’s share of fielded, intelligent tutors have been rule-based systems that are hard-coded by domain experts to give learners feedback for pre-defined scenarios (e.g., Koedinger et al. 1997, Brusilovsky and Peylo 2003, VanLehn et al. 2005, and Butz et al. 2006). The specificity of such systems is counterbalanced by their high development cost in terms of both time and money, which has limited their scalability and impact in practice.

In a fresh direction, recent progress has been made on applying *machine learning* algorithms to mine learner interaction data and educational content (see the overview articles by Romero and Ventura 2007 and Baker and Yacef 2009). In contrast to rule-based approaches, machine learning-based PLSs promise to be rapid and inexpensive to deploy, which will enhance their scalability and impact. Indeed, the dawning age of “big data” provides new opportunities to build PLSs based on data rather than rules. We conceptualize the architecture of a generic machine learning-based PLS to have three interlocking components:

- *Learning analytics*: Algorithms that estimate what each learner does and does not understand based on data obtained from tracking their interactions with learning content.
- *Content analytics*: Algorithms that organize learning content such as text, video, simulations, questions, and feedback hints.
- *Scheduling*: Algorithms that use the results of learning and content analytics to suggest to each learner at each moment what they should be doing in order to maximize their learning outcomes, in effect closing the learning feedback loop.

1.2 Sparse Factor Analysis (SPARFA)

In this paper, we develop a new model and a suite of algorithms for joint machine learning-based *learning analytics* and *content analytics*. Our model (developed in Section 2) represents the probability that a learner provides the correct response to a given question in

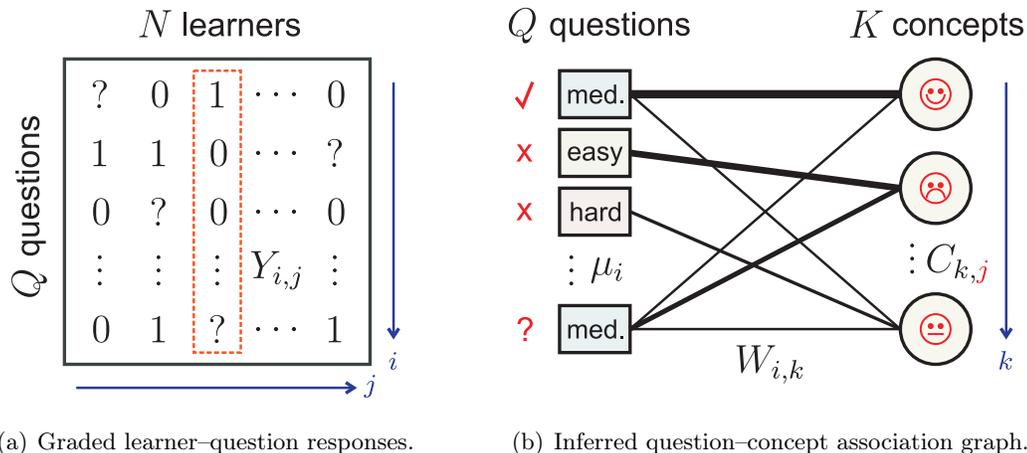


Figure 1: (a) The SPARFA framework processes a (potentially incomplete) binary-valued data set of graded learner-question responses to (b) estimate the underlying questions-concept association graph and the abstract conceptual knowledge of each learner (illustrated here by smiley faces for learner $j = 3$, the column in (a) selected by the red dashed box).

terms of three factors: their knowledge of the underlying concepts, the concepts involved in each question, and each question’s intrinsic difficulty.

Figure 1 provides a graphical depiction of our approach. As shown in Figure 1(a), we are provided with data related to the correctness of the learners’ responses to a collection of questions. We encode these graded responses in a “gradebook,” a source of information commonly used in the context of classical test theory (Norvick 1966). Specifically, the “gradebook” is a matrix with entry $Y_{i,j} = 1$ or 0 depending on whether learner j answers question i correctly or incorrectly, respectively. Question marks correspond to incomplete data due to unanswered or unassigned questions. Working left-to-right in Figure 1(b), we assume that the collection of questions (rectangles) is related to a small number of abstract concepts (circles) by a bipartite graph, where the edge weight $W_{i,k}$ indicates the degree to which question i involves concept k . We also assume that question i has intrinsic difficulty μ_i . Denoting learner j ’s knowledge of concept k by $C_{k,j}$, we calculate the probabilities that the learners answer the questions correctly in terms of $\mathbf{WC} + \mathbf{M}$, where \mathbf{W} and \mathbf{C} are matrix versions of $W_{i,k}$ and $C_{k,j}$, respectively, and \mathbf{M} is a matrix containing the intrinsic question difficulty μ_i on row i . We transform the probability of a correct answer to an actual 1/0 correctness via a standard probit or logit link function (see Rasmussen and Williams 2006).

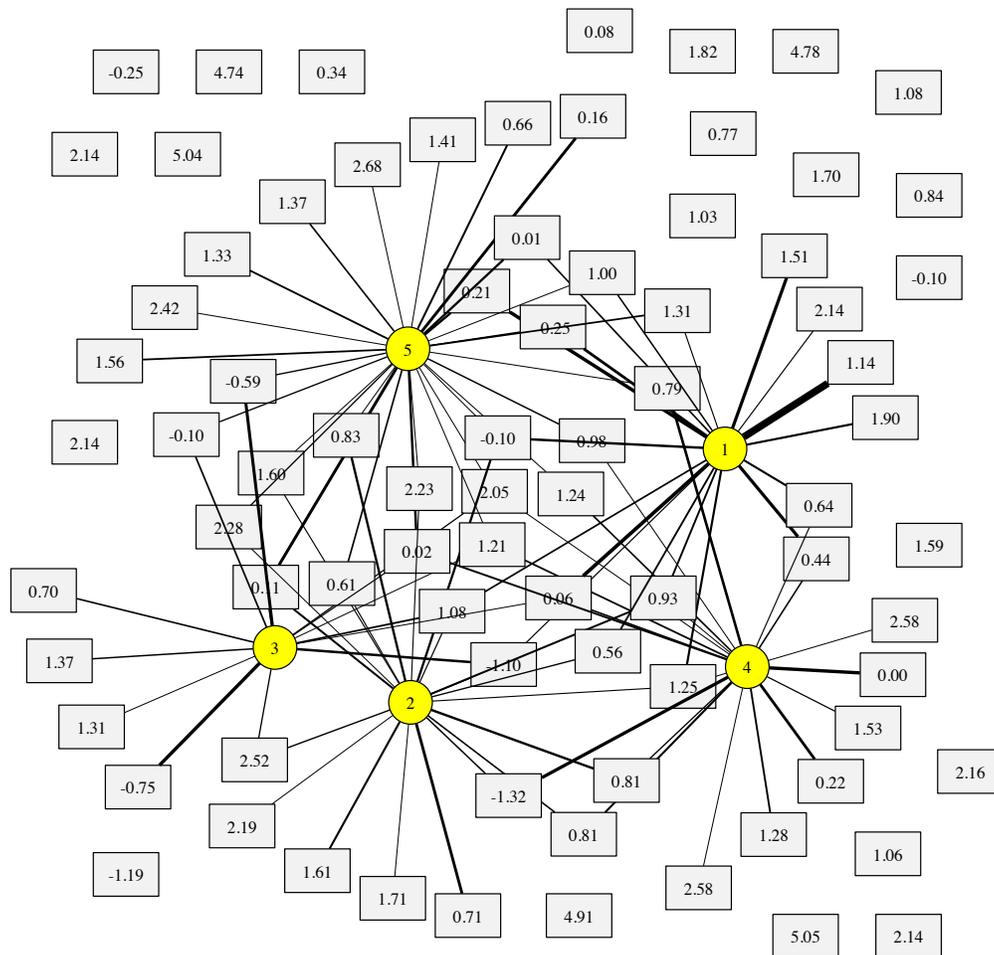
Armed with this model and given incomplete observations of the graded learner-question responses $Y_{i,j}$, our goal is to estimate the factors \mathbf{W} , \mathbf{C} , and \mathbf{M} . Such a *factor-analysis problem* is ill-posed in general, especially when each learner answers only a small subset of the collection of questions (see Harman 1976 for a factor analysis overview). Our first key observation that enables a well-posed solution is the fact that typical educational domains of interest involve only a small number of key concepts (i.e., we have $K \ll N, Q$ in Figure 1). Consequently, \mathbf{W} becomes a tall, narrow $Q \times K$ matrix that relates the questions to a small set of abstract concepts, while \mathbf{C} becomes a short, wide $K \times N$ matrix that relates learner knowledge to that same small set of abstract concepts. Note that the concepts are

“abstract” in that they will be estimated from the data rather than dictated by a subject matter expert. Our second key observation is that each question involves only a small subset of the abstract concepts. Consequently, the matrix \mathbf{W} is sparsely populated. Our third observation is that the entries of \mathbf{W} should be non-negative, since we postulate that having strong concept knowledge should never hurt a learner’s chances to answer questions correctly. This constraint on \mathbf{W} ensures that large positive values in \mathbf{C} represent strong knowledge of the associated abstract concepts, which is crucial for a PLS to generate human-interpretable feedback to learners on their strengths and weaknesses.

Leveraging these observations, we propose a suite of new algorithms for solving the *SPARse Factor Analysis* (SPARFA) problem. Section 3 develops SPARFA-M, a matrix factorization method which uses an efficient bi-convex optimization approach to produce point estimates of the factors. Section 4 develops SPARFA-B, a Bayesian factor analysis method to produce posterior distributions of the factors. SPARFA-M is computationally efficient and scales to large-scale applications, while SPARFA-B is more computationally intensive but also provide richer statistical information on the latent factors. Since the concepts are abstract mathematical quantities estimated by the SPARFA algorithms, we develop a *post-processing* step in Section 5 to facilitate interpretation of the estimated latent concepts by associating user-defined tags for each question with each abstract concept.

In Section 6, we report on a range of experiments with a variety of synthetic and real-world data that demonstrate the wealth of information provided by the estimates of \mathbf{W} , \mathbf{C} and \mathbf{M} . As an example, Figure 2 provides the results for a data set collected from learners using STEMscopes (2012), a science curriculum platform. The data set consists of 145 Grade 8 learners from a single school district answering a manually tagged set of 80 questions on Earth science; only 13.5% of all graded learner–question responses were observed. We applied the SPARFA-B algorithm to retrieve the factors \mathbf{W} , \mathbf{C} , and \mathbf{M} using 5 latent concepts. The resulting sparse matrix \mathbf{W} is displayed as a bipartite graph in Figure 2(a); circles denote the abstract concepts and boxes denote questions. Each question box is labeled with its estimated intrinsic difficulty μ_i , with large positive values denoting easy questions. Links between the concept and question nodes represent the active (non-zero) entries of \mathbf{W} , with thicker links denoting larger values $W_{i,k}$. Unconnected questions are those for which no concept explained the learners’ answer pattern; such questions typically have either very low or very high intrinsic difficulty, resulting in nearly all learners answering them correctly or incorrectly. The tags provided in Figure 2(b) enable human-readable interpretability of the estimated abstract concepts.

We envision a range of potential learning and content analytics applications for the SPARFA framework that go far beyond the standard practice of merely forming column sums of the “gradebook” matrix (with entries $Y_{i,j}$) to arrive at a final scalar numerical score for each learner (which is then often further quantized to a letter grade on a 5-point scale). Each column of the estimated \mathbf{C} matrix can be interpreted as a measure of the corresponding learner’s knowledge about the abstract concepts. Low values indicate concepts ripe for remediation, while high values indicate concepts ripe for enrichment. The sparse graph stemming from the estimated \mathbf{W} matrix automatically groups questions into similar types based on their concept association; this graph makes it straightforward to find a set of questions similar to a given target question. Finally, the estimated \mathbf{M} matrix (with entries μ_i on each row) provides an estimate of each question’s intrinsic difficulty. This



(a) Inferred question–concept association graph.

Concept 1		Concept 2		Concept 3	
Changes to land	(45%)	Evidence of the past	(74%)	Alternative energy	(76%)
Properties of soil	(28%)	Mixtures and solutions	(14%)	Environmental changes	(19%)
Uses of energy	(27%)	Environmental changes	(12%)	Changes from heat	(5%)
Concept 4		Concept 5			
Properties of soil	(77%)	Formulation of fossil fuels	(54%)		
Environmental changes	(17%)	Mixtures and solutions	(28%)		
Classifying matter	(6%)	Uses of energy	(18%)		

(b) Most important tags and relative weights for the estimated concepts.

Figure 2: (a) Sparse question–concept association graph and (b) most important tags associated with each concept for Grade 8 Earth science with $N = 135$ learners answering $Q = 80$ questions. Only 13.5% of all graded learner–question responses were observed.

property enables an instructor to assign questions in an orderly fashion as well as to prune out potentially problematic questions that are either too hard, too easy, too confusing, or unrelated to the concepts underlying the collection of questions.

In Section 7, we provide an overview of related work on machine learning-based personalized learning, and we conclude in Section 8. All proofs are relegated to three Appendices.

2. Statistical Model for Learning and Content Analytics

Our approach to learning and content analytics is based on a new statistical model that encodes the probability that a learner will answer a given question correctly in terms of three factors: (i) the learner’s knowledge of a set of latent, abstract concepts, (ii) how the question is related to each concept, and (iii) the intrinsic difficulty of the question.

2.1 Model for Graded Learner Response Data

Let N denote the total number of learners, Q the total number of questions, and K the number of latent abstract concepts. We define $C_{k,j}$ as the *concept knowledge* of learner j on concept k , with large positive values of $C_{k,j}$ corresponding to a better chance of success on questions related to concept k . Stack these values into the column vector $\mathbf{c}_j \in \mathbb{R}^K$, $j \in \{1, \dots, N\}$ and the $K \times N$ matrix $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_N]$. We further define $W_{i,k}$ as the *question–concept association* of question i with respect to concept k , with larger values denoting stronger involvement of the concept. Stack these values into the column vector $\bar{\mathbf{w}}_i \in \mathbb{R}^K$, $i \in \{1, \dots, Q\}$ and the $Q \times K$ matrix $\mathbf{W} = [\bar{\mathbf{w}}_1, \dots, \bar{\mathbf{w}}_Q]^T$. Finally, we define the scalar $\mu_i \in \mathbb{R}$ as the *intrinsic difficulty* of question i , with larger values representing easier questions. Stack these values into the column vector $\boldsymbol{\mu}$ and form the $Q \times N$ matrix $\mathbf{M} = \boldsymbol{\mu} \mathbf{1}_{1 \times N}$ as the product of $\boldsymbol{\mu} = [\mu_1, \dots, \mu_Q]^T$ with the N -dimensional all-ones row vector $\mathbf{1}_{1 \times N}$.

Given these definitions, we propose the following model for the binary-valued graded response variable $Y_{i,j} \in \{0, 1\}$ for learner j on question i , with 1 representing a correct response and 0 an incorrect response:

$$\begin{aligned} Z_{i,j} &= \bar{\mathbf{w}}_i^T \mathbf{c}_j + \mu_i, & \forall i, j, \\ Y_{i,j} &\sim \text{Ber}(\Phi(Z_{i,j})), & (i, j) \in \Omega_{\text{obs}}. \end{aligned} \tag{1}$$

Here, $\text{Ber}(z)$ designates a Bernoulli distribution with success probability z , and $\Phi(z)$ denotes an *inverse link function*¹ that maps a real value z to the success probability of a binary random variable. Thus, the slack variable $\Phi(Z_{i,j}) \in [0, 1]$ governs the probability of learner j answering question i correctly.

The set $\Omega_{\text{obs}} \subseteq \{1, \dots, Q\} \times \{1, \dots, N\}$ in (1) contains the indices associated with the observed graded learner response data. Hence, our framework is able to handle the case of incomplete or missing data (e.g., when the learners do not answer all of the questions).²

1. Inverse link functions are often called *response functions* in the generalized linear models literature (see, e.g., Guisan et al. 2002).

2. Two common situations lead to missing learner response data. First, a learner might not attempt a question because it was not assigned or available to them. In this case, we simply exclude their response from Ω_{obs} . Second, a learner might not attempt a question because it was assigned to them but was too difficult. In this case, we treat their response as incorrect, as is typical in standard testing settings.

Stack the values $Y_{i,j}$ and $Z_{i,j}$ into the $Q \times N$ matrices \mathbf{Y} and \mathbf{Z} , respectively. We can conveniently rewrite (1) in matrix form as

$$Y_{i,j} \sim \text{Ber}(\Phi(Z_{i,j})), (i, j) \in \Omega_{\text{obs}} \quad \text{with} \quad \mathbf{Z} = \mathbf{WC} + \mathbf{M}. \quad (2)$$

In this paper, we focus on the two most commonly used link functions in the machine learning literature. The *inverse probit* function is defined as

$$\Phi_{\text{pro}}(x) = \int_{-\infty}^x \mathcal{N}(t) dt = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt, \quad (3)$$

where $\mathcal{N}(t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$ is the probability density function (PDF) of the standard normal distribution (with mean zero and variance one). The *inverse logit* link function is defined as

$$\Phi_{\text{log}}(x) = \frac{1}{1 + e^{-x}}. \quad (4)$$

As noted in the Introduction, \mathbf{C} , \mathbf{W} , and $\boldsymbol{\mu}$ (or equivalently, \mathbf{M}) have natural interpretations in real education settings. Column j of \mathbf{C} can be interpreted as a measure of learner j 's knowledge about the abstract concepts, with larger $C_{k,j}$ values implying more knowledge. The non-zero entries in \mathbf{W} can be used to visualize the connectivity between concepts and questions (see Figure 1(b) for an example), with larger $W_{i,k}$ values implying stronger ties between question i and concept k . The values of $\boldsymbol{\mu}$ contains estimates of each question's intrinsic difficulty.

2.2 Joint Estimation of Concept Knowledge and Question–Concept Association

Given a (possibly partially observed) matrix of graded learner response data \mathbf{Y} , we aim to estimate the learner concept knowledge matrix \mathbf{C} , the question–concept association matrix \mathbf{W} , and the question intrinsic difficulty vector $\boldsymbol{\mu}$. In practice, the latent factors \mathbf{W} and \mathbf{C} , and the vector $\boldsymbol{\mu}$ will contain many more unknowns than we have observations in \mathbf{Y} ; hence, estimating \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ is, in general, an ill-posed inverse problem. The situation is further exacerbated if many entries in \mathbf{Y} are unobserved.

To regularize this inverse problem, prevent over-fitting, improve identifiability,³ and enhance interpretability of the entries in \mathbf{C} and \mathbf{W} , we appeal to the following three observations regarding education that are reasonable for typical exam, homework, and practice questions at all levels. We will exploit these observations extensively in the sequel as fundamental assumptions:

- (A1) *Low-dimensionality*: The number of latent, abstract concepts K is small relative to both the number of learners N and the number of questions Q . This implies that the questions are redundant and that the learners' graded responses live in a low-dimensional space. The parameter K dictates the concept *granularity*. Small K extracts just a few general, broad concepts, whereas large K extracts more specific and detailed concepts.⁴

3. If $\mathbf{Z} = \mathbf{WC}$, then for any orthonormal matrix \mathbf{H} with $\mathbf{H}^T \mathbf{H} = \mathbf{I}$, we have $\mathbf{Z} = \mathbf{WH}^T \mathbf{HC} = \widetilde{\mathbf{W}} \widetilde{\mathbf{C}}$. Hence, the estimation of \mathbf{W} and \mathbf{C} is, in general, non-unique up to a unitary matrix rotation.

4. Standard techniques like cross-validation (Hastie et al. 2010) can be used to select K . We provide the corresponding details in Section 6.3.

- (A2) *Sparsity*: Each question should be associated with only a small subset of the concepts in the domain of the course/assessment. In other words, we assume that the matrix \mathbf{W} is sparsely populated, i.e., contains mostly zero entries.
- (A3) *Non-negativity*: A learner’s knowledge of a given concept does not negatively affect their probability of correctly answering a given question, i.e., knowledge of a concept is not “harmful.” In other words, the entries of \mathbf{W} are non-negative, which provides a natural interpretation for the entries in \mathbf{C} : Large values $C_{k,j}$ indicate strong knowledge of the corresponding concept, whereas negative values indicate weak knowledge.

In practice, N can be larger than Q and vice versa, and hence, we do not impose any additional assumptions on their values. Assumptions (A2) and (A3) impose sparsity and non-negativity constraints on \mathbf{W} . Since these assumptions are likely to be violated under arbitrary unitary transforms of the factors, they help alleviate several well-known identifiability problems that arise in factor analysis.

We will refer to the problem of estimating \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$, given the observations \mathbf{Y} , under the assumptions (A1)–(A3) as the *SPARse Factor Analysis* (SPARFA) problem. We now develop two complementary algorithms to solve the SPARFA problem. Next, we detail SPARFA-M and SPARFA-B, a matrix-factorization approach and a Bayesian approach to estimate the quantities of interest.

3. SPARFA-M: Maximum Likelihood-based Sparse Factor Analysis

Our first algorithm, SPARFA-M, solves the SPARFA problem using maximum-likelihood-based probit or logistic regression.

3.1 Problem Formulation

To estimate \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$, we maximize the likelihood of the observed data $Y_{i,j}$, $(i, j) \in \Omega_{\text{obs}}$

$$p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) = \Phi(\bar{\mathbf{w}}_i^T \mathbf{c}_j)^{Y_{i,j}} (1 - \Phi(\bar{\mathbf{w}}_i^T \mathbf{c}_j))^{1-Y_{i,j}}$$

given \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ and subject to the assumptions (A1), (A2), and (A3) from Section 2.2. This likelihood yields the following optimization problem:

$$(P^*) \begin{cases} \underset{\mathbf{W}, \mathbf{C}}{\text{maximize}} & \sum_{(i,j) \in \Omega_{\text{obs}}} \log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) \\ \text{subject to} & \|\bar{\mathbf{w}}_i\|_0 \leq s \ \forall i, \quad \|\bar{\mathbf{w}}_i\|_2 \leq \kappa \ \forall i, \quad W_{i,k} \geq 0 \ \forall i, k, \quad \|\mathbf{C}\|_F = \xi. \end{cases}$$

Let us take a quick tour of the problem (P*) and its constraints. The intrinsic difficulty vector $\boldsymbol{\mu}$ is incorporated as an additional column of \mathbf{W} , and \mathbf{C} is augmented with an all-ones row accordingly. We impose sparsity on each vector $\bar{\mathbf{w}}_i$ to comply with (A2) by limiting its maximum number of nonzero coefficients using the constraint $\|\bar{\mathbf{w}}_i\|_0 \leq s$; here $\|\mathbf{a}\|_0$ counts the number of non-zero entries in the vector \mathbf{a} . The ℓ_2 -norm constraint on each vector $\bar{\mathbf{w}}_i$ with $\kappa > 0$ is required for our convergence proof below. We enforce non-negativity on each entry $W_{i,k}$ to comply with (A3). Finally, we normalize the Frobenius norm of the concept knowledge matrix \mathbf{C} to a given $\xi > 0$ to suppress arbitrary scalings between the entries in both matrices \mathbf{W} and \mathbf{C} .

Unfortunately, optimizing over the sparsity constraints $\|\bar{\mathbf{w}}_i\|_0 \leq s$ requires a combinatorial search over all K -dimensional support sets having no more than s non-zero entries. Hence, (P*) cannot be solved efficiently in practice for the typically large problem sizes of interest. In order to arrive at an optimization problem that can be solved with a reasonable computational complexity, we *relax* the sparsity constraints $\|\bar{\mathbf{w}}_i\|_0 \leq s$ in (P*) to ℓ_1 -norm constraints as in Chen et al. (1998) and move them, the ℓ_2 -norm constraints, and the Frobenius norm constraint, into the objective function via Lagrange multipliers:

$$(P) \quad \underset{\mathbf{W}, \mathbf{C}: W_{i,k} \geq 0 \forall i,k}{\text{minimize}} \quad \sum_{(i,j) \in \Omega_{\text{obs}}} -\log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) + \lambda \sum_i \|\bar{\mathbf{w}}_i\|_1 + \frac{\mu}{2} \sum_i \|\bar{\mathbf{w}}_i\|_2^2 + \frac{\gamma}{2} \|\mathbf{C}\|_F^2.$$

The first regularization term $\lambda \sum_i \|\bar{\mathbf{w}}_i\|_1$ induces sparsity on each vector $\bar{\mathbf{w}}_i$, with the single parameter $\lambda > 0$ controlling the sparsity level. Since one can arbitrarily increase the scale of the vectors $\bar{\mathbf{w}}_i$ while decreasing the scale of the vectors \mathbf{c}_j accordingly (and vice versa) without changing the likelihood, we gauge these vectors using the second and third regularization terms $\frac{\mu}{2} \sum_i \|\bar{\mathbf{w}}_i\|_2^2$ and $\frac{\gamma}{2} \|\mathbf{C}\|_F^2$ with the regularization parameters $\mu > 0$ and $\gamma > 0$, respectively.⁵ We emphasize that since $\|\mathbf{C}\|_F^2 = \sum_j \|\mathbf{c}_j\|_2^2$, we can impose a regularizer on each column rather than the entire matrix \mathbf{C} , which facilitates the development of the efficient algorithm detailed below.

3.2 The SPARFA-M Algorithm

Since the first negative log-likelihood term in the objective function of (P) is convex in the product \mathbf{WC} for both the probit and the logit functions (see, e.g., Hastie et al. 2010), and since the rest of the regularization terms are convex in either \mathbf{W} or \mathbf{C} while the non-negativity constraints on $W_{i,k}$ are with respect to a convex set, the problem (P) is *biconvex* in the individual factors \mathbf{W} and \mathbf{C} . More importantly, with respect to blocks of variables $\bar{\mathbf{w}}_i, \mathbf{c}_j$, the problem (P) is *block multi-convex* in the sense of Xu and Yin (2012).

SPARFA-M is an alternating optimization approach to (approximately) *solve* (P) that proceeds as follows. We initialize \mathbf{W} and \mathbf{C} with random entries and then iteratively optimize the objective function of (P) for both factors in an alternating fashion. Each outer iteration involves solving two kinds of inner subproblems. In the first subproblem, we hold \mathbf{W} constant and separately optimize each block of variables in \mathbf{c}_j ; in the second subproblem, we hold \mathbf{C} constant and separately optimize each block of variables $\bar{\mathbf{w}}_i$. Each subproblem is solved using an iterative method; see Section 3.3 for the respective algorithms. The outer loop is terminated whenever a maximum number of outer iterations I_{max} is reached, or if the decrease in the objective function of (P) is smaller than a certain threshold.

The two subproblems constituting the inner iterations of SPARFA-M correspond to the following convex ℓ_1/ℓ_2 -norm and ℓ_2 -norm regularized regression (RR) problems:

$$\begin{aligned} (\text{RR}_1^+) \quad & \underset{\bar{\mathbf{w}}_i: W_{i,k} \geq 0 \forall k}{\text{minimize}} \quad \sum_{j: (i,j) \in \Omega_{\text{obs}}} -\log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) + \lambda \|\bar{\mathbf{w}}_i\|_1 + \frac{\mu}{2} \|\bar{\mathbf{w}}_i\|_2^2, \\ (\text{RR}_2) \quad & \underset{\mathbf{c}_j}{\text{minimize}} \quad \sum_{i: (i,j) \in \Omega_{\text{obs}}} -\log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\gamma}{2} \|\mathbf{c}_j\|_2^2. \end{aligned}$$

5. The first ℓ_1 -norm regularization term in (RR₁⁺) already gauges the norm of the $\bar{\mathbf{w}}_i$. The ℓ_2 -norm regularizer $\frac{\mu}{2} \sum_i \|\bar{\mathbf{w}}_i\|_2^2$ is included only to aid in establishing the convergence results for SPARFA-M as detailed in Section 3.4.

We develop two novel first-order methods that efficiently solve (RR_1^+) and (RR_2) for both probit and logistic regression. These methods scale well to high-dimensional problems, in contrast to existing second-order methods. In addition, the probit link function makes the explicit computation of the Hessian difficult, which is only required for second-order methods. Therefore, we build our algorithm on the fast iterative soft-thresholding algorithm (FISTA) framework developed in Beck and Teboulle (2009), which enables the development of efficient first-order methods with accelerated convergence.

3.3 Accelerated First-Order Methods for Regularized Probit/Logistic Regression

The FISTA framework (Beck and Teboulle 2009) iteratively solves optimization problems whose objective function is given by $f(\cdot) + g(\cdot)$, where $f(\cdot)$ is a continuously differentiable convex function and $g(\cdot)$ is convex but potentially non-smooth. This approach is particularly well-suited to the inner subproblem (RR_1^+) due to the presence of the non-smooth ℓ_1 -norm regularizer and the non-negativity constraint. Concretely, we associate the log-likelihood function plus the ℓ_2 -norm regularizer $\frac{\mu}{2}\|\bar{\mathbf{w}}_i\|_2^2$ with $f(\cdot)$ and the ℓ_1 -norm regularization term with $g(\cdot)$. For the inner subproblem (RR_2) , we associate the log-likelihood function with $f(\cdot)$ and the ℓ_2 -norm regularization term with $g(\cdot)$.⁶

Each FISTA iteration consists of two steps: (i) a gradient-descent step in $f(\cdot)$ and (ii) a shrinkage step determined by $g(\cdot)$. For simplicity of exposition, we consider the case where all entries in \mathbf{Y} are observed, i.e., $\Omega_{\text{obs}} = \{1, \dots, Q\} \times \{1, \dots, N\}$; the extension to the case with missing entries in \mathbf{Y} is straightforward. We will derive the algorithm for the case of probit regression first and then point out the departures for logistic regression.

For (RR_1^+) , the gradients of $f(\bar{\mathbf{w}}_i)$ with respect to the i^{th} block of regression coefficients $\bar{\mathbf{w}}_i$ are given by

$$\nabla f_{\text{pro}}^i = \nabla_{\bar{\mathbf{w}}_i}^{\text{pro}}(-\sum_j \log p_{\text{pro}}(Y_{i,j}|\bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\mu}{2}\|\bar{\mathbf{w}}_i\|_2^2) = -\mathbf{C}\mathbf{D}^i(\bar{\mathbf{y}}^i - \mathbf{p}_{\text{pro}}^i) + \mu\bar{\mathbf{w}}_i, \quad (5)$$

where $\bar{\mathbf{y}}^i$ is an $N \times 1$ column vector corresponding to the transpose of the i^{th} row of \mathbf{Y} . $\mathbf{p}_{\text{pro}}^i$ is an $N \times 1$ vector whose j^{th} element equals the probability of $Y_{i,j}$ being 1; that is, $p_{\text{pro}}(Y_{i,j} = 1|\bar{\mathbf{w}}_i, \mathbf{c}_j) = \Phi_{\text{pro}}(\bar{\mathbf{w}}_i^T \mathbf{c}_j)$. The entries of the $N \times N$ diagonal matrix \mathbf{D} are given by

$$D_{j,j}^i = \frac{\mathcal{N}(\bar{\mathbf{w}}_i^T \mathbf{c}_j)}{\Phi_{\text{pro}}(\bar{\mathbf{w}}_i^T \mathbf{c}_j)(1 - \Phi_{\text{pro}}(\bar{\mathbf{w}}_i^T \mathbf{c}_j))}.$$

The gradient step in each FISTA iteration $\ell = 1, 2, \dots$ corresponds to

$$\hat{\bar{\mathbf{w}}}_i^{\ell+1} \leftarrow \bar{\mathbf{w}}_i^\ell - t_\ell \nabla f_{\text{pro}}^i, \quad (6)$$

where t_ℓ is a suitable step-size. To comply with (A3), the shrinkage step in (RR_1^+) corresponds to a non-negative soft-thresholding operation

$$\bar{\mathbf{w}}_i^{\ell+1} \leftarrow \max\{\hat{\bar{\mathbf{w}}}_i^{\ell+1} - \lambda t_\ell, 0\}. \quad (7)$$

6. Of course, both $f(\cdot)$ and $g(\cdot)$ are smooth for (RR_2) . Hence, we could also apply an accelerated gradient-descent approach instead, e.g., as described in Nesterov (2007).

For (RR₂), the gradient step becomes

$$\hat{\mathbf{c}}_j^{\ell+1} \leftarrow \mathbf{c}_j^\ell - t_\ell \nabla f_{\text{pro}}^i,$$

which is the same as (5) and (6) after replacing \mathbf{C} with \mathbf{W}^T and μ with γ . The shrinkage step for (RR₂) is the simple re-scaling

$$\mathbf{c}_j^{\ell+1} \leftarrow \frac{1}{1 + \gamma t_\ell} \hat{\mathbf{c}}_j^{\ell+1}. \quad (8)$$

In the logistic regression case, the steps (6), (7), and (8) remain the same but the gradient changes to

$$\nabla f_{\text{log}}^i = \nabla_{\bar{\mathbf{w}}_i}^{\text{log}}(-\sum_j \log p_{\text{log}}(Y_{i,j}|\bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\mu}{2}\|\bar{\mathbf{w}}_i\|_2^2) = -\mathbf{C}(\bar{\mathbf{y}}^i - \mathbf{p}_{\text{log}}^i) + \mu\bar{\mathbf{w}}_i, \quad (9)$$

where the $N \times 1$ vector $\mathbf{p}_{\text{log}}^i$ has elements $\mathbf{p}_{\text{log}}^i(Y_{i,j} = 1|\bar{\mathbf{w}}_i, \mathbf{c}_j) = \Phi_{\text{log}}(\bar{\mathbf{w}}_i^T \mathbf{c}_j)$.

The above steps require a suitable step-size t_ℓ to ensure convergence to the optimal solution. A common approach that guarantees convergence is to set $t_\ell = 1/L$, where L is the Lipschitz constant of $f(\cdot)$ (see Beck and Teboulle 2009 for the details). The Lipschitz constants for both the probit and logit cases are analyzed in Theorem 1 below. Alternatively, one can also perform backtracking, which—under certain circumstances—can be more efficient; see (Beck and Teboulle, 2009, p. 194) for more details.

3.4 Convergence Analysis of SPARFA-M

While the SPARFA-M objective function is guaranteed to be non-increasing over the outer iterations (Boyd and Vandenberghe 2004), the factors \mathbf{W} and \mathbf{C} do not necessarily converge to a global or local optimum due to its biconvex (or more generally, block multi-convex) nature. It is difficult, in general, to develop rigorous statements for the convergence behavior of block multi-convex problems. Nevertheless, we can establish the global convergence of SPARFA-M from any starting point to a critical point of the objective function using recent results developed in Xu and Yin (2012). The convergence results below appear to be novel for both sparse matrix factorization as well as dictionary learning.

3.4.1 CONVERGENCE ANALYSIS OF REGULARIZED REGRESSION USING FISTA

In order to establish the SPARFA-M convergence result, we first adapt the convergence results for FISTA in Beck and Teboulle (2009) to prove convergence on the two subproblems (RR₁⁺) and (RR₂). The following theorem is a consequence of (Beck and Teboulle, 2009, Thm. 4.4) combined with Lemmata 4 and 5 in Appendix A. If back-tracking is used to select step-size t_ℓ (Beck and Teboulle, 2009, p. 194), then let α correspond to the back-tracking parameter. Otherwise set $\alpha = 1$ and for (RR₁⁺) let $t_\ell = 1/L_1$ and for (RR₂) let $t_\ell = 1/L_2$. In Lemma 5, we compute that $L_1 = \sigma_{\max}^2(\mathbf{C}) + \mu$ and $L_2 = \sigma_{\max}^2(\mathbf{W})$ for the probit case, and $L_1 = \frac{1}{4}\sigma_{\max}^2(\mathbf{C}) + \mu$ and $L_2 = \frac{1}{4}\sigma_{\max}^2(\mathbf{W})$ for the logit case.

Theorem 1 (Linear convergence of RR using FISTA) *Given i and j , let*

$$F_1(\bar{\mathbf{w}}_i) = \sum_{j:(i,j) \in \Omega_{\text{obs}}} -\log p(Y_{i,j}|\bar{\mathbf{w}}_i, \mathbf{c}_j) + \lambda\|\bar{\mathbf{w}}_i\|_1 + \frac{\mu}{2}\|\bar{\mathbf{w}}_i\|_2^2, \quad W_{i,k} \geq 0 \quad \forall k,$$

$$F_2(\mathbf{c}_j) = \sum_{i:(i,j) \in \Omega_{\text{obs}}} -\log p(Y_{i,j}|\bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\gamma}{2}\|\mathbf{c}_j\|_2^2$$

be the cost functions of (RR_1^+) and (RR_2) , respectively. Then, we have

$$F_1(\bar{\mathbf{w}}_i^\ell) - F_1(\bar{\mathbf{w}}_i^*) \leq \frac{2\alpha L_1 \|\bar{\mathbf{w}}_i^0 - \bar{\mathbf{w}}_i^*\|^2}{(\ell + 1)^2},$$

$$F_2(\mathbf{c}_j^\ell) - F_2(\mathbf{c}_j^*) \leq \frac{2\alpha L_2 \|\mathbf{c}_j^0 - \mathbf{c}_j^*\|^2}{(\ell + 1)^2},$$

where $\bar{\mathbf{w}}_i^0$ and \mathbf{c}_j^0 are the initialization points of (RR_1^+) and (RR_2) , $\bar{\mathbf{w}}_i^\ell$ and \mathbf{c}_j^ℓ designate the solution estimates at the ℓ^{th} inner iteration, and $\bar{\mathbf{w}}_i^*$ and \mathbf{c}_j^* denote the optimal solutions.

In addition to establishing convergence, Theorem 1 reveals that the difference between the cost functions at the current estimates and the optimal solution points, $F_1(\bar{\mathbf{w}}_i^\ell) - F_1(\bar{\mathbf{w}}_i^*)$ and $F_2(\mathbf{c}_j^\ell) - F_2(\mathbf{c}_j^*)$, decrease as $O(\ell^{-2})$.

3.4.2 CONVERGENCE ANALYSIS OF SPARFA-M

We are now ready to establish global convergence of SPARFA-M to a critical point. To this end, we first define $\mathbf{x} = [\bar{\mathbf{w}}_1^T, \dots, \bar{\mathbf{w}}_Q^T, \mathbf{c}_1^T, \dots, \mathbf{c}_N^T]^T \in \mathbb{R}^{(N+Q)K}$ and rewrite the objective function (P) of SPARFA-M as follows:

$$F(\mathbf{x}) = \sum_{(i,j) \in \Omega_{\text{obs}}} -\log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\mu}{2} \sum_i \|\bar{\mathbf{w}}_i\|_2^2 + \lambda \sum_i \|\bar{\mathbf{w}}_i\|_1 + \sum_{i,k} \delta(W_{i,k} < 0) + \frac{\gamma}{2} \sum_j \|\mathbf{c}_j\|_2^2$$

with the indicator function $\delta(z < 0) = \infty$ if $z < 0$ and 0 otherwise. Note that we have re-formulated the non-negativity constraint as a set indicator function and added it to the objective function of (P). Since minimizing $F(\mathbf{x})$ is equivalent to solving (P), we can now use the results developed in Xu and Yin (2012) to establish the following convergence result for the SPARFA-M algorithm. The proof can be found in Appendix B.

Theorem 2 (Global convergence of SPARFA-M) *From any starting point \mathbf{x}^0 , let $\{\mathbf{x}^t\}$ be the sequence of estimates generated by the SPARFA-M algorithm with $t = 1, 2, \dots$ as the outer iteration number. Then, the sequence $\{\mathbf{x}^t\}$ converges to the finite limit point $\hat{\mathbf{x}}$, which is a critical point of (P). Moreover, if the starting point \mathbf{x}^0 is within a close neighborhood of a global optimum of (P), then SPARFA-M converges to this global optimum.*

Since the problem (P) is bi-convex in nature, we cannot guarantee that SPARFA-M always converges to a *global optimum* from an *arbitrary* starting point. Nevertheless, the use of multiple randomized initialization points can be used to increase the chance of being in the close vicinity of a global optimum, which improves the (empirical) performance of SPARFA-M (see Section 3.5 for details). Note that we do not provide the convergence rate of SPARFA-M, since the associated parameters in (Xu and Yin, 2012, Thm. 2.9) are difficult to determine for the model at hand; a detailed analysis of the convergence rate for SPARFA-M is part of ongoing work.

3.5 Algorithmic Details and Improvements for SPARFA-M

In this section, we outline a toolbox of techniques that improve the empirical performance of SPARFA-M and provide guidelines for choosing the key algorithm parameters.

3.5.1 REDUCING COMPUTATIONAL COMPLEXITY IN PRACTICE

To reduce the computational complexity of SPARFA-M in practice, we can improve the convergence rates of (RR_1^+) and (RR_2) . In particular, the regularizer $\frac{\mu}{2}\|\bar{\mathbf{w}}_i\|_2^2$ in (RR_1^+) has been added to (P) to facilitate the proof for Theorem 2. This term, however, typically slows down the (empirical) convergence of FISTA, especially for large values of μ . We therefore set μ to a small positive value (e.g., $\mu = 10^{-4}$), which leads to fast convergence of (RR_1^+) while still guaranteeing convergence of SPARFA-M.

Selecting the appropriate (i.e., preferably large) step-sizes t_ℓ in (6), (7), and (8) is also crucial for fast convergence. In Lemmata 4 and 5, we derive the Lipschitz constants L for (RR_1^+) and (RR_2) , which enables us to set the step-sizes t_ℓ to the constant value $t = 1/L$. In all of our experiments below, we exclusively use constant step-sizes, since we observed that backtracking (Beck and Teboulle 2009, p. 194) provided no advantage in terms of computational complexity for SPARFA-M.

To further reduce the computational complexity of SPARFA-M without degrading its empirical performance noticeably, we have found that instead of running the large number of inner iterations it typically takes to converge, we can run just a few (e.g., 10) inner iterations per outer iteration.

3.5.2 REDUCING THE CHANCE OF GETTING STUCK IN LOCAL MINIMA

The performance of SPARFA-M strongly depends on the initialization of \mathbf{W} and \mathbf{C} , due to the bi-convex nature of (P). We have found that running SPARFA-M multiple times with different starting points and picking the solution with the smallest overall objective function delivers excellent performance. In addition, we can deploy the standard heuristics used in the dictionary-learning literature (Aharon et al., 2006, Section IV-E) to further improve the convergence towards a global optimum. For example, every few outer iterations, we can evaluate the current \mathbf{W} and \mathbf{C} . If two rows of \mathbf{C} are similar (as measured by the absolute value of the inner product between them), then we re-initialize one of them as an i.i.d. Gaussian vector. Moreover, if some columns in \mathbf{W} contain only zero entries, then we re-initialize them with i.i.d. Gaussian vectors. Note that the convergence proof in Section 3.4 does not apply to implementations employing such trickery.

3.5.3 PARAMETER SELECTION

The input parameters to SPARFA-M include the number of concepts K and the regularization parameters γ and λ . The number of concepts K is a user-specified value. In practice, cross-validation can be used to select K if the task is to predict missing entries of \mathbf{Y} , (see Section 6.3). The sparsity parameter λ and the ℓ_2 -norm penalty parameter γ strongly affect the output of SPARFA-M; they can be selected using any of a number of criteria, including the Bayesian information criterion (BIC) or cross-validation, as detailed in Hastie et al. (2010). Both criteria resulted in similar performance in all of the experiments reported in Section 6.

3.6 Related Work on Maximum Likelihood-based Sparse Factor Analysis

Sparse logistic factor analysis has previously been studied in Lee et al. (2010) in the context of principal components analysis. There are three major differences with the SPARFA framework. First, Lee et al. (2010) do not impose the non-negativity constraint on \mathbf{W} that is critical for the interpretation of the estimated factors. Second, they impose an orthonormality constraint on \mathbf{C} that does not make sense in educational scenarios. Third, they optimize an upper bound on the negative log-likelihood function in each outer iteration, in contrast to SPARFA-M, which optimizes the exact cost functions in (RR_1^+) and (RR_2) .

The problem (P) shares some similarities with the method for missing data imputation outlined in (Mohamed et al., 2012, Eq. 7). However, the problem (P) studied here includes an additional non-negativity constraint on \mathbf{W} and the regularization term $\frac{\mu}{2} \sum_i \|\bar{\mathbf{w}}_i\|_2^2$ that are important for the interpretation of the estimated factors and the convergence analysis. Moreover, SPARFA-M utilizes the accelerated FISTA framework as opposed to the more straightforward but less efficient gradient descent method in Mohamed et al. (2012).

SPARFA-M is capable of handling both the inverse logit and inverse probit link functions. For the inverse logit link function, one could solve (RR_1^+) and (RR_2) using an iteratively reweighted second-order algorithm as in Hastie et al. (2010), Minka (2003), Lee et al. (2006), Park and Hastie (2008), or an interior-point method as in Koh et al. (2007). However, none of these techniques extend naturally to the inverse probit link function, which is essential for some applications, e.g., in noisy compressive sensing recovery from 1-bit measurements (e.g., Jacques et al. 2013 or Plan and Vershynin 2012, submitted). Moreover, second-order techniques typically do not scale well to high-dimensional problems due to the necessary computation of the Hessian. In contrast, SPARFA-M scales favorably due to the fact that it utilizes the accelerated first-order FISTA method, avoiding the computation of the Hessian.

4. SPARFA-B: Bayesian Sparse Factor Analysis

Our second algorithm, SPARFA-B, solves the SPARFA problem using a Bayesian method based on Markov chain Monte-Carlo (MCMC) sampling. In contrast to SPARFA-M, which computes point estimates for each of the parameters of interest, SPARFA-B computes full posterior distributions for \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$.

While SPARFA-B has a higher computational complexity than SPARFA-M, it has several notable benefits in the context of learning and content analytics. First, the full posterior distributions enable the computation of informative quantities such as credible intervals and posterior modes for all parameters of interest. Second, since MCMC methods explore the full posterior space, they are not subject to being trapped indefinitely in local minima, which is possible with SPARFA-M. Third, the hyperparameters used in Bayesian methods generally have intuitive meanings, in contrary to the regularization parameters of optimization-based methods like SPARFA-M. These hyperparameters can also be specially chosen to incorporate additional prior information about the problem.

4.1 Problem Formulation

As discussed in Section 2.2, we require the matrix \mathbf{W} to be both sparse (A2) and non-negative (A3). We enforce these assumptions through the following prior distributions that

are a variant of the well-studied spike-slab model (West, 2003; Ishwaran and Rao, 2005) adapted for non-negative factor loadings:

$$W_{i,k} \sim r_k \text{Exp}(\lambda_k) + (1 - r_k) \delta_0, \quad \lambda_k \sim \text{Ga}(\alpha, \beta), \quad \text{and} \quad r_k \sim \text{Beta}(e, f). \quad (10)$$

Here, $\text{Exp}(x|\lambda) \sim \lambda e^{-\lambda x}$, $x \geq 0$, and $\text{Ga}(x|\alpha, \beta) \sim \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}$, $x \geq 0$, δ_0 is the Dirac delta function, and α, β, e, f are hyperparameters. The model (10) uses the latent random variable r_k to control the sparsity via the hyperparameters e and f . This set of priors induces a conjugate form on the posterior that enables efficient sampling. We note that both the exponential rate parameters λ_k as well as the inclusion probabilities r_k are grouped per factor. The remaining priors used in the proposed Bayesian model are summarized as

$$\mathbf{c}_j \sim \mathcal{N}(0, \mathbf{V}), \quad \mathbf{V} \sim \text{IW}(\mathbf{V}_0, h), \quad \text{and} \quad \mu_i \sim \mathcal{N}(\mu_0, v_\mu), \quad (11)$$

where \mathbf{V}_0, h, μ_0 , and v_μ are hyperparameters.

4.2 The SPARFA-B Algorithm

We obtain posterior distribution estimates for the parameters of interest through an MCMC method based on the Gibbs' sampler. To implement this, we must derive the conditional posteriors for each of the parameters of interest. We note again that the graded learner-response matrix \mathbf{Y} will not be fully observed, in general. Thus, our sampling method must be equipped to handle missing data.

The majority of the posterior distributions follow from standard results in Bayesian analysis and will not be derived in detail here. The exception is the posterior distribution of $W_{i,k} \forall i, k$. The spike-slab model that enforces sparsity in \mathbf{W} requires first sampling $W_{i,k} \neq 0 | \mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}$ and then sampling $W_{i,k} | \mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}$, for all $W_{i,k} \neq 0$. These posterior distributions differ from previous results in the literature due to our assumption of an exponential (rather than a normal) prior on $W_{i,k}$. We next derive these two results in detail.

4.2.1 DERIVATION OF POSTERIOR DISTRIBUTION OF $W_{i,k}$

We seek both the probability that an entry $W_{i,k}$ is active (non-zero) and the distribution of $W_{i,k}$ when active given our observations. The following theorem states the final sampling results; the proof is given in Appendix C.

Theorem 3 (Posterior distributions for \mathbf{W}) *For all $i = 1, \dots, Q$ and all $k = 1, \dots, K$, the posterior sampling results for $W_{i,k} = 0 | \mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}$ and $W_{i,k} | \mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}, W_{i,k} \neq 0$ are given by*

$$\begin{aligned} \widehat{R}_{i,k} &= p(W_{i,k} = 0 | \mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}) = \frac{\mathcal{N}^r(0 | \widehat{M}_{i,k}, \widehat{S}_{i,k}, \lambda_k) \text{Exp}(0 | \lambda_k) (1 - r_k)}{\mathcal{N}^r(0 | \widehat{M}_{i,k}, \widehat{S}_{i,k}, \lambda_k) \text{Exp}(0 | \lambda_k) (1 - r_k) + r_k}, \\ W_{i,k} | \mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}, W_{i,k} \neq 0 &\sim \mathcal{N}^r(\widehat{M}_{i,k}, \widehat{S}_{i,k}, \lambda_k), \\ \widehat{M}_{i,k} &= \frac{\sum_{\{j:(i,j) \in \Omega_{obs}\}} ((Z_{i,j} - \mu_i) - \sum_{k' \neq k} W_{i,k'} C_{k',j}) C_{k,j}}{\sum_{\{j:(i,j) \in \Omega_{obs}\}} C_{k,j}^2}, \\ \widehat{S}_{i,k} &= \left(\sum_{\{j:(i,j) \in \Omega_{obs}\}} C_{k,j}^2 \right)^{-1}, \end{aligned}$$

where $\mathcal{N}^r(x|m, s, \lambda) = \frac{e^{\lambda m - \lambda^2 s/2}}{\sqrt{2\pi s} \Phi\left(\frac{m - \lambda s}{\sqrt{s}}\right)} e^{-(x-m)^2/2s - \lambda m}$ represents a rectified normal distribution (see Schmidt et al. 2009).

4.2.2 SAMPLING METHODOLOGY

SPARFA-B carries out the following MCMC steps to compute posterior distributions for all parameters of interest:

1. For all $(i, j) \in \Omega_{\text{obs}}$, draw $Z_{i,j} \sim \mathcal{N}((\mathbf{WC})_{i,j} + \mu_i, 1)$, truncating above 0 if $Y_{i,j} = 1$, and truncating below 0 if $Y_{i,j} = 0$.
2. For all $i = 1, \dots, Q$, draw $\mu_i \sim \mathcal{N}(m_i, v)$ with $v = (v_{\boldsymbol{\mu}}^{-1} + n')^{-1}$, $m_i = \mu_0 + v \sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} (Z_{i,j} - \bar{\mathbf{w}}_i^T \mathbf{c}_j)$, and n' the number of learners responding to question i .
3. For all $j = 1, \dots, N$, draw $\mathbf{c}_j \sim \mathcal{N}(\mathbf{m}_j, \mathbf{M}_j)$ with $\mathbf{M}_j = (\mathbf{V}^{-1} + \widetilde{\mathbf{W}}^T \widetilde{\mathbf{W}})^{-1}$, and $\mathbf{m}_j = \mathbf{M}_j \widetilde{\mathbf{W}}^T (\tilde{\mathbf{z}}_j - \tilde{\boldsymbol{\mu}})$. The notation $\widetilde{(\cdot)}$ denotes the restriction of the vector or matrix to the set of rows $i : (i, j) \in \Omega_{\text{obs}}$.
4. Draw $\mathbf{V} \sim IW(\mathbf{V}_0 + \mathbf{C}\mathbf{C}^T, N + h)$.
5. For all $i = 1, \dots, Q$ and $k = 1, \dots, K$, draw $W_{i,k} \sim \widehat{R}_{i,k} \mathcal{N}^r(\widehat{M}_{i,k}, \widehat{S}_{i,k}) + (1 - \widehat{R}_{i,k})\delta_0$, where $\widehat{R}_{i,k}$, $\widehat{M}_{i,k}$, and $\widehat{S}_{i,k}$ are as stated in Theorem 3.
6. For all $k = 1, \dots, K$, let b_k define the number of active (i.e., non-zero) entries of $\bar{\mathbf{w}}_k$. Draw $\lambda_k \sim Ga(\alpha + b_k, \beta + \sum_{i=1}^Q W_{i,k})$.
7. For all $k = 1, \dots, K$, draw $r_k \sim Beta(e + b_k, f + Q - b_k)$, with b_k defined as in Step 6.

4.3 Algorithmic Details and Improvements for SPARFA-B

Here we discuss some several practical issues for efficiently implementing SPARFA-B, selecting the hyperparameters, and techniques for easy visualization of the SPARFA-B results.

4.3.1 IMPROVING COMPUTATIONAL EFFICIENCY

The Gibbs sampling scheme of SPARFA-B enables efficient implementation in several ways. First, draws from the truncated normal in Step 1 of Section 4.2.2 are decoupled from one another, allowing them to be performed independently and, potentially, in parallel. Second, sampling of the elements in each column of \mathbf{W} can be carried out in parallel by computing the relevant factors of Step 5 in matrix form. Since $K \ll Q, N$ by assumption (A1), the relevant parameters are recomputed only a relatively small number of times. One taxing computation is the calculation of the covariance matrix \mathbf{M}_j for each $j = 1, \dots, N$ in Step 3. This computation is necessary, since we do not constrain each learner to answer the same set of questions which, in turn, changes the nature of the covariance calculation for each individual learner. For data sets where all learners answer the same set of questions, this covariance matrix is the same for all learners and, hence, can be carried out once per MCMC iteration.

4.3.2 PARAMETER SELECTION

The selection of the hyperparameters is performed at the discretion of the user. As is typical for Bayesian methods, non-informative (broad) hyperparameters can be used to avoid biasing results and to allow for adequate exploration of the posterior space. Tighter hyperparameters can be used when additional side information is available. For example, prior information from subject matter experts might indicate which concepts are related to which questions or might indicate the intrinsic difficulty of the questions.

Since SPARFA-M has a substantial speed advantage over SPARFA-B, it may be advantageous to first run SPARFA-M and then use its output to help in determining the hyperparameters or to initialize the SPARFA-B variables directly.

4.3.3 POST-PROCESSING FOR DATA VISUALIZATION

As discussed above, the generation of posterior statistics is one of the primary advantages of SPARFA-B. However, for many tasks, such as visualization of the retrieved knowledge base, it is often convenient to post-process the output of SPARFA-B to obtain point estimates for each parameter. For many Bayesian methods, simply computing the posterior mean is often sufficient. This is the case for most parameters computed by SPARFA-B, including \mathbf{C} and $\boldsymbol{\mu}$. The posterior mean of \mathbf{W} , however, is generally non-sparse, since the MCMC will generally explore the possibility of including each entry of \mathbf{W} . Nevertheless, we can easily generate a sparse \mathbf{W} by examining the posterior mean of the inclusion statistics contained in $\hat{R}_{i,k}, \forall i, k$. Concretely, if the posterior mean of $\hat{R}_{i,k}$ is small, then we set the corresponding entry of $W_{i,k}$ to zero. Otherwise, we set $W_{i,k}$ to its posterior mean. We will make use of this method throughout the experiments presented in Section 6.

4.4 Related Work on Bayesian Sparse Factor Analysis

Sparsity models for Bayesian factor analysis have been well-explored in the statistical literature (West, 2003; Tipping, 2001; Ishwaran and Rao, 2005). One popular avenue for promoting sparsity is to place a prior on the variance of each component in \mathbf{W} (see, e.g., Tipping 2001, Fokoue 2004, and Pournara and Wernisch 2007). In such a model, large variance values indicate active components, while small variance values indicate inactive components. Another approach is to model active and inactive components directly using a form of a spike-slab model due to West (2003) and used in Goodfellow et al. (2012), Mohamed et al. (2012), and Hahn et al. (2012):

$$W_{i,k} \sim r_k \mathcal{N}(0, v_k) + (1 - r_k) \delta_0, \quad v_k \sim IG(\alpha, \beta), \quad \text{and} \quad r_k \sim \text{Beta}(e, f).$$

The approach employed in (10) utilizes a spike-slab prior with an exponential distribution, rather than a normal distribution, for the active components of \mathbf{W} . We chose this prior for several reasons. First, it enforces the non-negativity assumption (A3). Second, it induces a posterior distribution that can be both computed in closed form and sampled efficiently. Third, its tail is slightly heavier than that of a standard normal distribution, which improves the exploration of quantities further away from zero.

A sparse factor analysis model with non-negativity constraints that is related to the one proposed here was discussed in Meng et al. (2010), although their methodology is quite

different from ours. Specifically, they impose non-negativity on the (dense) matrix \mathbf{C} rather than on the sparse factor loading matrix \mathbf{W} . Furthermore, they enforce non-negativity using a truncated normal⁷ rather than an exponential prior.

5. Tag Analysis: Post-Processing to Interpret the Estimated Concepts

So far we have developed SPARFA-M and SPARFA-B to estimate \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ (or equivalently, \mathbf{M}) in (2) given the partial binary observations in \mathbf{Y} . Both \mathbf{W} and \mathbf{C} encode a small number of latent concepts. As we initially noted, the concepts are “*abstract*” in that they are estimated from the data rather than dictated by a subject matter expert. In this section we develop a principled post-processing approach to interpret the meaning of the abstract concepts after they have been estimated from learner responses, which is important if our results are to be usable for learning analytics and content analytics in practice. Our approach applies when the questions come with a set of user-generated “tags” or “labels” that describe in a free-form manner what ideas underlie each question.

We develop a post-processing algorithm for the estimated matrices \mathbf{W} and \mathbf{C} that estimates the association between the latent concepts and the user-generated tags, enabling concepts to be interpreted as a “bag of tags.” Additionally, we show how to extract a personalized tag knowledge profile for each learner. The efficacy of our tag-analysis framework will be demonstrated in the real-world experiments in Section 6.2.

5.1 Incorporating Question–Tag Information

Suppose that a set of tags has been generated for each question that represent the topic(s) or theme(s) of each question. The tags could be generated by the course instructors, subject matter experts, learners, or, more broadly, by crowd-sourcing. In general, the tags provide a redundant representation of the true knowledge components, i.e., concepts are associated to a “bag of tags.”

Assume that there is a total number of M tags associated with the Q questions. We form a $Q \times M$ matrix \mathbf{T} , where each column of \mathbf{T} is associated to one of the M pre-defined tags. We set $T_{i,m} = 1$ if tag $m \in \{1, \dots, M\}$ is present in question i and 0 otherwise. Now, we postulate that the question association matrix \mathbf{W} extracted by SPARFA can be further factorized as $\mathbf{W} = \mathbf{TA}$, where \mathbf{A} is an $M \times K$ matrix representing the tags-to-concept mapping. This leads to the following additional assumptions:

- (A4) *Non-negativity*: The matrix \mathbf{A} is non-negative. This increases the interpretability of the result, since concepts should not be negatively correlated with any tags, in general.
- (A5) *Sparsity*: Each column of \mathbf{A} is sparse. This ensures that the estimated concepts relate to only a few tags.

7. One could alternatively employ a truncated normal distribution on the support $[0, \infty)$ for the active entries in \mathbf{W} . In experiments with this model, we found a slight, though noticeable, improvement in prediction performance on real-data experiments using the exponential prior.

5.2 Estimating the Concept–Tag Associations and Learner–Tag Knowledge

The assumptions (A4) and (A5) enable us to extract \mathbf{A} using ℓ_1 -norm regularized non-negative least-squares as described in Hastie et al. (2010) and Chen et al. (1998). Specifically, to obtain each column \mathbf{a}_k of \mathbf{A} , $k = 1, \dots, K$, we solve the following convex optimization problem, a non-negative variant of *basis pursuit denoising*:

$$(\text{BPDN}_+) \quad \underset{\mathbf{a}_k: A_{m,k} \geq 0 \forall m}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{w}_k - \mathbf{T}\mathbf{a}_k\|_2^2 + \eta \|\mathbf{a}_k\|_1.$$

Here, \mathbf{w}_k represents the k^{th} column of \mathbf{W} , and the parameter η controls the sparsity level of the solution \mathbf{a}_k .

We propose a first-order method derived from the FISTA framework in Beck and Teboulle (2009) to solve (BPDN₊). The algorithm consists of two steps: A gradient step with respect to the ℓ_2 -norm penalty function, and a projection step with respect to the ℓ_1 -norm regularizer subject to the non-negative constraints on \mathbf{a}_k . By solving (BPDN₊) for $k = 1, \dots, K$, and building $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_K]$, we can (i) assign tags to each concept based on the non-zero entries in \mathbf{A} and (ii) estimate a tag-knowledge profile for each learner.

5.2.1 ASSOCIATING TAGS TO EACH CONCEPT

Using the concept–tag association matrix \mathbf{A} we can directly associate tags to each concept estimated by the SPARFA algorithms. We first normalize the entries in \mathbf{a}_k such that they sum to one. With this normalization, we can then calculate percentages that show the proportion of each tag that contributes to concept k corresponding to the non-zero entries of \mathbf{a}_k . This concept tagging method typically will assign multiple tags to each concept, thus, enabling one to identify the coarse meaning of each concept (see Section 6.2 for examples using real-world data).

5.2.2 LEARNER TAG KNOWLEDGE PROFILES

Using the concept–tag association matrix \mathbf{A} , we can assess each learner’s knowledge of each tag. To this end, we form an $M \times N$ matrix $\mathbf{U} = \mathbf{A}\mathbf{C}$, where the $U_{m,j}$ characterizes the knowledge of learner j of tag m . This information could be used, for example, by a PLS to automatically inform each learner which tags they have strong knowledge of and which tags they do not. Course instructors can use the information contained in \mathbf{U} to extract measures representing the knowledge of all learners on a given tag, e.g., to identify the tags for which the entire class lacks strong knowledge. This information would enable the course instructor to select future learning content that deals with those specific tags. A real-world example demonstrating the efficacy of this framework is shown below in Section 6.2.1.

6. Experiments

In this section, we validate SPARFA-M and SPARFA-B on both synthetic and real-world educational data sets. First, using synthetic data, we validate that both algorithms can accurately estimate the underlying factors from binary-valued observations and characterize their performance under different circumstances. Specifically, we benchmark the factor estimation performance of SPARFA-M and SPARFA-B against a variant of the well-established

K-SVD algorithm (Aharon et al. 2006) used in dictionary-learning applications. Second, using real-world graded learner-response data we demonstrate the efficacy SPARFA-M (both probit and logit variants) and of SPARFA-B for learning and content analytics. Specifically, we showcase how the estimated learner concept knowledge, question–concept association, and intrinsic question difficulty can support machine learning-based personalized learning. Finally, we compare SPARFA-M against the recently proposed binary-valued collaborative filtering algorithm CF-IRT (Bergner et al. 2012) that predicts unobserved learner responses.

6.1 Synthetic Data Experiments

We first characterize the estimation performance of SPARFA-M and SPARFA-B using synthetic test data generated from a known ground truth model. We generate instances of \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ under pre-defined distributions and then generate the binary-valued observations \mathbf{Y} according to (2).

Our report on the synthetic experiments is organized as follows. In Section 6.1.1, we outline K-SVD₊, a variant of the well-established K-SVD dictionary-learning (DL) algorithm originally proposed in Aharon et al. (2006); we use it as a baseline method for comparison to both SPARFA algorithms. In Section 6.1.2 we detail the performance metrics. We compare SPARFA-M, SPARFA-B, and K-SVD₊ as we vary the problem size and number of concepts (Section 6.1.3), observation incompleteness (Section 6.1.4), and the sparsity of \mathbf{W} (Section 6.1.5). In the above-referenced experiments, we simulate the observation matrix \mathbf{Y} via the inverse probit link function and use only the probit variant of SPARFA-M in order to make a fair comparison with SPARFA-B. In a real-world situation, however, the link function is generally unknown. In Section 6.1.6 we conduct model-mismatch experiments, where we generate data from one link function but analyze assuming the other.

In all synthetic experiments, we average the results of all performance measures over 25 Monte-Carlo trials, limited primarily by the computational complexity of SPARFA-B, for each instance of the model parameters we control.

6.1.1 BASELINE ALGORITHM: K-SVD₊

Since we are not aware of any existing algorithms to solve (2) subject to the assumptions (A1)–(A3), we deploy a novel baseline algorithm based on the well-known K-SVD algorithm of Aharon et al. (2006), which is widely used in various dictionary learning settings but ignores the inverse probit or logit link functions. Since the standard K-SVD algorithm also ignores the non-negativity constraint used in the SPARFA model, we develop a variant of the non-negative K-SVD algorithm proposed in Aharon et al. (2005) that we refer to as K-SVD₊. In the sparse coding stage of K-SVD₊, we use the non-negative variant of orthogonal matching pursuit (OMP) outlined in Bruckstein et al. (2008); that is, we enforce the non-negativity constraint by iteratively picking the entry corresponding to the maximum inner product without taking its absolute value. We also solve a non-negative least-squares problem to determine the residual error for the next iteration. In the dictionary update stage of K-SVD₊, we use a variant of the rank-one approximation algorithm detailed in (Aharon et al., 2005, Figure 4), where we impose non-negativity on the elements in \mathbf{W} but not on the elements of \mathbf{C} .

K-SVD₊ has as input parameters the sparsity level of each row of \mathbf{W} . In what follows, we provide K-SVD₊ with the known *ground truth* for the number of non-zero components in order to obtain its best-possible performance. This will favor K-SVD₊ over both SPARFA algorithms, since, in practice, such oracle information is not available.

6.1.2 PERFORMANCE MEASURES

In each simulation, we evaluate the performance of SPARFA-M, SPARFA-B, and K-SVD₊ by comparing the fidelity of the estimates $\widehat{\mathbf{W}}$, $\widehat{\mathbf{C}}$, and $\widehat{\boldsymbol{\mu}}$ to the ground truth \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$. Performance evaluation is complicated by the facts that (i) SPARFA-B outputs posterior distributions rather than simple point estimates of the parameters and (ii) factor-analysis methods are generally susceptible to permutation of the latent factors. We address the first concern by post-processing the output of SPARFA-B to obtain point estimates for \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ as detailed in Section 4.3.3 using $\widehat{R}_{i,k} < 0.35$ for the threshold value. We address the second concern by normalizing the columns of \mathbf{W} , $\widehat{\mathbf{W}}$ and the rows of \mathbf{C} , $\widehat{\mathbf{C}}$ to unit ℓ_2 -norm, permuting the columns of $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{C}}$ to best match the ground truth, and then compare \mathbf{W} and \mathbf{C} with the estimates $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{C}}$. We also compute the Hamming distance between the support set of \mathbf{W} and that of the (column-permuted) estimate $\widehat{\mathbf{W}}$. To summarize, the performance measures used in the sequel are

$$\begin{aligned} E_{\mathbf{W}} &= \|\mathbf{W} - \widehat{\mathbf{W}}\|_F^2 / \|\mathbf{W}\|_F^2, & E_{\mathbf{C}} &= \|\mathbf{C} - \widehat{\mathbf{C}}\|_F^2 / \|\mathbf{C}\|_F^2, \\ E_{\boldsymbol{\mu}} &= \|\boldsymbol{\mu} - \widehat{\boldsymbol{\mu}}\|_2^2 / \|\boldsymbol{\mu}\|_2^2, & E_{\mathbf{H}} &= \|\mathbf{H} - \widehat{\mathbf{H}}\|_F^2 / \|\mathbf{H}\|_F^2, \end{aligned}$$

where $\mathbf{H} \in \{0, 1\}^{Q \times K}$ with $H_{i,k} = 1$ if $W_{i,k} > 0$ and $H_{i,k} = 0$ otherwise. The $Q \times K$ matrix $\widehat{\mathbf{H}}$ is defined analogously using $\widehat{\mathbf{W}}$.

6.1.3 IMPACT OF PROBLEM SIZE AND NUMBER OF CONCEPTS

In this experiment, we study the performance of SPARFA vs. KSVD₊ as we vary the number of learners N , the number of questions Q , and the number of concepts K .

Experimental setup We vary the number of learners N and the number of questions $Q \in \{50, 100, 200\}$, and the number of concepts $K \in \{5, 10\}$. For each combination of (N, Q, K) , we generate \mathbf{W} , \mathbf{C} , $\boldsymbol{\mu}$, and \mathbf{Y} according to (10) and (11) with $v_{\boldsymbol{\mu}} = 1$, $\lambda_k = 2/3 \forall k$, and $\mathbf{V}_0 = \mathbf{I}_K$. For each instance, we choose the number of non-zero entries in each row of \mathbf{W} as $DU(1, 3)$ where $DU(a, b)$ denotes the discrete uniform distribution in the range a to b . For each trial, we run the probit version of SPARFA-M, SPARFA-B, and K-SVD₊ to obtain the estimates $\widehat{\mathbf{W}}$, $\widehat{\mathbf{C}}$, $\widehat{\boldsymbol{\mu}}$, and calculate $\widehat{\mathbf{H}}$. For all of the synthetic experiments with SPARFA-M, we set the regularization parameters $\gamma = 0.1$ and select λ using the BIC (Hastie et al. 2010). For SPARFA-B, we set the hyperparameters to $h = K + 1$, $v_{\boldsymbol{\mu}} = 1$, $\alpha = 1$, $\beta = 1.5$, $e = 1$, and $f = 1.5$; moreover, we burn-in the MCMC for 30,000 iterations and take output samples over the next 30,000 iterations.

Results and discussion Figure 3 shows box-and-whisker plots for the three algorithms and the four performance measures. We observe that the performance of all of the algorithms generally improves as the problem size increases. Moreover, SPARFA-B has superior performance for $E_{\mathbf{W}}$, $E_{\mathbf{C}}$, and $E_{\boldsymbol{\mu}}$. We furthermore see that both SPARFA-B and SPARFA-M outperform K-SVD₊ on $E_{\mathbf{W}}$, $E_{\mathbf{C}}$, and especially $E_{\boldsymbol{\mu}}$. K-SVD₊ performs very well in terms of

$E_{\mathbf{H}}$ (slightly better than both SPARFA-M and SPARFA-B) due to the fact that we provide it with the oracle sparsity level, which is, of course, not available in practice. SPARFA-B’s improved estimation accuracy over SPARFA-M comes at the price of significantly higher computational complexity. For example, for $N = Q = 200$ and $K = 5$, SPARFA-B requires roughly 10 minutes on a 3.2 GHz quad-core desktop PC, while SPARFA-M and K-SVD₊ require only 6 s.

In summary, SPARFA-B is well-suited to small problems where solution accuracy or the need for confidence statistics are the key factors; SPARFA-M, in contrast, is destined for analyzing large-scale problems where low computational complexity (e.g., to generate immediate learner feedback) is important.

6.1.4 IMPACT OF THE NUMBER OF INCOMPLETE OBSERVATIONS

In this experiment, we study the impact of the number of observations in \mathbf{Y} on the performance of the probit version of SPARFA-M, SPARFA-B, and K-SVD₊.

Experimental setup We set $N = Q = 100$, $K = 5$, and all other parameters as in Section 6.1.3. We then vary the percentage P_{obs} of entries in \mathbf{Y} that are observed as 100%, 80%, 60%, 40%, and 20%. The locations of missing entries are generated i.i.d. and uniformly over the entire matrix.

Results and discussion Figure 4 shows that the estimation performance of all methods degrades gracefully as the percentage of missing observations increases. Again, SPARFA-B outperforms the other algorithms on $E_{\mathbf{W}}$, $E_{\mathbf{C}}$, and $E_{\boldsymbol{\mu}}$. K-SVD₊ performs worse than both SPARFA algorithms except on $E_{\mathbf{H}}$, where it achieves comparable performance. We conclude that SPARFA-M and SPARFA-B can both reliably estimate the underlying factors, even in cases of highly incomplete data.

6.1.5 IMPACT OF SPARSITY LEVEL

In this experiment, we study the impact of the sparsity level in \mathbf{W} on the performance of the probit version of SPARFA-M, SPARFA-B, and K-SVD₊.

Experimental setup We choose the active entries of \mathbf{W} i.i.d. $\text{Ber}(q)$ and vary $q \in \{0.2, 0.4, 0.6, 0.8\}$ to control the number of non-zero entries in each row of \mathbf{W} . All other parameters are set as in Section 6.1.3. This data-generation method allows for scenarios in which some rows of \mathbf{W} contain no active entries, as well as scenarios with all active entries. We set the hyperparameters for SPARFA-B to $h = K + 1 = 6$, $v_{\boldsymbol{\mu}} = 1$, and $e = 1$, and $f = 1.5$. For $q = 0.2$ we set $\alpha = 2$ and $\beta = 5$. For $q = 0.8$ we set $\alpha = 5$ and $\beta = 2$. For all other cases, we set $\alpha = \beta = 2$.

Results and discussion Figure 5 shows that sparser \mathbf{W} lead to lower estimation errors. This demonstrates that the SPARFA algorithms are well-suited to applications where the underlying factors have a high level of sparsity. SPARFA-B outperforms SPARFA-M across all metrics. The performance of K-SVD₊ is worse than both SPARFA algorithms except on the support estimation error $E_{\mathbf{H}}$, which is due to the fact that K-SVD₊ is aware of the oracle sparsity level.

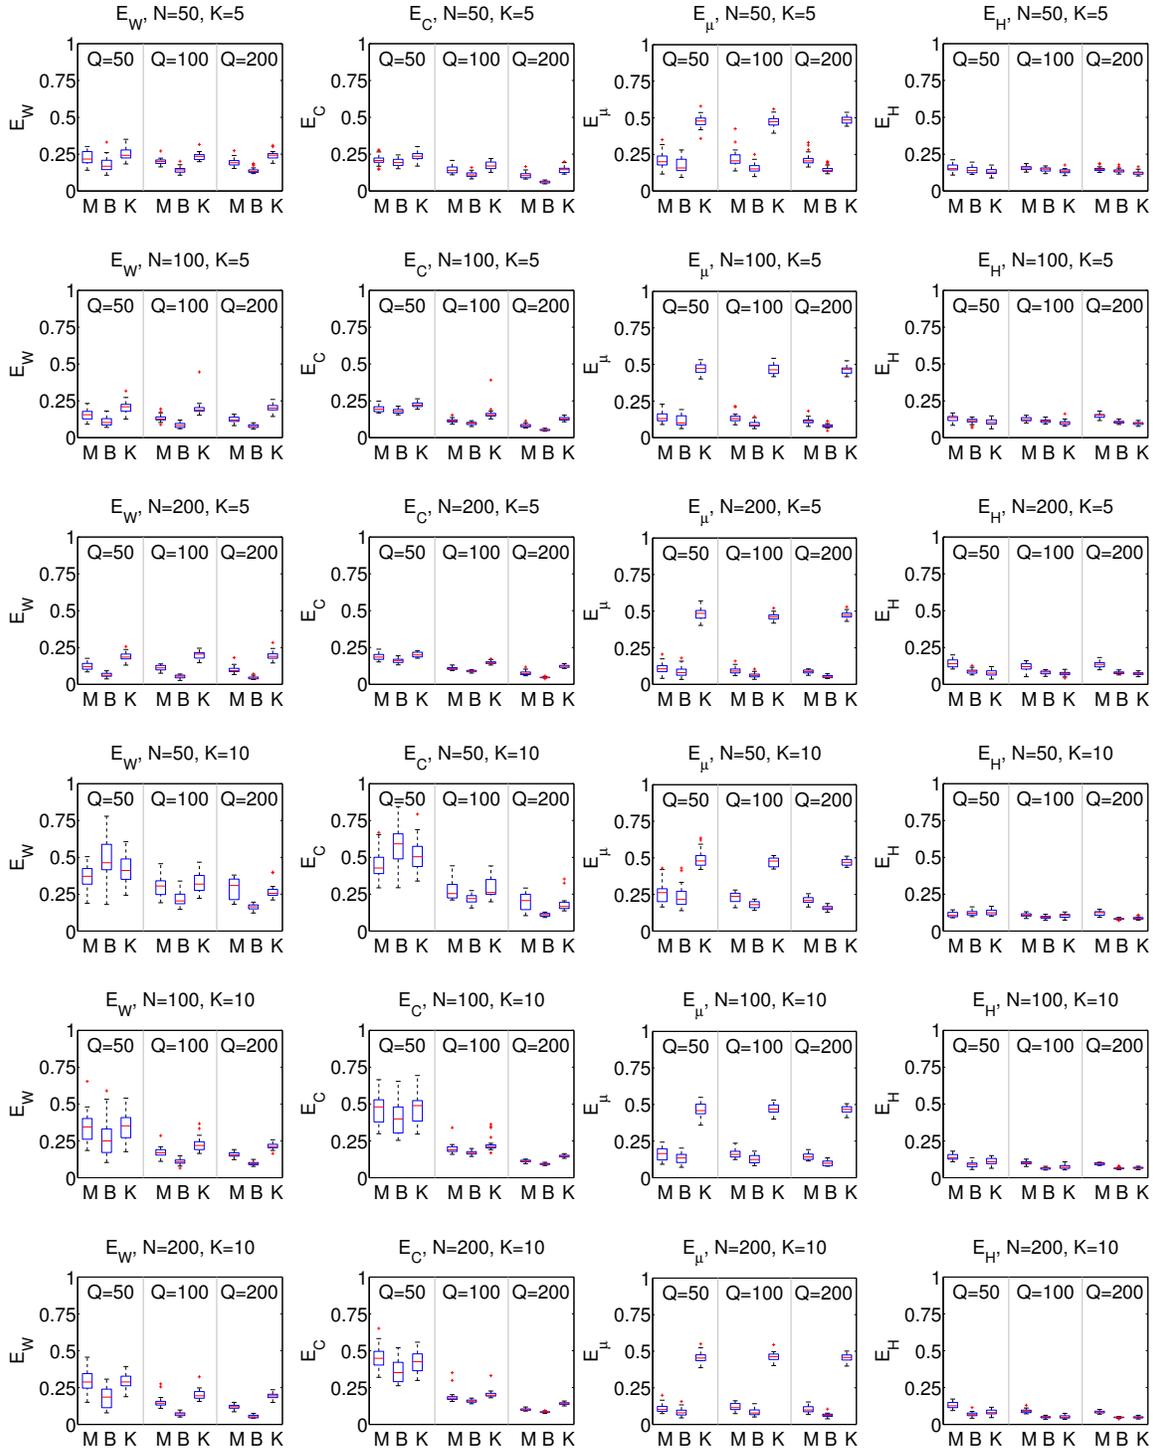


Figure 3: Performance comparison of SPARFA-M, SPARFA-B, and K-SVD₊ for different problem sizes $Q \times N$ and number of concepts K . The performance naturally improves as the problem size increases, while both SPARFA algorithms outperform K-SVD₊. M denotes SPARFA-M, B denotes SPARFA-B, and K denotes K-SVD₊.

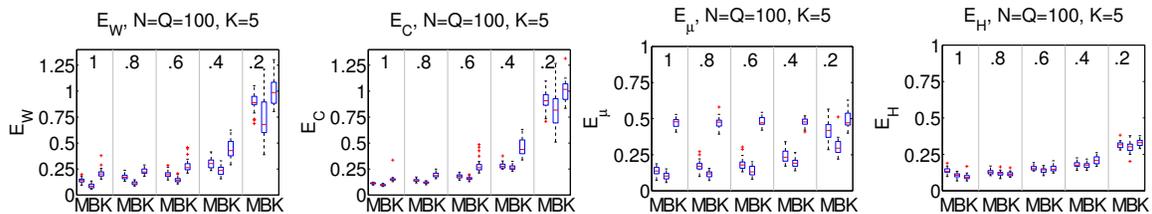


Figure 4: Performance comparison of SPARFA-M, SPARFA-B, and K-SVD₊ for different percentages of observed entries in \mathbf{Y} . The performance degrades gracefully as the number of observations decreases, while the SPARFA algorithms outperform K-SVD₊.

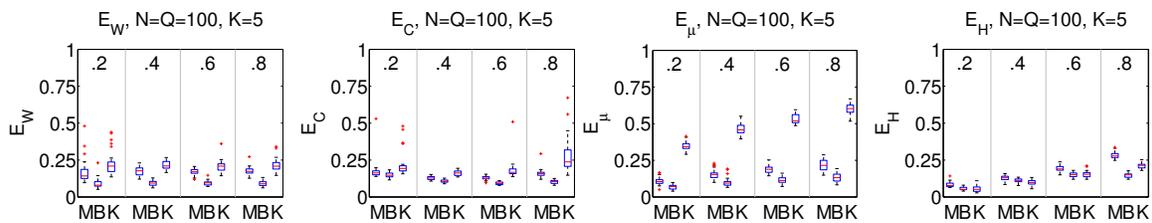


Figure 5: Performance comparison of SPARFA-M, SPARFA-B, and K-SVD₊ for different sparsity levels in the rows in \mathbf{W} . The performance degrades gracefully as the sparsity level increases, while the SPARFA algorithms outperform K-SVD₊.

6.1.6 IMPACT OF MODEL MISMATCH

In this experiment, we examine the impact of model mismatch by using a link function for estimation that does not match the true link function from which the data is generated.

Experimental setup We fix $N = Q = 100$ and $K = 5$, and set all other parameters as in Section 6.1.3. Then, for each generated instance of \mathbf{W} , \mathbf{C} and $\boldsymbol{\mu}$, we generate \mathbf{Y}_{pro} and \mathbf{Y}_{log} according to both the inverse probit link and the inverse logit link, respectively. We then run SPARFA-M (both the probit and logit variants), SPARFA-B (which uses only the probit link function), and K-SVD₊ on both \mathbf{Y}_{pro} and \mathbf{Y}_{log} .

Results and discussion Figure 6 shows that model mismatch does not severely affect $E_{\mathbf{W}}$, $E_{\mathbf{C}}$, and $E_{\mathbf{H}}$ for both SPARFA-M and SPARFA-B. However, due to the difference in the functional forms between the probit and logit link functions, model mismatch does lead to an increase in $E_{\boldsymbol{\mu}}$ for both SPARFA algorithms. We also see that K-SVD₊ performs worse than both SPARFA methods, because it ignores the link function.

6.2 Real Data Experiments

We next test the SPARFA algorithms on three real-world educational data sets. Since all variants of SPARFA-M and SPARFA-B obtained similar results in the synthetic data experiments in Section 6.1, for the sake of brevity, we will show the results for only one of the algorithms for each data set. In what follows, we select the sparsity penalty param-

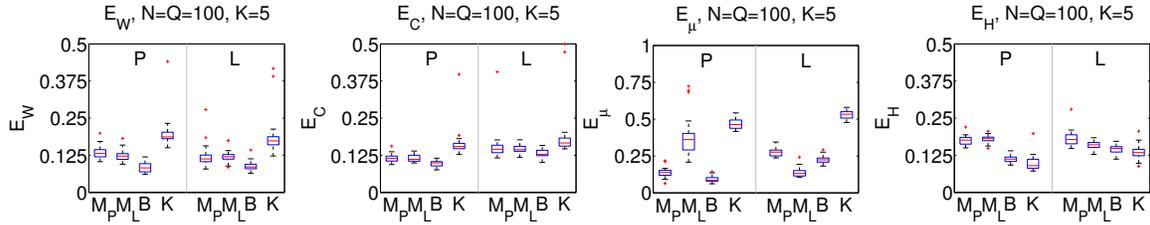


Figure 6: Performance comparison of SPARFA-M, SPARFA-B, and K-SVD₊ with probit/logit model mismatch; M_P and M_L indicate probit and logit SPARFA-M, respectively. In the left/right halves of each box plot, we generate \mathbf{Y} according to the inverse probit/logit link functions. The performance degrades only slightly with mismatch, while both SPARFA algorithms outperform K-SVD₊.

eter λ in SPARFA-M using the BIC as described in Hastie et al. (2010) and choose the hyperparameters for SPARFA-B to be largely non-informative.

6.2.1 UNDERGRADUATE DSP COURSE

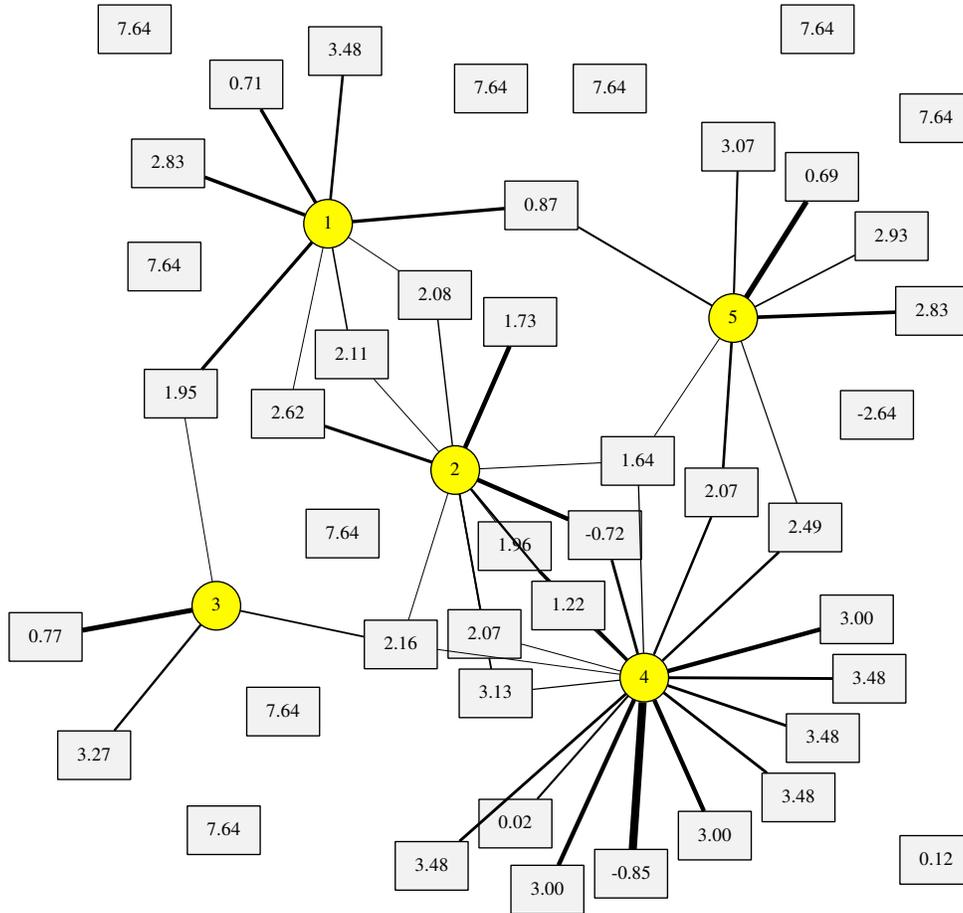
Data set We analyze a very small data set consisting of $N = 15$ learners answering $Q = 44$ questions taken from the final exam of an introductory course on digital signal processing (DSP) taught at Rice University in Fall 2011 (ELEC 301, Rice University 2011). There is no missing data in the matrix \mathbf{Y} .

Analysis We estimate \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ from \mathbf{Y} using the logit version of SPARFA-M assuming $K = 5$ concepts to achieve a concept granularity that matches the complexity of the analyzed data set. Since the questions had been manually tagged by the course instructor, we deploy the tag-analysis approach proposed in Section 5. Specifically, we form a 44×12 matrix \mathbf{T} using the $M = 12$ available tags and estimate the 12×5 concept-tag association matrix \mathbf{A} in order to interpret the meaning of each retrieved concept. For each concept, we only show the top 3 tags and their relative contributions. We also compute the 12×15 learner tag knowledge profile matrix \mathbf{U} .

Results and discussion Figure 7(a) visualizes the estimated question-concept association matrix $\widehat{\mathbf{W}}$ as a bipartite graph consisting of question and concept nodes.⁸ In the graph, circles represent the estimated concepts and squares represent questions, with thicker edges indicating stronger question-concept associations (i.e., larger entries $\widehat{W}_{i,k}$). Questions are also labeled with their estimated intrinsic difficulty μ_i , with larger positive values of μ_i indicating easier questions. Note that ten questions are not linked to any concept. All $Q = 15$ learners answered these questions correctly; as a result nothing can be estimated about their underlying concept structure. Figure 7(b) provides the concept-tag association (top 3 tags) for each of the 5 estimated concepts.

Table 1 provides Learner 1’s knowledge of the various tags relative to other learners. Large positive values mean that Learner 1 has strong knowledge of the tag, while large negative values indicate a deficiency in knowledge of the tag. Table 2 shows the average tag

8. To avoid the scaling identifiability problem that is typical in factor analysis, we normalize each row of \mathbf{C} to unit ℓ_2 -norm and scale each column of \mathbf{W} accordingly prior to visualizing the bipartite graph. This enables us to compare the strength of question-concept associations across different concepts.



(a) Question–concept association graph. Circles correspond to concepts and rectangles to questions; the values in each rectangle corresponds to that question’s intrinsic difficulty.

Concept 1		Concept 2		Concept 3	
Frequency response	(46%)	Fourier transform	(40%)	z -transform	(66%)
Sampling rate	(23%)	Laplace transform	(36%)	Pole/zero plot	(22%)
Aliasing	(21%)	z -transform	(24%)	Laplace transform	(12%)
Concept 4		Concept 5			
Fourier transform	(43%)	Impulse response	(74%)		
Systems/circuits	(31%)	Transfer function	(15%)		
Transfer function	(26%)	Fourier transform	(11%)		

(b) Most important tags and relative weights for the estimated concepts.

Figure 7: (a) Question–concept association graph and (b) most important tags associated with each concept for an undergraduate DSP course with $N = 15$ learners answering $Q = 44$ questions.

z -transform	Impulse response	Transfer function	Fourier transform	Laplace transform
1.09	-1.80	-0.50	0.99	-0.77

Table 1: Selected tag knowledge of Learner 1.

z -transform	Impulse response	Transfer function	Fourier transform	Laplace transform
0.04	-0.03	-0.10	0.11	0.03

Table 2: Average tag knowledge of all learners.

knowledge of the entire class, computed by averaging the entries of each row in the learner tag knowledge matrix \mathbf{U} as described in Section 5.2.2. Table 1 indicates that Learner 1 has particularly weak knowledge of the tag “Impulse response.” Armed with this information, a PLS could automatically suggest remediation about this concept to Learner 1. Table 2 indicates that the entire class has (on average) weak knowledge of the tag “Transfer function.” With this information, a PLS could suggest to the class instructor that they provide remediation about this concept to the entire class.

6.2.2 GRADE 8 SCIENCE COURSE

Data set The STEMscopes data set was introduced in Section 1.2. There is substantial missing data in the matrix \mathbf{Y} , with only 13.5% of its entries observed.

Analysis We compare the results of SPARFA-M and SPARFA-B on this data set to highlight the pros and cons of each approach. For both algorithms, we select $K = 5$ concepts. For SPARFA-B, we fix reasonably broad (non-informative) values for all hyperparameters. For μ_0 we calculate the average rate of correct answers p_s on observed graded responses of all learners to all questions and use $\mu_0 = \Phi_{\text{pro}}^{-1}(p_s)$. The variance v_{μ} is left sufficiently broad to enable adequate exploration of the intrinsic difficulty for each questions. Point estimates of \mathbf{W} , \mathbf{C} and $\boldsymbol{\mu}$ are generated from the SPARFA-B posterior distributions using the methods described in Section 4.3.3. Specifically, an entry $\widehat{W}_{i,k}$ that has a corresponding active probability $\widehat{R}_{i,k} < 0.55$ is thresholded to 0. Otherwise, we set $\widehat{W}_{i,k}$ to its posterior mean. On a 3.2 GHz quad-core desktop PC, SPARFA-M converged to its final estimates in 4s, while SPARFA-B required 10 minutes.

Results and discussion Both SPARFA-M and SPARFA-B deliver comparable factorizations. The estimated question–concept association graph for SPARFA-B is shown in Figure 2(a), with the accompanying concept–tag association in Figure 2(b). Again we see a sparse relationship between questions and concepts. The few outlier questions that are not associated with any concept are generally those questions with very low intrinsic difficulty or those questions with very few responses.

One advantage of SPARFA-B over SPARFA-M is its ability to provide not only point estimates of the parameters of interest but also reliability information for those estimates. This reliability information can be useful for decision making, since it enables one to tailor

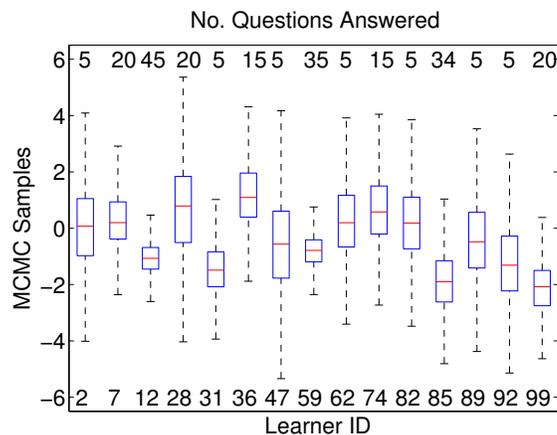


Figure 8: Concept 5 knowledge estimates generated by SPARFA-B for the STEMscopes data for a randomly selected subset of learners. The box-whisker plot shows the posterior variance of the MCMC samples, with each box-whisker plot corresponding to a different learner in the data set. Anonymized learner IDs are shown on the bottom, while the number of relevant questions answered by each learner answered is indicated on the top of the plot.

actions according to the associated uncertainty. If there is considerable uncertainty regarding learner mastery of a particular concept, for example, it may be a more appropriate use of time of the learner to ask additional questions that reduce the uncertainty, rather than assigning new material for which the learner may not be adequately prepared.

We demonstrate the utility of SPARFA-B’s posterior distribution information on the learner concept knowledge matrix \mathbf{C} . Figure 8 shows box-whisker plots of the MCMC output samples over 30,000 iterations (after a burn-in period of 30,000 iterations) for a set of learners for Concept 5. Each box-whisker plot corresponds to the posterior distribution for a different learner. These plots enable us to visualize both the posterior mean and variance associated with the concept knowledge estimates \hat{c}_j . As one would expect, the estimation variance tends to decrease as the number of answered questions increases (shown in the top portion of Figure 8).

The exact set of questions answered by a learner also affects the posterior variance of our estimate, as different questions convey different levels of information regarding a learner’s concept mastery. An example of this phenomenon is observed by comparing Learners 7 and 28. Each of these two learners answered 20 questions and had a nearly equal number of correct answers (16 and 17, respectively). A conventional analysis that looked only at the percentage of correct answers would conclude that both learners have similar concept mastery. However, the actual set of questions answered by each learner is not the same, due to their respective instructors assigning different questions. While SPARFA-B finds a similar posterior mean for Learner 7 and Learner 28, it finds very different posterior variances, with considerably more variance for Learner 28. The SPARFA-B posterior samples shed additional light on the situation at hand. Most of the questions answered by Learner 28 are deemed easy (defined as having intrinsic difficulties $\hat{\mu}_i$ larger than one). Moreover, the remaining, more difficult questions answered by Learner 28 show stronger affinity to concepts

		C1	C2	C3	C4	C5
Q3 (27 responses)	M	Yes	No	No	No	Yes
	B	94%	36%	48%	18%	80%
Q56 (5 responses)	M	No	No	No	No	No
	B	30%	30%	26%	31%	31%
Q72 (6 responses)	M	No	No	No	No	Yes
	B	61%	34%	29%	36%	58%

Table 3: Comparison of SPARFA-M and SPARFA-B for three selected questions and the $K = 5$ estimated concepts in the STEMscopes data set. For SPARFA-M, the labels “Yes” and “No” indicate whether a particular concept was detected in the question. For SPARFA-B, we show the posterior inclusion probability (in percent), which indicates the percentage of iterations in which a particular concept was sampled.

other than Concept 5. In contrast, roughly half of the questions answered by Learner 7 are deemed hard and all of these questions have stronger affinity to Concept 5. Thus, the questions answered by Learner 28 convey only weak information about the knowledge of Concept 5, while those answered by Learner 7 convey strong information. Thus, we cannot determine from Learner 28’s responses whether they have mastered Concept 5 well or not. Such SPARFA-B posterior data would enable a PLS to quickly assess this scenario and tailor the presentation of future questions to Learner 28—in this case, presenting more difficult questions related to Concept 5 would reduce the estimation variance on their concept knowledge and allow a PLS to better plan future educational tasks for this particular learner.

Second, we demonstrate the utility of SPARFA-B’s posterior distribution information on the question–concept association matrix \mathbf{W} . Accurate estimation of \mathbf{W} enables course instructors and content authors to validate the extent to which problems measure knowledge across various concepts. In general, there is a strong degree of commonality between the results of SPARFA-M and SPARFA-B, especially as the number of learners answering a question grow. We present some illustrative examples of support estimation on \mathbf{W} for both SPARFA algorithms in Table 3. We use the labels “Yes”/“No” to indicate inclusion of a concept by SPARFA-M and show the posterior inclusion probabilities for each concept by SPARFA-B. Here, both SPARFA-M and SPARFA-B agree strongly on both Question 3 and Question 56. Question 72 is answered by only 6 learners, and SPARFA-M discovers a link between this question and Concept 5. SPARFA-B proposes Concept 5 in 58% of all MCMC iterations, but also Concept 1 in 60% of all MCMC iterations. Furthermore, the proposals of Concept 1 and Concept 5 are nearly mutually exclusive; in most iterations only one of the two concepts is proposed, but both are rarely proposed jointly. This behavior implies that SPARFA-B has found two competing models that explain the data associated with Question 72. To resolve this ambiguity, a PLS would need to gather more learner responses.

6.2.3 ALGEBRA TEST ADMINISTERED ON AMAZON MECHANICAL TURK

For a final demonstration of the capabilities the SPARFA algorithms, we analyze a data set from a high school algebra test carried out by Daniel Calderón of Rice University on Amazon Mechanical Turk, a crowd-sourcing marketplace (Amazon Mechanical Turk 2012). *Data set* The data set consists of $N = 99$ learners answering $Q = 34$ questions covering topics such as geometry, equation solving, and visualizing function graphs. Calderón manually labeled the questions from a set of $M = 10$ tags. The data set is fully populated, with no missing entries.

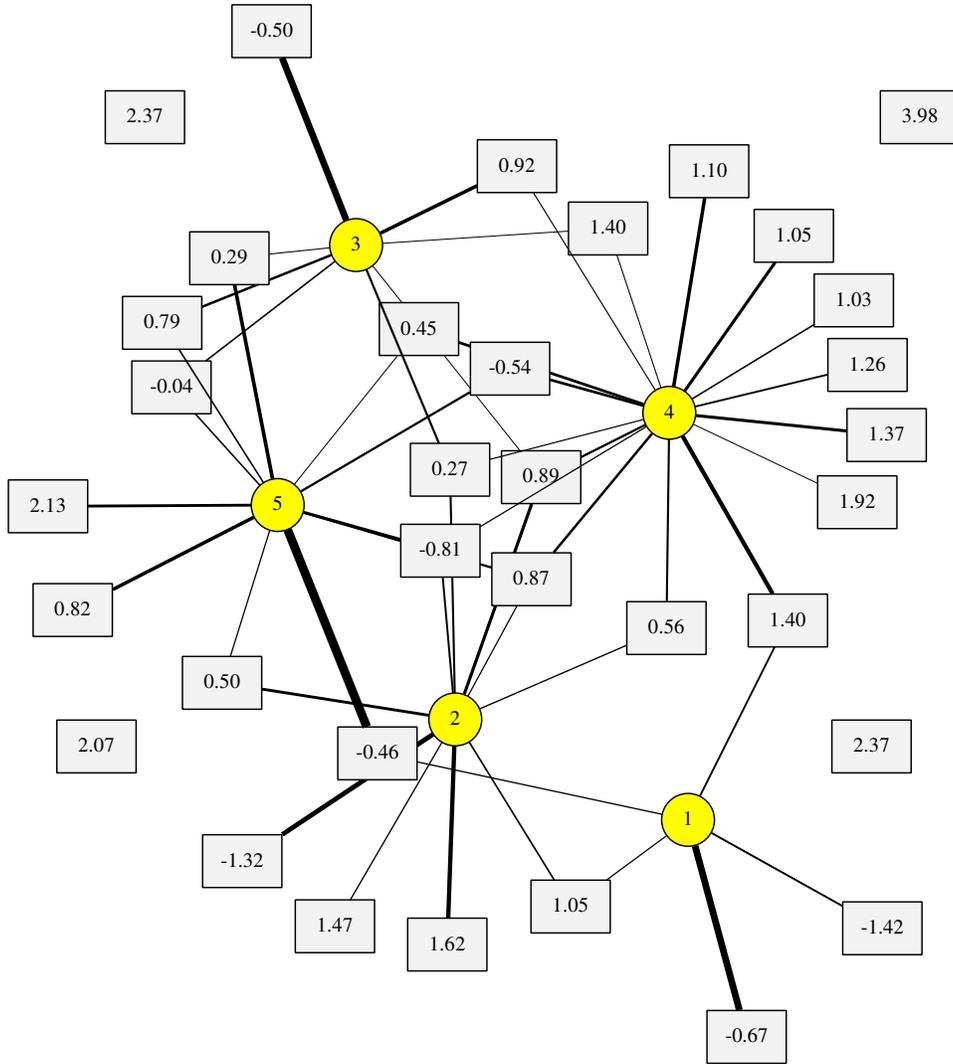
Analysis We estimate \mathbf{W} , \mathbf{C} , and $\boldsymbol{\mu}$ from the fully populated 34×99 binary-valued matrix \mathbf{Y} using the logit version of SPARFA-M assuming $K = 5$ concepts. We deploy the tag-analysis approach proposed in Section 5 to interpret each concept. Additionally, we calculate the likelihoods of the responses using (1) and the estimates $\widehat{\mathbf{W}}$, $\widehat{\mathbf{C}}$ and $\widehat{\boldsymbol{\mu}}$. The results from SPARFA-M are summarized in Figure 9. We detail the results of our analysis for Questions 19–26 in Table 4 and for Learner 1 in Table 5.

Results and discussion With the aid of SPARFA, we can analyze the strengths and weaknesses of each learner’s concept knowledge both individually and relative to other users. We can also detect outlier responses that are due to guessing, cheating, or carelessness. The values in the estimated concept knowledge matrix $\widehat{\mathbf{C}}$ measure each learner’s concept knowledge relative to all other learners. The estimated intrinsic difficulties of the questions $\widehat{\boldsymbol{\mu}}$ provide a relative measure that summarizes how all users perform on each question.

Let us now consider an example in detail; see Table 4 and Table 5. Learner 1 incorrectly answered Questions 21 and 26 (see Table 4), which involve Concepts 1 and 2. Their knowledge of these concepts is not heavily penalized, however (see Table 5), due to the high intrinsic difficulty of these two questions, which means that most other users also incorrectly answered them. User 1 also incorrectly answered Questions 24 and 25, which involve Concepts 2 and 4. Their knowledge of these concepts is penalized, due to the low intrinsic difficulty of these two questions, which means that most other users correctly answered them. Finally, Learner 1 correctly answered Questions 19 and 20, which involve Concepts 1 and 5. Their knowledge of these concepts is boosted, due to the high intrinsic difficulty of these two questions.

SPARFA can also be used to identify each user’s individual strengths and weaknesses. Continuing the example, Learner 1 needs to improve their knowledge of Concept 4 (associated with the tags “Simplifying expressions”, “Trigonometry,” and “Plotting functions”) significantly, while their deficiencies on Concepts 2 and 3 are relatively minor.

Finally, by investigating the likelihoods of the graded responses, we can detect outlier responses, which would enables a PLS to detect guessing and cheating. By inspecting the concept knowledge of Learner 1 in Table 5, we can identify insufficient knowledge of Concept 4. Hence, Learner 1’s correct answer to Question 22 is likely due to a random guess, since the predicted likelihood of providing the correct answer is estimated at only 0.21.



(a) Question–concept association graph.

Concept 1		Concept 2		Concept 3	
Fractions	(57%)	Plotting functions	(64%)	Geometry	(63%)
Solving equations	(42%)	System of equations	(27%)	Simplifying expressions	(27%)
Arithmetic	(1%)	Simplifying expressions	(9%)	Trigonometry	(10%)
Concept 4		Concept 5			
Simplifying expressions	(64%)	Trigonometry	(53%)		
Trigonometry	(21%)	Slope	(40%)		
Plotting Functions	(15%)	Solving equations	(7%)		

(b) Most important tags and relative weights for the estimated concepts.

Figure 9: (a) Question–concept association graph and (b) most important tags associated with each concept for a high-school algebra test carried out on Amazon Mechanical Turk with $N = 99$ users answering $Q = 34$ questions.

Question number	19	20	21	22	23	24	25	26
Learner’s graded response $Y_{i,j}$	1	1	0	1	1	0	0	0
Correct answer likelihood $p(Y_{i,j} = 1 \bar{\mathbf{w}}_i, \mathbf{c}_j, \mu_i)$	0.79	0.71	0.11	0.21	0.93	0.23	0.43	0.00
Underlying concepts	1	1, 5	1	2, 3, 4	3, 5	2, 4	1, 4	2, 4
Intrinsic difficulty μ_i	-1.42	-0.46	-0.67	0.27	0.79	0.56	1.40	-0.81

Table 4: Graded responses and their underlying concepts for Learner 1 (1 designates a correct response and 0 an incorrect response).

Concept number	1	2	3	4	5
Concept knowledge	0.46	-0.35	0.72	-1.67	0.61

Table 5: Estimated concept knowledge for Learner 1.

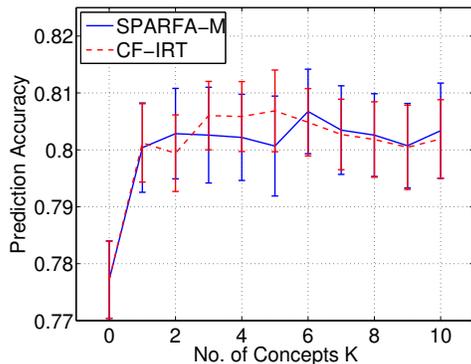
6.3 Predicting Unobserved Learner Responses

We now compare SPARFA-M against the recently proposed binary-valued collaborative filtering algorithm CF-IRT (Bergner et al. 2012) in an experiment to predict unobserved learner responses.

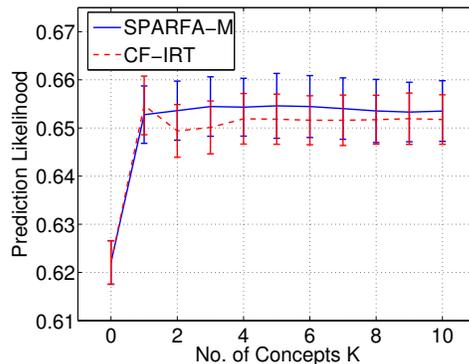
Data set and experimental setup In this section, we study both the Mechanical Turk algebra test data set and a portion of the ASSISTment data set (Pardos and Heffernan 2010). The ASSISTment data set consists of $N = 403$ learners answering $Q = 219$ questions, with 25% of the responses observed (see Vats et al. 2013 for additional details on the data set). In each of the 25 trials we run for both data sets, we hold out 20% of the observed learner responses as a test set, and train both the logistic variant of SPARFA-M⁹ and CF-IRT on the rest. The regularization parameters of both algorithms are selected using 4-fold cross-validation on the training set. We use two performance metrics to evaluate the performance of these algorithms, namely (i) the prediction accuracy, which corresponds to the percentage of correctly predicted unobserved responses, and (ii) the average prediction likelihood $\frac{1}{|\bar{\Omega}_{\text{obs}}|} \sum_{i,j:(i,j) \in \bar{\Omega}_{\text{obs}}} p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j)$ of the unobserved responses, as proposed in González-Brenes and Mostow (2012), for example.

Results and discussion Figure 10 shows the prediction accuracy and prediction likelihood for both the Mechanical Turk algebra test data set and the ASSISTment data set. We see that SPARFA-M delivers comparable (sometimes slightly superior) prediction performance to CF-IRT in predicting unobserved learner responses.

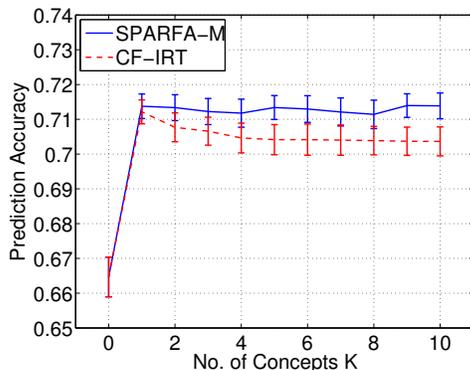
9. In order to arrive at a fair comparison, we choose to use the logistic variant of SPARFA-M, since CF-IRT also relies on a logistic model.



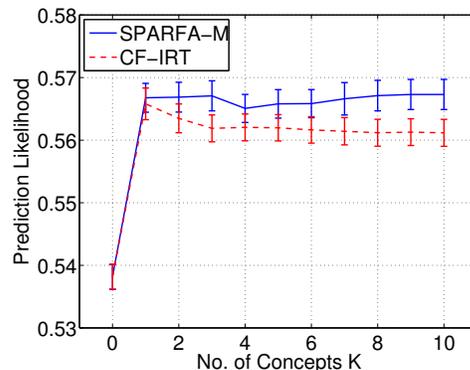
(a) Prediction accuracy for the Mechanical Turk algebra test data set.



(b) Average prediction likelihood for the Mechanical Turk algebra test data set.



(c) Prediction accuracy for the ASSISTment data set.



(d) Average prediction likelihood for the ASSISTment data set.

Figure 10: Performance comparison of SPARFA-M and CF-IRT on (a) prediction accuracy and (b) average prediction likelihood for the Mechanical Turk algebra test data set, (c) prediction accuracy and (d) average prediction likelihood for the ASSISTment data set. SPARFA-M achieves comparable or better performance than CF-IRT while enabling interpretability of the estimated latent concepts

Furthermore, we see from Figure 10 that the prediction performance varies little over different values of K , meaning that the specific choice of K has little influence on the prediction performance within a certain range. This phenomenon agrees with other collaborative filtering results (see, e.g., Koren et al. 2009; Koren and Sill 2011). Consequently, the choice of K essentially dictates the granularity of the abstract concepts we wish to estimate. We choose $K = 5$ in the real data experiments of Section 6.2 when we visualize the question–concept associations as bi-partite graphs, as it provides a desirable granularity of the estimated concepts in the data sets. We emphasize that SPARFA-M is able to provide interpretable estimated factors while achieving comparable (or slightly superior) prediction performance than that achieved by CF-IRT, which does not provide interpretability. This

feature of SPARFA is key for the development of PLSs, as it enables an automated way of generating interpretable feedback to learners in a purely data-driven fashion.

7. Related Work on Machine Learning-based Personalized Learning

A range of different machine learning algorithms have been applied in educational contexts. Bayesian belief networks have been successfully used to probabilistically model and analyze learner response data (e.g., Krudysz et al. 2006; Woolf 2008; Krudysz and McClellan 2011). Such models, however, rely on predefined question–concept dependencies (that are not necessarily the true dependencies governing learner responses) and primarily only work for a single concept. In contrast, SPARFA discovers question–concept dependencies from solely the graded learner responses to questions and naturally estimates multi-concept question dependencies.

Modeling question–concept associations has been studied in Barnes (2005), Thai-Nghe et al. (2011a), Thai-Nghe et al. (2011b), and Desmarais (2011). The approach in Barnes (2005) characterizes the underlying question–concept associations using binary values, which ignore the relative strengths of the question–concept associations. In contrast, SPARFA differentiates between strong and weak relationships through the real-valued weights $W_{i,k}$. The matrix and tensor factorization methods proposed in Barnes (2005), Thai-Nghe et al. (2011a), and Thai-Nghe et al. (2011b) treat graded learner responses as real but deterministic values. In contrast, the probabilistic framework underlying SPARFA provides a statistically principled model for graded responses; the likelihood of the observed graded responses provides even more explanatory power.

Existing intelligent tutoring systems capable of modeling question–concept relations probabilistically include Khan Academy (Dijksman and Khan 2011; Hu 2011) and the system of Bachrach et al. (2012). Both approaches, however, are limited to dealing with a single concept. In contrast, SPARFA is built from the ground up to deal with multiple latent concepts.

A probit model for graded learner responses is used in Desmarais (2011) without exploiting the idea of low-dimensional latent concepts. In contrast, SPARFA leverages multiple latent concepts and therefore can create learner concept knowledge profiles for personalized feedback. Moreover, SPARFA-M is compatible with the popular logit model.

The recent results developed in Beheshti et al. (2012) and Bergner et al. (2012) address the problem of *predicting* the missing entries in a binary-valued graded learner response matrix. Both papers use low-dimensional latent factor techniques specifically developed for collaborative filtering, as, e.g., discussed in Linden et al. (2003) and Herlocker et al. (2004). While predicting missing correctness values is an important task, these methods do not take into account the sparsity and non-negativity of the matrix \mathbf{W} ; this inhibits the interpretation of the relationships among questions and concepts. In contrast, SPARFA accounts for both the sparsity and non-negativity of \mathbf{W} , which enables the interpretation of the value $C_{k,j}$ as learner j 's knowledge of concept k .

There is a large body of work on item response theory (IRT), which uses statistical models to analyze and score graded question response data (see, e.g., Lord 1980, Baker and Kim 2004, and Reckase 2009 for overview articles). The main body of the IRT literature builds on the model developed by Rasch (1993) and has been applied mainly in the

context of adaptive testing, e.g., in the graduate record examination (GRE) and graduate management (GMAT) tests (Chang and Ying 2009, Thompson 2009, and Linacre 1999). While the SPARFA model shares some similarity to the model in Rasch (1993) by modeling question–concept association strengths and intrinsic difficulties of questions, it also models each learner in terms of a multi-dimensional concept knowledge vector. This capability of SPARFA is in stark contrast to the Rasch model, where each learner is characterized by a single, scalar ability parameter. Consequently, the SPARFA framework is able to provide stronger explanatory power in the estimated factors compared to that of the conventional Rasch model. We finally note that multi-dimensional variants of IRT have been proposed in McDonald (2000), Yao (2003), and Reckase (2009). We emphasize, however, that the design of these algorithms leads to poor interpretability of the resulting parameter estimates.

8. Conclusions

In this paper, we have formulated a new approach to learning and content analytics, which is based on a new statistical model that encodes the probability that a learner will answer a given question correctly in terms of three factors: (i) the learner’s knowledge of a set of latent concepts, (ii) how the question relates to each concept, and (iii) the intrinsic difficulty of the question. We have proposed two algorithms, SPARFA-M and SPARFA-B, to estimate the above three factors given incomplete observations of graded learner question responses. SPARFA-M uses an efficient Maximum Likelihood-based bi-convex optimization approach to produce point estimates of the factors, while SPARFA-B uses Bayesian factor analysis to produce posterior distributions of the factors. In practice, SPARFA-M is beneficial in applications where timely results are required; SPARFA-B is favored in situations where posterior statistics are required. We have also introduced a novel method for incorporating user-defined tags on questions to facilitate the interpretability of the estimated factors. Experiments with both synthetic and real world education data sets have demonstrated both the efficacy and robustness of the SPARFA algorithms.

The quantities estimated by SPARFA can be used directly in a range of PLS functions. For instance, we can identify the knowledge level of learners on particular concepts and diagnose why a given learner has incorrectly answered a particular question or type of question. Moreover, we can discover the hidden relationships among questions and latent concepts, which is useful for identifying questions that do and do not aid in measuring a learner’s conceptual knowledge. Outlier responses that are either due to guessing or cheating can also be detected. In concert, these functions can enable a PLS to generate personalized feedback and recommendation of study materials, thereby enhancing overall learning efficiency.

Various extensions and refinements to the SPARFA framework developed here have been proposed recently. Most of these results aim at improving interpretability of the SPARFA model parameters. In particular, a variant of SPARFA-M that analyzes ordinal rather than binary-valued responses and directly utilizes tag information in the probabilistic model has been detailed in Lan et al. (2013a). Another variant of SPARFA-M that further improves the interpretability of the underlying concepts via the joint analysis of graded learner responses and question/response text has been proposed in Lan et al. (2013b). A nonparametric Bayesian variant of SPARFA-B that estimates both the number of concepts K as well as the

reliability of each learner from data has been developed in Fronczyk et al. (2013, submitted). The results of this nonparametric method confirm our choice of $K = 5$ concepts for the real-world educational data sets considered in Section 6.2.

Before closing, we would like to point out a connection between SPARFA and *dictionary learning* that is of independent interest. This connection can be seen by noting that (2) for both the probit and inverse logit functions is statistically equivalent to (see Rasmussen and Williams 2006):

$$Y_{i,j} = [\text{sign}(\mathbf{WC} + \mathbf{M} + \mathbf{N})]_{i,j}, \quad (i, j) \in \Omega_{\text{obs}},$$

where $\text{sign}(\cdot)$ denotes the entry-wise sign function and the entries of \mathbf{N} are i.i.d. and drawn from either a standard Gaussian or standard logistic distribution. Hence, estimating \mathbf{W} , \mathbf{C} , and \mathbf{M} (or equivalently, $\boldsymbol{\mu}$) is equivalent to learning a (possibly overcomplete) dictionary from the data \mathbf{Y} . The key departures from the dictionary-learning literature (Aharon et al. 2006; Mairal et al. 2010) and algorithm variants capable of handling missing observations (Studer and Baraniuk 2012) are the binary-valued observations and the non-negativity constraint on \mathbf{W} . Note that the algorithms developed in Section 3 to solve the sub-problems by holding one of the factors \mathbf{W} or \mathbf{C} fixed and solving for the other variable can be used to solve noisy binary-valued (or 1-bit) compressive sensing or sparse signal recovery problems, e.g., as studied in Boufounos and Baraniuk (2008), Jacques et al. (2013), and Plan and Vershynin (2012, submitted). Thus, the proposed SPARFA algorithms can be applied to a wide range of applications beyond education, including the analysis of survey data, voting patterns, gene expression, and signal recovery from noisy 1-bit compressive measurements.

Acknowledgments

Thanks to Wotao Yin and Yangyang Xu for helpful discussions on the convergence proof of SPARFA-M, Genevera Allen for insights into probit regression, Marina Vannucci for helpful discussions regarding Bayesian factor analysis, Daniel Calderón for organizing and administering the Amazon Mechanical Turk experiments in Section 6.2.3, and Carlos Monroy and Reid Whitaker for providing the STEMscopes data. We furthermore thank the anonymous reviewers for their valuable comments which improved the exposition of our results.

This work was supported by the National Science Foundation under Cyberlearning grant IIS-1124535, the Air Force Office of Scientific Research under grant FA9550-09-1-0432, the Google Faculty Research Award program.

Please see our website www.sparfa.com, where you can learn more about the project and purchase SPARFA t-shirts and other merchandise.

Appendix A. Proof of Theorem 1

We now establish the convergence of the FISTA algorithms that solve the SPARFA-M sub-problems $(\text{RR})_1^+$ and $(\text{RR})_2$. We start by deriving the relevant Lipschitz constants.

Lemma 4 (Scalar Lipschitz constants) *Let $g_{\text{pro}}(x) = \frac{\Phi'_{\text{pro}}(x)}{\Phi_{\text{pro}}(x)}$ and $g_{\text{log}}(x) = \frac{\Phi'_{\text{log}}(x)}{\Phi_{\text{log}}(x)}$, $x \in \mathbb{R}$, where $\Phi_{\text{pro}}(x)$ and $\Phi_{\text{log}}(x)$ are the inverse probit and logit link functions defined*

in (3) and (4), respectively. Then, for $y, z \in \mathbb{R}$ we have

$$|g_{pro}(y) - g_{pro}(z)| \leq L_{pro}|y - z|, \tag{12}$$

$$|g_{log}(y) - g_{log}(z)| \leq L_{log}|y - z|, \tag{13}$$

with the constants $L_{pro} = 1$ for the probit case and $L_{log} = 1/4$ for the logit case.

Proof For simplicity of exposition, we omit the subscripts designating the probit and logit cases in what follows. We first derive L_{pro} for the probit case by computing the derivative of $g(x)$ and bounding its derivative from below and above. The derivative of $g(x)$ is given by

$$g'(x) = -\frac{\mathcal{N}(x)}{\Phi(x)} \left(x + \frac{\mathcal{N}(x)}{\Phi(x)} \right). \tag{14}$$

where $\mathcal{N}(t) = \frac{1}{\sqrt{2\pi}}e^{-t^2/2}$ is the PDF of the standard normal distribution.

We first bound this derivative for $x \leq 0$. To this end, we individually bound the first and second factor in (14) using the following bounds listed in Olver (2010):

$$-\frac{x}{2} + \sqrt{\frac{x^2}{4} + \frac{2}{\pi}} \leq \frac{\mathcal{N}(x)}{\Phi(x)} \leq -\frac{x}{2} + \sqrt{\frac{x^2}{4} + 1}, \quad x \leq 0$$

and

$$\frac{x}{2} + \sqrt{\frac{x^2}{4} + \frac{2}{\pi}} \leq x + \frac{\mathcal{N}(x)}{\Phi(x)} \leq \frac{x}{2} + \sqrt{\frac{x^2}{4} + 1}, \quad x \leq 0.$$

Multiplying the above inequalities leads to the bounds

$$-1 \leq g'(x) \leq -\frac{2}{\pi}, \quad x \leq 0. \tag{15}$$

We next bound the derivative of (14) for $x > 0$. For $x > 0$, $\mathcal{N}(x)$ is a positive decreasing function and $\Phi(x)$ is a positive increasing function; hence $\frac{\mathcal{N}(x)}{\Phi(x)}$ is a decreasing function and $\frac{\mathcal{N}(x)}{\Phi(x)} \leq \frac{\mathcal{N}(0)}{\Phi(0)} = \sqrt{2/\pi}$. Thus, we arrive at

$$g'(x) = -\frac{\mathcal{N}(x)}{\Phi(x)} \left(x + \frac{\mathcal{N}(x)}{\Phi(x)} \right) \geq -\frac{\mathcal{N}(x)}{\Phi(x)} \left(x + \sqrt{2/\pi} \right),$$

where we have used the facts that $\Phi(x) \geq 1/2$ and $\mathcal{N}(x) \leq \frac{1}{\sqrt{2\pi}}$ for $x > 0$. According to (3) and the bound of Chu (1955), we have

$$\Phi(x) = \frac{1}{2} + \int_0^x \mathcal{N}(t | 0, 1) dt \geq \frac{1}{2} + \frac{1}{2} \sqrt{1 - e^{-x^2/2}} \geq 1 - \frac{1}{2} e^{-x^2/2}, \tag{16}$$

where the second inequality follows from the fact that $(1 - e^{-x^2/2}) \in [0, 1]$. Using (16) we can further bound $g'(x)$ from below as

$$g'(x) \geq -\frac{\mathcal{N}(x)}{1 - \frac{1}{2}e^{-x^2/2}} \left(x + \sqrt{2/\pi} \right).$$

Let us now assume that

$$-\frac{\mathcal{N}(x)}{1 - \frac{1}{2}e^{-x^2/2}} \left(x + \sqrt{2/\pi}\right) \geq -1.$$

In order to prove that this assumption is true, we rearrange terms to obtain

$$\left(\frac{x}{\sqrt{2\pi}} + (1/\pi + 1/2)\right) e^{-x^2/2} \leq 1. \tag{17}$$

Now we find the maximum of the LHS of (17) for $x > 0$. To this end, we observe that $\frac{x}{\sqrt{2\pi}} + (1/\pi + 1/2)$ is monotonically increasing and that $e^{-x^2/2}$ monotonically decreasing for $x > 0$; hence, this function has a unique maximum in this region. By taking its derivative and setting it to zero, we obtain

$$x^2 + \sqrt{2/\pi} + \sqrt{\pi/2} - 1 = 0$$

Substituting the result of this equation, i.e., $\hat{x} \approx 0.4068$, into (17) leads to

$$\left(\frac{\hat{x}}{\sqrt{2\pi}} + (1/\pi + 1/2)\right) e^{-\hat{x}^2/2} \approx 0.9027 \leq 1,$$

which certifies our assumption. Hence, we have

$$-1 \leq g'(x) \leq 0, \quad x > 0.$$

Combining this result with the one for $x \leq 0$ in (15) yields

$$-1 \leq g'(x) \leq 0, \quad x \in \mathbb{R}.$$

We finally obtain the following bound on the scalar Lipschitz constant (12):

$$|g_{\text{pro}}(y) - g_{\text{pro}}(z)| \leq \left| \int_y^z |g'_{\text{pro}}(x)| dx \right| \leq \left| \int_y^z 1 dx \right| = |y - z|,$$

which concludes the proof for the probit case.

We now develop the bound L_{\log} for the logit case. To this end, we bound the derivative of $g_{\log}(x) = \frac{1}{1+e^x}$ as follows:

$$0 \geq g'_{\log}(x) = -\frac{e^x}{(1+e^x)^2} = -\frac{1}{e^x + e^{-x} + 2} \geq -\frac{1}{4}.$$

where we used the inequality of arithmetic and geometric means. Consequently, we have the following bound on the scalar Lipschitz constant (13):

$$|g_{\log}(y) - g_{\log}(z)| \leq \left| \int_y^z |g'_{\log}(x)| dx \right| \leq \left| \int_y^z \frac{1}{4} dx \right| = \frac{1}{4}|y - z|,$$

which concludes the proof for the logit case. ■

The following lemma establishes a bound on the (vector) Lipschitz constants for the individual regularized regression problems (RR_1^+) and (RR_2) for both the probit and the logit case, using the results in Lemma 4. We work out in detail the analysis of (RR_1^+) for $\bar{\mathbf{w}}_i$, i.e., the transpose of the i^{th} row of \mathbf{W} . The proofs for the remaining subproblems for other rows of \mathbf{W} and all columns of \mathbf{C} follow analogously.

Lemma 5 (Lipschitz constants) For a given i and j , let

$$f_w(\bar{\mathbf{w}}_i) = - \sum_j \log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\mu}{2} \|\bar{\mathbf{w}}_i\|_2^2 = - \sum_j \log \Phi((2Y_{i,j} - 1)\bar{\mathbf{w}}_i^T \mathbf{c}_j) + \frac{\mu}{2} \|\bar{\mathbf{w}}_i\|_2^2,$$

$$f_c(\mathbf{c}_j) = - \sum_i \log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) = - \sum_i \log \Phi((2Y_{i,j} - 1)\bar{\mathbf{w}}_i^T \mathbf{c}_j),$$

where $Y_{i,j}$, $\bar{\mathbf{w}}_i$, and \mathbf{c}_j are defined as in Section 2.1. Here, $\Phi(x)$ designates the inverse link function, which can either be (3) or (4). Then, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^K$, we have

$$\|\nabla f_w(\mathbf{x}) - \nabla f_w(\mathbf{y})\|_2 \leq (L\sigma_{\max}^2(\mathbf{C}) + \mu)\|\mathbf{x} - \mathbf{y}\|_2,$$

$$\|\nabla f_c(\mathbf{x}) - \nabla f_c(\mathbf{y})\|_2 \leq L\sigma_{\max}^2(\mathbf{W})\|\mathbf{x} - \mathbf{y}\|_2,$$

where $L = L_{\text{pro}} = 1$ and $L = L_{\text{log}} = 1/4$ are the scalar Lipschitz constants for the probit and logit cases from Lemma 4, respectively.

Proof For the sake of brevity, we only show the proof for $f_w(\mathbf{x})$ in the probit case. The logit cases and the cases for $f_c(\mathbf{x})$ follow analogously. In what follows, the PDF $\mathcal{N}(x)$ and CDF $\Phi(x)$ of the standard normal density (the inverse probit link function) defined in (3) are assumed to operate element-wise on the vector $\mathbf{x} \in \mathbb{R}^K$.

In order to simplify the derivation of the proof, we define the following effective matrix associated to $\bar{\mathbf{w}}_i$ as

$$\mathbf{C}_{\text{eff},i} = [(2Y_{i,1} - 1)\mathbf{c}_1, \dots, (2Y_{i,N} - 1)\mathbf{c}_N],$$

which is equivalent to a right-multiplication $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_N]$ with a diagonal matrix containing the binary-valued response variables $(2Y_{i,j} - 1) \in \{-1, +1\} \forall j$. We can now establish an upper bound of the ℓ_2 -norm of the difference between the gradients at two arbitrary points \mathbf{x} and \mathbf{y} as follows:

$$\begin{aligned} \|\nabla f_w(\mathbf{x}) - \nabla f_w(\mathbf{y})\|_2 &= \left\| \mathbf{C}_{\text{eff},i} \frac{\mathcal{N}(\mathbf{C}_{\text{eff},i}^T \mathbf{x})}{\Phi(\mathbf{C}_{\text{eff},i}^T \mathbf{x})} - \mathbf{C}_{\text{eff},i} \frac{\mathcal{N}(\mathbf{C}_{\text{eff},i}^T \mathbf{y})}{\Phi(\mathbf{C}_{\text{eff},i}^T \mathbf{y})} + \mu\mathbf{x} - \mu\mathbf{y} \right\|_2 \\ &\leq \sigma_{\max}(\mathbf{C}_{\text{eff},i}) \left\| \frac{\mathcal{N}(\mathbf{C}_{\text{eff},i}^T \mathbf{x})}{\Phi(\mathbf{C}_{\text{eff},i}^T \mathbf{x})} - \frac{\mathcal{N}(\mathbf{C}_{\text{eff},i}^T \mathbf{y})}{\Phi(\mathbf{C}_{\text{eff},i}^T \mathbf{y})} \right\|_2 + \mu\|\mathbf{x} - \mathbf{y}\|_2 \quad (18) \\ &\leq L\sigma_{\max}(\mathbf{C}_{\text{eff},i}) \|\mathbf{C}_{\text{eff},i}^T \mathbf{x} - \mathbf{C}_{\text{eff},i}^T \mathbf{y}\|_2 + \mu\|\mathbf{x} - \mathbf{y}\|_2 \quad (19) \\ &\leq L\sigma_{\max}^2(\mathbf{C}_{\text{eff},i}) \|\mathbf{x} - \mathbf{y}\|_2 + \mu\|\mathbf{x} - \mathbf{y}\|_2 \quad (20) \\ &= (L\sigma_{\max}^2(\mathbf{C}) + \mu)\|\mathbf{x} - \mathbf{y}\|_2. \quad (21) \end{aligned}$$

Here, (18) uses the triangle inequality and the Rayleigh-Ritz theorem of Horn and Johnson (1991), where $\sigma_{\max}(\mathbf{C}_{\text{eff},i})$ denotes the principal singular value of $\mathbf{C}_{\text{eff},i}$. The bound (19) follows from Lemma 4, and (20) is, once more, a consequence of the Rayleigh-Ritz theorem. The final equality (21) follows from the fact that flipping the signs of the columns of a matrix (as we did to arrive at $\mathbf{C}_{\text{eff},i}$) does not affect its singular values, which concludes the proof. Note that the proof for $f_c(\cdot)$ follows by omitting μ and substitute \mathbf{C} by \mathbf{W} in (21). \blacksquare

Note that in all of the above proofs we only considered the case where the observation matrix \mathbf{Y} is fully populated. Our proofs easily adapt to the case of missing entries in \mathbf{Y} , by replacing the matrix \mathbf{C} to $\mathbf{C}_{\mathcal{I}}$, where $\mathbf{C}_{\mathcal{I}}$ corresponds to the matrix containing the columns of \mathbf{C} corresponding to the observed entries indexed by the set $\mathcal{I} = \{j : (i, j) \in \Omega_{\text{obs}}\}$. We omit the details for the sake of brevity.

Appendix B. Proof of Theorem 2

Minimizing $F(\mathbf{x})$ as defined in Theorem 2 using SPARFA-M corresponds to a multi-block coordinate descent problem, where the subproblems $(\text{RR})_1^+$ and $(\text{RR})_2$ correspond to (Xu and Yin, 2012, Problem. 1.2b and 1.2a), respectively. Hence, we can use the results of (Xu and Yin, 2012, Lemma 2.6, Corollary 2.7, and Theorem 2.8) to establish global convergence of SPARFA-M. To this end, we must verify that the problem (P) satisfies all of the assumptions in (Xu and Yin, 2012, Assumption 1, Assumption 2, and Lemma 2.6).

B.1 Prerequisites

We first show that the smooth part of the cost function in (P), i.e., the negative log-likelihood plus both ℓ_2 -norm regularization terms, is Lipschitz continuous on any bounded set in $\mathbb{R}^{(N+Q)K}$. Then, we show that the probit log-likelihood function is real analytic. Note that the logit log-likelihood function is real analytic as shown in (Xu and Yin, 2012, Section 2.3). Finally, we combine both results to prove Theorem 2, which establishes the global convergence of SPARFA-M.

Lemma 6 (Lipschitz continuity) *Define $\mathbf{x} = [\bar{\mathbf{w}}_1^T, \dots, \bar{\mathbf{w}}_Q^T, \mathbf{c}_1^T, \dots, \mathbf{c}_N^T]^T$, and let*

$$f(\mathbf{x}) = - \sum_{(i,j) \in \Omega_{\text{obs}}} \log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\mu}{2} \sum_i \|\bar{\mathbf{w}}_i\|_2^2 + \frac{\gamma}{2} \sum_j \|\mathbf{c}_j\|_2^2.$$

Then, $f(\mathbf{x})$ is Lipschitz continuous on any bounded set $\mathcal{D} = \{\mathbf{x} : \|\mathbf{x}\|_2 \leq D\}$.

Proof Let $\mathbf{y}, \mathbf{z} \in \mathcal{D}$, recall the notation of Lemma 5, and let $\bar{\mathbf{w}}_i^y, \bar{\mathbf{w}}_i^z, \mathbf{c}_j^y$, and \mathbf{c}_j^z denote the blocks of variables $\bar{\mathbf{w}}_i$ and \mathbf{c}_j in \mathbf{y} and \mathbf{z} , respectively. We now have

$$\begin{aligned} \|\nabla f(\mathbf{y}) - \nabla f(\mathbf{z})\|_2 &= \left(\sum_{i,j} ((\nabla f_w(\bar{\mathbf{w}}_i^y) - \nabla f_w(\bar{\mathbf{w}}_i^z))^2 + (\nabla f_c(\mathbf{c}_j^y) - \nabla f_c(\mathbf{c}_j^z))^2 + \gamma^2 \|\mathbf{c}_j^y - \mathbf{c}_j^z\|_2^2) \right)^{\frac{1}{2}} \\ &\leq \left(\sum_{i,j} ((L\sigma_{\max}^2(\mathbf{C}) + \mu)^2 \|\bar{\mathbf{w}}_i^y - \bar{\mathbf{w}}_i^z\|_2^2 + (L^2\sigma_{\max}^4(\mathbf{W}) + \gamma^2) \|\mathbf{c}_j^y - \mathbf{c}_j^z\|_2^2) \right)^{\frac{1}{2}} \end{aligned} \tag{22}$$

$$\begin{aligned} &\leq (L(\|\mathbf{W}\|_F^2 + \|\mathbf{C}\|_F^2) + \max\{\mu, \gamma\}) \|\mathbf{y} - \mathbf{z}\|_2 \tag{23} \\ &\leq (LD^2 + \max\{\mu, \gamma\}) \|\mathbf{y} - \mathbf{z}\|_2, \end{aligned}$$

where (22) follows from Lemma 5, and (23) follows from the fact that the maximum singular value of a matrix is no greater than its Frobenius norm (Horn and Johnson 1991), which is bounded by D for $\mathbf{y}, \mathbf{z} \in \mathcal{D}$. We furthermore have $L = 1$ for the probit case and $L = 1/4$

for the logit case, shown in Lemma 4. Thus, $f(\mathbf{x})$ is Lipschitz continuous on any bounded set. \blacksquare

Lemma 7 (Real analyticity) Define $\mathbf{x} = [\bar{\mathbf{w}}_1^T, \dots, \bar{\mathbf{w}}_Q^T, \mathbf{c}_1^T, \dots, \mathbf{c}_N^T]^T$, and let

$$g(\mathbf{x}) = - \sum_{(i,j) \in \Omega_{obs}} \log p(Y_{i,j} | \bar{\mathbf{w}}_i, \mathbf{c}_j) = - \sum_{(i,j) \in \Omega_{obs}} \log \Phi_{pro}((2Y_{i,j} - 1)\bar{\mathbf{w}}_i^T \mathbf{c}_j),$$

where $\Phi_{pro}(\cdot)$ is the inverse probit link function defined in (3). Then, $g(\mathbf{x})$ is real analytic.

Proof Recall the important property established by Krantz and Parks (2002) that compositions of real analytic functions are real analytic. Therefore, the standard normal density $\mathcal{N}(x)$ is real analytic, since the exponential function and x^2 are both real analytical functions. Consequently, let $\mathcal{N}^{(k)}(x)$ denote the k^{th} derivative of $\mathcal{N}(x)$, then $\left(\frac{\mathcal{N}^{(k)}(x)}{k!}\right)^{\frac{1}{k}}$ is bounded for all k , according to the definition of real analytic functions.

Now we show that $\Phi_{pro}(x) = \int_{-\infty}^x \mathcal{N}(t)dt$ is also real analytic. Its k^{th} derivative is given by $\Phi_{pro}^{(k)}(x) = \mathcal{N}^{(k-1)}(x)$, and therefore $\left(\frac{\Phi_{pro}^{(k)}(x)}{k!}\right)^{\frac{1}{k}}$ is obviously bounded for all k , since $\left(\frac{\mathcal{N}^{(k)}(x)}{k!}\right)^{\frac{1}{k}}$ is bounded for all k as we have just shown. Thus, $\Phi_{pro}(x)$ is real analytic.

Given that $\Phi_{pro}(x)$ is real-analytic, it follows that the negative log-probit-likelihood $-\log\Phi_{pro}(x)$ is real analytic, since both the logarithm function and the inverse probit link function are real analytic. Finally, extending the proof from scalar functions to vector functions preserves analyticity according to (Xu and Yin, 2012, Section 2.2). \blacksquare

B.2 Proof of Theorem 2

We are finally armed to prove Theorem 2. We begin by showing that our problem (P) meets (Xu and Yin, 2012, Assumptions 1 and 2). Then, we show that (P) meets all the additional assumptions needed for the convergence results in (Xu and Yin, 2012, Lemma 2.6), through which we can establish convergence of the sequence $\{\mathbf{x}^t\}$ from certain starting points to some finite limit point. Finally, we use (Xu and Yin, 2012, Theorem 2.8) to show global convergence of SPARFA-M from any starting point.

B.2.1 ASSUMPTION 1

We start by showing that (P) meets (Xu and Yin, 2012, Assumption 1). Since every term in our objective function in (P) is non-negative, we have $F(\mathbf{x}) > -\infty$. It is easy to verify that (P) is also block multi-convex in the variable \mathbf{x} , with the rows of \mathbf{W} and columns of \mathbf{C} forming the blocks. Consequently, the problem (P) has at least one critical point, since $F(\mathbf{x})$ is lower bounded by 0. Therefore, Assumption 1 is met.

B.2.2 ASSUMPTION 2

Problem (P) also meets (Xu and Yin, 2012, Assumption 2) regarding the strong convexity of the individual subproblems. Due to the presence of the quadratic terms $\frac{\mu}{2}\|\bar{\mathbf{w}}_i\|_2^2$ and $\frac{\gamma}{2}\|\mathbf{c}_j\|_2^2$, the smooth part of the objective functions of the individual subproblems (RR_1^+) and (RR_2) are strongly convex with parameters μ and γ , respectively. Consequently, Assumption 2 is satisfied.

B.2.3 ASSUMPTIONS IN (XU AND YIN, 2012, LEM. 2.6)

Problem (P) also meets the assumptions in (Xu and Yin, 2012, Lem. 2.6) regarding the Lipschitz continuity of the subproblems and the Kurdyka-Łojasiewicz inequality. Lemma 6 shows that $f(\mathbf{x}) = -\sum_{(i,j)\in\Omega_{\text{obs}}}\log p(Y_{i,j}|\bar{\mathbf{w}}_i, \mathbf{c}_j) + \frac{\mu}{2}\sum_i\|\bar{\mathbf{w}}_i\|_2^2 + \frac{\gamma}{2}\sum_j\|\mathbf{c}_j\|_2^2$, satisfies the Lipschitz continuous requirement in (Xu and Yin, 2012, Lemma 2.6). As shown in Lemma 7 for the probit case and as shown in (Xu and Yin, 2012, Section 2.2) for the logit case, the negative log-likelihood term in (P) is real analytic, therefore also sub-analytic. All the regularizer functions in $F(\mathbf{x})$ defined in Theorem 2 are semi-algebraic and therefore sub-analytic, a consequence of (Bolte et al., 2006, Section 2.1) and (Xu and Yin, 2012, Section 2.2). Using (Fischer, 2008, Theorems 1.1 and 1.2), the objective function $F(\mathbf{x})$ is also sub-analytic, since all of its parts are sub-analytic and bounded below (non-negative), therefore satisfying the Kurdyka-Łojasiewicz inequality at any point \mathbf{x} , as shown in (Bolte et al., 2006, Theorem 3.1). Finally, the SPARFA-M algorithm uses $\omega_i^{k-1} \equiv 0$ and $\ell = \min\{\mu, \gamma\}$ where ω_i^{k-1} and ℓ as defined in (Xu and Yin, 2012, Lemma 2.6).

Up to this point, we have shown that (P) satisfies all assumptions and requirements in (Xu and Yin, 2012, Lemma 2.6). Now, SPARFA-M follows (Xu and Yin, 2012, Lemma 2.6) in the sense that, if \mathbf{x}^0 is sufficiently close to some critical point $\hat{\mathbf{x}}$ of (P), (more specifically, $\mathbf{x}^0 \in \mathcal{B}$ for some $\mathcal{B} \subset \mathcal{U}$ where \mathcal{U} is a neighborhood of $\hat{\mathbf{x}}$ in which Kurdyka-Łojasiewicz inequality holds), then $\{\mathbf{x}^t\}$ converges to a point in \mathcal{B} . This establishes the convergence of SPARFA-M to a local minimum point from certain starting points.

B.2.4 GLOBAL CONVERGENCE

Finally, we can use (Xu and Yin, 2012, Lemma 2.6) to establish global convergence of SPARFA-M. It is obvious that the objective function (P) is bounded on any bounded set. Hence, the sequence $\{\mathbf{x}^k\}$ will always have a finite limit point and meet the assumptions in (Xu and Yin, 2012, Theorem 2.8). The final statement of Theorem 2 now directly follows from (Xu and Yin, 2012, Theorem 2.8). Moreover, if the starting point is in close proximity to a global minimum, then SPARFA-M is guaranteed to converge to a global minimum. This is a consequence of (Xu and Yin, 2012, Corollary 2.7).

Appendix C. Proof of Theorem 3

Proof To prove Theorem 3, we first define some notation. Let $\mathcal{N}(x|m, s) = \frac{1}{\sqrt{2\pi s}}e^{-(x-m)^2/2s}$ define the normal PDF with mean m and variance s . Furthermore, let $\text{Exp}(m|\lambda) = \lambda e^{-\lambda m}$, $m \geq 0$ define the PDF of the exponential distribution with rate parameter λ .

We are ultimately concerned with identifying whether the factor $W_{i,k}$ is active given our current beliefs about all other parameters in the model. Given the probit model, this

is equivalent to determining whether or not an exponential random variable is present in Gaussian noise. Let $x|m, s \sim \mathcal{N}(0|m, s)$ with $m \sim r \text{Exp}(m|\lambda) + (1-r)\delta_0$ and δ_0 the Dirac delta function located at 0. The posterior distribution $p(m=0|x)$ can be derived via Bayes' rule as follows:

$$\begin{aligned}
 p(m=0|x) &= \frac{\mathcal{N}(x|m=0, s)(1-r)}{\mathcal{N}(x|m=0, s)(1-r) + r \int \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) dm}, \\
 &= \frac{\frac{\mathcal{N}(x|0, s)}{\int \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) dm} (1-r)}{\frac{\mathcal{N}(x|0, s)}{\int \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) dm} (1-r) + r}, \\
 &= \frac{\frac{\text{Exp}(0|\lambda) \mathcal{N}(x|0, s)}{\int \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) dm} (1-r)}{\frac{\text{Exp}(0|\lambda) \mathcal{N}(x|0, s)}{\int \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) dm} (1-r) + r \text{Exp}(0|\lambda)}. \tag{24}
 \end{aligned}$$

Here, it is important to recognize that $\frac{\text{Exp}(m|\lambda) \mathcal{N}(x|m, s)}{\int \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) dm}$ denotes the posterior under the continuous portion of the prior (i.e., $m \neq 0$). Since the exponential prior we have chosen is not conjugate to the normal likelihood, we must compute this distribution in closed form. To this end, let $\mathcal{N}^r(x|m, s, \lambda) \propto \mathcal{N}(x|m, s) \text{Exp}(m|\lambda) = C_0 e^{-(x-m)^2/2s - \lambda m}$ denote a rectified normal distribution with normalization constant C_0 . Completing the square and carrying out the integration, we find $C_0 = \frac{e^{\lambda m - \lambda^2 s/2}}{\sqrt{2\pi s} \Phi\left(\frac{m - \lambda s}{\sqrt{s}}\right)}$, which leads to

$$\mathcal{N}^r(x|m, s, \lambda) = \frac{e^{\lambda m - \lambda^2 s/2}}{\sqrt{2\pi s} \Phi\left(\frac{m - \lambda s}{\sqrt{s}}\right)} e^{-(x-m)^2/2s - \lambda m}.$$

We can now rewrite (24) as

$$p(m=0|x) = \frac{\frac{\mathcal{N}^r(0|\hat{m}, \hat{s}, \lambda)}{\text{Exp}(0|\lambda)} (1-r)}{\frac{\mathcal{N}^r(0|\hat{m}, \hat{s}, \lambda)}{\text{Exp}(0|\lambda)} (1-r) + r}$$

or, alternatively, as

$$\hat{r} = p(m \neq 0|x) = 1 - p(m=0|x) = \frac{\frac{\text{Exp}(0|\lambda)}{\mathcal{N}^r(0|\hat{m}, \hat{s}, \lambda)}}{\frac{\text{Exp}(0|\lambda)}{\mathcal{N}^r(0|\hat{m}, \hat{s}, \lambda)} + \frac{1-r}{r}}. \tag{25}$$

All that remains now is to determine \hat{m} and \hat{s} in (25) for our full factor analysis scenario. Recall that our probabilistic model corresponds to $\mathbf{Z} = \mathbf{WC} + \mathbf{M}$. Further recall our definition of the observation set $\Omega_{\text{obs}} = \{(i, j) : Y_{i,j} \text{ is observed}\}$. We can now calculate the

posterior on each coefficient $W_{i,k} \neq 0$ as follows:

$$\begin{aligned}
 p(W_{i,k}|\mathbf{Z}, \mathbf{C}, \boldsymbol{\mu}) &\propto p(W_{i,k}) p(\mathbf{Z}|\mathbf{W}_{-(i,k)}, \mathbf{C}, \boldsymbol{\mu}) \\
 &\propto e^{-\lambda W_{i,k}} e^{-\frac{1}{2\sigma^2} \sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} ((Z_{i,j} - \mu_i) - \sum_{k=1}^K W_{i,k} C_{k,j})^2} \\
 &= e^{-\lambda W_{i,k}} e^{-\frac{1}{2\sigma^2} \sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} ((Z_{i,j} - \mu_i) - \sum_{k' \neq k} W_{i,k'} C_{k',j} - W_{i,k} C_{k,j})^2} \\
 &\propto e^{-\lambda W_{i,k}} e^{-\frac{1}{2\sigma^2} \sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} (W_{i,k}^2 C_{k,j}^2 - 2((Z_{i,j} - \mu_i) - \sum_{k' \neq k} W_{i,k'} C_{k',j}) W_{i,k} C_{k,j})} \\
 &\propto e^{-\lambda W_{i,k}} e^{-\frac{\sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} C_{k,j}^2}{2\sigma^2} \left(W_{i,k} - \frac{\sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} ((Z_{i,j} - \mu_i) - \sum_{k' \neq k} W_{i,k'} C_{k',j}) C_{k,j}}{\sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} C_{k,j}^2} \right)^2},
 \end{aligned} \tag{26}$$

where the last step is obtained by completing the square in $W_{i,k}$.

The final result in (26) implies that $W_{i,k} \sim \mathcal{N}^r(\hat{m}, \hat{s}, \lambda)$, where

$$\hat{m} = \frac{\sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} ((Z_{i,j} - \mu_i) - \sum_{k' \neq k} W_{i,k'} C_{k',j}) C_{k,j}}{\sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} C_{k,j}^2}$$

and $\hat{s} = \frac{\sigma^2}{\sum_{\{j:(i,j) \in \Omega_{\text{obs}}\}} C_{k,j}^2}$. Combining the results of (25) and (26), recognizing that $\sigma^2 = 1$ in the standard probit model, and adopting the notation $\hat{R}_{i,k}$, $\hat{M}_{i,k}$ and $\hat{S}_{i,k}$ for the values of \hat{r} , \hat{m} and \hat{s} corresponding to each $\widehat{W}_{i,k}$, furnishes the final sampling result. \blacksquare

References

- M. Aharon, M. Elad, and A. M. Bruckstein. K-SVD and its non-negative variant for dictionary design. In *Proc. SPIE Conf. on Wavelets*, volume 5914, pages 327–339, July 2005.
- M. Aharon, M. Elad, and A. M. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, Dec. 2006.
- Amazon Mechanical Turk, Sep. 2012. URL <https://www.mturk.com/mturk/welcome>.
- Y. Bachrach, T. P. Minka, J. Guiver, and T. Graepel. How to grade a test without knowing the answers – a Bayesian graphical model for adaptive crowdsourcing and aptitude testing. In *Proc. 29th Intl. Conf. on Machine Learning*, pages 1183–1190, June 2012.
- F. B. Baker and S. H. Kim. *Item Response Theory: Parameter Estimation Techniques*. Marcel Dekker Inc., 2nd edition, 2004.
- R. Baker and K. Yacef. The state of educational data mining in 2009: A review and future visions. *Journal of Educational Data Mining*, 1(1):3–17, Oct. 2009.
- T. Barnes. The Q-matrix method: Mining student response data for knowledge. In *Proc. AAAI Workshop Educational Data Mining*, July 2005.

- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Science*, 2(1):183–202, Mar. 2009.
- B. Beheshti, M. Desmarais, and R. Naceur. Methods to find the number of latent skills. In *Proc. 5th Intl. Conf. on Educational Data Mining*, pages 81–86, June 2012.
- Y. Bergner, S. Droschler, G. Kortemeyer, S. Rayyan, D. Seaton, and D. Pritchard. Model-based collaborative filtering analysis of student response data: Machine-learning item response theory. In *Proc. 5th Intl. Conf. on Educational Data Mining*, pages 95–102, June 2012.
- J. Bolte, A. Daniilidis, and A. Lewis. The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems. *SIAM Journal on Optimization*, 17(4):1205–1223, Dec. 2006.
- P. T. Boufounos and R. G. Baraniuk. 1-bit compressive sensing. In *Proc. Conf. on Information Science and Systems (CISS)*, Mar. 2008.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- A. M. Bruckstein, M. Elad, and M. Zibulevsky. On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations. *IEEE Transactions on Information Theory*, 54(11):4813–4820, Nov. 2008.
- P. Brusilovsky and C. Peylo. Adaptive and intelligent web-based educational systems. *Intl. Journal of Artificial Intelligence in Education*, 13(2-4):159–172, Apr. 2003.
- C. J. Butz, S. Hua, and R. B. Maguire. A web-based Bayesian intelligent tutoring system for computer programming. *Web Intelligence and Agent Systems*, 4(1):77–97, Nov. 2006.
- H. Chang and Z. Ying. Nonlinear sequential designs for logistic item response theory models with applications to computerized adaptive tests. *Annals of Statistics*, 37(3):1466–1488, June 2009.
- S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, Mar. 1998.
- J. T. Chu. On bounds for the normal integral. *IEEE Transactions on Signal Processing*, 42(1/2):263–265, June 1955.
- M. Desmarais. Conditions for effectively deriving a Q-matrix from data with non-negative matrix factorization. In *Proc. 4th Intl. Conf. on Educational Data Mining*, pages 41–50, July 2011.
- J. A. Dijkstra and S. Khan. Khan Academy: the world’s free virtual school. In *APS Meeting Abstracts*, page 14006, Mar. 2011.
- ELEC 301, Rice University. Introduction to signals and systems, May 2011. URL <http://dsp.rice.edu/courses/elec301>.
- A. Fischer. On sums of subanalytic functions. *Preprint*, 2008.

- E. Fokoue. Stochastic determination of the intrinsic structure in Bayesian factor analysis. Technical report, Statistical and Applied Mathematical Sciences Institute, June 2004.
- K. Fronczyk, A. E. Waters, M. Guindani, R. G. Baraniuk, and M. Vannucci. A Bayesian infinite factor model for learning and content analytics. *Computational Statistics and Data Analysis*, June 2013, submitted.
- J. P. González-Brenes and J. Mostow. Dynamic cognitive tracing: Towards unified discovery of student and cognitive models. In *Proc. 5th Intl. Conf. on Educational Data Mining*, pages 49–56, June 2012.
- I. Goodfellow, A. Courville, and Y. Bengio. Large-scale feature learning with spike-and-slab sparse coding. In *Proc. 29th Intl. Conf. on Machine Learning*, pages 1439–1446, July 2012.
- A. Guisan, T. C. Edwards Jr, and T. Hastie. Generalized linear and generalized additive models in studies of species distributions: setting the scene. *Ecological Modelling*, 157 (2–3):89–100, Nov. 2002.
- P. R. Hahn, C. M. Carvalho, and J. G. Scott. A sparse factor-analytic probit model for congressional voting patterns. *Journal of the Royal Statistical Society*, 61(4):619–635, Aug. 2012.
- H. H. Harman. *Modern Factor Analysis*. The University of Chicago Press, 1976.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, 2010.
- J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1):5–53, Jan. 2004.
- R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- D. Hu. How Khan academy is using machine learning to assess student mastery. *online: <http://david-hu.com/>*, Nov. 2011.
- H. Ishwaran and J. S. Rao. Spike and slab variable selection: frequentist and Bayesian strategies. *Annals of Statistics*, 33(2):730–773, Apr. 2005.
- L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transaction on Info Theory*, 59(4), Apr. 2013.
- Knewton. Knewton adaptive learning: Building the world’s most powerful recommendation engine for education. *online: <http://www.knewton.com/adaptive-learning-white-paper/>*, June 2012.

- K. R. Koedinger, J. R. Anderson, W. H. Hadley, and M. A. Mark. Intelligent tutoring goes to school in the big city. *Intl. Journal of Artificial Intelligence in Education*, 8(1):30–43, 1997.
- K. Koh, S. Kim, and S. Boyd. An interior-point method for large-scale ℓ_1 -regularized logistic regression. *Journal of Machine Learning Research*, 8:1519–1555, 2007.
- Y. Koren and J. Sill. OrdRec: an ordinal model for predicting personalized item rating distributions. In *Proc. of the 5th ACM Conf. on Recommender Systems*, pages 117–124, Oct. 2011.
- Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, Aug. 2009.
- S. G. Krantz and H. R. Parks. *A Primer of Real Analytic Functions*. Birkhauser, 2002.
- G. A. Krudysz and J. H. McClellan. Collaborative system for signal processing education. In *IEEE Intl. Conf. on Acoustics, Speech and Signal Processing*, pages 2904–2907, May 2011.
- G. A. Krudysz, J. S. Li, and J. H. McClellan. Web-based Bayesian tutoring system. In *12th Digital Signal Processing Workshop - 4th Signal Processing Education Workshop*, pages 129–134, Sep. 2006.
- A. S. Lan, C. Studer, A. E. Waters, and R. G. Baraniuk. Tag-aware ordinal sparse factor analysis for learning and content analytics. In *Proc. 6th Intl. Conf. on Educational Data Mining*, pages 90–97, July 2013a.
- A. S. Lan, C. Studer, A. E. Waters, and R. G. Baraniuk. Joint topic modeling and factor analysis of textual information and graded response data. In *Proc. 6th Intl. Conf. on Educational Data Mining*, pages 324–325, July 2013b.
- S. Lee, H. Lee, P. Abbeel, and A. Y. Ng. Efficient ℓ_1 regularized logistic regression. In *Proc. National Conf. on Artificial Intelligence*, volume 21, pages 401–408, 2006.
- S. Lee, J. Z. Huang, and J. Hu. Sparse logistic principal components analysis for binary data. *Annals of Applied Statistics*, 4(3):1579–1601, Sept. 2010.
- N. Li, W. W. Cohen, and K. R. Koedinger. A machine learning approach for automatic student model discovery. In *Proc. 4th Intl. Conf. on Educational Data Mining*, pages 31–40, July 2011.
- J. M. Linacre. Understanding Rasch measurement: Estimation methods for Rasch measures. *Journal of Outcome Measurement*, 3(4):382–405, 1999.
- G. Linden, B. Smith, and J. York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, Jan. 2003.
- W. J. V. D. Linden and editors Glas, C. A. W. *Computerized Adaptive Testing: Theory and Practice*. Kluwer Academic Publishers, 2000.

- F. M. Lord. *Applications of Item Response Theory to Practical Testing Problems*. Erlbaum Associates, 1980.
- J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11:19–60, 2010.
- R. P. McDonald. A basis for multidimensional item response theory. *Applied Psychological Measurement*, 24(2):99–114, June 2000.
- J. Meng, J. Zhang, Y. Qi, Y. Chen, and Y. Huang. Uncovering transcriptional regulatory networks by sparse Bayesian factor model. *EURASIP Journal on Advances in Signal Processing*, 2010(3):1–18, Mar. 2010.
- T. P. Minka. A comparison of numerical optimizers for logistic regression. Technical report, 2003. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.85.7017&rep=rep1&type=pdf>.
- S. Mohamed, K. Heller, and Z. Ghahramani. Bayesian and ℓ_1 approaches for sparse unsupervised learning. In *Proc. 29th Intl. Conf. on Machine Learning*, pages 751–758, July 2012.
- R. C. Murray, K. VanLehn, and J. Mostow. Looking ahead to select tutorial actions: A decision-theoretic approach. *Intl. Journal of Artificial Intelligence in Education*, 14(3–4): 235–278, Dec. 2004.
- Y. Nesterov. Gradient methods for minimizing composite objective function. Technical report, Université catholique de Louvain, Sep. 2007.
- M. R. Norvick. The axioms and principal results of classical test theory. *Journal of Mathematical Psychology*, 3(1):1–18, Feb. 1966.
- F. W. J. Olver, editor. *NIST Handbook of Mathematical Functions*. Cambridge University Press, 2010.
- Z. A. Pardos and N. T. Heffernan. Modeling individualization in a Bayesian networks implementation of knowledge tracing. In *User Modeling, Adaptation, and Personalization*, volume 6075, pages 255–266. Springer, June 2010.
- M. Y. Park and T. Hastie. Penalized logistic regression for detecting gene interactions. *Biostatistics*, 9(1):30–50, Jan. 2008.
- Y. Plan and R. Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. 59(1), Feb. 2012, submitted.
- I. Pournara and L. Wernisch. Factor analysis for gene regulatory networks and transcription factor activity profiles. *BMC Bioinformatics*, 8(1):61, Feb. 2007.
- J. Psotka, L. D. Massey, and editors Mutter, S. A. *Intelligent Tutoring Systems: Lessons Learned*. Lawrence Erlbaum Associates, 1988.

- A. N. Rafferty, E. Brunskill, T. L. Griffiths, and P. Shafto. Faster teaching by POMDP planning. In *Proc. 15th Intl. Conf. on Artificial Intelligence in Education*, pages 280–287, June 2011.
- G. Rasch. *Probabilistic Models for Some Intelligence and Attainment Tests*. MESA Press, 1993.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Process for Machine Learning*. MIT Press, 2006.
- M. D. Reckase. *Multidimensional Item Response Theory*. Springer Publishing Company, Incorporated, 1st edition, 2009.
- C. Romero and S. Ventura. Educational data mining: A survey from 1995 to 2005. *Expert Systems with Applications*, 33(1):135–146, July 2007.
- M. N. Schmidt, O. Winther, and L. K. Hansen. Bayesian non-negative matrix factorization. In *Independent Component Analysis and Signal Separation*, volume 5441, pages 540–547, Mar. 2009.
- J. C. Stamper, T. Barnes, and M. Croy. Extracting student models for intelligent tutoring systems. In *Proc. National Conf. on Artificial Intelligence*, volume 22, pages 113–147, July 2007.
- STEMscopes. STEMscopes science education, Sep. 2012. URL <http://stemscopes.com/>.
- C. Studer and R. G. Baraniuk. Dictionary learning from sparsely corrupted or compressed signals. In *IEEE Intl. Conf. on Acoustics, Speech and Signal Processing*, pages 3341–3344, Mar. 2012.
- N. Thai-Nghe, L. Drumond, T. Horvath, and L. Schmidt-Thieme. Multi-relational factorization models for predicting student performance. *KDD Workshop on Knowledge Discovery in Educational Data (KDDinED)*, Aug. 2011a.
- N. Thai-Nghe, T. Horvath, and L. Schmidt-Thieme. Factorization models for forecasting student performance. In *Proc. 4th Intl. Conf. on Educational Data Mining*, pages 11–20, July 2011b.
- N. A. Thompson. Item selection in computerized classification testing. *Educational and Psychological Measurement*, 69(5):778–793, Oct. 2009.
- M. E. Tipping. Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, 2001.
- K. VanLehn, C. Lynch, K. Schulze, J. A. Shapiro, R. Shelby, L. Taylor, D. Treacy, A. Weinstein, and M. Wintersgill. The Andes physics tutoring system: Lessons learned. *Intl. Journal of Artificial Intelligence in Education*, 15(3):147–204, Sep. 2005.
- D. Vats, C. Studer, A. S. Lan, L. Carin, and R. G. Baraniuk. Test size reduction for concept estimation. In *Proc. 6th Intl. Conf. on Educational Data Mining*, pages 292–295, July 2013.

- M. West. Bayesian factor regression models in the “large p , small n ” paradigm. *Bayesian Statistics*, 7:723–732, Sep. 2003.
- B. P. Woolf. *Building Intelligent Interactive Tutors: Student-centered Strategies for Revolutionizing E-learning*. Morgan Kaufman Publishers, 2008.
- Y. Xu and W. Yin. A block coordinate descent method for multi-convex optimization with applications to nonnegative tensor factorization and completion. Technical report, Rice University CAAM, Sep. 2012.
- L. Yao. *BMIRT: Bayesian Multivariate Item Response Theory*. CTb/McGraw-Hill, 2003.