

Tag-Aware Ordinal Sparse Factor Analysis for Learning and Content Analytics

Andrew S. Lan, Christoph Studer, Andrew E. Waters, Richard G. Baraniuk

Rice University, TX, USA

{mr.lan, studer, waters, richb}@sparfa.com

ABSTRACT

Machine learning offers novel ways and means to design *personalized learning systems* wherein each student’s educational experience is customized in real time depending on their background, learning goals, and performance to date. SPARse Factor Analysis (SPARFA) is a novel framework for machine learning-based *learning analytics*, which estimates a learner’s knowledge of the concepts underlying a domain, and *content analytics*, which estimates the relationships among a collection of questions and those concepts. SPARFA jointly learns the associations among the questions and the concepts, learner concept knowledge profiles, and the underlying question difficulties, solely based on the correct/incorrect graded responses of a population of learners to a collection of questions. In this paper, we extend the SPARFA framework significantly to enable: (i) the analysis of graded responses on an ordinal scale (partial credit) rather than a binary scale (correct/incorrect); (ii) the exploitation of tags/labels for questions that partially describe the question–concept associations. The resulting *Ordinal SPARFA-Tag* framework greatly enhances the interpretability of the estimated concepts. We demonstrate using real educational data that Ordinal SPARFA-Tag outperforms both SPARFA and existing collaborative filtering techniques in predicting missing learner responses.

Keywords

Factor analysis, ordinal regression, matrix factorization, personalized learning, block coordinate descent

1. INTRODUCTION

Today’s education system typically provides only a “one-size-fits-all” learning experience that does not cater to the background, interests, and goals of individual learners. Modern machine learning (ML) techniques provide a golden opportunity to reinvent the way we teach and learn by making it more personalized and, hence, more efficient and effective. The last decades have seen a great acceleration in the development of *personalized learning systems* (PLSs), which can be grouped into two broad categories: (i) high-quality, but labor-intensive rule-based systems designed by domain experts that are hard-coded to give feedback in pre-defined scenarios [8], and (ii) more affordable and scalable ML-based systems that mine various forms of learner data in order to make performance predictions for each learner [15, 18, 30].

1.1 Learning and content analytics

Learning analytics (LA, estimating what a learner understands based on data obtained from tracking their inter-

actions with learning content) and *content analytics* (CA, organizing learning content such as questions, instructional text, and feedback hints) enable a PLS to generate automatic, targeted feedback to learners, their instructors, and content authors [23]. Recently we proposed a new framework for LA and CA based on SPARse Factor Analysis (SPARFA) [24]. SPARFA consists of a statistical model and convex-optimization-based inference algorithms for analytics that leverage the fact that the knowledge in a given subject can typically be decomposed into a small set of latent knowledge components that we term *concepts* [24]. Leveraging the latent concepts and based only on the graded binary-valued responses (i.e., correct/incorrect) to a set of questions, SPARFA jointly estimates (i) the associations among the questions and the concepts (via a “concept graph”), (ii) learner concept knowledge profiles, and (iii) the underlying question difficulties.

1.2 Contributions

In this paper, we develop *Ordinal SPARFA-Tag*, a significant extension to the SPARFA framework that enables the exploitation of the additional information that is often available in educational settings. First, Ordinal SPARFA-Tag exploits the fact that responses are often graded on an ordinal scale (partial credit), rather than on a binary scale (correct/incorrect). Second, Ordinal SPARFA-Tag exploits tags/labels (i.e., keywords characterizing the underlying knowledge component related to a question) that can be attached by instructors and other users to questions. Exploiting pre-specified tags within the estimation procedure provides significantly more interpretable question–concept associations. Furthermore, our statistical framework can discover new concept–question relationships that would not be in the pre-specified tag information but, nonetheless, explain the graded learner–response data.

We showcase the superiority of Ordinal SPARFA-Tag compared to the methods in [24] via a set of synthetic “ground truth” simulations and on a variety of experiments with real-world educational datasets. We also demonstrate that Ordinal SPARFA-Tag outperforms existing state-of-the-art collaborative filtering techniques in terms of predicting missing ordinal learner responses.

2. STATISTICAL MODEL

We assume that the learners’ knowledge level on a set of abstract latent *concepts* govern the responses they provide to a set of questions. The SPARFA statistical model charac-

terizes the probability of learners’ binary (correct/incorrect) graded responses to questions in terms of three factors: (i) question–concept associations, (ii) learners’ concept knowledge, and (iii) intrinsic question difficulties; details can be found in [24, Sec. 2]. In this section, we will first extend the SPARFA framework to characterize *ordinal* (rather than binary-valued) responses, and then impose additional structure in order to model real-world educational behavior more accurately.

2.1 Model for ordinal learner response data

Suppose that we have N learners, Q questions, and K underlying concepts. Let $Y_{i,j}$ represent the graded response (i.e., score) of the j^{th} learner to the i^{th} question, which are from a set of P ordered labels, i.e., $Y_{i,j} \in \mathcal{O}$, where $\mathcal{O} = \{1, \dots, P\}$. For the i^{th} question, with $i \in \{1, \dots, Q\}$, we propose the following model for the learner–response relationships:

$$\begin{aligned} Z_{i,j} &= \mathbf{w}_i^T \mathbf{c}_j + \mu_i, \quad \forall (i,j), \\ Y_{i,j} &= \mathcal{Q}(Z_{i,j} + \epsilon_{i,j}), \quad \epsilon_{i,j} \sim \mathcal{N}(0, 1/\tau_{i,j}), \quad (i,j) \in \Omega_{\text{obs}}, \end{aligned} \quad (1)$$

where the column vector $\mathbf{w}_i \in \mathbb{R}^K$ models the *concept associations*; i.e., it encodes how question i is related to each concept. Let the column vector $\mathbf{c}_j \in \mathbb{R}^K$, $j \in \{1, \dots, N\}$, represent the latent *concept knowledge* of the j^{th} learner, with its k^{th} component representing the j^{th} learner’s knowledge of the k^{th} concept. The scalar μ_i models the *intrinsic difficulty* of question i , with large positive value of μ for an easy question. The quantity $\epsilon_{i,j}$ models the uncertainty of learner j answering question i correctly/incorrectly and $\mathcal{N}(0, 1/\tau_{i,j})$ denotes a zero-mean Gaussian distribution with precision parameter $\tau_{i,j}$, which models the *reliability* of the observation of learner j answering question i . We will further assume $\tau_{i,j} = \tau$, meaning that all the observations have the same reliability.¹ The slack variable $Z_{i,j}$ in (1) governs the probability of the observed grade $Y_{i,j}$. The set $\Omega_{\text{obs}} \subseteq \{1, \dots, Q\} \times \{1, \dots, N\}$ contains the indices associated to the observed learner–response data, in case the response data is not fully observed.

In (1), $\mathcal{Q}(\cdot) : \mathbb{R} \rightarrow \mathcal{O}$ is a scalar quantizer that maps a real number into P ordered labels according to

$$\mathcal{Q}(x) = p \quad \text{if } \omega_{p-1} < x \leq \omega_p, \quad p \in \mathcal{O},$$

where $\{\omega_0, \dots, \omega_P\}$ is the set of quantization bin boundaries satisfying $\omega_0 < \omega_1 < \dots < \omega_{P-1} < \omega_P$, with ω_0 and ω_P denoting the lower and upper bound of the domain of the quantizer $\mathcal{Q}(\cdot)$.² This quantization model leads to the equivalent input–output relation

$$\begin{aligned} Z_{i,j} &= \mathbf{w}_i^T \mathbf{c}_j + \mu_i, \quad \forall (i,j), \quad \text{and} \\ p(Y_{i,j} = p \mid Z_{i,j}) &= \int_{\omega_{p-1}}^{\omega_p} \mathcal{N}(s \mid Z_{i,j}, 1/\tau_{i,j}) \, ds \\ &= \Phi(\tau(\omega_p - Z_{i,j})) - \Phi(\tau(\omega_{p-1} - Z_{i,j})), \quad (i,j) \in \Omega_{\text{obs}}, \end{aligned} \quad (2)$$

where $\Phi(x) = \int_{-\infty}^x \mathcal{N}(s \mid 0, 1) \, ds$ denotes the *inverse probit function*, with $\mathcal{N}(s \mid 0, 1)$ representing the value of a standard normal distribution evaluated at s .³

¹Accounting for learner/question-varying reliabilities is straightforward and omitted for the sake of brevity.

²In most situations, we have $\omega_0 = -\infty$ and $\omega_P = \infty$.

³Space limitations preclude us from discussing a corresponding logistic-based model; the extension is straightforward.

We can conveniently rewrite (1) and (2) in matrix form as

$$\begin{aligned} \mathbf{Z} &= \mathbf{W}\mathbf{C}, \quad \forall (i,j), \quad \text{and} \\ p(Y_{i,j} \mid Z_{i,j}) &= \Phi(\tau(U_{i,j} - Z_{i,j})) \\ &\quad - \Phi(\tau(L_{i,j} - Z_{i,j})), \quad (i,j) \in \Omega_{\text{obs}}, \end{aligned} \quad (3)$$

where \mathbf{Y} and \mathbf{Z} are $Q \times N$ matrices. The $Q \times (K+1)$ matrix \mathbf{W} is formed by concatenating $[\mathbf{w}_1, \dots, \mathbf{w}_Q]^T$ with the intrinsic difficulty vector $\boldsymbol{\mu}$ and \mathbf{C} is a $(K+1) \times N$ matrix formed by concatenating the $K \times N$ matrix $[\mathbf{c}_1, \dots, \mathbf{c}_N]$ with an all-ones row vector $\mathbf{1}_{1 \times N}$. We furthermore define the $Q \times N$ matrices \mathbf{U} and \mathbf{L} to contain the upper and lower bin boundaries corresponding to the observations in \mathbf{Y} , i.e., we have $U_{i,j} = \omega_{Y_{i,j}}$ and $L_{i,j} = \omega_{Y_{i,j}-1}$, $\forall (i,j) \in \Omega_{\text{obs}}$.

We emphasize that the statistical model proposed above is significantly more general than the original SPARFA model proposed in [24], which is a special case of (1) with $P = 2$ and $\tau = 1$. The precision parameter τ does not play a central role in [24] (it has been set to $\tau = 1$), since the observations are binary-valued with bin boundaries $\{-\infty, 0, \infty\}$. For ordinal responses (with $P > 2$), however, the precision parameter τ significantly affects the behavior of the statistical model and, hence, we estimate the precision parameter τ directly from the observed data.

2.2 Fundamental assumptions

Estimating \mathbf{W} , $\boldsymbol{\mu}$ and \mathbf{C} from \mathbf{Y} is an ill-posed problem, in general, since there are more unknowns than observations and the observations are ordinal (and not real-valued). To ameliorate the illposedness, [24] proposed three assumptions accounting for real-world educational situations:

- (A1) *Low-dimensionality*: Redundancy exists among the questions in an assessment, and the observed graded learner responses live in a low-dimensional space, i.e., $K \ll N, Q$.
- (A2) *Sparsity*: Each question measures the learners’ knowledge on only a few concepts (relative to N and Q), i.e., the question–concept association matrix \mathbf{W} is sparse.
- (A3) *Non-negativity*: The learners’ knowledge on concepts does not reduce the chance of receiving good score on any question, i.e., the entries in \mathbf{W} are non-negative. Therefore, large positive values of the entries in \mathbf{C} represent good concept knowledge, and vice versa.

Although these assumptions are reasonable for a wide range of educational contexts (see [24] for a detailed discussion), they are hardly complete. In particular, additional information is often available regarding the questions and the learners in some situations. Hence, we impose one additional assumption:

- (A4) *Oracle support*: Instructor-provided tags on questions provide prior information on some question–concept associations. In particular, associating each tag with a *single concept* will partially (or fully) determine the locations of the non-zero entries in \mathbf{W} .

As we will see, assumption (A4) significantly improves the limited interpretability of the estimated factors \mathbf{W} and \mathbf{C} over the conventional SPARFA framework [24], which relies on a (somewhat ad-hoc) post-processing step to associate instructor provided tags with concepts. In contrast, we utilize the tags as “oracle” support information on \mathbf{W} *within* the

model, which enhances the explanatory performance of the statistical framework, i.e., it enables to associate each concept directly with a predefined tag. Note that user-specified tags might not be precise or complete. Hence, the proposed estimation algorithm must be capable of discovering new question–concept associations and removing predefined associations that cannot be explained from the observed data.

3. ALGORITHM

We start by developing *Ordinal SPARFA-M*, a generalization of SPARFA-M from [24] to ordinal response data. Then, we detail *Ordinal SPARFA-Tag*, which considers prespecified question tags as oracle support information of \mathbf{W} , to estimate \mathbf{W} , \mathbf{C} , and τ , from the ordinal response matrix \mathbf{Y} while enforcing the assumptions (A1)–(A4).

3.1 Ordinal SPARFA-M

To estimate \mathbf{W} , \mathbf{C} , and τ in (3) given \mathbf{Y} , we maximize the log-likelihood of \mathbf{Y} subject to (A1)–(A4) by solving

$$(P) \quad \underset{\mathbf{W}, \mathbf{C}, \tau}{\text{minimize}} \quad -\sum_{i,j:(i,j) \in \Omega_{\text{obs}}} \log p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j) \\ + \lambda \sum_i \|\mathbf{w}_i\|_1 \quad \text{subject to} \quad \mathbf{W} \geq 0, \tau > 0, \|\mathbf{C}\| \leq \eta.$$

Here, the likelihood of each response $p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j)$ is given by (2). The regularization term $\lambda \sum_i \|\mathbf{w}_i\|_1$ imposes sparsity on each vector \mathbf{w}_i to account for (A2). To prevent arbitrary scaling between \mathbf{W} and \mathbf{C} , we gauge the norm of the matrix \mathbf{C} by applying a matrix norm constraint $\|\mathbf{C}\| \leq \eta$. For example, the Frobenius norm constraint $\|\mathbf{C}\|_F \leq \eta$ can be used. Alternatively, the nuclear norm constraint $\|\mathbf{C}\|_* \leq \eta$ can also be used, promoting low-rankness of \mathbf{C} [9], motivated by the facts that (i) reducing the number of degrees-of-freedom in \mathbf{C} helps to prevent overfitting to the observed data and (ii) learners can often be clustered into a few groups due to their different demographic backgrounds and learning preferences.

The log-likelihood of the observations in (P) is concave in the product $\tau \mathbf{w}_i^T \mathbf{c}_j$ [36]. Consequently, the problem (P) is *tri-concave*, in the sense that the problem obtained by holding two of the three factors \mathbf{W} , \mathbf{C} , and τ constant and optimizing the third one is convex. Therefore, to arrive at a practicable way of solving (P), we propose the following computationally efficient block coordinate descent approach, with \mathbf{W} , \mathbf{C} , and τ as the different blocks of variables.

The matrices \mathbf{W} and \mathbf{C} are initialized as i.i.d. standard normal random variables, and we set $\tau = 1$. We then iteratively optimize the objective of (P) for all three factors in round-robin fashion. Each (outer) iteration consists of three phases: first, we hold \mathbf{W} and τ constant and optimize \mathbf{C} ; second, we hold \mathbf{C} and τ constant and separately optimize each row vector \mathbf{w}_i ; third, we hold \mathbf{W} and \mathbf{C} fixed and optimize over the precision parameter τ . These three phases form the outer loop of Ordinal SPARFA-M.

The sub-problems for estimating \mathbf{W} and \mathbf{C} correspond to the following *ordinal regression* (OR) problems [12]:

$$(OR-W) \quad \underset{\mathbf{w}_i: W_{i,k} \geq 0 \forall k}{\text{minimize}} \quad -\sum_j \log p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j) + \lambda \|\mathbf{w}_i\|_1,$$

$$(OR-C) \quad \underset{\mathbf{C}: \|\mathbf{C}\| \leq \eta}{\text{minimize}} \quad -\sum_{i,j} \log p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j).$$

To solve (OR-W) and (OR-C), we deploy the iterative first-order methods detailed below. To optimize the precision parameter τ , we compute the solution to

$$\underset{\tau > 0}{\text{minimize}} \quad -\sum_{i,j:(i,j) \in \Omega_{\text{obs}}} \log \left(\Phi \left(\tau \left(\mathbf{U}_{i,j} - \mathbf{w}_i^T \mathbf{c}_j \right) \right) - \Phi \left(\tau \left(\mathbf{L}_{i,j} - \mathbf{w}_i^T \mathbf{c}_j \right) \right) \right),$$

via the secant method [26].

Instead of fixing the quantization bin boundaries $\{\omega_0, \dots, \omega_P\}$ introduced in Sec. 2 and optimizing the precision and intrinsic difficulty parameters, one can fix $\tau = 1$ and optimize the bin boundaries instead, an approach used in, e.g., [21]. We emphasize that optimization of the bin boundaries can also be performed straightforwardly via the secant method, iteratively optimizing each bin boundary while keeping the others fixed. We omit the details for the sake of brevity. Note that we have also implemented variants of Ordinal SPARFA-M that directly optimize the bin boundaries, while keeping τ constant; the associated prediction performance is shown in Sec. 4.3.

3.2 First-order methods for regularized ordinal regression

As in [24], we solve (OR-W) using the FISTA framework [4]. (OR-C) also falls into the FISTA framework, by re-writing the convex constraint $\|\mathbf{C}\| \leq \eta$ as a penalty term $\delta(\mathbf{C} : \|\mathbf{C}\| > \eta)$ and treat it as a non-smooth regularizer, where $\delta(\mathbf{C} : \|\mathbf{C}\| > \eta)$ is the delta function, equaling 0 if $\|\mathbf{C}\| \leq \eta$ and ∞ otherwise. Each iteration of both algorithms consists of two steps: A gradient-descent step and a shrinkage/projection step. Take (OR-W), for example, and let $f(\mathbf{w}_i) = -\sum_j \log p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j)$. Then, the gradient step is given by⁴

$$\nabla f = \nabla_{\mathbf{w}_i} \left(-\sum_j \log p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j) \right) = -\mathbf{C}\mathbf{p}. \quad (4)$$

Here, \mathbf{p} is a $N \times 1$ vector, with the j^{th} element equal to

$$\frac{\mathcal{N}(\tau(U_{i,j} - \mathbf{w}_i^T \mathbf{c}_j)) - \mathcal{N}(\tau(L_{i,j} - \mathbf{w}_i^T \mathbf{c}_j))}{\Phi(\tau(U_{i,j} - \mathbf{w}_i^T \mathbf{c}_j)) - \Phi(\tau(L_{i,j} - \mathbf{w}_i^T \mathbf{c}_j))},$$

where $\Phi(\cdot)$ is the inverse probit function. The gradient step and the shrinkage step for \mathbf{w}_i corresponds to

$$\hat{\mathbf{w}}_i^{\ell+1} \leftarrow \mathbf{w}_i^\ell - t_\ell \nabla f, \quad (5)$$

and

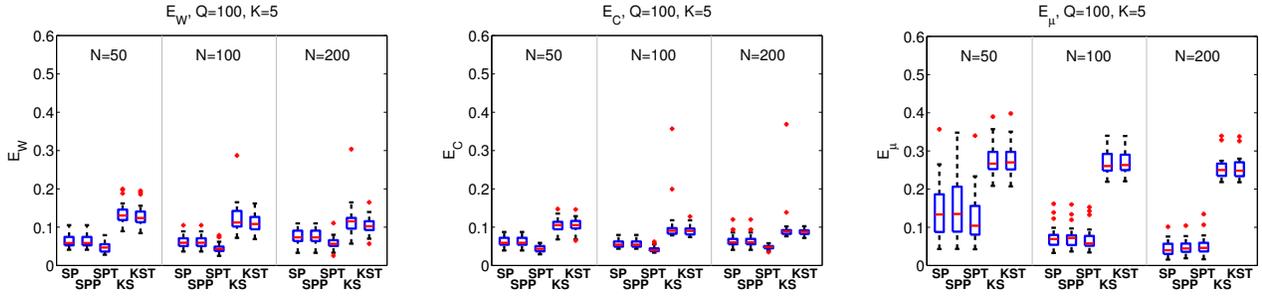
$$\mathbf{w}_i^{\ell+1} \leftarrow \max\{\hat{\mathbf{w}}_i^{\ell+1} - \lambda t_\ell, 0\}, \quad (6)$$

respectively, where t_ℓ is a suitable step-size. For (OR-C), the gradient with respect to each column \mathbf{c}_j is given by substituting \mathbf{W}^T for \mathbf{C} and \mathbf{c}_j for \mathbf{w}_i in (4). Then, the gradient for \mathbf{C} is formed by aggregating all these individual gradient vectors for \mathbf{c}_j into a corresponding gradient matrix.

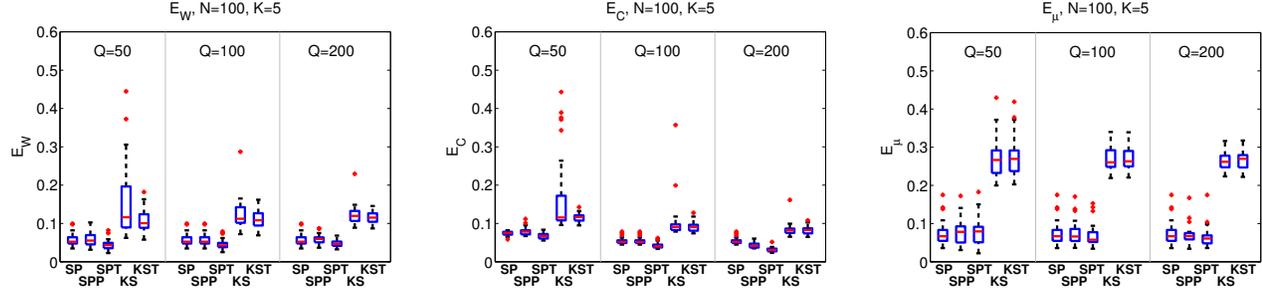
For the Frobenius norm constraint $\|\mathbf{C}\|_F \leq \eta$, the projection step is given by [7]

$$\mathbf{C}^{\ell+1} \leftarrow \begin{cases} \hat{\mathbf{C}}^{\ell+1} & \text{if } \|\hat{\mathbf{C}}^{\ell+1}\|_F \leq \eta \\ \eta \frac{\hat{\mathbf{C}}^{\ell+1}}{\|\hat{\mathbf{C}}^{\ell+1}\|_F} & \text{otherwise.} \end{cases} \quad (7)$$

⁴Here, we assume $\Omega_{\text{obs}} = \{1, \dots, Q\} \times \{1, \dots, N\}$ for simplicity; a generalization to the case of missing entries in \mathbf{Y} is straightforward.



(a) Impact of the number of learners, $N \in \{50, 100, 200\}$, with the number of questions Q fixed.



(b) Impact of the number of questions, $Q \in \{50, 100, 200\}$, with the number of learners N fixed.

Figure 1: Performance comparison of Ordinal SPARFA-M vs. K-SVD₊. “SP” denotes Ordinal SPARFA-M without given support Γ of \mathbf{W} , “SPP” denotes the variant with estimated precision τ , and “SPT” denotes Ordinal SPARFA-Tag. “KS” stands for K-SVD₊, and “KST” denotes its variant with given support Γ .

For the nuclear-norm constraint $\|\mathbf{C}\|_* \leq \eta$, the projection step is given by

$$\mathbf{C}^{\ell+1} \leftarrow \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}^T, \text{ with } \mathbf{s} = \text{Proj}_\eta(\text{diag}(\mathbf{S})), \quad (8)$$

where $\hat{\mathbf{C}}^{\ell+1} = \mathbf{U} \mathbf{S} \mathbf{V}^T$ denotes the singular value decomposition, and $\text{Proj}_\eta(\cdot)$ is the projection onto the ℓ_1 -ball with radius η (see, e.g., [16] for the details).

The update steps (5), (6), and (7) (or (8)) require a suitable step-size t_ℓ to ensure convergence. We consider a constant step-size and set t_ℓ to the reciprocal of the Lipschitz constant [4]. The Lipschitz constants correspond to $\tau^2 \sigma_{\max}^2(\mathbf{C})$ for (OR-W) and $\tau^2 \sigma_{\max}^2(\mathbf{W})$ for (OR-C), with $\sigma_{\max}(\mathbf{X})$ representing the maximum singular value of \mathbf{X} .

3.3 Ordinal SPARFA-Tag

We now develop the Ordinal SPARFA-Tag algorithm that incorporates (A4). Assume that the total number of tags associated with the Q questions equal K (each of the K concepts correspond to a tag), and define the set $\Gamma = \{(i, k) : \text{question } i \text{ has tag } k\}$ as the set of indices of entries in \mathbf{W} identified by pre-defined tags, and $\bar{\Gamma}$ as the set of indices not in Γ , we can re-write the optimization problem (P) as:

$$\begin{aligned} (\text{P}_\Gamma) \quad & \underset{\mathbf{W}, \mathbf{C}, \tau}{\text{minimize}} && - \sum_{i,j:(i,j) \in \Omega_{\text{obs}}} \log p(Y_{i,j} | \tau \mathbf{w}_i^T \mathbf{c}_j) \\ & && + \lambda \sum_i \|\mathbf{w}_i^{(\bar{\Gamma})}\|_1 + \gamma \sum_i \frac{1}{2} \|\mathbf{w}_i^{(\Gamma)}\|_2^2 \\ & \text{subject to} && \mathbf{W} \geq 0, \tau > 0, \|\mathbf{C}\|_* \leq \eta. \end{aligned}$$

Here, $\mathbf{w}_i^{(\Gamma)}$ is a vector of those entries in \mathbf{w}_i belonging to the set Γ , while $\mathbf{w}_i^{(\bar{\Gamma})}$ is a vector of entries in \mathbf{w}_i not belonging to Γ . The ℓ_2 -penalty term on $\mathbf{w}_i^{(\Gamma)}$ regularizes the entries in \mathbf{W} that are part of the (predefined) support of \mathbf{W} ; we set $\gamma = 10^{-6}$ in all our experiments. The ℓ_1 -penalty term

on $\mathbf{w}_i^{(\bar{\Gamma})}$ induces sparsity on the entries in \mathbf{W} that are not predefined but might be in the support of \mathbf{W} . Reducing the parameter λ enables one to discover new question–concept relationships (corresponding to new non-zero entries in \mathbf{W}) that were not contained in Γ .

The problem (P_Γ) is solved analogously to the approach described in Sec. 3.2, except that we split the \mathbf{W} update step into two parts that operate separately on the entries indexed by Γ and $\bar{\Gamma}$. For the entries in Γ , the projection step corresponds to

$$\mathbf{w}_i^{(\Gamma), \ell+1} \leftarrow \max\{\hat{\mathbf{w}}_i^{(\Gamma), \ell+1} / (1 + \gamma t_\ell), 0\}.$$

The step for the entries indexed by $\bar{\Gamma}$ is given by (6). Since Ordinal SPARFA-Tag is tri-convex, it does not necessarily converge to a global optimum. Nevertheless, we can leverage recent results in [24, 35] in order to show that Ordinal SPARFA-Tag converges to a local optimum from an arbitrary starting point. Furthermore, if the starting point is within a close neighborhood of a global optimum of (P), then Ordinal SPARFA-Tag converges to this global optimum.

4. EXPERIMENTS

We first showcase the performance of Ordinal SPARFA-Tag on synthetic data to demonstrate its convergence to a known ground truth. We then demonstrate the ease of interpretation of the estimated factors by leveraging instructor provided tags in combination with a Frobenius or nuclear norm constraint for two real educational datasets. We finally compare the performance of Ordinal SPARFA-M to state-of-the-art collaborative filtering techniques on predicting unobserved ordinal learner responses.

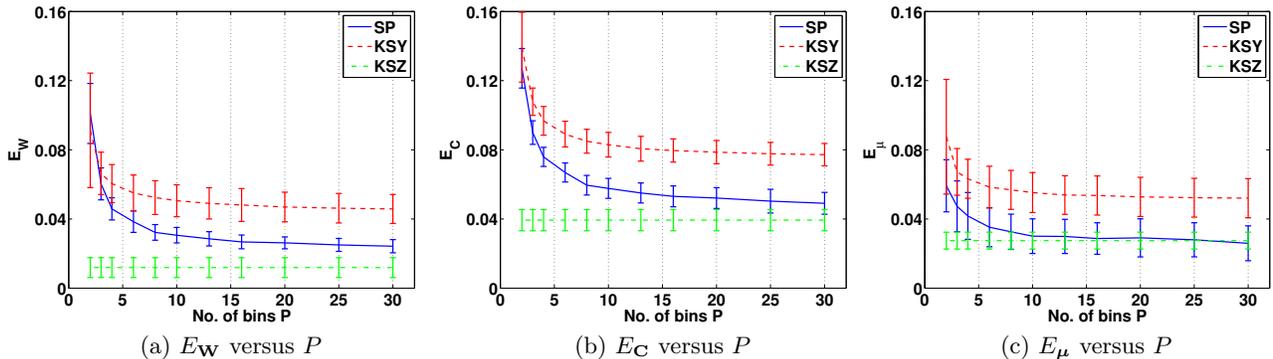


Figure 2: Performance comparison of Ordinal SPARFA-M vs. K-SVD₊ by varying the number of quantization bins. “SP” denotes Ordinal SPARFA-M, “KSY” denotes K-SVD₊ operating on \mathbf{Y} , and “KSZ” denotes K-SVD₊ operating on \mathbf{Z} in (3) (the unquantized data).

4.1 Synthetic data

In order to show that Ordinal SPARFA-Tag is capable of estimating latent factors based on binary observations, we compare the performance of Ordinal SPARFA-Tag to a non-negative variant of the popular K-SVD dictionary learning algorithm [1], referred to as K-SVD₊, which we have detailed in [24]. We consider both the case when the precision τ is known a-priori and also when it must be estimated. In all synthetic experiments, the algorithm parameters λ and η are selected according to Bayesian information criterion (BIC) [17]. All experiments are repeated for 25 Monte-Carlo trials.

In all synthetic experiments, we retrieve estimates of all factors, $\widehat{\mathbf{W}}$, $\widehat{\mathbf{C}}$, and $\widehat{\boldsymbol{\mu}}$. For Ordinal SPARFA-M and K-SVD₊, the estimates $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{C}}$ are re-scaled and permuted as in [24]. We consider the following error metrics:

$$E_{\mathbf{W}} = \frac{\|\mathbf{W} - \widehat{\mathbf{W}}\|_F^2}{\|\mathbf{W}\|_F^2}, \quad E_{\mathbf{C}} = \frac{\|\mathbf{C} - \widehat{\mathbf{C}}\|_F^2}{\|\mathbf{C}\|_F^2}, \quad E_{\boldsymbol{\mu}} = \frac{\|\boldsymbol{\mu} - \widehat{\boldsymbol{\mu}}\|_2^2}{\|\boldsymbol{\mu}\|_2^2}.$$

We generate the synthetic test data \mathbf{W} , \mathbf{C} , $\boldsymbol{\mu}$ as in [24, Eq. 10] with $K = 5$, $\mu_0 = 0$, $v_\mu = 1$, $\lambda_k = 0.66 \forall k$, and $\mathbf{V}_0 = \mathbf{I}_K$. \mathbf{Y} is generated according to (3), with $P = 5$ bins and $\{\omega_0, \dots, \omega_5\} = \{-\infty, -2.1, -0.64, 0.64, 2.1, \infty\}$, such that the entries of \mathbf{Z} fall evenly into each bin. The number of concepts K for each question is chosen uniformly in $\{1, 2, 3\}$. We first consider the impact of problem size on estimation error in Fig. 2. To this end, we fix $Q = 100$ and sweep $N \in \{50, 100, 200\}$ for $K = 5$ concepts, and then fix $N = 100$ and sweep $Q \in \{50, 100, 200\}$.

Impact of problem size: We first study the performance of Ordinal SPARFA-M versus K-SVD₊ while varying the problem size parameters Q and N . The corresponding box-and-whisker plots of the estimation error for each algorithm are shown in Fig. 1. In Fig. 1(a), we fix the number of questions Q and plot the errors $E_{\mathbf{W}}$, $E_{\mathbf{C}}$ and $E_{\boldsymbol{\mu}}$ for the number of learners $N \in \{50, 100, 200\}$. In Fig. 1(b), we fix the number of learners N and plot the errors $E_{\mathbf{W}}$, $E_{\mathbf{C}}$ and $E_{\boldsymbol{\mu}}$ for the number of questions $Q \in \{50, 100, 200\}$. It is evident that $E_{\mathbf{W}}$, $E_{\mathbf{C}}$, and $E_{\boldsymbol{\mu}}$ decrease as the problem size increases for all considered algorithms. Moreover, Ordinal SPARFA-M has superior performance to K-SVD₊ in all cases and for all error metrics. Ordinal SPARFA-Tag and the oracle support provided versions of K-SVD outperform Ordinal SPARFA-

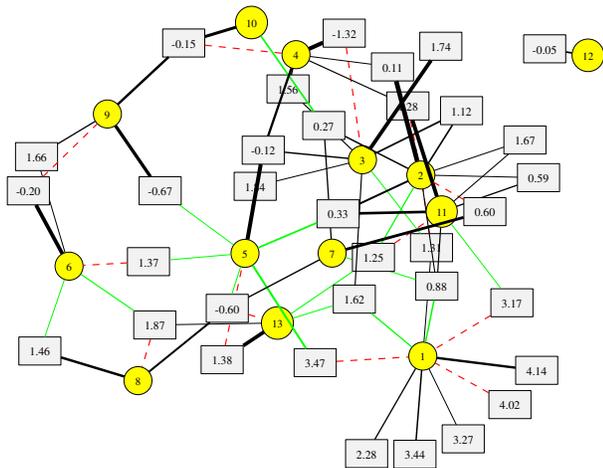
M and K-SVD₊. We furthermore see that the variant of Ordinal SPARFA-M without knowledge of the precision τ performs as well as knowing τ ; this implies that we can accurately learn the precision parameter directly from data.

Impact of the number of quantization bins: We now consider the effect of the number of quantization bins P in the observation matrix \mathbf{Y} on the performance of our algorithms. We fix $N = Q = 100$, $K = 5$ and generate synthetic data as before up to \mathbf{Z} in (3). For this experiment, a different number of bins P is used to quantize \mathbf{Z} into \mathbf{Y} . The quantization boundaries are set to $\{\Phi^{-1}(0), \Phi^{-1}(1/P), \dots, \Phi^{-1}(1)\}$. To study the impact of the number of bins needed for Ordinal SPARFA-M to provide accurate factor estimates that are comparable to algorithms operating with real-valued observations, we also run K-SVD₊ directly on the \mathbf{Z} values (recall (3)) as a base-line. Figure 2 shows that the performance of Ordinal SPARFA-M consistently outperforms K-SVD₊. We furthermore see that all error measures decrease by about half when using 6 bins, compared to 2 bins (corresponding to binary data). Hence, ordinal SPARFA-M clearly outperforms the conventional SPARFA model [24], when ordinal response data is available. As expected, Ordinal SPARFA-M approaches the performance of K-SVD₊ operating directly on \mathbf{Z} (unquantized data) as the number of quantization bins P increases.

4.2 Real-world data

We now demonstrate the superiority of Ordinal SPARFA-Tag compared to regular SPARFA as in [24]. In particular, we show the advantages of using tag information directly within the estimation algorithm and of imposing a nuclear norm constraint on the matrix \mathbf{C} . For all experiments, we apply Ordinal SPARFA-Tag to the graded learner response matrix \mathbf{Y} with oracle support information obtained from instructor-provided question tags. The parameters λ and η are selected via cross-validation.

Algebra test: We analyze a dataset from a high school algebra test carried out on Amazon Mechanical Turk [2], a crowd-sourcing marketplace. The dataset consists of $N = 99$ users answering $Q = 34$ multiple-choice questions covering topics such as geometry, equation solving, and visualizing function graphs. The questions were manually labeled with a set of 13 tags. The dataset is fully populated, with no



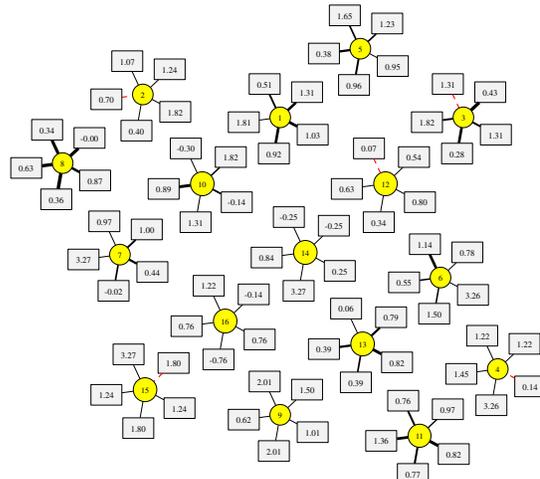
Concept 1	Concept 2	Concept 3	Concept 4	Concept 5
Arithmetic	Simplifying expressions	Solving equations	Fractions	Quadratic functions
Concept 6	Concept 7	Concept 8	Concept 9	Concept 10
Geometry	Inequality	Slope	Trigonometry	Limits
Concept 11	Concept 12	Concept 13		
Polynomials	System equations	Plotting functions		

Figure 3: Question–concept association graph for a high-school algebra test with $N = 99$ users answering $Q = 34$ questions. Boxes represent questions; circles represent concepts. We furthermore show the unique tag associated with each concept.

missing entries. A domain expert manually mapped each possible answer to one of $P = 4$ bins, i.e., assigned partial credit to each choice as follows: totally wrong ($p = 1$), wrong ($p = 2$), mostly correct ($p = 3$), and correct ($p = 4$).

Figure 3 shows the question–concept association map estimated by Ordinal SPARFA-Tag using the Frobenius norm constraint $\|\mathbf{C}\|_F \leq \eta$. Circles represent concepts, and squares represent questions (labelled by their intrinsic difficulties μ_i). Large positive values of μ_i indicate easy questions; negative values indicate hard questions. Connecting lines indicate whether a concept is present in a question; thicker lines represent stronger question–concept associations. Black lines represent the question–concept associations estimated by Ordinal SPARFA-Tag, corresponding to the entries in \mathbf{W} as specified by Γ . Red, dashed lines represent the “mislabeled” associations (entries of \mathbf{W} in Γ) that are estimated to be zero. Green solid lines represent new discovered associations, i.e., entries in \mathbf{W} that were not in Γ that were discovered by Ordinal SPARFA-Tag.

By comparing Fig. 3 with [24, Fig. 9], we can see that Ordinal SPARFA-Tag provides unique concept labels, i.e., one tag is associated with one concept; this enables precise interpretable feedback to individual learners, as the values in \mathbf{C} represent directly the tag knowledge profile for each learner. This tag knowledge profile can be used by a PLS to provide targeted feedback to learners. The estimated question–concept association matrix can also serve as useful tool to domain experts or course instructors, as they indicate missing and inexistent tag–question associations.



Concept 1	Concept 2	Concept 3	Concept 4	Concept 5	
Classifying matter	Properties of water	Mixtures and solutions	Changes from heat	Uses of energy	
Concept 6	Concept 7	Concept 8	Concept 9	Concept 10	
Circuits and electricity	Forces and motion	Formation of fossil fuels	Changes to land	Evidence of the past	
Concept 11	Concept 12	Concept 13	Concept 14	Concept 15	Concept 16
Earth, sun and moon	Alternative energy	Properties of soil	Earth’s forces	Food webs	Environmental changes

Figure 4: Question–concept association graph for a grade 8 Earth Science course with $N = 145$ learners answering $Q = 80$ questions (Y is highly incomplete with only 13.5% entries observed). We furthermore show the unique tag associated with each concept.

Grade 8 Earth Science course: As a second example of Ordinal SPARFA-Tag, we analyze a Grade 8 Earth Science course dataset [31]. This dataset contains $N = 145$ learners answering $Q = 80$ questions and is highly incomplete (only 13.5% entries of \mathbf{Y} are observed). The matrix \mathbf{Y} is binary-valued; domain experts labeled all questions with 16 tags.

The result of Ordinal SPARFA-Tag with the nuclear norm constraint $\|\mathbf{C}\|_* \leq \eta$ on \mathbf{Y} is shown in Fig. 4. The estimated question–concept associations mostly matches those pre-defined by domain experts. Note that our algorithm identified some question–concept associations to be non-existent (indicated with red dashed lines). Moreover, no new associations have been discovered, verifying the accuracy of the pre-specified question tags from domain experts. Comparing to the question–concept association graph of the high school algebra test in Fig. 3, we see that for this dataset, the pre-specified tags represent disjoint knowledge components, which is indeed the case in the underlying question set. Interestingly, the estimated concept matrix \mathbf{C} has rank 3; note that we are estimating $K = 16$ concepts. This observation suggests that all learners can be accurately represented by a linear combination of only 3 different “eigen-learner” vectors. Further investigation of this clustering phenomenon is part of on-going research.

4.3 Predicting unobserved learner responses

We now compare the prediction performance of ordinal SPARFA-M on unobserved learner responses against state-of-the-art collaborative filtering techniques: (i) SVD++ in [20], which treats ordinal values as real numbers, and

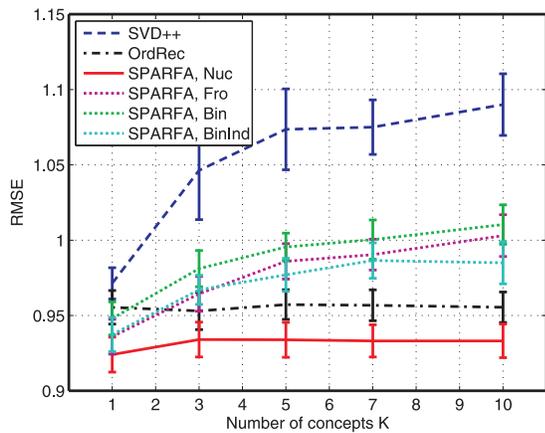


Figure 5: Prediction performance on the Mechanical Turk algebra test dataset. We compare the collaborative filtering methods SVD++ and OrdRec to various Ordinal SPARFA-M based methods: “Nuc” uses the nuclear norm constraint, “Fro” uses the Frobenius norm constraint, “Bin” and “BinInd” learn the bin boundaries, whereas “Bin” learns one set of bin boundaries for the entire dataset and “BinInd” learns individual bin boundaries for each question.

(ii) OrdRec in [21], which relies on an ordinal logit model. We compare different variants of Ordinal SPARFA-M: (i) optimizing the precision parameter, (ii) optimizing a set of bins for all learners, (iii) optimizing a set of bins for each question, and (iv) using the nuclear norm constraint on \mathbf{C} . We consider the Mechanical Turk algebra test, hold out 20% of the observed learner responses as test sets, and train all algorithms on the rest. The regularization parameters of all algorithms are selected using 4-fold cross-validation on the training set. Figure 5 shows the root mean square error (RMSE) $\sqrt{\frac{1}{|\Omega_{\text{obs}}|} \sum_{i,j:(i,j) \in \Omega_{\text{obs}}} \|\hat{Y}_{i,j} - Y_{i,j}\|_2^2}$ where $\hat{Y}_{i,j}$ is the predicted score for $Y_{i,j}$, averaged over 50 trials.

Figure 5 demonstrates that the nuclear norm variant of Ordinal SPARFA-M outperforms OrdRec, while the performance of other variants of ordinal SPARFA are comparable to OrdRec. SVD++ performs worse than all compared methods, suggesting that the use of a probabilistic model considering ordinal observations enables accurate predictions on unobserved responses. We furthermore observe that the variants of Ordinal SPARFA-M that optimize the precision parameter or bin boundaries deliver almost identical performance.

We finally emphasize that Ordinal SPARFA-M not only delivers superior prediction performance over the two state-of-the-art collaborative filtering techniques in predicting learner responses, but it also provides interpretable factors, which is key in educational applications.

5. RELATED WORK

A range of different ML algorithms have been applied in educational contexts. Bayesian belief networks have been successfully used to probabilistically model and analyze learner response data in order to trace learner concept knowledge and estimate question difficulty (see, e.g., [13, 22, 33, 34]).

Such models, however, rely on predefined question–concept dependencies (that are not necessarily accurate), in contrast to the framework presented here that estimates the dependencies solely from data.

Item response theory (IRT) uses a statistical model to analyze and score graded question response data [25, 29]. Our proposed statistical model shares some similarity to the Rasch model [28], the additive factor model [10], learning factor analysis [19, 27], and the instructional factors model [11]. These models, however, rely on pre-defined question features, do not support disciplined algorithms to estimate the model parameters solely from learner response data, or do not produce interpretable estimated factors. Several publications have studied factor analysis approaches on learner responses [3, 14, 32], but treat learner responses as real and deterministic values rather than ordinal values determined by statistical quantities. Several other results have considered probabilistic models in order to characterize learner responses [5, 6], but consider only binary-valued responses and cannot be generalized naturally to ordinal data.

While some ordinal factor analysis methods, e.g., [21], have been successful in predicting missing entries in datasets from ordinal observations, our model enables interpretability of the estimated factors, due to (i) the additional structure imposed on the learner–concept matrix (non-negativity combined with sparsity) and (ii) the fact that we associate unique tags to each concept within the estimation algorithm.

6. CONCLUSIONS

We have significantly extended the SPARse Factor Analysis (SPARFA) framework of [24] to exploit (i) ordinal learner question responses and (ii) instructor generated tags on questions as oracle support information on the question–concept associations. We have developed a computationally efficient new algorithm to compute an approximate solution to the associated ordinal factor-analysis problem. Our proposed *Ordinal SPARFA-Tag* framework not only estimates the strengths of the pre-defined question–concept associations provided by the instructor but can also discover new associations. Moreover, the algorithm is capable of imposing a nuclear norm constraint on the learner concept knowledge matrix, which achieves better prediction performance on unobserved learner responses than state-of-the-art collaborative filtering techniques, while improving the interpretability of the estimated concepts relative to the user-defined tags.

The Ordinal SPARFA-Tag framework enables a PLS to provide readily interpretable feedback to learners about their latent concept knowledge. The tag-knowledge profile can, for example, be used to make personalized recommendations to learners, such as recommending remedial or enrichment material to learners according to their tag (or concept) knowledge status. Instructors also benefit from the capability to discover new question–concept associations underlying their learning materials.

7. ACKNOWLEDGMENTS

Thanks to Daniel Calderon for administering the Algebra test on Amazon’s Mechanical Turk. This work was supported by the National Science Foundation under Cyber-learning grant IIS-1124535, the Air Force Office of Scientific

Research under grant FA9550-09-1-0432, the Google Faculty Research Award program, and the Swiss National Science Foundation under grant PA00P2-134155.

8. REFERENCES

- [1] M. Aharon, M. Elad, and A. M. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans Sig. Proc.*, 54(11):4311–4322, Dec. 2006.
- [2] Amazon Mechanical Turk. <http://www.mturk.com/mturk/welcome>, Sep. 2012.
- [3] T. Barnes. The Q-matrix method: Mining student response data for knowledge. In *Proc. AAAI EDM Workshop*, July 2005.
- [4] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. on Imaging Science*, 2(1):183–202, Mar. 2009.
- [5] B. Beheshti, M. Desmarais, and R. Naceur. Methods to find the number of latent skills. In *Proc. 5th Intl. Conf. on EDM*, pages 81–86, June 2012.
- [6] 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 EDM*, pages 95–102, June 2012.
- [7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [8] P. Brusilovsky and C. Peylo. Adaptive and intelligent web-based educational systems. *Intl. J. of Artificial Intelligence in Education*, 13(2-4):159–172, Apr. 2003.
- [9] J. F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM J. on Optimization*, 20(4):1956–1982, Mar. 2010.
- [10] H. Cen, K. R. Koedinger, and B. Junker. Learning factors analysis—a general method for cognitive model evaluation and improvement. In M. Ikeda, K. D. Ashley, and T. W. Chan, editors, *Intelligent Tutoring Systems*, volume 4053 of *Lecture Notes in Computer Science*, pages 164–175. Springer, June 2006.
- [11] M. Chi, K. Koedinger, G. Gordon, and P. Jordan. Instructional factors analysis: A cognitive model for multiple instructional interventions. In *Proc. 4th Intl. Conf. on EDM*, pages 61–70, July 2011.
- [12] W. Chu and Z. Ghahramani. Gaussian processes for ordinal regression. *J. of Machine Learning Research*, 6:1019–1041, July 2005.
- [13] A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4(4):253–278, Dec. 1994.
- [14] M. Desmarais. Conditions for effectively deriving a Q-matrix from data with non-negative matrix factorization. In *Proc. 4th Intl. Conf. on EDM*, pages 41–50, July 2011.
- [15] J. A. Dijkstra and S. Khan. Khan Academy: The world’s free virtual school. In *APS Meeting Abstracts*, page 14006, Mar. 2011.
- [16] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proc. 25th Intl. Conf. on ML*, pages 272–279, July 2008.
- [17] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, 2010.
- [18] D. Hu. How Khan academy is using machine learning to assess student mastery. *Online: http://david-hu.com*, Nov. 2011.
- [19] K. R. Koedinger, E. A. McLaughlin, and J. C. Stamper. Automated student model improvement. In *Proc. 5th Intl. Conf. on EDM*, pages 17–24, June 2012.
- [20] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, Aug. 2009.
- [21] 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.
- [22] G. A. Krudysz and J. H. McClellan. Collaborative system for signal processing education. In *Proc. IEEE ICASSP*, pages 2904–2907, May 2011.
- [23] J. A. Kulik. Meta-analytic studies of findings on computer-based instruction. *Technology assessment in education and training*, pages 9–33, 1994.
- [24] A. S. Lan, A. E. Waters, C. Studer, and R. G. Baraniuk. Sparse factor analysis for learning and content analytics. Oct. 2012, submitted.
- [25] F. M. Lord. *Applications of Item Response Theory to Practical Testing Problems*. Erlbaum Associates, 1980.
- [26] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Verlag, 1999.
- [27] P. I. Pavlik, H. Cen, and K. R. Koedinger. Learning factors transfer analysis: Using learning curve analysis to automatically generate domain models. In *Proc. 2nd Intl. Conf. on EDM*, pages 121–130, July 2009.
- [28] G. Rasch. *Probabilistic Models for Some Intelligence and Attainment Tests*. MESA Press, 1993.
- [29] M. D. Reckase. *Multidimensional Item Response Theory*. Springer Publishing Company Incorporated, 2009.
- [30] C. Romero and S. Ventura. Educational data mining: A survey from 1995 to 2005. *Expert Systems with Applications*, 33(1):135–146, July 2007.
- [31] STEMscopes Science Education. <http://stemscopes.com>, Sep. 2012.
- [32] N. Thai-Nghe, T. Horvath, and L. Schmidt-Thieme. Factorization models for forecasting student performance. In *Proc. 4th Intl. Conf. on EDM*, pages 11–20, July 2011.
- [33] K. Wauters, P. Desmet, and W. Van Den Noortgate. Acquiring item difficulty estimates: a collaborative effort of data and judgment. In *Proceedings of the 4th Intl. Conf. on EDM*, pages 121–128, July 2011.
- [34] B. P. Woolf. *Building Intelligent Interactive Tutors: Student-centered Strategies for Revolutionizing E-learning*. Morgan Kaufman Publishers, 2008.
- [35] 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.
- [36] A. Zymnis, S. Boyd, and E. Candès. Compressed sensing with quantized measurements. *IEEE Sig. Proc. Letters*, 17(2):149–152, Feb. 2010.