

1-Bit Tensor Completion

Anastasia Aidini^{1,2}, Grigorios Tsagakatakis¹, Panagiotis Tsakalides^{1,2}

¹ ICS - Foundation for Research and Technology - Hellas (FORTH), Crete, Greece

² Department of Computer Science, University of Crete, Greece

Abstract

Higher-order tensor structured data arise in many imaging scenarios, including hyperspectral imaging and color video. The recovery of a tensor from an incomplete set of its entries, known as tensor completion, is crucial in applications like compression. Furthermore, in many cases observations are not only incomplete, but also highly quantized. Quantization is a critical step for high dimensional data transmission and storage in order to reduce storage requirements and power consumption, especially for energy-limited systems. In this paper, we propose a novel approach for the recovery of low-rank tensors from a small number of binary (1-bit) measurements. The proposed method, called 1-bit Tensor Completion, relies on the application of 1-bit matrix completion over different matricizations of the underlying tensor. Experimental results on hyperspectral images demonstrate that directly operating with the binary measurements, rather than treating them as real values, results in lower recovery error.

Introduction

Massive multiway data emerge in many fields. Multidimensional arrays, known as Tensors, are higher-order generalizations of matrices and vectors, and provide a natural way to represent data objects whose entries are indexed by several variables. For instance, a hyperspectral image can be modeled as a third-order tensor defined by two indices for spatial variables and one index for the spectral dimension while a color video corresponds a fourth-order tensor with an additional index for the temporal dimension. Employing therefore tensors to process high dimensional observations has become increasingly popular [1, 2].

One issue which arises frequently is that a tensor may have a significant number of entries missing. Fortunately, in high-dimensional spaces, the corresponding tensors can be well approximated by lower dimensional structures. In such case, we would like to fill in those missing entries based on the available observations. The problem of recovering a low-rank tensor from an incomplete set of its entries, is known as low-rank tensor completion [18]. In many cases, not only are many observations missing, but the available also quantized to a predefined number of bits. Quantization is an integral part of data acquisition, especially for remote sensing scenario like communing measurements from airborne and space-borne platform where channel bandwidth is very limited, as well as application involving energy-limited systems such as Wireless Sensor Networks and Internet of Things platforms, where full data transmission is directly related to an increase in power consumption and a subsequent reduction in the network lifetime. As a result, quantization of a signal is a critical step for high dimensional data transmission but also for data storage.

The question we address in this work is whether it is pos-

sible to recover the real-valued entries of tensors from a small number of highly quantized (binary) observations. This problem significantly differs from currently available tensor completion algorithms which are challenged when the observations are highly quantized, since they treat them as real values instead of discrete. More specifically, we introduce a novel approach for the recovery of a low-rank tensor from a small number of binary measurements, where given a 3-order tensor where only a small number of extremely quantized entries are available, we unfold it into 3 matrices and we apply the 1-bit matrix completion algorithm to the all-mode matricizations of the tensor as illustrated in Figure 1. To the best of our knowledge, this is the first work that examines the interaction between quantization and sampling in high-order structured data. Furthermore, by considering binary encoding, the proposed scheme is directly applicable to tensor data compression, a new and underexplored research topic. In short, the key novelties of this work are

- Explore the recovery of missing and real-valued entries of highly quantized high-order tensors.
- Propose a formally approach for the recovery of a tensor from a small number of binary observations.
- Investigate the performance of the proposed method on publicly available hyperspectral images.

Related Work

The problem of low-rank tensor completion can be regarded as an extension of low rank Matrix Completion. According to MC, the recovery is possible provided the matrix is characterized by a small rank (compared to its dimensions) and enough randomly selected entries of the matrix are acquired [3]. Unfortunately, rank minimization is an NP-hard problem therefore cannot be used for reasonably size data. Since the rank of a matrix corresponds to the number of nonzero singular values, the rank minimization can be replaced by the minimization of the nuclear norm, which is the sum of the singular values of the matrix [4]. The nuclear norm is a convex function which makes the above problem easy to solve, while under mild conditions, the nuclear norm minimization can estimate the same matrix as the rank minimization with high probability [5]. To solve the nuclear norm minimization problem, various approaches have been proposed including the Singular Value Thresholding [7] and the Augmented Lagrange Multiplier Method [8].

There are many methods which extent the problem of LRMC to low-rank matrix completion. Specifically, in [17, 18] is employed matrix nuclear-norm minimization and is used the singular value decomposition (SVD) in the algorithms. A non-convex approach of this problem presented in [20], where the authors apply low-rank matrix factorization [9] to each mode unfolding of

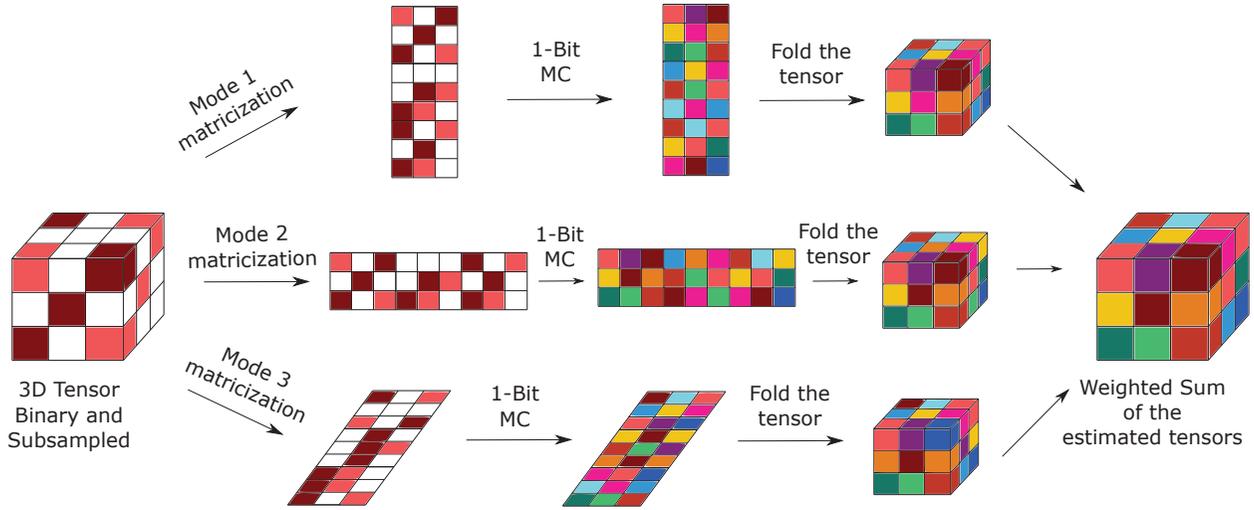


Figure 1. The proposed method applied to a binary and subsampled 3D tensor. Specifically, we have a tensor with two possible values (two colors) and missing entries (white cubes), we unfold it into 3 matrices and we apply the 1-bit matrix completion algorithm to its all-mode matricizations. Then, we fold each of the recovered matrices, which have all their real-valued entries (different colors) and we take the weighted sum of the estimated tensors, with dynamic weights which depend on the fitting error.

the tensor in order to enforce low-rankness and update the matrix factors alternatively. This approach has better performance than the model that performs low-rank matrix factorization to only one mode unfolding [19].

Although numerous algorithms have been presented for the recovery of the missing entries of a tensor, no prior work has been presented for the recovery of a tensor from quantized observations, especially to a single bit. Unlike tensors however, existing methods for quantized matrix completion in [10, 11, 12] have developed statistical models to solve the convex optimization problem of the recovery of a low-rank matrix from quantized and possibly corrupted measurements, either estimating the set of quantization bin boundaries from observed data using an alternative optimization procedure, or assuming that the quantization bin boundaries are known. For the extreme case of noisy 1-bit observations, recent work has estimated the matrix with all its real entries via solving a constrained maximum likelihood optimization problem. Under the assumption that the matrix is low-rank these works have used convex relaxations for the rank via the nuclear norm [14] or max-norm [15], assuming that the entries are sampled according to a uniform distribution, or a non-uniform distribution respectively. In addition, in [16] is considered constrained maximum likelihood estimation of the underlying matrix, under a constraint on its entry-wise infinity-norm and an exact rank constraint.

Quantization and Statistical Model

We aim at the recovery of an unknown, low-rank tensor $\mathcal{M} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ from partial binary observations $\mathcal{Y} = \mathcal{P}_\Omega(Q(\mathcal{M}))$, where $\Omega \subseteq \{1, \dots, I_1\} \times \dots \times \{1, \dots, I_N\}$ is the index set of observed entries and \mathcal{P}_Ω is a random sampling operator which keeps the entries in Ω and zeros out others, retaining only a small number

of entries from the tensor. In addition, the function $Q(\cdot) : \mathbb{R} \rightarrow \mathcal{F}$ corresponds to a non-uniform scalar quantizer that maps a real number to a set of two ordered labels $\mathcal{F} = \{1, 2\}$ according to

$$Q(x) = \begin{cases} 1, & \text{if } w_0 < x \leq w_1 \\ 2, & \text{if } w_1 < x \leq w_2 \end{cases} \quad (1)$$

where $\{w_0, w_1, w_2\}$ represents the set of quantization bin boundaries of all measurements, which satisfies $w_0 \leq w_1 \leq w_2$. We will assume that the set of quantization bin boundaries is known a priori.

Let $\mathcal{Y}_{i_1 \dots i_N} \in \mathcal{F}$ represents the binary measurement of the (i_1, \dots, i_N) -th entry of the (unknown) tensor \mathcal{M} . We use the following model for the binary measurements $\mathcal{Y}_{i_1 \dots i_N}$:

$$\begin{aligned} \mathcal{Y}_{i_1 \dots i_N} &= Q(\mathcal{M}_{i_1 \dots i_N} + \varepsilon_{i_1 \dots i_N}), \quad (i_1, \dots, i_N) \in \Omega \\ \varepsilon_{i_1 \dots i_N} &\sim \text{Logistic}(0, 1) \text{ or } \varepsilon_{i_1 \dots i_N} \sim \mathcal{N}(0, 1) \end{aligned} \quad (2)$$

The quantities $\varepsilon_{i_1 \dots i_N}$ model the uncertainty on each measurement of $\mathcal{M}_{i_1 \dots i_N}$. $\text{Logistic}(0, 1)$ denotes a logistic distribution with zero mean and unit scale and $\mathcal{N}(0, 1)$ denotes the standard normal distribution.

In terms of the likelihood of the observations $\mathcal{Y}_{i_1 \dots i_N}$, the model in (2) can be written equivalently as

$$p(\mathcal{Y}_{i_1 \dots i_N} | \mathcal{M}_{i_1 \dots i_N}) = \Phi(\mathcal{U}_{i_1 \dots i_N} - \mathcal{M}_{i_1 \dots i_N}) - \Phi(\mathcal{L}_{i_1 \dots i_N} - \mathcal{M}_{i_1 \dots i_N}), \quad (3)$$

where the $I_1 \times \dots \times I_N$ tensors \mathcal{U} and \mathcal{L} contain the upper and lower bin boundaries corresponding to the measurements $\mathcal{Y}_{i_1 \dots i_N}$, i.e., we have $\mathcal{U}_{i_1 \dots i_N} = w_{\mathcal{Y}_{i_1 \dots i_N}}$ and $\mathcal{L}_{i_1 \dots i_N} = w_{\mathcal{Y}_{i_1 \dots i_N} - 1}$. Furthermore, the function $\Phi(x)$ corresponds to an inverse link function. For the logistic model (logistic noise), we use the inverse

logit link function $\Phi_{\log}(x) = \frac{1}{1+e^{-x}}$, and for the probit model (standard normal noise), we use the inverse probit link function $\Phi_{\text{pro}}(x) = \int_{-\infty}^x \mathcal{N}(s | 0, 1) ds$. The proposed algorithm can be formulated for both noise models.

1-Bit Tensor Completion

In order to recover the missing entries, but also recover the real values of the low-rank tensor \mathcal{M} from partial binary observations, we unfold the measurement tensor $\mathcal{Y} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ into N matrices and for each of them, we apply the 1-bit matrix completion algorithm. Formally, the n -th of these matrices is called the mode- n matricization or unfolding of the tensor \mathcal{Y} , and is denoted as $\text{unfold}_n(\mathcal{Y}) = \mathbf{Y}_{(n)} \in \mathbb{R}^{I_n \times \prod_{j \neq n} I_j}$ and corresponds to a matrix with columns being the vectors obtained by fixing all indices of \mathcal{Y} except the n -th index. The estimated tensors $\mathcal{Z}_n = \text{fold}_n(\mathbf{Z}_{(n)})$, $n = 1, \dots, N$, is produced by folding each of the recovered matrices $\mathbf{Z}_{(n)}$ such that:

$$\mathcal{M} \approx \sum_{n=1}^N a_n \cdot \mathcal{Z}_n \quad (4)$$

where a_n , $n = 1, \dots, N$, are weights, which depend on the fitting error, and satisfy $\sum_n a_n = 1$.

1-Bit Matrix Completion Algorithm

Our model can be regarded as an extension of the quantized matrix completion [11] to the case of quantized tensor completion, for the extreme case of 1-bit observations. In particular, in order to recover the low-rank mode- n matricization $\mathbf{M}_{(n)}$ from binary measurements, one seeks to minimize the negative log-likelihood of $\mathbf{Y}_{(n)}(j,k)$, $(j,k) \in \Omega_n$ (where Ω_n is the index set of observed entries of $\mathbf{M}_{(n)}$), given by (3), subject to a low-rank constraint on $\mathbf{M}_{(n)}$, i.e., we seek to solve the following constrained optimization problem:

$$\begin{aligned} & \text{minimize } \mathbf{M}_{(n)} && - \sum_{j,k:(j,k) \in \Omega_n} \log p(\mathbf{Y}_{(n)}(j,k) | \mathbf{M}_{(n)}(j,k)) \\ & \text{subject to} && \|\mathbf{M}_{(n)}\|_* \leq \lambda. \end{aligned} \quad (5)$$

The nuclear norm constraint $\|\mathbf{M}_{(n)}\|_* \leq \lambda$ is a convex relaxation of the low-rank constraint which promotes low-rankness of $\mathbf{M}_{(n)}$ [6] and the parameter $\lambda > 0$ is used to control its rank.

Since the gradient of the negative log-likelihood of the inverse logit and probit link functions are convex in $\mathbf{M}_{(n)}$ when keeping the quantization bin boundaries w_0, w_1, w_2 fixed, the optimization problem in (5) can be solved efficiently. Starting with an initialization of the estimated matrix $\mathbf{Z}_{(n)}$ as the measurement matrix $\mathbf{Y}_{(n)}$, i.e., $\mathbf{Z}_{(n)}^1 = \mathbf{Y}_{(n)}$, the algorithm performs two steps at each iteration $l = 1, 2, \dots$. Both steps are repeated until a maximum number of iteration l_{\max} is reached or the change in $\mathbf{Z}_{(n)}$ between consecutive iterations is below a given threshold.

The first step aims at reducing the objective function $f(\mathbf{Z}_{(n)}) = -\sum_{j,k:(j,k) \in \Omega_n} \log p(\mathbf{Y}_{(n)}(j,k) | \mathbf{Z}_{(n)}(j,k))$ of (5) and is given by

$$\hat{\mathbf{Z}}_{(n)}^{l+1} \leftarrow \mathbf{Z}_{(n)}^l - s_l \cdot \nabla f, \quad (6)$$

where s_l is the step-size at iteration l . For simplicity, we use a constant step-size $s_l = \frac{1}{L}$, where L is the Lipschitz constant, which

is given by $L_{\log} = \frac{1}{4}$ for the logistic model and $L_{\text{pro}} = 1$ for the probit model. The gradient of the objective function $f(\mathbf{Z}_{(n)})$, with respect to $\mathbf{Z}_{(n)}$, is given by

$$[\nabla f]_{j,k} = \begin{cases} \frac{\Phi'(\mathbf{L}_{(n)}(j,k) - \mathbf{Z}_{(n)}(j,k)) - \Phi'(\mathbf{U}_{(n)}(j,k) - \mathbf{Z}_{(n)}(j,k))}{\Phi(\mathbf{U}_{(n)}(j,k) - \mathbf{Z}_{(n)}(j,k)) - \Phi(\mathbf{L}_{(n)}(j,k) - \mathbf{Z}_{(n)}(j,k))} & \text{if } (j,k) \in \Omega_n \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathbf{L}_{(n)}$ and $\mathbf{U}_{(n)}$ are the mode- n matricizations of \mathcal{L} and \mathcal{U} and contain the lower and the upper bin boundaries of the observations $\mathbf{Y}_{(n)}(j,k)$ respectively, i.e., $\mathbf{L}_{(n)} = \text{unfold}_n(\mathcal{L})$ and $\mathbf{U}_{(n)} = \text{unfold}_n(\mathcal{U})$. The derivative of the inverse link function $\Phi'(x)$ can be calculated as $\Phi'_{\log}(x) = \frac{1}{2+e^{-x}+e^x}$ and $\Phi'_{\text{pro}}(x) = \mathcal{N}(x | 0, 1)$.

The second step aims to impose low-rankness on $\mathbf{Z}_{(n)}$ in order to make the solution satisfy the constraint $\|\mathbf{Z}_{(n)}\|_* \leq \lambda$. In order to achieve this, we apply the Augmented Lagrangian Multipliers (ALM) method [8]. Specifically, we solve the following optimization problem

$$\min_{\mathbf{Z}_{(n)}^{l+1}} \|\mathbf{Z}_{(n)}^{l+1}\|_* \quad \text{subject to } \mathbf{Z}_{(n)}^{l+1} + \mathbf{E}^{l+1} = \hat{\mathbf{Z}}_{(n)}^{l+1}, \quad \mathcal{P}_{\Omega}(\mathbf{E}^{l+1}) = 0 \quad (7)$$

As \mathbf{E}^l will compensate for the unknown entries of $\hat{\mathbf{Z}}_{(n)}^l$ are simply set as zeros. So, the partial augmented Lagrangian function of (7) is

$$\begin{aligned} L(\mathbf{Z}_{(n)}^{l+1}, \mathbf{E}^{l+1}, \mathbf{A}^{l+1}, \mu) &= \|\mathbf{Z}_{(n)}^{l+1}\|_* + \\ &< \mathbf{A}^{l+1}, \hat{\mathbf{Z}}_{(n)}^{l+1} - \mathbf{Z}_{(n)}^{l+1} - \mathbf{E}^{l+1} > + \frac{\mu}{2} \|\hat{\mathbf{Z}}_{(n)}^{l+1} - \mathbf{Z}_{(n)}^{l+1} - \mathbf{E}^{l+1}\|_F^2 \end{aligned} \quad (8)$$

Then we can have the inexact ALM approach for the matrix completion problem, where for updating \mathbf{E}^l the constraint $\mathcal{P}_{\Omega}(\mathbf{E}^l) = 0$ should be enforced when minimizing the Lagrangian function $L(\mathbf{Z}_{(n)}^l, \mathbf{E}^l, \mathbf{A}^l, \mu)$. The inexact ALM approach is described with more details in [8].

Dynamic weights

The weights a_1, \dots, a_N in (4) can uniformly set to $\frac{1}{N}$. But, in some cases, the recovery in one unfolding maybe better than others. Therefore, instead of fixed weights, we use dynamic weights which depend on the fitting error

$$\text{fit}_n(\mathbf{Z}_{(n)}) = \|\mathcal{P}_{\Omega}(\text{fold}_n(\mathbf{Z}_{(n)}) - \mathcal{Y})\|_F, \quad (9)$$

where $\|\mathcal{X}\|_F = \sqrt{\langle \mathcal{X}, \mathcal{X} \rangle}$ is the Frobenius norm of \mathcal{X} (with $\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{i_1=1}^{I_1} \dots \sum_{i_N=1}^{I_N} x_{i_1 \dots i_N} y_{i_1 \dots i_N}$ is denoted to be the inner product of $\mathcal{X}, \mathcal{Y} \in \mathbb{R}^{I_1 \times \dots \times I_N}$). The smaller $\text{fit}_n(\mathbf{Z}_{(n)})$ is, the larger a_n should be. Specifically, we set

$$a_n = \frac{[\text{fit}_n(\mathbf{Z}_{(n)})]^{-1}}{\sum_{i=1}^N [\text{fit}_i(\mathbf{Z}_{(i)})]^{-1}}, \quad n = 1, \dots, N. \quad (10)$$

As demonstrated below, the dynamic weights a_n can improve the recovery quality of the recovered tensor.

Experimental results

In this section, we present experimental results on hyperspectral Earth Observation images taken from airborne or satellites which are publicly available [21]. Specifically, we considered the hyperspectral images over Indian Pines, Botswana, Pavia Center, Pavia University and Kennedy Space Center, in order to validate the efficacy of the 1-bit tensor completion algorithm (1BTC). The first two images use 14 bits per pixel, the other two use 13 bits per pixel and the last one uses 16 bits per pixel. We quantized the images to a single bit, we subsampled their entries and we recovered each image applying the proposed algorithm. To assess the recovery performance of our algorithm for different sampling percentages, we use the Spectral Angle Mapper (SAM) [22]. SAM is spectral technique that measures the similarity of image pixel spectra to the spectra of the reconstructed image.

Figure 2 presents the results of the 1-bit tensor completion algorithm to each mode matricization, using the probit model on the hyperspectral image over Indian Pines. In this figure, is demonstrated that the dynamic weights improve the recovery quality of the recovered tensor, as the mode 1 and mode 2 matricizations have better performance than the mode 3 matricization.

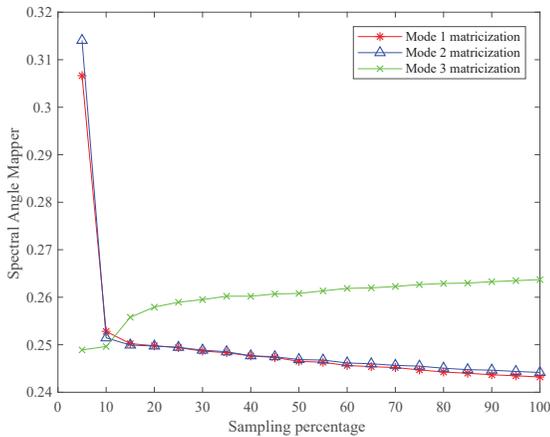


Figure 2. Reconstruction error to each mode matricization, using the probit model and the hyperspectral image over Indian Pines.

In Figure 3, we can see the original and the reconstructed image for 20% sampling percentage, using the same hyperspectral image and the same model.

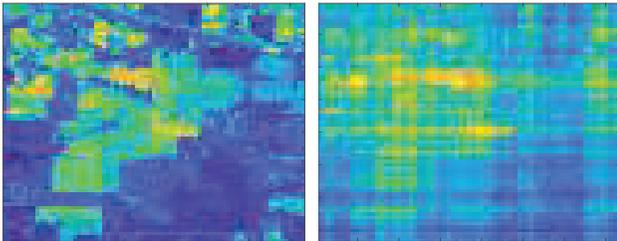


Figure 3. The original and the reconstructed image of Indian Pines for 20% sampling percentage, using the probit model.

Figure 4 shows the proposed algorithm applied on the test image of Figure 3, using the logistic and the probit model, in

comparison with the linear interpolation method for the recovery of the unknown tensor. As we can see, our method outperforms the linear interpolation even for low sampling percentage. In addition, the logistic and the probit model have the same performance.

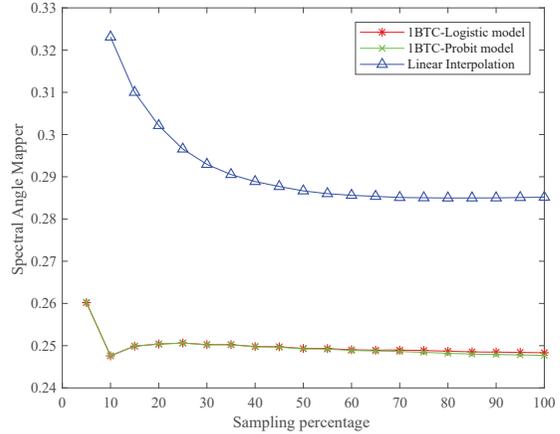


Figure 4. Comparison of 1BTC algorithm using the logistic and the probit model versus linear interpolation on the hyperspectral image over Indian Pines.

For Figure 5, we quantized a part of the hyperspectral image over Indian Pines to 1, 4, 8 and 14 bits and we used these images as the original image in each case. Applying the 1-bit tensor completion algorithm using the probit model, the results are presented in Figure 5. As it was expected, the fewer the bits per pixel of the original image, the better the performance of reconstruction.

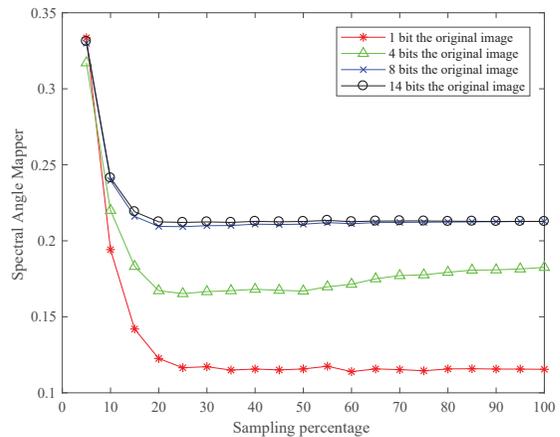


Figure 5. Reconstruction error for each number of bits per pixel of the original image, using the probit model.

Finally, in Figure 6 we applied the 1-bit tensor completion algorithm on each hyperspectral image that was described above and we measured the performance by computing the peak signal to noise ratio (PSNR) in decibels, between the original and the estimated image. Higher PSNR represents better the quality of the recovered image. Specifically, PSNR is computed using the

equation

$$PSNR = 10 \cdot \log_{10}\left(\frac{R^2}{MSE}\right) \quad (11)$$

where R is the maximum fluctuation in the input image data type and MSE is denoted as the mean square error, which is the average of the squares of the differences between the original and the estimated signal.

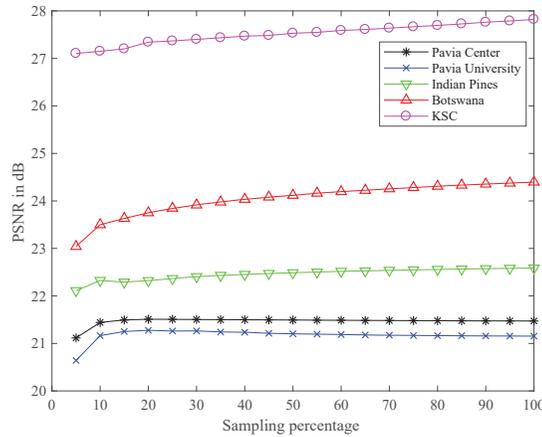


Figure 6. PSNR for each hyperspectral image, using the probit model.

In Figures 7 and 8, we can see the original and the reconstructed image for 50% sampling percentage of the hyperspectral images over Pavia Center and Pavia University, using the probit model.

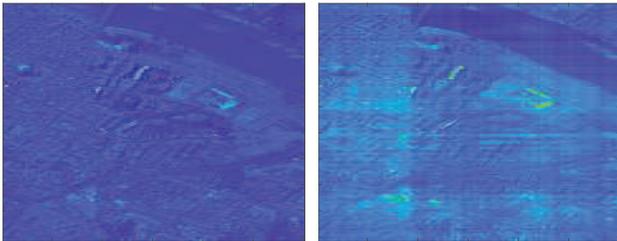


Figure 7. The original and the reconstructed image of Pavia Center for 50% sampling percentage, using the probit model.

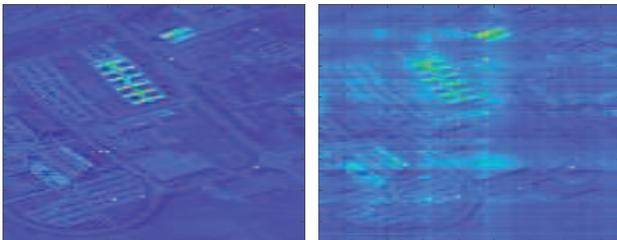


Figure 8. The original and the reconstructed image of Pavia University for 50% sampling percentage, using the probit model.

Conclusions

In this work, we presented a novel approach for the recovery of a low-rank tensor from an incomplete set of its binary entries. This problem is crucial especially in compression while, in many applications of tensor completion are considered discrete observations, often in the form of binary measurements. However, a simple method like linear interpolation, is extremely ill-posed, even if one collects a binary measurement for each of the tensor entries. Experimental results on real data demonstrate that it is better to take into account that the quantization of the measurements compared to treating them as actual observed values. An other issue that would also be interesting to study is how the proposed algorithm performs with measurements that are quantized to more than 2 (but still a small number) of different values, but we leave such investigations for future work.

Acknowledgments

This work was funded by the General Secretariat for Research and Technology (GSRT) and by the Hellenic Foundation for Research and Innovation (HFRI) within the scholarships for doctoral candidates.

References

- [1] Giannopoulos, Michalis, et al., Application of Tensor and Matrix Completion on Environmental Sensing Data.
- [2] Savvaki, Sofia, et al., Matrix and Tensor Completion on a Human Activity Recognition Framework, IEEE journal of biomedical and health informatics 21.6 (2017): 1554-1561.
- [3] Candès, Emmanuel J., and Benjamin Recht, Exact matrix completion via convex optimization, Foundations of Computational mathematics 9.6 (2009): 717.
- [4] Recht, Benjamin, Maryam Fazel, and Pablo A. Parrilo, Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization, SIAM review 52.3, pp. 471-501. (2010).
- [5] Candès, Emmanuel J., and Terence Tao, The power of convex relaxation: Near-optimal matrix completion, IEEE Transactions on Information Theory 56.5, pp. 2053-2080. (2010).
- [6] Fazel, Maryam, Haitham Hindi, and Stephen P. Boyd, A rank minimization heuristic with application to minimum order system approximation, American Control Conference, 2001, Proceedings of the 2001, Vol. 6, IEEE, pp. 4734-4739. (2001).
- [7] Cai, Jian-Feng, Emmanuel J. Candès, and Zuowei Shen, A singular value thresholding algorithm for matrix completion, SIAM Journal on Optimization 20.4, pp. 1956-1982. (2010).
- [8] Lin, Zhouchen, Minming Chen, and Yi Ma, The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices, arXiv preprint arXiv:1009.5055 (2010).
- [9] Z. Wen, W. Yin, and Y. Zhang, Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm, Mathematical Programming Computation, pp. 1-29.(2012).
- [10] A. S. Lan, C. Studer, and R. G. Baraniuk, Matrix recovery from quantized and corrupted measurements, In Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on, IEEE, pp. 4973-4977. (2014).
- [11] Lan, Andrew S., Christoph Studer, and Richard G. Baraniuk, Quantized matrix completion for personalized learning, arXiv preprint arXiv:1412.5968 (2014).
- [12] Bhaskar, Sonia A., Probabilistic low-rank matrix completion from

- quantized measurements, *Journal of Machine Learning Research* 17.60, pp. 1–34. (2016).
- [13] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, Efficient projections onto the l_1 -ball for learning in high dimensions, *Proceedings of the 25th international conference on Machine learning*, ACM, pp. 272–279. (2008).
- [14] M. A. Davenport, Y. Plan, E. van den Berg, and M. Wooters, 1-bit matrix completion, *arXiv preprint:1209.3672* (2012).
- [15] Cai, Tony, and Wen-Xin Zhou, A max-norm constrained minimization approach to 1-bit matrix completion, *The Journal of Machine Learning Research* 14.1, pp. 3619–3647. (2013).
- [16] Bhaskar, Sonia A., and Adel Javanmard, 1-bit matrix completion under exact low-rank constraint, *Information Sciences and Systems (CISS), 2015 49th Annual Conference on*. IEEE, pp. 1–6. (2015).
- [17] J. Liu, P. Musialski, P. Wonka, and J. Ye, Tensor completion for estimating missing values in visual data, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 208–220. (2013).
- [18] S. Gandy, B. Recht, and I. Yamada, Tensor completion and low-rank tensor recovery via convex optimization, *Inverse Problems*, 27, pp. 1–19. (2011).
- [19] S. Ma, D. Goldfarb, and L. Chen, Fixed point and Bregman iterative methods for matrix rank minimization, *Mathematical Programming*, 128, pp. 321–353. (2011).
- [20] Yangyang Xu, Ruru Hao, Wotao Yin, and Zhixun Su, Parallel matrix factorization for low-rank tensor completion, *arXiv preprint arXiv:1312.1254* (2013).
- [21] www.ehu.eus/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes
- [22] Yang, Chenghai, James H. Everitt, and Joe M. Bradford, Yield estimation from hyperspectral imagery using spectral angle mapper (SAM), *Transactions of the ASABE* 51.2, pp. 729–737. (2008).

Author Biography

Anastasia Aidini is currently pursuing the PhD degree at the Computer Science Department of University of Crete. She received her M.S. degree in Mathematical Foundations of Informatics from the Department of Mathematics and Applied Mathematics of University of Crete and B.S. degree from the Department of Mathematics of University of Crete in 2016 and 2014 respectively. Her main research interests lie in the fields of Matrix and Tensor Completion, Compressed Sensing and Coding.

Grigorios Tsagkatakis received his Diploma and M.S. degrees in Electronics and Computer Engineering from the Technical University of Crete, Greece in 2005 and 2007 respectively, and his Ph.D. in Imaging Science from the Rochester Institute of Technology, New York, in 2011. Currently, he is a research associate with the Signal Processing Laboratory at ICS-FORTH and is working on the intersection of imaging science and computational intelligence. He has co-authored more than 40 peer-reviewed conference papers, journals and book chapters, has been awarded one US patent and two best paper awards in the 2010 Western New York Image Processing Workshop and in 2016 Image Sensors and Imaging Systems Conference, Electronic Imaging Symposium.

Panagiotis Tsakalides received the Diploma degree from Aristotle University of Thessaloniki, Greece, and the Ph.D. degree from the University of Southern California, Los Angeles, USA, in 1990 and 1995, respectively, both in electrical engineering. He is a Professor of Computer Science, the Vice Rector for Finance at the University of Crete, and the Head of the Signal Processing Laboratory at FORTH-ICS. He has coauthored over 150 technical publications, including 35+ journal papers. He has been the Project Coordinator in seven European Commission and nine

national projects. His research interests include statistical signal processing with emphasis in non-Gaussian estimation and detection theory, sparse representations, and applications in sensor networks, audio, imaging, and multimedia systems.