INSTITUTE OF COMPUTER SCIENCE

H.F.R.I.
Hellenic Foundation for Research & Innovation

GENERAL SECRETARIAT FOR RESEARCH AND TECHNOLOGY

QUANTIZED TENSOR ROBUST PRINCIPAL COMPONENT ANALYSIS

Anastasia Aidini[1,2], Grigorios Tsagkatakis[1], Panagiotis Tsakalides[1,2]

{aidini, greg, tsakalid}@ics.forth.gr

[1]Institute of Computer Science FORTH, Greece

[2]Computer Science Department, UOC

ICASSP, 2020

Motivation

- Continuous growth of data
- Multidimensional observations known as **tensors**

  Hyperspectral images
  Time series of images
  Color video

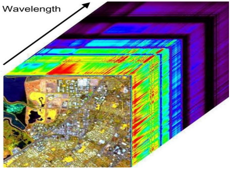

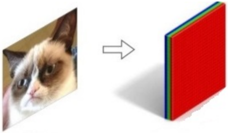

Challenges:

- *Quantized* measurements for compression purposes
- *Lost* measurements due to communication failures
- *Corrupted* measurements due to noise
- *Anomalies* in the data

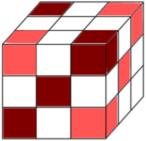

Problem

Recovery of all the real-valued entries of a high-dimensional signal from a number of quantized and sparsely corrupted measurements.

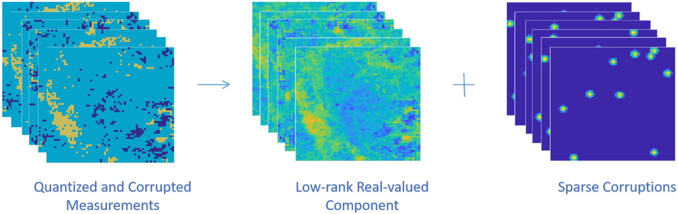

Proposed solution:
A novel *quantized tensor robust principal component analysis* algorithm is used to recover the tensor as a sum of a low-rank and a sparse tensor, through matricizations in each mode.

Related Work

Tensor Robust Principal Component Analysis using

- tensor singular value decomposition
- tensor decomposition
- tensor unfoldings

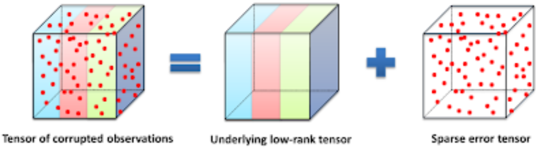

Tensor completion from binary measurements using

- nuclear norm constraint on the different matricizations
- tensor nuclear norm
- tensor decomposition

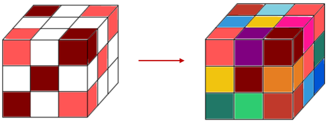

From Tensor to Matrices

$\mathscr{M} \in \mathbb{R}^{I_1 \times \cdots \times I_N}$ is a N-way array, the unknown tensor, such that

$$\mathscr{M} = \mathscr{L} + \mathscr{S},$$

where $\mathscr{L}$ is a low-rank tensor and $\mathscr{S}$ a sparse tensor of the same dimensions.

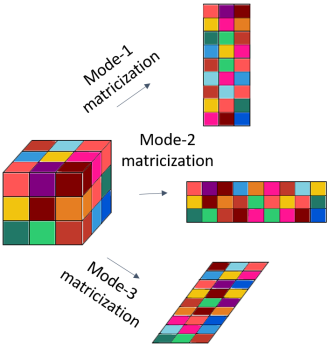

The *mode-n unfolded matrix* $\mathbf{M}_{(n)} \in \mathbb{R}^{I_n \times \prod_{i \neq n} I_i}$ corresponds to a matrix with columns being the vectors obtained by fixing all indices of $\mathscr{M}$ except the n-th index.

Robust Principal Component Analysis

If $\mathscr{P}_\Omega$ is a random sampling operator with sampling set Ω, then the optimization problem on real-valued matrices is:

$$\begin{aligned} \text{minimize} \quad & \|\mathbf{Z}_{(n)}\|_* + \lambda \|\mathbf{S}_{(n)}\|_1 \\ \text{subject to} \quad & \mathscr{P}_\Omega(\mathbf{Z}_{(n)} + \mathbf{S}_{(n)}) = \mathscr{P}_\Omega(\mathbf{M}_{(n)}) \end{aligned}$$

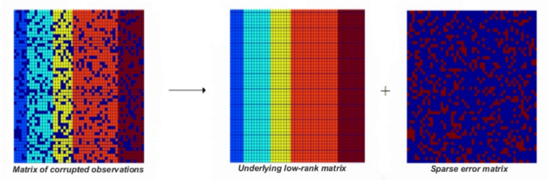

$$\mathbf{M}_{(n)} \quad = \quad \mathbf{Z}_{(n)} \quad + \quad \mathbf{S}_{(n)}$$

Quantization Model

The quantized measurement of the $(i_1, ..., i_N) - th$ entry of $\mathscr{M}$ is

$$\mathscr{Y}_{i_1..i_N} = Q(\mathscr{Z}_{i_1..i_N} + \mathscr{S}_{i_1..i_N} + \epsilon_{i_1..i_N}),\ (i_1, .., i_N) \in \Omega$$

$$\epsilon_{i_1...i_N} \sim \text{Logistic}(0, 1) \ \text{ or } \ \epsilon_{i_1...i_N} \sim \mathscr{N}(0, 1)$$

where $\mathscr{Q} : \mathbb{R} \to \{1, .., K\}$ is a uniform scalar quantizer that

$$\mathscr{Q}(x) = l \text{ if } w_{l-1} < x \leq w_l,\ l \in \{1, ..., K\},$$

where $\{w_0, w_1, \cdots, w_K\}$ represents the set of quantization bin boundaries of all measurements (we assume that is known a priori).

Quantized Robust Principal Component Analysis

We solve the following constrained optimization problem:

$$\begin{array}{ll} \min_{\mathbf{Z}_{(n)},\mathbf{S}_{(n)}} & -\sum_{(j,k)\in\Omega_n} \log p(\mathbf{Y}_{(n)_{j,k}} \mid \mathbf{Z}_{(n)_{j,k}} + \mathbf{S}_{(n)_{j,k}}) \\ \text{subject to} & \|\mathbf{Z}_{(n)}\|_* \leq \lambda \text{ and } \|\mathbf{S}_{(n)}\|_1 \leq \sigma \end{array}$$

If the tensors $\mathscr{U}, \mathscr{L} \in \mathbb{R}^{I_1 \times \cdots \times I_N}$ contain the upper and lower bin boundaries corresponding to the measurements, then $p(\mathbf{Y}_{(n)_{j,k}}|\mathbf{Z}_{(n)_{j,k}} + \mathbf{S}_{(n)_{j,k}}) = \Phi(\mathbf{U}_{(n)_{j,k}} - \mathbf{Z}_{(n)_{j,k}} - \mathbf{S}_{(n)_{j,k}}) - \Phi(\mathbf{L}_{(n)_{j,k}} - \mathbf{Z}_{(n)_{j,k}} - \mathbf{S}_{(n)_{j,k}})$.

The function $\Phi(x)$ corresponds to an inverse link function.

- Logistic model: $\Phi_{\text{log}}(x) = \frac{1}{1+e^{-x}}$,
- Probit model: $\Phi_{\text{pro}}(x) = \int_{-\infty}^{x} \mathcal{N}(s \mid 0, 1)\, ds$.

Algorithm (1/2)

Starting with a random low-rank component $\mathbf{Z}_{(n)}$ and a zero sparse component $\mathbf{S}_{(n)}$, the algorithm optimize iteratively each individual variable while holding the other fixed. At each iteration q:

1. Optimize $\mathbf{Z}_{(n)}$ iteratively by holding $\mathbf{S}_{(n)}$ fixed. At each iteration l:

    i) Gradient step to optimize the low-rank component $\mathbf{Z}_{(n)}$:

$$\hat{\mathbf{Z}}_{(n)}^{l+1} \leftarrow \mathbf{Z}_{(n)}^{l} - c \cdot \nabla F,$$

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

where $\mathbf{X} = \mathbf{Z}_{(n)}^{l} + \mathbf{S}_{(n)}^{q}$ and $c = \frac{1}{L}$ is the step-size ($L_{\log} = \frac{1}{4}$, $L_{\text{pro}} = 1$).

Algorithm (2/2)

ii) Projection step to impose low-rankness on $\hat{\mathbf{Z}}_{(n)}^{l+1}$, using the projection B_λ onto the l_1-ball with radius λ:

$$\begin{aligned}(\tilde{\mathbf{U}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T) &\leftarrow \mathrm{svd}(\hat{\mathbf{Z}}_{(n)}^{l+1}) \\ \tilde{\mathbf{s}} &\leftarrow B_\lambda(\mathrm{diag}(\tilde{\mathbf{S}})) \\ \mathbf{Z}_{(n)}^{l+1} &\leftarrow \tilde{U} \cdot \mathrm{diag}(\tilde{\mathbf{s}}) \cdot \tilde{V}^T\end{aligned}$$

2. Optimize $\mathbf{S}_{(n)}$ by holding $\mathbf{Z}_{(n)}$ fixed:

i) Gradient step to optimize the sparse component $\mathbf{S}_{(n)}$:

$$\hat{\mathbf{S}}_{(n)}^{q+1} \leftarrow \mathbf{S}_{(n)}^{q} - c \cdot \nabla F$$

ii) Projection step to impose sparsity on $\hat{\mathbf{S}}_{(n)}^{q+1}$, using the soft-thresholding operator $H_\sigma(x) = \mathrm{sign}(x) \cdot \max(|x| - \sigma, 0)$:

$$\mathbf{S}_{(n)}^{q+1} \leftarrow H_\sigma(\hat{\mathbf{S}}_{(n)}^{q+1}).$$

Dynamic Weights

If the estimated tensor in each unfolding is $\mathscr{M}_n = \text{fold}_n(\mathbf{Z}_{(n)} + \mathbf{S}_{(n)})$, the recovered tensor is calculated as

$$\mathscr{M} \approx \sum_{n=1}^{N} a_n \cdot \mathscr{M}_n,$$

where

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

and the fitting error is given by

$$\text{fit}_n(\mathbf{Z}_{(n)} + \mathbf{S}_{(n)}) = \|\mathscr{P}_\Omega(\mathscr{Q}(\text{fold}_n(\mathbf{Z}_{(n)} + \mathbf{S}_{(n)}))) - \mathscr{Y})\|_F.$$

The dynamic weights a_n can improve the recovery quality of the recovered tensor.

Experiments

- Data: Time-series of images of the land surface temperature, acquired by the MODIS satellite over the region of Brazil.
- Size: $64 \times 64 \times 22$
  The last dimension indicates 22 days of July and August of 2019.
- We selected this region and time period because of the extensive fires that affected a large portion of the Amazonian rain-forest.
- Original images use 8 bits per pixel per band.
- The recovery performance is measured in terms of the Peak-Signal-to-Noise-Ratio (PSNR) between the original and the estimated images of the time-series.

Sparsity Parameter σ

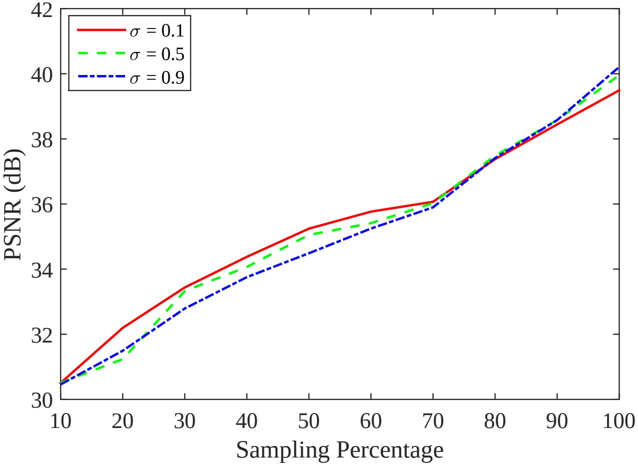

Figure: Recovery error for different sparsity values and sampling percentages, using 4 bits of quantization.

Number of Quantization Bits

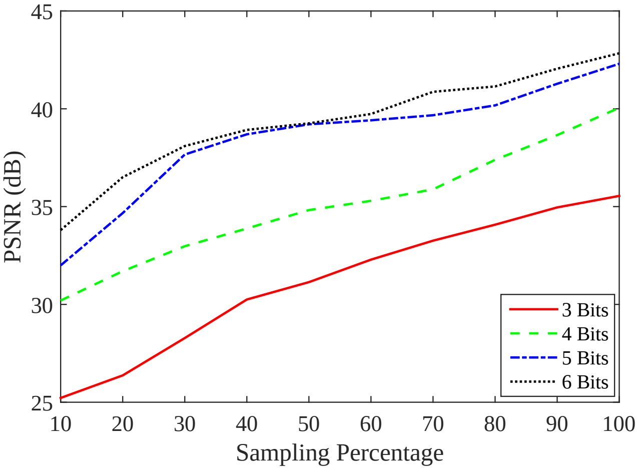

Figure: Recovery error for different sampling percentages and bits of quantization.

Tensor Unfoldings

Table: Recovery error for different sampling percentages on each mode matricization and on the weighted sum of them, using 3 bits of quantization.

PSNR	Sampling Percentage				
	10	30	50	70	100
Mode-1	24.43	26.62	28.43	33.15	35.32
Mode-2	23.38	25.96	31.33	33.17	35.35
Mode-3	25.92	28.52	30.46	32.35	34.89
Weighted sum	25.21	28.28	31.13	33.25	35.54

Anomaly Detection

$$\text{True Positive Rate} = \frac{\#\text{True predicted positions of the anomalies}}{\#\text{True positions of the anomalies}}$$

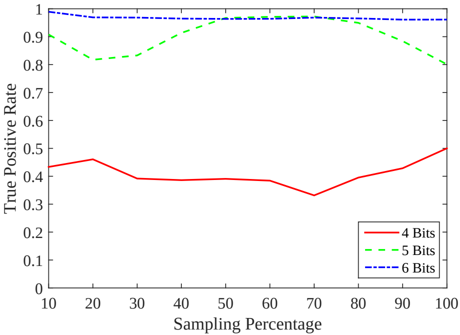

Figure: True Positive Rate for different sampling percentages and bits of quantization measured only in the sparse component and the available observations.

Recovered Images

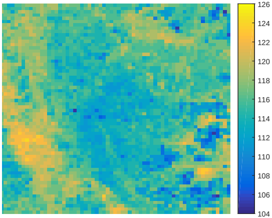

(a) Original image

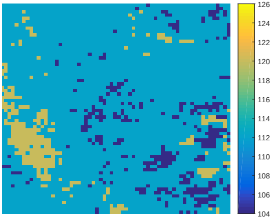

(b) Quantized image

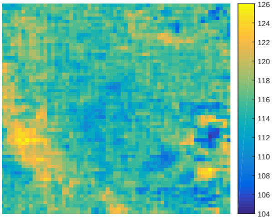

(c) Low-rank component

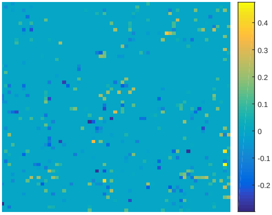

(d) Sparse component

Conclusion

- A formal approach is presented for the recovery of a tensor from a number of quantized and sparsely corrupted measurements.
- Investigation of the interaction between quantization and sampling in high-order structured data.
- Detection of sparse outliers in the data (e.g., temperature anomalies).
- Evaluation on time-series of satellite derived images of land surface temperature.

Acknowledgments

This research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) and the General Secretariat for Research and Technology (GSRT), under HFRI PhD Fellowship grant no. 1509 and HFRI Faculty grant no. 1725.