SAND2024-10878C

Sandia National Laboratories

Exceptional service in the national interest

PROBABILISTIC GUARANTEES FOR LOW-RANK TENSOR DECOMPOSITIONS

Speaker:
Carlos Llosa (Sandia)

Other Team Members:
Danny Dunlavy (Sandia)
Rich Lehoucq (Sandia)
Arvind Prasadan (Sandia)
Oscar Lopez (Florida Atlantic University)

September 2024, ML/DL

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2023-14378C

U.S. DEPARTMENT OF ENERGY

NNSA

GOAL: PROVIDE TRUST IN LOW-RANK TENSOR DECOMPOSITIONS

Low-rank tensor models/decompositions are useful in many data analysis applications

Scientific computing, cybersecurity, remote sensing, text analysis, …

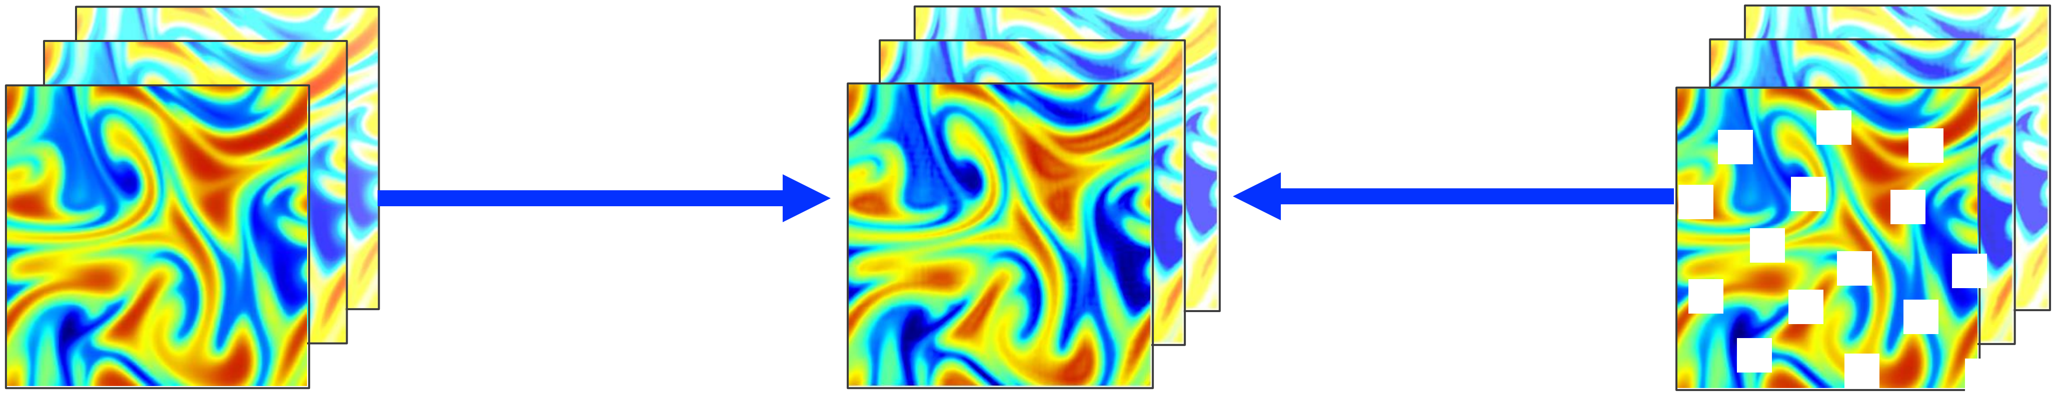

Combustion Simulation Tensor Data

Compression: ~7700X
Model Error: ϵ = 1e-1

Subset/Sample of Data

How good is a model given the data?
Cramer-Rao Bounds

How good is a model Given a sample of the data?
Sampling complexity

Current state-of-the-art:

- 😀 Efficient algorithms for several data models (Canonical Polyadic [CP], Tucker, tensor train, …)
- 😀 Some models (e.g., CP) provide useful interpretations of latent features/patterns/signals in data
- 😥 No general approaches exist for assessing how trustworthy these models are

PROJECT ACCOMPLISHMENTS TO DATE

- Zero-truncated Poisson regression for sparse multiway count data corrupted by false zeros
  - *Advance:* Understand the impact of untrusted/incomplete data in low-rank matrix and tensor modeling
  - *Benefits:* Faster, scalable low-rank decompositions with guaranteed minimal introduction of modeling error
- Spectral gap-based deterministic tensor completion
  - *Advance:* Improved error bounds on tensor completion using deterministic sampling
  - *Benefits:* Near-optimal error analysis for problems where sampling may be constrained (e.g., sensor placement problems)
- Minimax Rates in Constrained Poisson Tensors: Upper and Lower Bounds
  - *Advance:* Extension of existing low-rank matrix modeling with data constraints to tensor data
  - *Benefits:* Provides both lower (Cramer-Rao, best expected error) and upper (worst case, in expectation) error bounds
- On a Latent-Variable Formulation of the Poisson Canonical Polyadic Tensor Model
  - *Advance:* Formulation of existing low-rank matrix (NMF) and tensor (CP-APR) decompositions in Expectation Maximization framework
  - *Benefits:* Statistical error analyses, Fisher information matrix, sensitivity analyses, new model fitting algorithms

KEY RESULT: CRAMER-RAO BOUNDS FOR POISSON CP MODEL

Poisson CP model

$$\mathcal{X}_{i,j,k} \overset{\text{indep.}}{\sim} \text{Poisson}(\mathcal{M}_{i,j,k})$$

$$\sum_{r=1}^{R} A_{i,r} B_{j,r} C_{k,r}$$

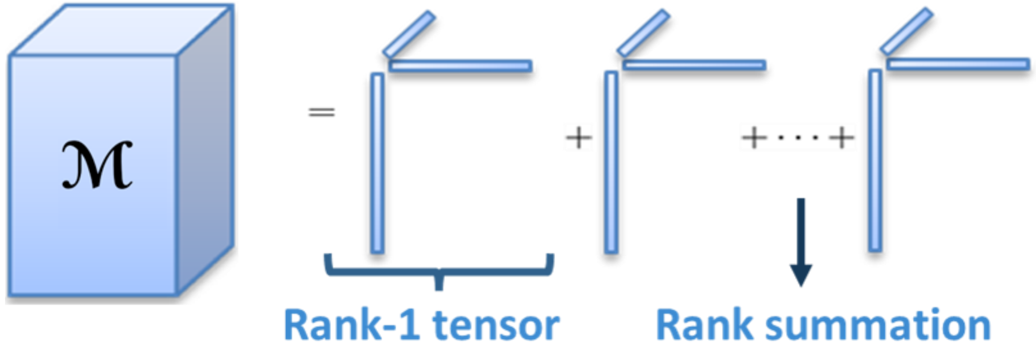

$$\boldsymbol{\theta} = [\text{vec}(A)' \text{vec}(B)' \text{vec}(C)']'$$

$$\ell(\boldsymbol{\theta}) = \sum_{i,j,k} [\mathcal{X}_{i,j,k} * \log(\mathcal{M}_{i,j,k}) - \mathcal{M}_{i,j,k}]$$

Cramer-Rao Lower Bound

$$\left(\frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E}(\hat{\boldsymbol{\theta}})\right) \mathcal{I}^{\dagger}(\boldsymbol{\theta}) \left(\frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E}(\hat{\boldsymbol{\theta}})\right)^{\top} \leq Var(\hat{\boldsymbol{\theta}})$$

Fisher Information

$$\mathbb{E}\left[-\frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \ell(\boldsymbol{x}|\boldsymbol{\theta})\right]$$

Challenges:
- Well posedness (identifiability)
- Inversion

Jacobian identity

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E}(\hat{\boldsymbol{\theta}}) = \text{Cov}\left(\hat{\boldsymbol{\theta}}, \frac{\partial}{\partial \boldsymbol{\theta}} \ell(\boldsymbol{x}|\boldsymbol{\theta})\right)$$

Challenges:
- biased estimators (MLE)

Simulation Setup

Visualize the bias-variance trade-off for tensors $\mathcal{M}(\boldsymbol{\theta}_o)$ with varying:

- Entry-wise mean $s = 1, 2, 3, 4$
  $s = \text{mean}(\mathcal{M})$
- Sizes $N = 10, 15, 20$
  $\mathcal{M} \in \mathbb{R}^{N \times N \times N}$
- Rank $R = 1, 2, \ldots, 16$

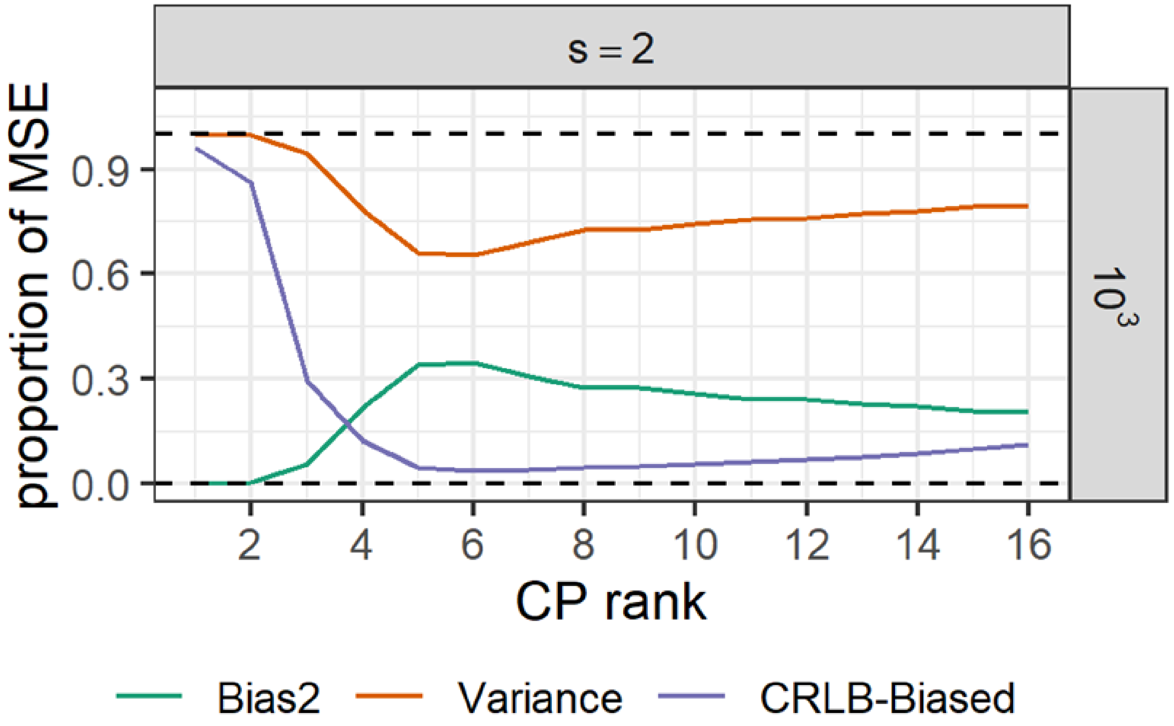

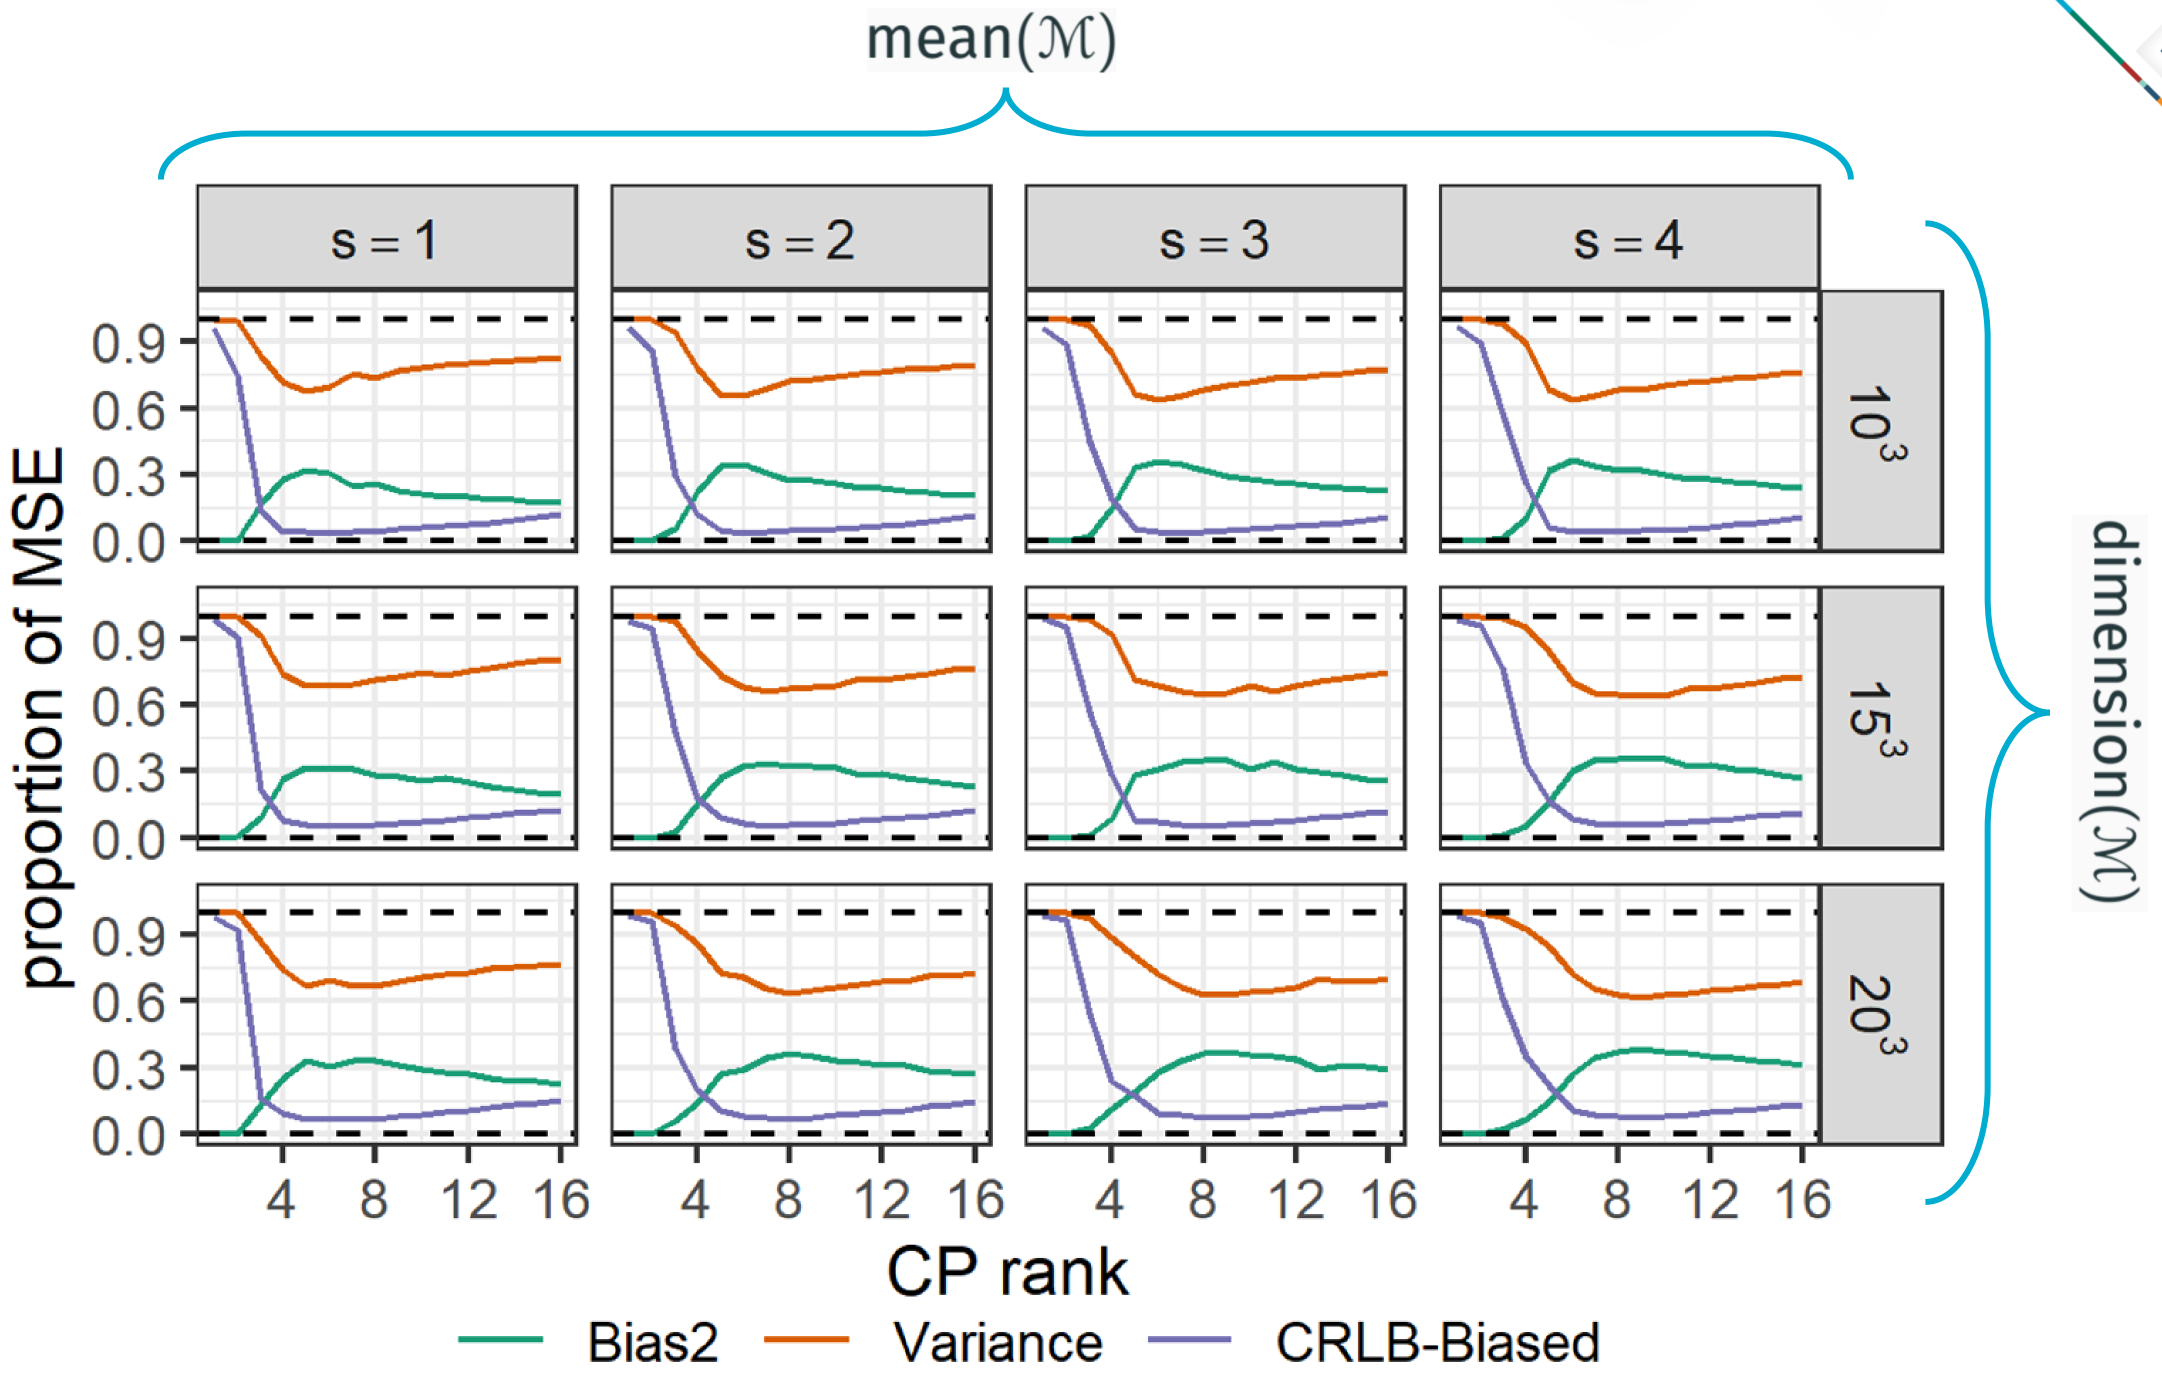
mean(ℳ)
s = 1
s = 2
s = 3
s = 4
10^3
15^3
20^3
dimension(ℳ)
proportion of MSE
0.9
0.6
0.3
0.0
4
8
12
16
CP rank
Bias2
Variance
CRLB-Biased

NEXT STEPS

- Analysis of biased estimators for (variety of) tensor models (different tensor models are used for answering a variety of questions in different applications)
  - Extend our results from CP to Tucker, Tensor Train, etc.
  - Extend our results from Poisson to GCP
- Bridging signal processing (sampling complexity) and statistics (FIM/CRLB) view of tensor modeling
- (Near) Optimal sampling complexity
- Produce software to allow others to use these new analysis techniques
  - As part of the pyttb python software

THANK YOU!

Probabilistic Guarantees for Low-Rank Tensor Decompositions

Carlos Llosa cjllosa@sandia.gov