

# Accelerating the Tucker Decomposition with Compressed Sparse Tensors

**Shaden Smith** and George Karypis

Department of Computer Science & Engineering, University of Minnesota  
{shaden, karypis}@cs.umn.edu

Euro-Par 2017

# Outline

---

Tensor Background

Computing the Tucker Decomposition

TTMc with a Compressed Sparse Tensor

Utilizing Multiple Compressed Tensors

Experiments

Conclusions

# Table of Contents

---

Tensor Background

Computing the Tucker Decomposition

TTMc with a Compressed Sparse Tensor

Utilizing Multiple Compressed Tensors

Experiments

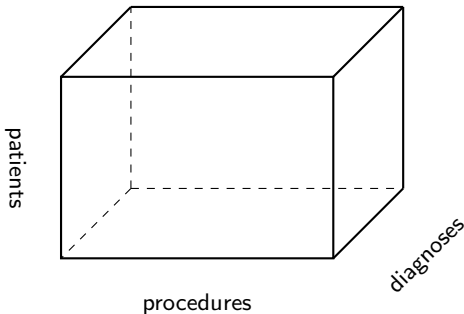
Conclusions

# Tensors

---

Tensors are the generalization of matrices to higher dimensions.

- ▶ Allow us to represent and analyze multi-dimensional data
- ▶ Applications in precision healthcare, cybersecurity, recommender systems, ...

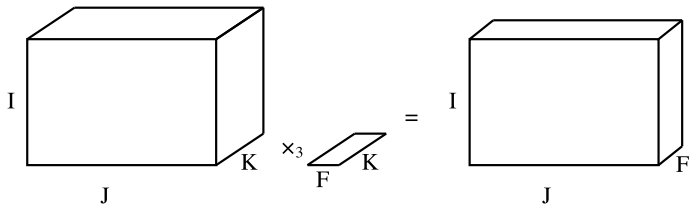


# Essential operation: tensor-matrix multiplication

---

Tensor-matrix multiplication (TTM; also called the  $n$ -way product)

- ▶ Given: tensor  $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$  and matrix  $\mathbf{M} \in \mathbb{R}^{F \times K}$ .
- ▶ Operation:  $\mathcal{X} \times_3 \mathbf{M}$
- ▶ Output:  $\mathcal{Y} \in \mathbb{R}^{I \times J \times F}$



Elementwise:

$$\mathcal{Y}(i, j, f) = \sum_{k=1}^K \mathcal{X}(i, j, k) \mathbf{M}(f, k).$$

# Chained tensor-matrix multiplication (TTMc)

---

Tensor-matrix multiplications are often performed in sequence (*chained*).

$$\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$$

## Notation

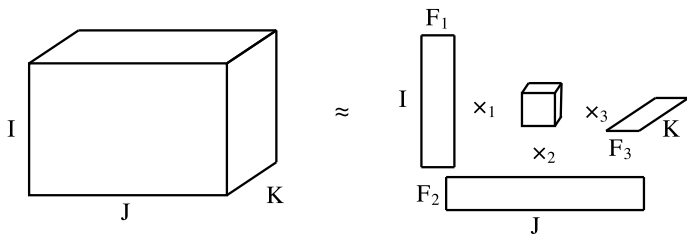
Tensors can be *unfolded* along one mode to matrix form:  $\mathbf{Y}_{(n)}$ .

- ▶ Mode  $n$  forms the rows and the remaining modes become columns.

# Tucker decomposition

---

The Tucker decomposition models a tensor  $\mathcal{X}$  as a set of orthogonal factor matrices and a core tensor.



## Notation

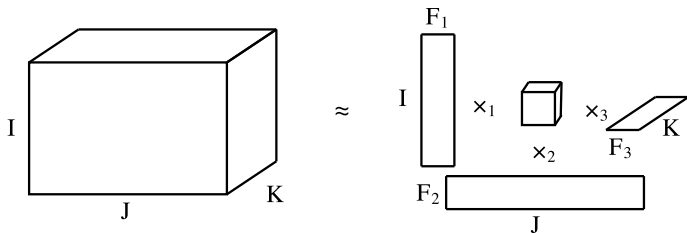
$\mathbf{A} \in \mathbb{R}^{I \times F_1}$ ,  $\mathbf{B} \in \mathbb{R}^{J \times F_2}$ , and  $\mathbf{C} \in \mathbb{R}^{K \times F_3}$  denote the factor matrices.

$\mathcal{G} \in \mathbb{R}^{F_1 \times F_2 \times F_3}$  denotes the core tensor.

# Tucker decomposition

---

The core tensor,  $\mathcal{G}$ , can be viewed as weights for the interactions between the low-rank factor matrices.



Elementwise:

$$\mathcal{X}(i, j, k) \approx \sum_{f_1=1}^{F_1} \sum_{f_2=1}^{F_2} \sum_{f_3=1}^{F_3} \mathcal{G}(f_1, f_2, f_3) \mathbf{A}(i, f_1) \mathbf{B}(j, f_2) \mathbf{C}(k, f_3)$$



# Example Tucker applications

---

## Dense: data compression

- ▶ The Tucker decomposition has long been used to compress (dense) tensor data (think truncated SVD).
- ▶ Folks at Sandia have had huge successes in compressing large simulation outputs<sup>1</sup>.

## Sparse: unstructured data analysis

- ▶ More recently, used to discover relationships in unstructured data.
- ▶ The resulting tensors are sparse and high-dimensional.
  - ▶ These large, sparse tensors are the focus of this talk.

---

<sup>1</sup>Woody Austin, Grey Ballard, and Tamara G Kolda. "Parallel tensor compression for large-scale scientific data". In: *International Parallel & Distributed Processing Symposium (IPDPS'16)*. IEEE. 2016, pp. 912–922.

## Example: dimensionality reduction for clustering

---

Factor interpretation:

- ▶ Each row of a factor matrix represents an object from the original data.
- ▶ The  $i$ th object is a point in low-dimensional space:  $\mathbf{A}(i, :)$ .
- ▶ These points can be clustered, etc.

# Example: dimensionality reduction for clustering

---

Factor interpretation:

- ▶ Each row of a factor matrix represents an object from the original data.
- ▶ The  $i$ th object is a point in low-dimensional space:  $\mathbf{A}(i, :)$ .
- ▶ These points can be clustered, etc.

## Application: citation network analysis [Kolda & Sun, ICDM '08]

- ▶ A citation network forms an *author*  $\times$  *conference*  $\times$  *keyword* sparse tensor.
- ▶ The rows of the resulting factors are clustered with  $k$ -means to reveal relationships.

**Authors:** Jiawei Han, Christos Faloutsos, ...

**Conferences:** KDD, ICDM, PAKDD, ...

**Keywords:** knowledge, learning, reasoning

# Table of Contents

---

Tensor Background

Computing the Tucker Decomposition

TTMc with a Compressed Sparse Tensor

Utilizing Multiple Compressed Tensors

Experiments

Conclusions

# Optimization problem

---

The resulting optimization problem is non-convex:

$$\begin{aligned} & \underset{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathcal{G}}{\text{minimize}} && \frac{1}{2} \left\| \boldsymbol{\mathcal{X}} - \left( \mathcal{G} \times_1 \mathbf{A}^T \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T \right) \right\|_F^2 \\ & \text{subject to} && \mathbf{A}^T \mathbf{A} = \mathbf{I} \\ & && \mathbf{B}^T \mathbf{B} = \mathbf{I} \\ & && \mathbf{C}^T \mathbf{C} = \mathbf{I} \end{aligned}$$

# Higher-Order Orthogonal Iterations (HOOI)

---

HOOI is an alternating optimization algorithm.

---

## Tucker Decomposition with HOOI

---

- 1: **while** not converged **do**
  - 2:      $\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$
  - 3:      $\mathbf{A} \leftarrow F_1$  leading left singular vectors of  $\mathbf{Y}_{(1)}$
  - 4:
  - 5:      $\mathcal{Y}_2 \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_3 \mathbf{C}^T$
  - 6:      $\mathbf{B} \leftarrow F_2$  leading left singular vectors of  $\mathbf{Y}_{(2)}$
  - 7:
  - 8:      $\mathcal{Y}_3 \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_2 \mathbf{B}^T$
  - 9:      $\mathbf{C} \leftarrow F_3$  leading left singular vectors of  $\mathbf{Y}_{(3)}$
  - 10:
  - 11:      $\mathcal{G} \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$
  - 12: **end while**
-

# Higher-Order Orthogonal Iterations (HOOI)

---

TTMc is the most expensive kernel in the HOOI algorithm.

---

## Tucker Decomposition with HOOI

---

- 1: **while** not converged **do**
  - 2:      $\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$
  - 3:      $\mathbf{A} \leftarrow F_1$  leading left singular vectors of  $\mathbf{Y}_{(1)}$
  - 4:
  - 5:      $\mathcal{Y}_2 \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_3 \mathbf{C}^T$
  - 6:      $\mathbf{B} \leftarrow F_2$  leading left singular vectors of  $\mathbf{Y}_{(2)}$
  - 7:
  - 8:      $\mathcal{Y}_3 \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_2 \mathbf{B}^T$
  - 9:      $\mathbf{C} \leftarrow F_3$  leading left singular vectors of  $\mathbf{Y}_{(3)}$
  - 10:
  - 11:      $\mathcal{G} \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$
  - 12: **end while**
-

# Intermediate memory blowup

---

A first step is to optimize a single TTM kernel and apply in sequence:

$$\mathcal{Y}_1 \leftarrow \left( \left( \mathcal{X} \times_2 \mathbf{B}^T \right) \times_3 \mathbf{C}^T \right)$$

Challenge:

- ▶ Intermediate results become more **dense** after each TTM.
- ▶ **Memory overheads** are dependent on sparsity pattern and factorization rank, but can be **several orders of magnitude**.

---

Tamara Kolda and Jimeng Sun. "Scalable tensor decompositions for multi-aspect data mining". In: *International Conference on Data Mining (ICDM)*. 2008.



# Intermediate memory blowup

---

$$\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$$

Solutions:

---

<sup>2</sup>Tamara Kolda and Jimeng Sun. “Scalable tensor decompositions for multi-aspect data mining”. In: *International Conference on Data Mining (ICDM)*. 2008.

<sup>3</sup>Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. Tech. rep. RR-8801. Inria-Research Centre Grenoble–Rhône-Alpes, 2015.

# Intermediate memory blowup

---

$$\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$$

Solutions:

1. Tile over  $\mathcal{Y}_1$  to constrain blowup<sup>2</sup>.
  - ▶ Requires multiple passes over the input tensor and many FLOPs.

---

<sup>2</sup>Tamara Kolda and Jimeng Sun. “Scalable tensor decompositions for multi-aspect data mining”. In: *International Conference on Data Mining (ICDM)*. 2008.

<sup>3</sup>Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. Tech. rep. RR-8801. Inria-Research Centre Grenoble–Rhône-Alpes, 2015.

# Intermediate memory blowup

---

$$\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T$$

Solutions:

1. Tile over  $\mathcal{Y}_1$  to constrain blowup<sup>2</sup>.
  - ▶ Requires multiple passes over the input tensor and many FLOPs.
2. Instead, fuse the TTMs and use a formulation based on non-zeros<sup>3</sup>.
  - ▶ Only a single pass over the tensor!

---

<sup>2</sup>Tamara Kolda and Jimeng Sun. “Scalable tensor decompositions for multi-aspect data mining”. In: *International Conference on Data Mining (ICDM)*. 2008.

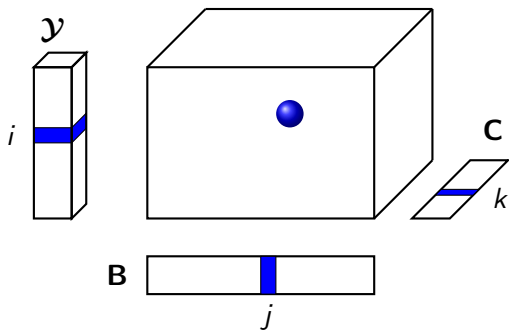
<sup>3</sup>Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. Tech. rep. RR-8801. Inria-Research Centre Grenoble–Rhône-Alpes, 2015.

# Elementwise formulation

---

Processing each non-zero individually has cost  $\mathcal{O}(\text{nnz}(\mathcal{X})F_2F_3)$  and  $\mathcal{O}(F_2F_3)$  memory overhead.

$$\mathcal{Y}_1(i, :, :) += \mathcal{X}(i, j, k) [\mathbf{B}(j, :) \circ \mathbf{C}(k, :)]$$

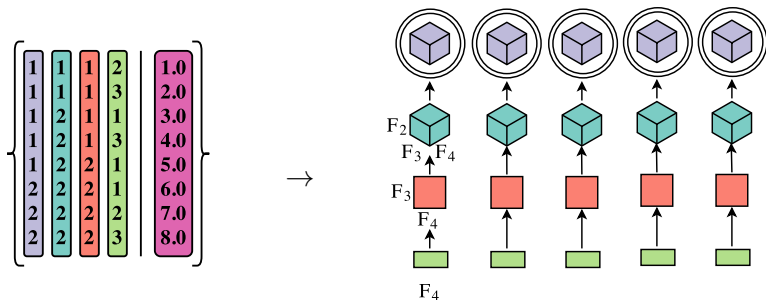


---

Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. Tech. rep. RR-8801. Inria-Research Centre Grenoble–Rhône-Alpes, 2015.

# TTMc with coordinate form

The elementwise formulation of TTMc naturally lends itself to a coordinate storage format:



# Memoization

---

Some of the intermediate results across TTMc kernels can be reused:

$$\mathcal{Y}_1 \leftarrow \mathcal{X} \times_2 \mathbf{B}^T \times_3 \mathbf{C}^T \times_4 \mathbf{D}^T$$

$$\mathcal{Y}_2 \leftarrow \mathcal{X} \times_1 \mathbf{A}^T \times_3 \mathbf{C}^T \times_4 \mathbf{D}^T$$

becomes:

$$\mathcal{Z} \leftarrow \mathcal{X} \times_3 \mathbf{C}^T \times_4 \mathbf{D}^T$$

$$\mathcal{Y}_1 \leftarrow \mathcal{Z} \times_2 \mathbf{B}^T$$

$$\mathcal{Y}_2 \leftarrow \mathcal{Z} \times_1 \mathbf{A}^T$$

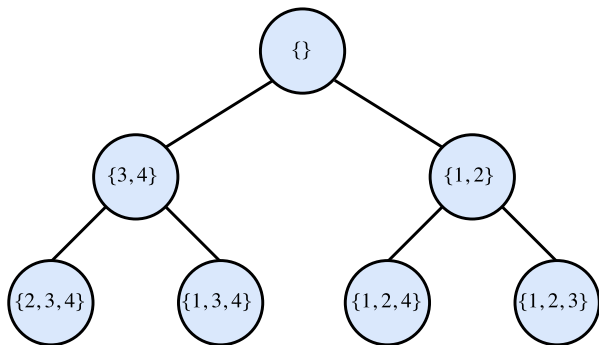
---

Muthu Baskaran et al. "Efficient and scalable computations with sparse tensors". In: *High Performance Extreme Computing (HPEC)*. 2012.

# TTMc with dimension trees

---

State-of-the-art TTMc:



Each node in the tree stores intermediate results from a set of modes.

---

Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. Tech. rep. RR-8801. Inria-Research Centre Grenoble–Rhône-Alpes, 2015.

# TTMc with dimension trees

---

Parallelism:

- ▶ Independent units of work within each node are identified.
- ▶ For flat dimension trees, this equates to parallelizing over  $\mathcal{Y}_1(i, :, :)$  slices.

---

Oguz Kaya and Bora Uçar. *High-performance parallel algorithms for the Tucker decomposition of higher order sparse tensors*. Tech. rep. RR-8801. Inria-Research Centre Grenoble–Rhône-Alpes, 2015.



# Table of Contents

---

Tensor Background

Computing the Tucker Decomposition

**TTMc with a Compressed Sparse Tensor**

Utilizing Multiple Compressed Tensors

Experiments

Conclusions

# Motivation

---

Existing algorithms either:

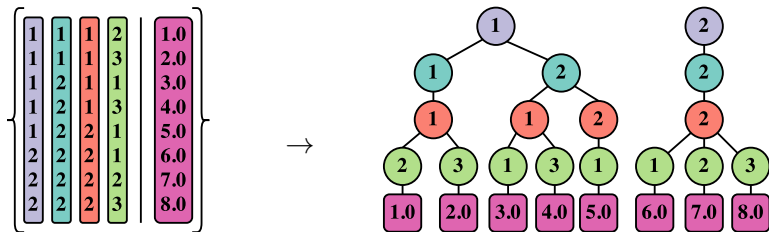
- ▶ have intermediate data blowup
- ▶ perform many operations
- ▶ trade memory for performance (i.e., memoization)
  - ▶ Overheads depend on the sparsity pattern and factorization rank

Can we accelerate TTMC without memory overheads?

# Compressed Sparse Fiber (CSF)

CSF encodes a sparse tensor as a forest.

- ▶ Each path from root to leaf encodes a non-zero.
- ▶ CSF can be viewed as a generalization of CSR.



Shaden Smith and George Karypis. "Tensor-Matrix Products with a Compressed Sparse Tensor". In: *5th Workshop on Irregular Applications: Architectures and Algorithms*. 2015.

# Arithmetic redundancies in TTMc

---

Going back to the non-zero formulation:

$$\mathcal{Y}_1(i, :, :) += \mathcal{X}(i, j, k) [\mathbf{B}(j, :) \circ \mathbf{C}(k, :)]$$

There are two arithmetic redundancies we can exploit:

1. Distributive outer products
2. Redundant outer products

# Distributive outer products

---

Consider two non-zeros in the same fiber  $\mathcal{X}(i, j, :)$

$$\mathcal{Y}_1(i, :, :) += \mathcal{X}(i, j, k_1) [\mathbf{B}(j, :) \circ \mathbf{C}(k_1, :)]$$

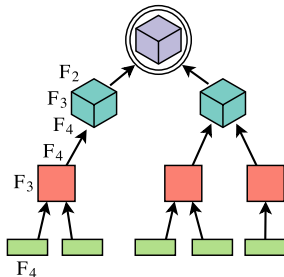
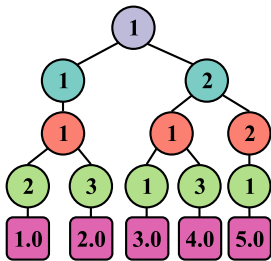
$$\mathcal{Y}_1(i, :, :) += \mathcal{X}(i, j, k_2) [\mathbf{B}(j, :) \circ \mathbf{C}(k_2, :)]$$

We can factor out  $\mathbf{B}(j, :)$

$$\mathcal{Y}_1(i, :, :) += \mathbf{B}(j, :) \circ [\mathcal{X}(i, j, k_1)\mathbf{C}(k_1, :) + \mathcal{X}(i, j, k_2)\mathbf{C}(k_2, :)]$$

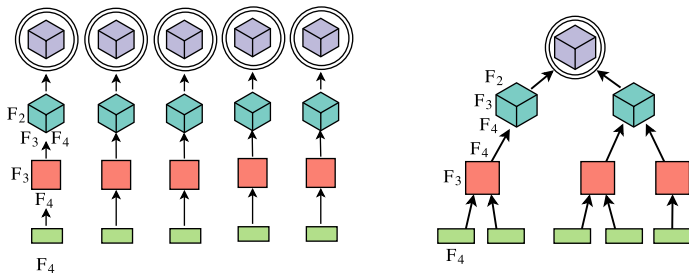
# Distributive outer products with CSF

- ▶ Children nodes are summed and then expanded with an outer product.
- ▶ Intermediate memory is kept negligible with a post-order depth-first traversal.
  - ▶ We only need to keep a single stack of intermediate results.



# Distributive outer products with CSF

Compare to computing with coordinate format:



**Savings:** The cost of each non-zero (leaf) is now linear in the rank.

## Redundant outer products

---

Suppose we're now computing  $\mathcal{Y}_3$ :

$$\mathcal{Y}_3(:, :, k_1) += \mathcal{X}(i, j, k_1) [\mathbf{A}(i, :) \circ \mathbf{B}(j, :)],$$

$$\mathcal{Y}_3(:, :, k_2) += \mathcal{X}(i, j, k_2) [\mathbf{A}(i, :) \circ \mathbf{B}(j, :)].$$

The outer product can be reused:

$$\mathbf{S} \leftarrow \mathbf{A}(i, :) \circ \mathbf{B}(j, :)$$

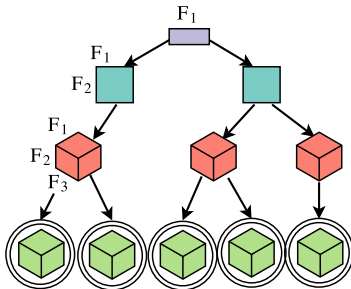
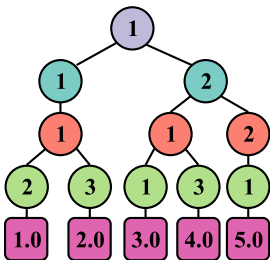
$$\mathcal{Y}_3(:, :, k_1) += \mathcal{X}(i, j, k_1) \cdot \mathbf{S}$$

$$\mathcal{Y}_3(:, :, k_2) += \mathcal{X}(i, j, k_2) \cdot \mathbf{S}$$



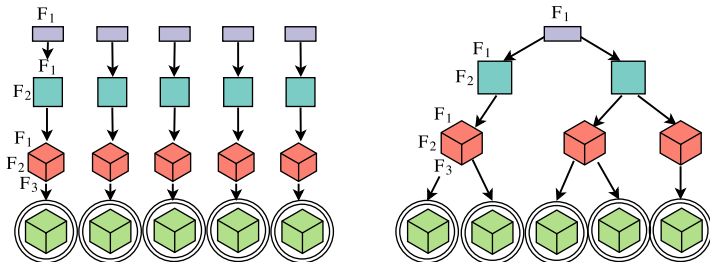
# Redundant outer products with CSF

- ▶ Each child node uses its parent in an outer product.
- ▶ Use a pre-order depth-first traversal to manage memory.



# Redundant outer products with CSF

Compare to computing with coordinate format:



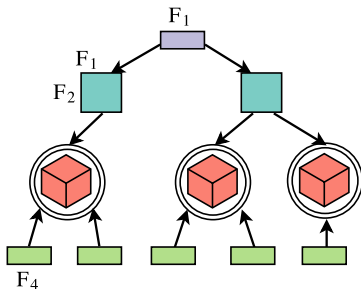
**Savings:** outer products are constructed less often, but non-zeros still have the same asymptotic cost as coordinate form.

# Putting it all together

---

These two optimizations can be combined within the same kernel.

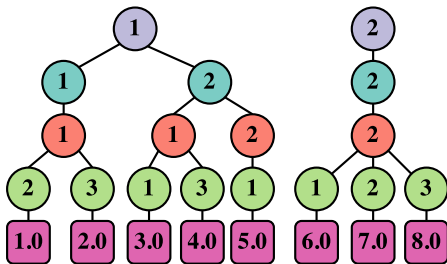
- ▶ Traversal is still depth-first
  - ▶ Use a pre-order traversal for levels above the mode of interest.
  - ▶ Use a post-order traversal for levels below the mode of interest.



# Parallelism with CSF

---

We distribute trees to threads and use dynamic load balancing.



Race conditions are dependent on the mode of interest:

- ▶ Root nodes are unique, so no race conditions
- ▶ Otherwise, use a mutex to lock the slice of  $\mathcal{Y}$

# Table of Contents

---

Tensor Background

Computing the Tucker Decomposition

TTMc with a Compressed Sparse Tensor

**Utilizing Multiple Compressed Tensors**

Experiments

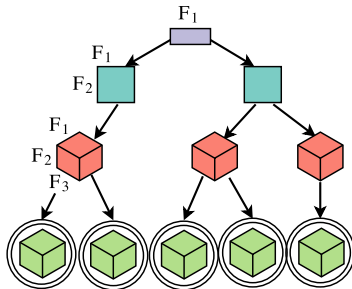
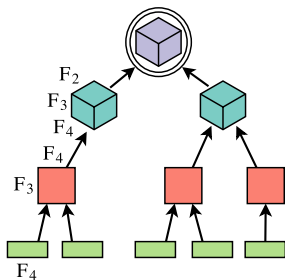
Conclusions

# Motivation

---

TTMc is significantly more expensive when computing for the lower-level modes.

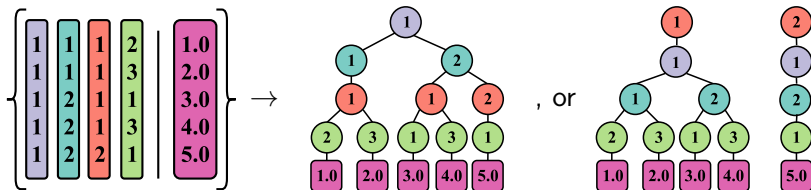
- ▶ This is due to FLOPs and synchronization costs.



# Multiple CSF representations

We can reorder the modes of  $\mathcal{X}$  and store additional copies of the tensor.

- ▶ Selectively place modes near the top which were previously expensive.



# CSF selection

---

Given storage for  $K$  copies of the tensor, how do we select them from the  $N!$  mode possible orderings?

Greedy, heuristic algorithm:

- ▶ The first CSF always sorts the modes based on their lengths.
- ▶ For each of the  $K - 1$  remaining CSF representations:
  - ▶ Select the mode which is estimated to be the most expensive based on FLOPs.
  - ▶ Place that mode at the top of the next CSF, with the remaining modes sorted by length.
  - ▶ Examine the new CSF and update the new best cost estimates.



# Table of Contents

---

Tensor Background

Computing the Tucker Decomposition

TTMc with a Compressed Sparse Tensor

Utilizing Multiple Compressed Tensors

**Experiments**

Conclusions

# Experimental Setup

---

Source code:

- ▶ Part of the **S**urprisingly **P**arallel **s**pArse **T**ensor **T**oolkit<sup>4</sup>
- ▶ Written in C and parallelized with OpenMP
- ▶ Compiled with `icc v16.0.3` and linked with Intel MKL

HyperTensor

- ▶ Implements dimension tree-based methods
- ▶ Written in C++ and parallelized with OpenMP

Machine specifications:

- ▶ 2x 12-core Intel E5-2680v3 processors (Haswell)
- ▶ Double-precision floats and 32-bit integers

---

<sup>4</sup>SPLATT: <https://github.com/ShadenSmith/splatt>

# Experimental Setup

---

## Experiments

- ▶ All measurements are for a sequence of TTMC kernels forming one iteration.
- ▶ We fix  $F_1 = F_2 = \dots = 20$ .

# Datasets

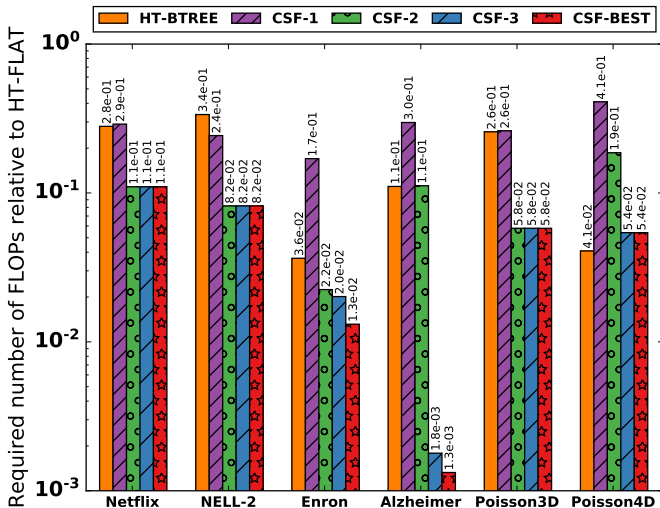
---

<b>Dataset</b>	<b>Non-zeros</b>	<b>Modes</b>	<b>Dimensions</b>
NELL-2	77M	3	12K, 9K, 29K
Netflix	100M	3	480K, 18K, 2K
Enron	54M	4	6K, 6K, 244K, 1K
Alzheimer	6.27M	5	5, 1K, 156, 1K, 396
Poisson3D, Poisson4D	100M	3,4	3K, ..., 3K

**K** and **M** stand for thousand and million, respectively.

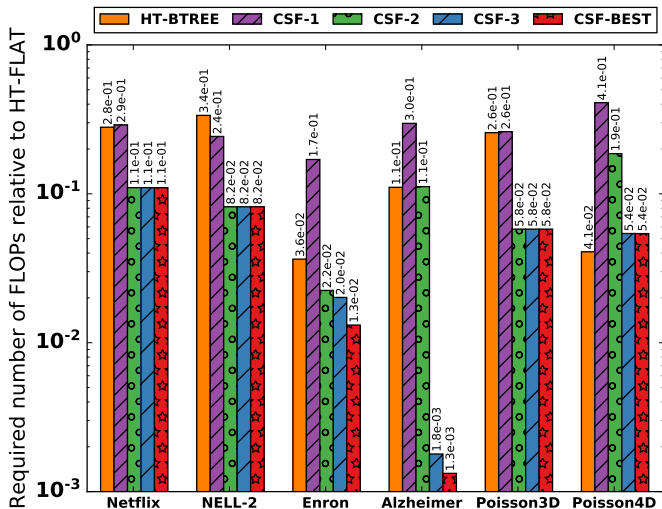
# Relative FLOP cost of TTMc

The greedy algorithm usually matches or gets close to the best possible CSF configuration.



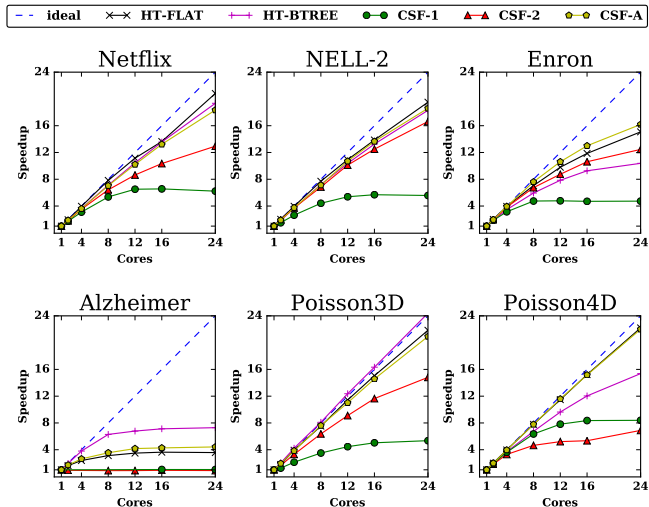
# Relative FLOP cost of TTMc

CSF benefits more as dimensionality increases.



# Parallel Scalability

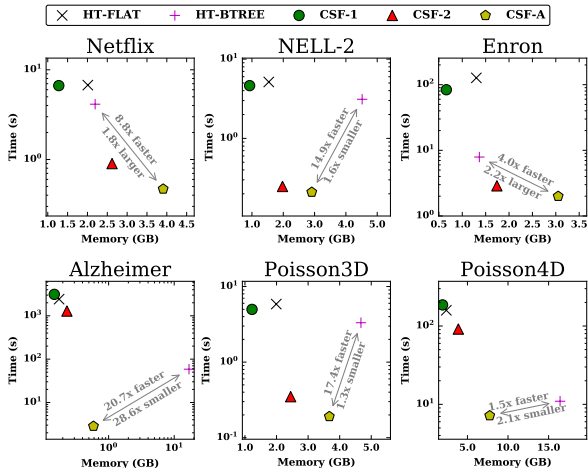
Adding CSF representations improves scalability due to fewer and smaller critical regions.



# Performance tradeoffs

Selecting the number of CSFs provides tuning for memory vs. speed.

- ▶ CSF always provides the options for the smallest and fastest executions.





# Table of Contents

---

Tensor Background

Computing the Tucker Decomposition

TTMc with a Compressed Sparse Tensor

Utilizing Multiple Compressed Tensors

Experiments

Conclusions

# Wrapping up

---

## Contributions:

- ▶ We optimized TTMc kernels via a compressed tensor representation
  - ▶ CSF naturally exposes arithmetic redundancies in TTMc
- ▶ Multiple CSF tensors can further accelerate computation
  - ▶ Up to  $20\times$  speedup over the state-of-the-art while using  $28\times$  less memory
  - ▶ Choosing the number of data copies offers *tunable* computation/memory tradeoff

# Wrapping up

---

## Contributions:

- ▶ We optimized TTMc kernels via a compressed tensor representation
  - ▶ CSF naturally exposes arithmetic redundancies in TTMc
- ▶ Multiple CSF tensors can further accelerate computation
  - ▶ Up to  $20\times$  speedup over the state-of-the-art while using  $28\times$  less memory
  - ▶ Choosing the number of data copies offers *tunable* computation/memory tradeoff

## Future work:

- ▶ Alternative decompositions to reduce synchronization costs
- ▶ Memoization is also applicable to CSF formulation!
  - ▶ *Li et al., IPDPS '17*

# Reproducibility

---

All of our work is open source (to be updated soon):

`https://github.com/ShadenSmith/splatt`

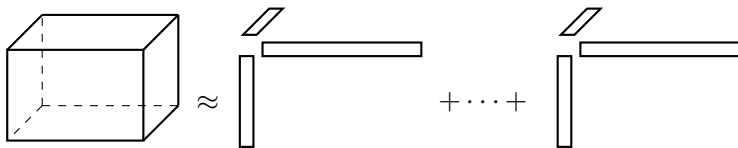
Datasets are freely available:

`http://frostdt.io/`

# Backup

# Canonical polyadic decomposition (CPD)

The CPD models a tensor as the summation of rank-1 tensors.



$$\underset{\mathbf{A}, \mathbf{B}, \mathbf{C}}{\text{minimize}} \quad \mathcal{L}(\mathcal{X}, \mathbf{A}, \mathbf{B}, \mathbf{C}) = \left\| \mathcal{X} - \sum_{f=1}^F \mathbf{A}(:, f) \circ \mathbf{B}(:, f) \circ \mathbf{C}(:, f) \right\|_F^2$$

## Notation

$\mathbf{A} \in \mathbb{R}^{I \times F}$ ,  $\mathbf{B} \in \mathbb{R}^{J \times F}$ , and  $\mathbf{C} \in \mathbb{R}^{K \times F}$  denote the factor matrices for a 3D tensor.

# TTMc with a CSF tensor (1)

---

## Algorithm 1 TTMc()

---

```
1: function TTMc( $\mathcal{X}$ , mode)
2:   for  $i_1 = 1, \dots, l_N$  in parallel do
3:     CONSTRUCT( $\mathcal{X}(i_1, :, \dots, :)$ , mode, 1)
4:   end for
5: end function
```

---

# TTMc with a CSF tensor (2)

---

## Algorithm 2 CONSTRUCT()

---

```
1:           ▷ Construct Kronecker products and push them down to level mode−1.
2: function CONSTRUCT(node, mode, above)
3:    $d \leftarrow \text{level}(\textit{node})$            ▷ The level in the tree (i.e., distance from the root).
4:    $i_d \leftarrow \text{node\_id}(\textit{node})$        ▷ The partial coordinate of a non-zero.
5:
6:   if  $d < \textit{mode}$  then
7:      $\textit{above} \leftarrow \textit{above} \otimes \mathbf{A}^{(d)}(i_d, :)$ 
8:     for  $c \in \text{children}(\textit{node})$  do
9:       CONSTRUCT( $c$ , mode, above)
10:    end for
11:
12:   else if  $d = \textit{mode}$  then
13:      $\textit{below} \leftarrow \sum_{c \in \text{children}(\textit{node})} \text{ACCUMULATE}(c)$ 
14:     Lock mutex  $i_d$ .
15:      $\mathbf{Y}_{(d)}(i_d, :) \leftarrow \mathbf{Y}_{(d)}(i_d, :) + (\textit{above} \otimes \textit{below})$            ▷ Update  $\mathbf{Y}_{(d)}$ .
16:     Unlock mutex  $i_d$ .
17:   end if
18: end function
```

---



# TTMc with a CSF tensor (3)

---

## Algorithm 3 ACCUMULATE()

---

```
1:                                     ▷ Pull Kronecker products up from the leaf nodes.
2: function ACCUMULATE(node)
3:    $i_d \leftarrow \text{node\_id}(\textit{node})$ 
4:   if level(node) =  $N$  then
5:     return  $\mathcal{X}(i_1, \dots, i_d) \cdot \mathbf{A}^{(N)}(i_d, :)$ 
6:   else
7:     return  $\mathbf{A}^{(d)}(i_d, :) \otimes \sum_{c \in \text{children}(\textit{node})} \text{ACCUMULATE}(c)$ 
8:   end if
9: end function
```

---