CS 598 EVS: Tensor Computations

Tensor Decomposition

Edgar Solomonik

University of Illinois at Urbana-Champaign

CP Decomposition Rank

- The *canonical polyadic or CANDECOMP/PARAFAC (CP) decomposition* expresses an order d tensor in terms of d factor matrices

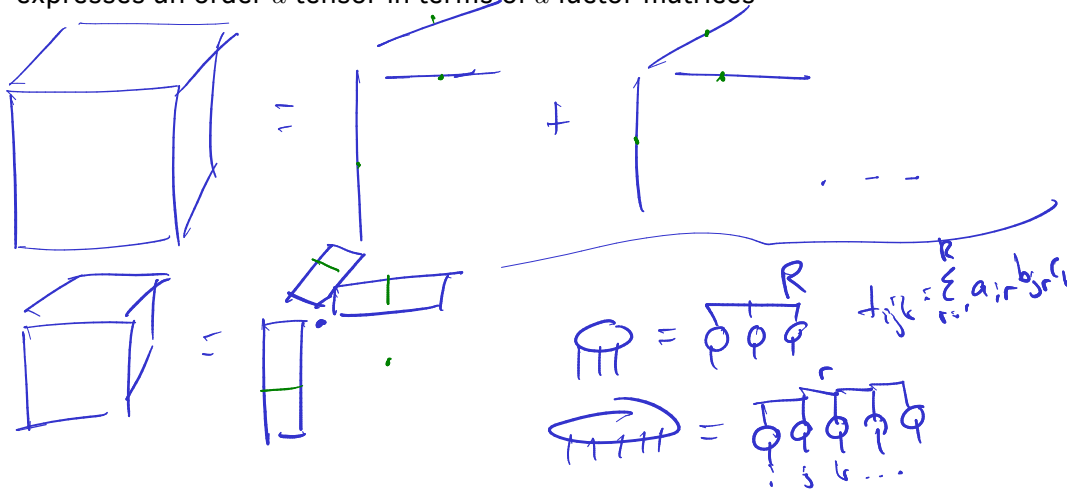

CP – CANDECOMP / PARAFAC

- canonical polyadic

1927 Hitchcock

Fogo

$n \times n \times n$

Applications

- bilinear algorithms (exact CP, $R > n$)
- data analysis (high-order clustering)
  - modelling (CP completion)
  - performance of applications for different parameters
- low-order methods (compression)
  - quantum chemistry

Tucker

# of processors

block size

mol. dir.

Tensor Rank Properties

- Tensor rank does not satisfy many of the properties of matrix rank

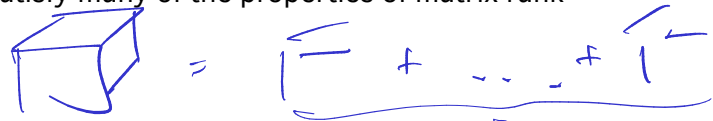

$$R \le \lceil (\hat{d})/n \rceil$$

generic rank over $\mathbb{C}$

$$T = \left[\begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} -1 & 1 \end{bmatrix}\right]$$

rank over $\mathbb{R}$ is 3

rank over $\mathbb{C}$ is 2

$$T = [\![A, B, C]\!] = \left[\begin{bmatrix} \cdot & \cdot \\ \cdot & \cdot \end{bmatrix}, \begin{bmatrix} \cdot & \cdot \\ \cdot & \cdot \end{bmatrix}, \begin{bmatrix} \cdot & \cdot \\ \cdot & \cdot \end{bmatrix}\right]$$

multiple typical ranks over $\mathbb{R}$

- e.g. $n=2$, $2\times 2\times 2$ tensors
  - over $\mathbb{R}$, rank is $\underline{2}$ with 79% prob
    $\underline{3}$ with 21% prob

- for $n=2$, $d=3$, there are 2 typical ranks
- if there is only one typical rank, it is the <u>generic</u> rank
- over $\mathbb{C}$, for all n, for symmetric tensors, there is a generic rank

$$R = \left\lceil \frac{\binom{n+d-1}{d}}{n} \right\rceil$$

Typical Rank and Generic Rank

- When there is only a single typical tensor rank, it is the *generic rank*

- for $n_1 \times n_2 \times n_3$ tensor

$$R \leq \min(n_1 n_2,\ n_2 n_3,\ n_1 n_3)$$

Uniqueness Sufficient Conditions

- Unlike the low-rank matrix case, the CP decomposition can be unique

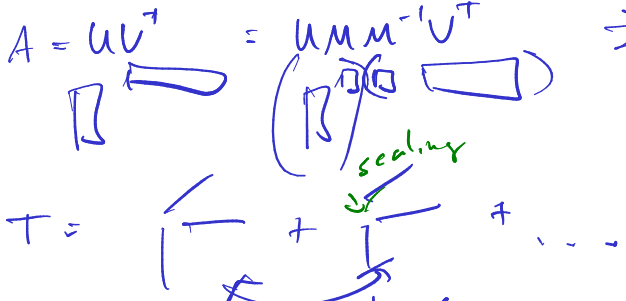

modulo

Sufficient condition for uniqueness (scaling, permutation) is

$k_A + k_B + k_C \geq 2R + 2$ $\quad k_A, k_B, k_C \leq n$ $\quad R \leq \frac{3n-2}{2}$

k_X Kruskal-rank of X, largest k_X such that any subset of k_X columns of X is linearly independent

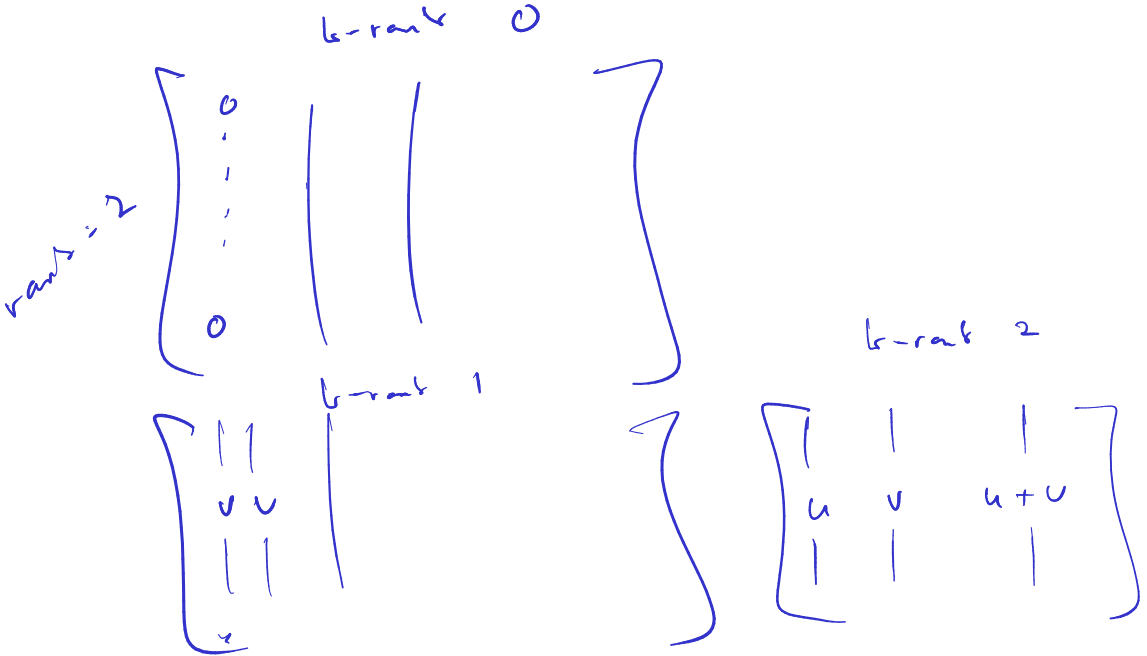
b-rank 0
rank = 2
0
0
b-rank 1
v
v
b-rank 2
u
v
u + v

Uniqueness Necessary Conditions

- Necessary conditions for uniqueness of the CP decomposition also exist

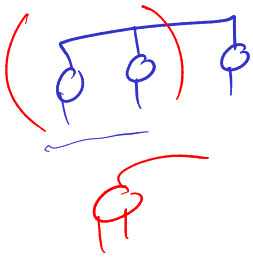

$T = [\![U, V, W]\!]$

$\text{rank}(U \odot V) = R$

$\text{rank}(U \odot W)$

$\text{rank}(V \odot W)$

$\text{rank}(U)\,\text{rank}(V) \geq R$

full rank

$\text{rank}(U \odot V) \leq \text{rank}(U \otimes V)$

$= \text{rank}(U)\,\text{rank}(V)$

Degeneracy

- The best rank-k approximation may not exist, a problem known as *degeneracy* of a tensor

$$T = a_1 \circ b_1 \circ c_2 + a_1 \circ b_2 \circ c_1 + a_2 \circ b_1 \circ c_1$$

CP rank of $T = 3$

$$T \approx \alpha\left(a_1 + \frac{a_2}{\alpha}\right) \circ \alpha\left(b_1 + \frac{b_2}{\alpha}\right) \circ \alpha\left(c_1 + \frac{c_2}{\alpha}\right) - \alpha\, a_1 \circ b_1 \circ c_1$$

$$= T + O\left(\frac{1}{\alpha}\right) = \tilde{T}_\alpha$$

$$\lim_{\alpha \to \infty} \| T - \tilde{T}_\alpha \| = 0$$

Border Rank

- Degeneracy motivates an approximate notion of rank, namely *border rank*

- border rank R implies we can find a sequence T_λ

$$\text{CP-rank}(T_\lambda) = R$$

$$\lim_{\lambda \to 0} \|T - T_\lambda\| = 0$$

- for bilinear problem tensor T representing 3×3 times 3×3 mat.-mul.

$9\times9\times9$

  - border rank is in $[14, 21]$
  - rank is in $[19, 23]$

Approximation by CP Decomposition

- Approximation via CP decomposition is a nonlinear optimization problem

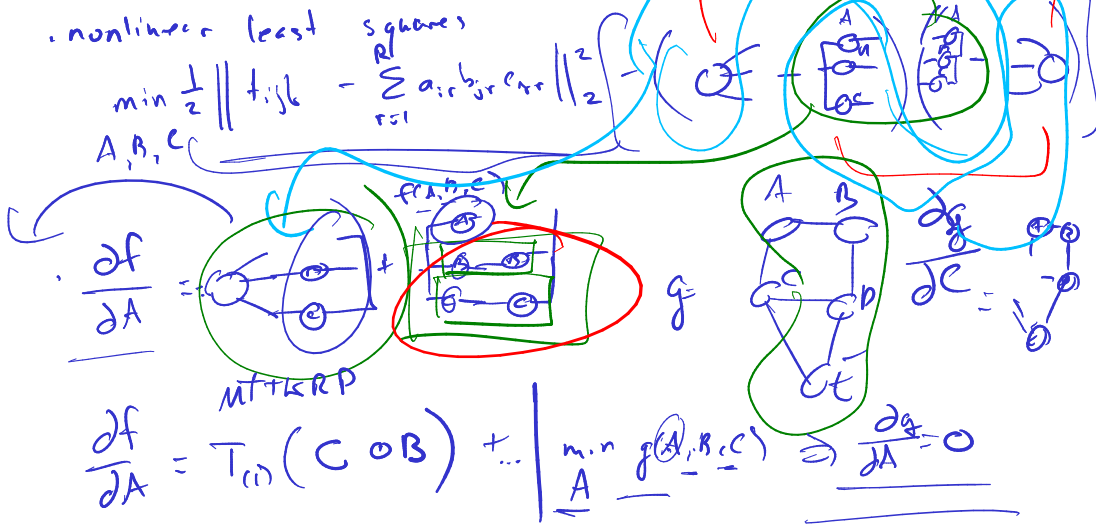

Alternating Least Squares Algorithm

- The standard approach for finding an approximate or exact CP decomposition of a tensor is the *alternating least squares (ALS) algorithm*

$= AB^T$

$BA^T \simeq T_{(1)}$

$A^T = (B^T B)^{-1} B^T T_{(1)}$

MTTKRP

ALS

linear

alternating least squares

tensor-network product

defines equations

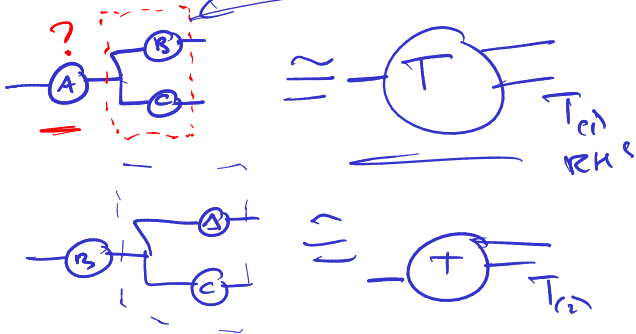

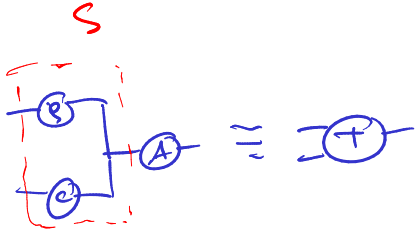

$S^T S\ A = S^T T_{cin}$

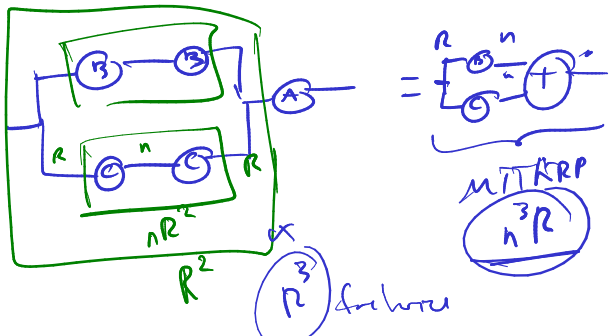

$R^2 n$

Alternating Least Squares for Tucker Decomposition

- For Tucker decomposition, an analogous optimization procedure to ALS is referred to as *high-order orthogonal iteration (HOOI)*

Dimension Trees for ALS

- The cost of ALS can be reduced by amortizing computation common terms

Fast Residual Norm Calculation

- Calculating the norm of the residual has cost $2ds^dR$, but can be done more cheaply within ALS

Pairwise Perturbation Algorithm

- A route to further reducing the cost of ALS is to perform it approximately via *pairwise perturbation*

Pairwise Perturbation Second Order Correction

- When approximating a tensor using CP, the partially converged CP factors can sometimes be used in place of the tensor to accelerate cost

Approximate CP ALS using Random Sampling

- Another approach to approximating ALS is to sample the least-squares equations[1]

[1] C. Battaglino, G. Ballard, T. G. Kolda, 2018

Gauss-Newton Algorithm

- ALS generally achieves linear convergence, while Newton-based methods can converge quadratically

Gauss-Newton for CP Decomposition

- CP decomposition for order $d = 3$ tensors ($d > 3$ is similar) minimizes

Gauss-Newton for CP Decomposition

- A step of Gauss-Newton requires solving a linear system with $\boldsymbol{H}$

```
u = []
for q in range(d):
  u.append(zeros((n,R)))
  for p in range(d):
    if q == p:
      u[q] += einsum("rz,kz->kr",G[q,p],v[p])
    else:
      u[q] += einsum("kz,lr,rz,lz->kr", \
                     U[q],U[p],G[q,p],v[p])
```

Tensor Completion

- The *tensor completion* problem seeks to build a model (e.g., CP decomposition) for a partially-observed tensor

- The problem was partially popularized by the Netflix prize collaborative filtering problem

CP Tensor Completion Gradient and Hessian

- The gradient of the tensor completion objective function is sparsified according to the set of observed entries

- ALS for tensor decomposition solves quadratic optimization problem for each row of each factor matrix, in the completion case, Newton's method on these subproblems yields different Hessians

Methods for CP Tensor Completion

- ALS for tensor completion with CP decomposition incurs additional cost

- Alternative methods for tensor completion include coordinate descent and stochastic gradient descent

Coordinate Descent for CP Tensor Completion

- Coordinate descent avoids the need to solve linear systems of equations

Sparse Tensor Contractions

- Tensor completion and sparse tensor decomposition require operations on sparse tensors

- Sparse tensor contractions often correspond to products of *hypersparse* matrices, i.e., matrices with mostly zero rows

Sparse Tensor Formats

- The overhead of transposition, and non-standard nature of the arising sparse matrix products, motivates sparse data structures for tensors that are suitable for tensor contractions of interest

- The *compressed sparse fiber (CSF)* format provides an effective representation for sparse tensors

Operations in Compressed Format

- CSF permits efficient execution of important sparse tensor kernels
  - Analogous to CSR format, which enables efficient implementation of the sparse matrix vector product
  - where `row[i]` stores a list of column indices and nonzeros in the ith row of $\boldsymbol{A}$

```
for i in range(n):
  for (a_ij,j) in row[i]:
    y[i] += a_ij * x[j]
```

  - In CSF format, a multilinear function evaluation $\boldsymbol{f}^{(\mathcal{T})}(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{T}_{(1)}(\boldsymbol{x} \odot \boldsymbol{y})$ can be implemented as

```
for (i,T_i) in T_CSF:
  for (j,T_ij) in T_i:
    for (k,t_ijk) in T_ij:
      z[i] += t_ijk * x[j] * y[k]
```

MTTKRP in Compressed Format

- MTTKRP and CSF pose additional implementation opportunities and challenges
  - MTTKRP $u_{ir} = \sum_{j,k} t_{ijk} v_{jr} w_{kr}$ can be implemented by adding a loop over r to our code for $f^{(\mathcal{T})}$, but would then require $3mr$ operations if m is the number of nonzeros in $\mathcal{T}$, can reduce to $2mr$ by amortization

```
for (i,T_i) in T_CSF:
  for (j,T_ij) in T_i:
    for r in range(R):
      f_ij = 0
      for (k,t_ijk) in T_ij:
        f_ij += t_ijk * w[k,r]
      u[i,r] = f_ij * v[j,r]
```

  - However, this amortization is harder (requires storage or iteration overheads) if the index i is a leaf node in the CSF tree
  - Similar challenges in achieving good reuse and obtaining good arithmetic intensity arise in implementation of other kernels, such as TTMc

All-at-once Contraction

- When working with sparse tensors, it is often more efficient to contract multiple operands in an all-at-once fashion

Constrained Tensor Decomposition

- Many applications of tensor decomposition in data science, feature additional structure, which can be enforced by constraints

Nonnegative Tensor Factorization

- *Nonnegative tensor factorization (NTF)*, such as CP decomposition with $\boldsymbol{\mathcal{T}} \geqslant 0$ and $\boldsymbol{U}, \boldsymbol{V}, \boldsymbol{W} \geqslant 0$ are widespread and a few classes of algorithms have been developed

Nonnegative Matrix Factorization

- NTF algorithms with alternating updates have a close correspondence with alternating update algorithms for *Nonnegative matrix factorization (NMF)*

Coordinate Descent for NMF and NTF

- Coordinate descent gives optimal closed-form updates for variables in NMF and NTF

Generalized Tensor Decomposition

- Aside from addition of constraints, the objective function may be modified by using different elementwise loss functions

- Some loss function admit ALS-like algorithms, while others may require gradient-based optimization