

## About the tutorial

- Introduce matrix and tensor tools through real mining applications
- Goal: find patterns, rules, clusters, outliers, ...
- in matrices and
- in tensors



## What is this tutorial NOT about?

- Classification methods
- Kernel methods
- Discriminative models
- Linear Discriminant Analysis (LDA)
- Canonical Correlation Analysis (CCA)
- Probabilistic latent variable models
- Probabilistic PCA
- Probabilistic latent semantic indexing
- Latent Dirichlet allocation





## Examples of Matrices: Market basket

- market basket as in Association Rules

| John <br> Peter | milk | bread | choc. | wine | ... |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | 13 | 11 | 22 | 55 | ... |
| Peter | 5 | 4 | 6 | 7 | ... |
| Mary | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| Nick | $\ldots$ | $\ldots$ | $\ldots$ | ... | $\ldots$ |
|  | ... | ... | ... | ... | ... |





## Motivation 2: Why tensor?

- Q: what is a tensor?


1-12


Motivation 2: Why tensor?

- A: N-D generalization of matrix:



CMU SCS

## Motivating Applications

- Why matrices are important?
- Why tensors are useful?
- P1: environmental sensors
- P2: data center monitoring ('autonomic')
- P3: social networks
- P4: network forensics
- P5: web mining
- P6: face recognition



CMU SCS

## P4: Network forensics

- Directional netsors
- A large ISPD Data in three modes Aink capacity (time, source, destination)
Task: Identify ant
cause


source
Faloutsos, Kolda, Sun




| CMU SCS <br> Dynamic Data model <br> - Tensor Streams <br> - A sequence of Mth order tensor <br> $\mathcal{X}_{1} \ldots \mathcal{X}_{\boldsymbol{t}}$ where $\mathcal{X}_{i} \in \mathrm{R}^{N_{1} \times \ldots \times N_{M}}$ $t$ is increasing over time |  |  |  |
| :---: | :---: | :---: | :---: |
| Order | 1st | $2^{\text {nd }}$ | $3{ }^{\text {rd }}$ |
| Correspondence | Multiple streams | Time evolving graphs | 3D arrays |
| Example |  |  |  |






## SVD - Interpretation

'documents', 'terms’ and 'concepts':
Q : if $\mathbf{A}$ is the document-to-term matrix, what is $\mathbf{A}^{\mathrm{T}} \mathbf{A}$ ?
A: term-to-term ([m x m]) similarity matrix
Q: $\mathbf{A ~}^{\mathrm{T}}$ ?
A: document-to-document ([n x n]) similarity matrix


## PCA interpretation

- best axis to project on: ('best' = min sum of squares of projection errors)



CMU SCS
風

## Roadmap

- Motivation
- Matrix tools
- Tensor basics
- Tensor extensions
- Software demo
$\left\{\begin{array}{l}\cdot \text { SVD, PCA } \\ \cdot \\ \text { - } H I T S, \text { PageRank }\end{array}\right.$
- Case studies
- Co-clustering
- Nonnegative Matrix factorization




## Kleinberg's algorithm HITS

- Step 1: expand by one move forward and backward



## Kleinberg＇s algorithm HITS

－on the resulting graph，give high score（＝ ＇authorities＇）to nodes that many important nodes point to
－give high importance score（＇hubs＇）to nodes that point to good＇authorities＇

hubs


## Kleinberg＇s algorithm HITS

observations
－recursive definition！
－each node（say，＇$i$＇－th node）has both an authoritativeness score $a_{i}$ and a hubness score $h_{i}$


## Kleinberg's algorithm: HITS

Then:

$$
a_{i}=h_{k}+h_{l}+h_{m}
$$

m

i
i
that is
$a_{i}=\operatorname{Sum}\left(h_{j}\right) \quad$ over all $j$ that (j,i) edge exists
or
$\mathbf{a}=\mathbf{A}^{\mathrm{T}} \mathbf{h}$


## Kleinberg's algorithm: HITS

In conclusion, we want vectors $h$ and a such that:

$$
\begin{aligned}
\mathbf{h} & =\mathbf{A} \mathbf{a} \\
\mathbf{a} & =\mathbf{A}^{\mathrm{T}} \mathbf{h}
\end{aligned}
$$

That is:

$$
\mathbf{a}=\mathbf{A}^{\mathrm{T}} \mathbf{A} \mathbf{a}
$$



## CMU SCS

## Kleinberg's algorithm - discussion

- 'authority' score can be used to find 'similar pages' (how?)
- closely related to 'citation analysis’, social networks / 'small world' phenomena



## Motivating problem: PageRank

Given a directed graph, find its most interesting/central node


A node is important, if it is connected with important nodes (recursive, but OK!)

## Motivating problem－PageRank solution

Given a directed graph，find its most interesting／central node
Proposed solution：Random walk；spot most ＇popular＇node（－＞steady state prob．（ssp））


A node has high ssp， if it is connected with high ssp nodes （recursive，but OK！）
（Simplified）PageRank algorithm
－Let $\mathbf{A}$ be the transition matrix（＝adjacency matrix）；let $\mathbf{A}$ become row－normalized－then



## (Simplified) PageRank algorithm

- $\mathbf{A p}=1^{*} \mathbf{p}$
- thus, $\mathbf{p}$ is the eigenvector that corresponds to the highest eigenvalue $(=1$, since the matrix is row-normalized)
- Why does it exist such a $\mathbf{p}$ ?
- $\mathbf{p}$ exists if A is nxn, nonnegative, irreducible [Perron-Frobenius theorem]


## （Simplified）PageRank algorithm

－In short：imagine a particle randomly moving along the edges
－compute its steady－state probabilities（ssp）

Full version of algo：with occasional random jumps
Why？To make the matrix irreducible

## Full Algorithm

－With probability 1－c，fly－out to a random node
－Then，we have

$$
\begin{aligned}
& \mathbf{p}=\mathrm{c} \mathbf{A} \mathbf{p}+(1-\mathrm{c}) / \mathrm{n} \mathbf{1}=> \\
& \mathbf{p}=(1-\mathrm{c}) / \mathrm{n}[\mathbf{I}-\mathrm{c} \mathbf{A}]^{-1} \mathbf{1}
\end{aligned}
$$




CMU SCS

## Motivation of CUR or CMD

- SVD, PCA all transform data into some abstract space (specified by a set basis)
- Interpretability problem
- Loss of sparsity


## Interpretability problem

- Each column of projection matrix $\mathrm{U}_{\mathrm{i}}$ is a linear combination of all dimensions along certain mode $U_{i}(:, 1)=[0.5 ;-0.5 ; 0.5 ; 0.5]$
- All the data are projected onto the span of $U_{i}$
- It is hard to interpret the projections

The sparsity problem - pictorially:



CMU SCS

## CUR (cont.)

- Key question:
- How to select/sample the columns and rows?
- Uniform sampling [Williams \& Seeger NIPS '00]
- Biased sampling
- CUR w/ absolute error bound
- CUR w/ relative error bound


## Reference:

1. Tutorial: Randomized Algorithms for Matrices and Massive Datasets, SDM'06
2. Drineas et al. Subspace Sampling and Relative-error Matrix Approximation: Column-Row-Based Methods, ESA2006
3. Drineas et al., Fast Monte Carlo Algorithms for Matrices III: Computing a Compressed Approximate Matrix Decomposition, SIAM Journal on Computing, 2006.



- CMD uses much smaller space to achieve the same accuracy
- CUR limitation: duplicate columns and rows
- SVD limitation: orthogonal projection densifies the data

Sun et al. Less is More: Compact Matrix Decomposition for Large Sparse Graphs, SDM'07


CMU SCS

## Co-clustering

- Given data matrix and the number of row and column groups $k$ and $l$
- Simultaneously
- Cluster rows of $p(X, Y)$ into $k$ disjoint groups
- Cluster columns of $p(X, Y)$ into $l$ disjoint groups


Faloutsos, Kolda, Sun



## Problem with Information Theoretic Co-clustering

- Number of row and column groups must be specified

Desiderata:
$\checkmark$ Simultaneously discover row and column groups
※ Fully Automatic: No "magic numbers"
$\checkmark$ Scalable to large graphs








CMU SCS

## Tensor Basics








Observe: For two vectors $\mathbf{a}$ and $\mathbf{b}, \mathbf{a} \circ \mathbf{b}$ and $\mathbf{a} \otimes \mathbf{b}$ have the same elements, but one is shaped into a matrix and the other into a vector.

## Specially Structured Tensors



## Specially Structured Tensors

- Tucker Tensor

$$
\begin{aligned}
& \boldsymbol{x}=\mathbf{S} \times{ }_{1} \mathbf{U} \times_{2} \mathbf{V} \times_{3} \mathbf{W} \\
& =\sum_{r} \sum_{s} \sum_{t} g_{r s t} \mathbf{u}_{r} \circ \mathbf{v}_{s} \circ \mathbf{W}_{t} \\
& \equiv[9 ; \mathbf{U}, \mathrm{V}, \mathrm{~W}]
\end{aligned}
$$

In matrix form:

$$
\begin{aligned}
\mathbf{x}_{(1)} & =\mathbf{U G} \mathbf{G}_{(1)}(\mathbf{W} \otimes \mathbf{V})^{\top} \\
\mathbf{x}_{(2)} & =\mathbf{V G}_{(2)}(\mathbf{W} \otimes \mathbf{U})^{\top} \\
\mathbf{x}_{(3)} & =\mathbf{W G}_{(3)}(\mathbf{V} \otimes \mathbf{U})^{\top} \\
& \\
\operatorname{vec}(\boldsymbol{X}) & =(\mathbf{W} \otimes \mathbf{V} \otimes \mathbf{U}) \operatorname{vec}(\boldsymbol{\mathcal { S }})
\end{aligned}
$$

- Kruskal Tensor

$$
\begin{aligned}
\boldsymbol{X} & =\sum_{r} \boldsymbol{\lambda}_{r} \mathbf{u}_{r} \circ \mathbf{v}_{r} \circ \mathbf{w}_{r} \\
& \equiv \llbracket \boldsymbol{\lambda} ; \mathbf{U}, \mathbf{V}, \mathbf{W}\rfloor
\end{aligned}
$$

In matrix form:
Let $A=\operatorname{diag}(\lambda)$ $\mathbf{X}_{(1)}=U \Lambda(W \in V)^{\top}$ $\mathbf{X}_{(2)}=\mathbf{V} \Lambda(\mathbf{W} \odot \mathbf{U})^{\top}$ $\mathbf{X}_{(3)}=\mathbf{W} \wedge(V \subset U)^{\top}$ $\operatorname{vec}(\boldsymbol{X})=(\mathbf{W} \odot \mathbf{V} \odot \mathbf{U}) \boldsymbol{\lambda}$

## What is the HO Analogue of the Matrix SVD?

Matrix SVD:

$$
\mathbf{X}=\mathbf{U} \mathbf{\Sigma} \mathbf{V}^{\mathbf{\top}}=
$$

Tucker Tensor (finding bases for each subspace):

$$
\mathbf{X}=\mathbf{\Sigma} \times_{1} \mathbf{U} \times_{2} \mathbf{V}=\llbracket \mathbf{\Sigma} ; \mathbf{U}, \mathbf{V} \rrbracket
$$

Kruskal Tensor (sum of rank-1 components):

$$
\mathbf{X}=\sum_{r=1}^{R} \sigma_{r} \mathbf{u}_{r} \circ \mathbf{v}_{r}=\llbracket \sigma ; \mathbf{U}, \mathbf{V} \rrbracket
$$

## Tensor Decompositions

Recall the equations for converting a tensor to a matrix

$$
\mathbf{x}_{(1)}=\mathbf{A} \mathbf{G}_{(1)}(\mathbf{C} \otimes B)^{\top}
$$

$$
\mathbf{x}_{(2)}=\mathbf{B G}_{(2)}(\mathbf{C} \otimes \mathbf{A})^{\top}
$$

$$
\mathbf{x}_{(3)}=\mathbf{C G}_{(3)}(\mathbf{B} \otimes A)^{\top}
$$

$$
\operatorname{vec}(\boldsymbol{X})=(\mathbf{C} \otimes \mathbf{B} \otimes \mathbf{A}) \operatorname{vec}(\boldsymbol{9})
$$

$$
\mathbf{S}=\left[\boldsymbol{X} ; \mathbf{A}^{\dagger}, \mathbf{B}^{\dagger}, \mathbf{C}^{\dagger}\right]
$$





## Tucker-Alternating Least Squares (ALS)

Successively solve for each component (A,B,C).


- Initialize
- Choose R, S, T
- Calculate A, B, C via HO-SVD
- Until converged do...
- A = R leading left singular vectors of $\mathbf{X}_{(1)}(\mathbf{C} \otimes \mathbf{B})$
$-\mathbf{B}=\mathrm{S}$ leading left singular vectors of $\mathbf{X}_{(2)}(\mathbf{C} \otimes \mathbf{A})$
- $\mathbf{C}=\mathrm{T}$ leading left singular vectors of $\mathbf{X}_{(3)}(\mathbf{B} \otimes \mathbf{A})$
- Solve for core:

$$
\mathbf{S}=\llbracket \boldsymbol{X} ; \mathbf{A}^{\top}, \mathbf{B}^{\top}, \mathbf{C}^{\top} \rrbracket
$$

## Tucker in Not Unique



Tucker decomposition is not unique. Let Y be an RxR orthogonal matrix. Then...

$$
\begin{gathered}
X \approx \mathcal{S} \times{ }_{1} \mathbf{A} \times_{2} \mathbf{B} \times{ }_{3} \mathbf{C}=\left(\mathcal{S} \times{ }_{1} \mathbf{Y}^{\top}\right) \times_{1}(\mathbf{A Y}) \times{ }_{2} \mathbf{B} \times{ }_{3} \mathbf{C} \\
\mathbf{X}_{(1)} \approx A \mathbf{G}_{(1)}(\mathbf{C} \otimes \mathbf{B})^{\top}=\mathbf{A Y Y} \mathbf{G}_{(1)}(\mathbf{C} \otimes \mathbf{B})^{\top}
\end{gathered}
$$




## Tucker vs. PARAFAC Decompositions

- Tucker
- Variable transformation in each mode
- Core G may be dense
- A, B, C generally orthonormal
- Not unique

- PARAFAC
- Sum of rank-1 components
- No core, i.e., superdiagonal core
- A, B, C may have linearly dependent columns
- Generally unique




## Other Tensor Decompositions

Combining Tucker \& PARAFAC

CMU SCS

## 2-Way DEDICOM



- 2-way DEDICOM introduced by Harshman (1978)
- X is a matrix of interactions between N entities
- Interactions can be nonsymmetric
- Assumes there are " M " roles
- Each entity has a weight for each role in A
- $\mathrm{R}_{\mathrm{ij}}=$ interaction weight for roles i \& j

- 3-way DEDICOM due to Kiers (1993)
- Once again, X captures interactions among entities
- Third dimension can correspond to time
- Diagonal slices capture participation of each role at each time
- See Bader et al., SAND2006-7744 , for application to Enron email data



CMU SCS

## Non-negative 3-Way PARAFAC Factorization

$$
\|\boldsymbol{X}-\llbracket \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket\| \leftarrow \underset{\text { A, B and } \mathrm{C} \text { being positive. }}{ }
$$

Lee-Seung-like update formulas can be derived for 3D and higher:
$\mathbf{A}=\mathbf{A} *\left(\mathbf{X}_{(1)}(\mathbf{C} \odot \mathbf{B})\right) \oslash\left(\mathbf{A}\left(\mathbf{C}^{\top} \mathbf{C} * \mathbf{B}^{\top} \mathbf{B}\right)\right)$
$B=B *\left(\mathbf{X}_{(2)}(\mathbf{C} \odot A)\right) \oslash\left(B\left(C^{\top} C * A^{\top} A\right)\right)$
$\mathbf{C}=\mathbf{C} *\left(\mathbf{X}_{(3)}(\mathbf{B} \odot \mathbf{A})\right) \oslash\left(\mathbf{C}\left(\mathbf{B}^{\top} \mathbf{B} * \mathbf{A}^{\top} \mathbf{A}\right)\right)$

Elementwise multiply Elementwise divide
(Hadamard product)
M. Mørup, L. K. Hansen, J. Parnas, S. M. Arnfred, Decomposing the time-frequency representation of EEG using non-negative 4-8 matrix and multi-way factorization, 2006


CMU SCS
A Quick Overview on Handling Missing Data

- Consider sparse PARAFAC where $\boldsymbol{X}$ is missing data:

$$
\mathcal{X} \approx \llbracket \lambda ; \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket
$$

- Typically, missing values are just set to zero
- More sophisticated approaches for handling missing values:
- Weighted least squares loss function
- Ignore missing values
- Data imputation
- Estimate missing values
- See, e.g., Kiers, Psychometrika, 1997 and Srebro \& Jaakkola, ICML 2003
Weighted
Least Squares

$$
\sum_{i} \sum_{j} \sum_{k} w_{i j k}\left(x_{i j k}-\sum_{r} \lambda_{r} a_{i r} b_{j r} c_{k r}\right)^{2}
$$

- But this problem is often too hard to solve directly!


## Missing Value Imputation

- Use the current estimate to fill in the missing values

$$
\boldsymbol{\mathcal { E }}=\llbracket \boldsymbol{\lambda} ; \mathbf{A}, \mathbf{B}, \mathbf{C} \rrbracket \quad \text { Current Estimate }
$$

- The tensor for the next iteration of the algorithm is:

$$
\begin{aligned}
\hat{\boldsymbol{X}} & =\overbrace{\mathbf{N} * \boldsymbol{X}+}^{\text {Known }} \underbrace{\mathbf{1}-\mathbf{N}, \mathbf{N}) * \boldsymbol{E}}_{\text {Sparse! }} \\
= & \underbrace{\text { Estimates of Unknowns }}_{\text {Kruskal Tensor }}
\end{aligned}
$$

- Challenge is finding a good initial estimate



## Tensor Toolbox for MATLAB

http://csmr.ca.sandia.gov/~tgkolda/TensorToolbox

- Six object-oriented tensor classes
- Working with tensors is easy
- Most comprehensive set of kernel operations in any language
- E.g., arithmetic, logical, multiplication operations
- Sparse tensors are unique
- Speed-ups of two orders of magnitude for smaller problems
- Larger problems than ever before

- Free for research or evaluations purposes
- 297 unique registered users from all over the world (as of January 17, 2006)

Bader \& Kolda, ACM TOMS 2006 \& SAND2006-7592

CMU SCS

## Dense Tensors

- Largest tensor that can be stored on a laptop is 200 x 200 x 200
- Typically, tensor operations are reduced to matrix operations
- Requires permuting and reshaping the tensor
- Example: Mode-n tensormatrix multiply


Example: Mode-1 Matrix Multiply

$$
\underset{M \times J \times K}{y}=\mathbf{X} \times \mathbf{1}_{M \times J \times K} \mathbf{U}
$$

$\mathbf{Y}_{(\boldsymbol{n})}=\mathbf{U X X}_{M \times 1} \mathbf{X}_{(\boldsymbol{n})}$
$M \times J K \quad I \times J K$




## Incremental Tensor Decomposition

- Dynamic data model
- Tensor Streams
- Dynamic Tensor Decomposition (DTA)
- Streaming Tensor Decomposition (STA)
- Window-based Tensor Decomposition (WTA)

- Streams come with structure
- (time, source, destination, port)
- (time, author, keyword)
- How to summarize tensor streams effectively and incrementally?



## Incremental Tensor Decomposition

© © Dynamic data model

- Tensor Streams
- Dynamic Tensor Decomposition (DTA)
- Streaming Tensor Decomposition (STA)
- Window-based Tensor Decomposition (WTA)

1. Jimeng Sun, Spiros Papadimitriou, Philip Yu. Window-based Tensor Analysis on High-dimensional and Multi-aspect Streams, ICDM 2006
2. Jimeng Sun, Dacheng Tao, Christos Faloutsos. Beyond Streams and Graphs: Dynamic Tensor Analysis, KDD 2006



## 1st order Dynamic tensoi Anaiysis



Input: new data vector $x \in \mathrm{R}^{\mathrm{N}}$, old variance matrix $\mathrm{C} \in \mathrm{R}^{\mathrm{N} \times \mathrm{N}}$
Output: new projection matrix $\mathrm{U} \in \mathrm{R}^{\mathrm{N} \times \mathrm{R}}$ Algorithm:

1. update variance matrix $C_{n e w}=x^{T} X+C$
2. Diagonalize $\mathrm{U} \Lambda \mathrm{U}^{\mathrm{T}}=\mathrm{C}_{\text {new }}$
3. Determine the rank R and return U


Diagonalization has to be done for every new $\mathbf{x}$ !
Faloutsos, Kolda, Sun


CMU SCS

M $^{\text {th }}$ order DTA - complexity

## Storage:

$O\left(\prod N_{i}\right)$, i.e., size of an input tensor at a single timestamp
Computation:
$\sum \mathrm{N}_{\mathrm{i}}^{3}$ (or $\sum \mathrm{N}_{\mathrm{i}}{ }^{2}$ ) diagonalization of C
$+\sum \mathrm{N}_{\mathrm{i}} \Pi \mathrm{N}_{\mathrm{i}} \quad$ matrix multiplication $\mathrm{X}_{(\mathrm{d})}{ }^{\mathrm{T}} \mathrm{X}_{(\mathrm{d})}$
For low order tensor(<3), diagonalization is the main cost
For high order tensor, matrix multiplication is the main cost

## Incremental Tensor Decomposition

Dynamic data model

- Tensor Streams
© Dynamic Tensor Decomposition (DTA)
- Streaming Tensor Decomposition (STA)
- Window-based Tensor Decomposition (WTA)

1. Jimeng Sun, Spiros Papadimitriou, Philip Yu. Window-based Tensor Analysis on High-dimensional and Multi-aspect Streams, ICDM 2006
2. Jimeng Sun, Dacheng Tao, Christos Faloutsos. Beyond Streams and Graphs: Dynamic Tensor Analysis, KDD 2006



CMU SCS

## Incremental Tensor Decomposition

© - Dynamic data model

- Tensor Streams
© Dynamic Tensor Decomposition (DTA)
© Streaming Tensor Decomposition (STA)
- Window-based Tensor Decomposition (WTA)

1. Jimeng Sun, Spiros Papadimitriou, Philip Yu. Window-based Tensor Analysis on High-dimensional and Multi-aspect Streams, ICDM 2006
2. Jimeng Sun, Dacheng Tao, Christos Faloutsos. Beyond Streams and Graphs: Dynamic Tensor Analysis, KDD 2006




CMU SCS
(17)

## Roadmap

- Motivation
- Matrix tools
- Tensor basics
- Tensor extensions
- Software demo
- Case studies



- $2^{\text {nd }}$ factor captures an atypical trend:
- Uniformly across all time
- Concentrating on 3 locations
- Mainly due to voltage
- Interpretation: two sensors have low battery, and the other one has high battery.













| Sumpary |  |  |  |
| :--- | :--- | :--- | :--- |
| Methods | Pros | Cons | Applications |
| SVD, PCA | Optimal in L2 <br> and Frobenius | Dense representation, <br> Negative entries | LSI, PageRank, <br> HITS |
| CUR, CMD | Interpretability, <br> sparse bases | Not optimal like SVD, <br> dense core | DNA SNP data, <br> network forensics |
| Co-clustering | Interpretability | Local minimum | Social networks, <br> microarray data |
| Tucker | Flexible <br> representation | Interpretability, non- <br> uniqueness, dense core | TensorFaces |
| PARAFAC | Interpretability, <br> efficient parse <br> computation | Slow convergence | TOPHISTS |
| Incrementalization | Efficiency | Non-optimal | Tensor Streams |
| Nonnegativity | Interpretability, <br> sparse results | Local minimum, non- <br> uniqueness | Image <br> segmentation |

## Conclusion

－Real data are often in high dimensions with multiple aspects（modes）
－Matrix and tensor provide elegant theory and algorithms for such data
－However，many problems are still open
－skew distribution，anomaly detection，streaming algorithm，distributed／parallel algorithms， efficient out－of－core processing

## Thank you！

－Christos Faloutsos www．cs．cmu．edu／～christos
－Tamara Kolda csmr．ca．sandia．gov／～tgkolda
－Jimeng Sun www．cs．cmu．edu／～jimeng


