


Tensor Analysis – Applications and Algorithms


Christos Faloutsos

CMU


Roadmap




- Applications – pattern discovery
 - ➔ – Brain scans – coupled matrix-tensor factorization
 - Power grid
- Applications – anomaly detection
- Algorithms
- Conclusions



Roadmap

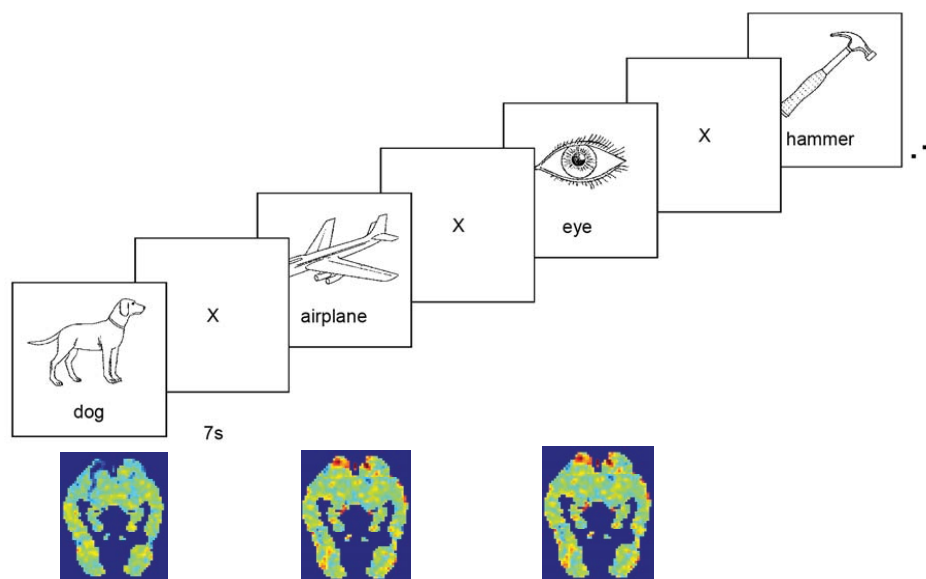


- Applications – pattern discovery
 - ➔ – Brain scans – coupled matrix-tensor factorization
 - Power grid
- Applications – anomaly detection
- Algorithms
- Conclusions



Neuro-semantic


- **Brain Scan Data***
 - 9 persons
 - 60 nouns
- **Questions**
 - 218 questions
 - ‘is it alive?’, ‘can you eat it?’



*Mitchell et al. *Predicting human brain activity associated with the meanings of nouns*. Science, 2008. Data@ www.cs.cmu.edu/afs/cs/project/theo-73/www/science2008/data.html

Neuro-semantic

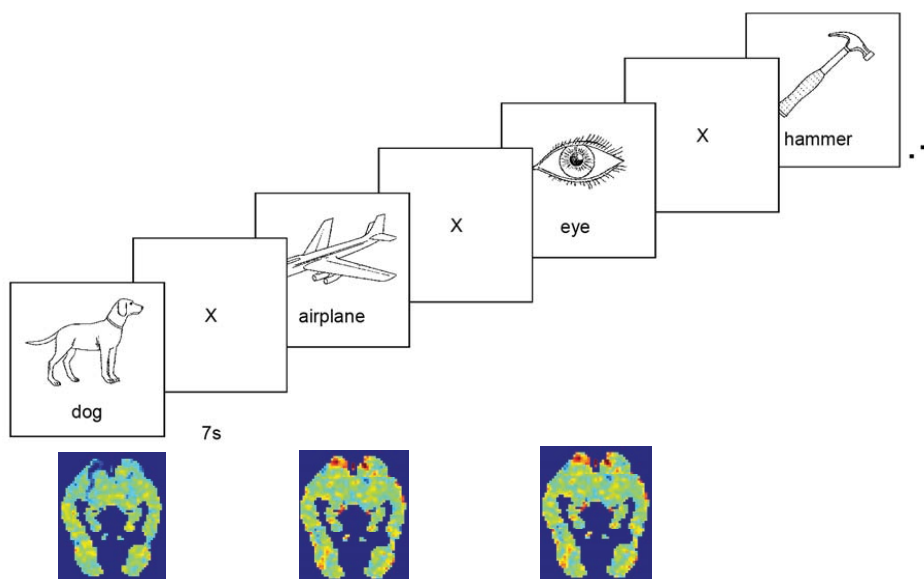
- **Brain Scan Data***
 - 9 persons
 - 60 nouns
- **Questions**
 - 218 questions
 - ‘is it alive?’, ‘can you eat it?’




Patterns?

Neuro-semantic

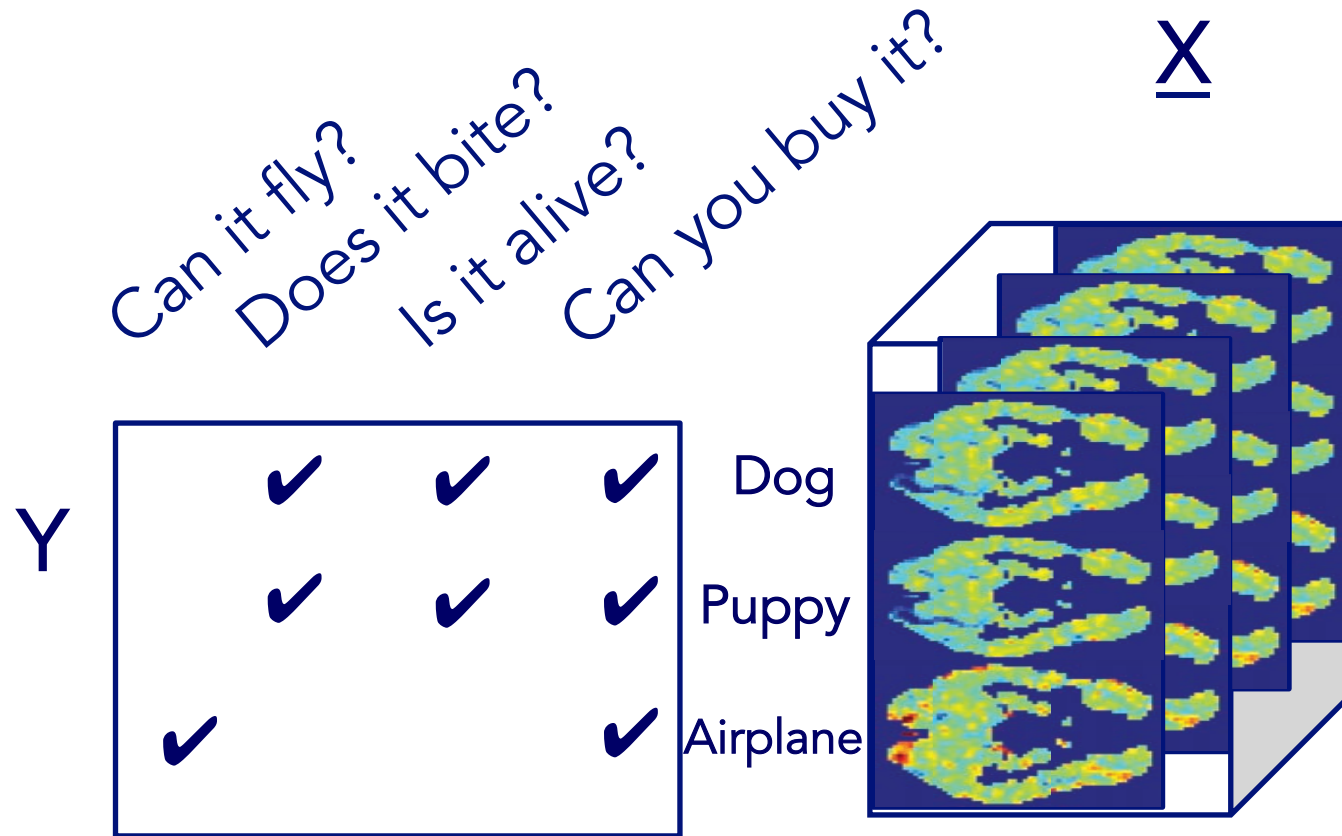
- **Brain Scan Data***
 - 9 persons
 - 60 nouns
- **Questions**
 - 218 questions
 - ‘is it alive?’, ‘can you eat it?’




Patterns?




Neuro-semantic



Coupled Matrix-Tensor Factorization (CMTF)



Neuro-semantic




Nouns

beetle
pants
bee

Questions

can it cause you pain?
do you see it daily?
is it conscious?




Group 1

Nouns

bear
cow
coat

Questions

does it grow?
is it alive?
was it ever alive?




Group 2

Nouns

glass
tomato
bell

Questions

can you pick it up?
can you hold it in one hand?
is it smaller than a golfball?




Group 3

Nouns

bed
house
car


Questions

does it use electricity?
can you sit on it?
does it cast a shadow?



Group 4

Neuro-semantic



**Small items ->
Premotor cortex**


Nouns

glass
tomato
bell

Questions

can you pick it up?
can you hold it in one hand?
is it smaller than a golfball?

- ✓ Unsupervised
- ✓ Matches intuition



Group 3

Neuro-semantic


Small items ->
Premotor cortex

Nouns

glass
tomato
bell

Questions

can you pick it up?
can you hold it in one hand?
is it smaller than a golfball?




Group 3



Evangelos Papalexakis, Tom Mitchell, Nicholas Sidiropoulos,
Christos Faloutsos, Partha Pratim Talukdar, Brian Murphy,
*Turbo-SMT: Accelerating Coupled Sparse Matrix-Tensor
Factorizations by 200x*, SDM 2014




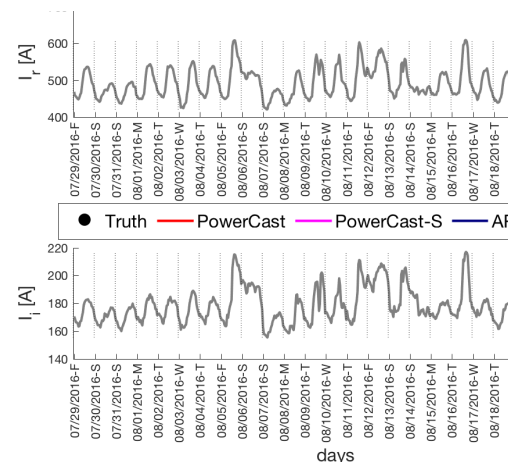
Roadmap



- Applications – pattern discovery
 - Brain scans – coupled matrix-tensor factorization
 - – Power grid
- Applications – anomaly detection
- Algorithms
- Conclusions

PowerCast

Mining electric power data



Hyun Ah Song, Bryan Hooi, Marko Jereminov, Amritanshu Pandey, Larry Pileggi, and CF, *PowerCast: Mining and Forecasting Power Grid Sequences*, PKDD'17, Skopje, FYROM

Problem definition

- **Given:** real and imaginary current and voltage
- **Forecast:** power demand in the future, and
- **Guess:** how the forecasts will change under various scenarios (e.g. population drops in half, etc)



Domain knowledge

$$I_r(t) = G(t)V_r(t) - B(t)V_i(t) + \alpha_r(t)$$


$$I_i(t) = B(t)V_r(t) + G(t)V_i(t) + \alpha_i(t)$$

resistors

Coils/capacitors

I_r

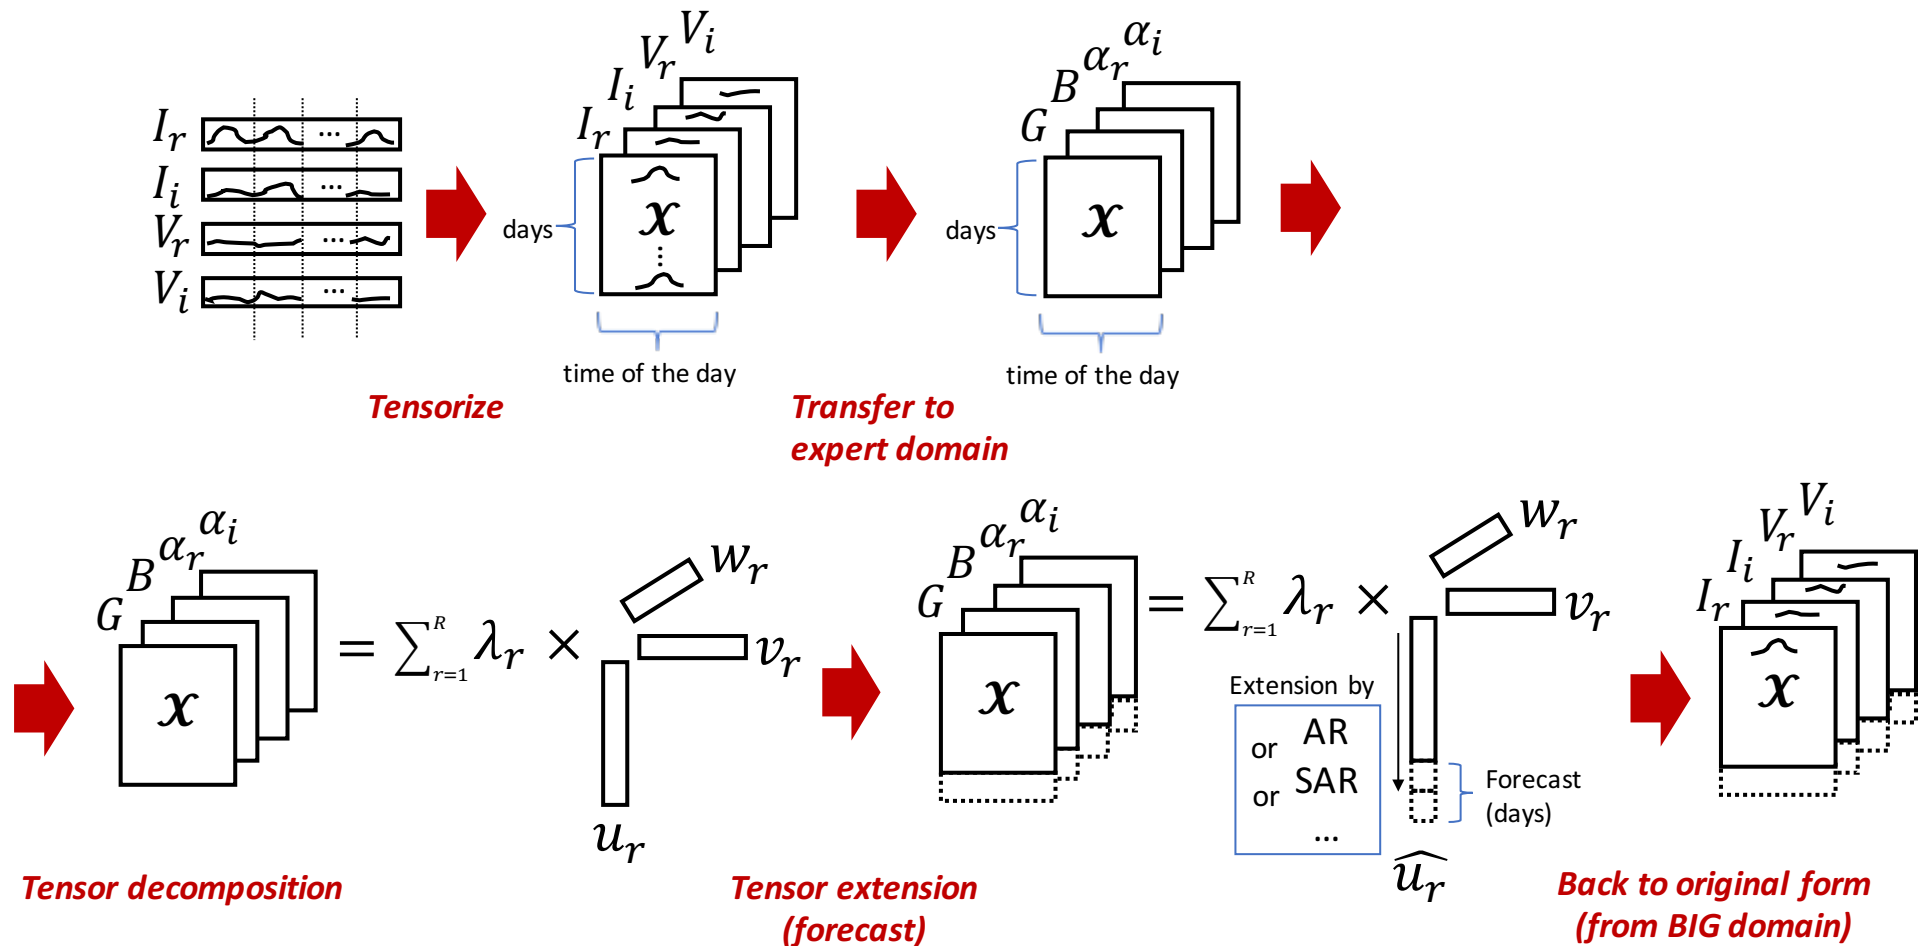
I_i



PowerCast

$$I_r(t) = G(t)V_r(t) - B(t)V_i(t) + \alpha_r(t)$$

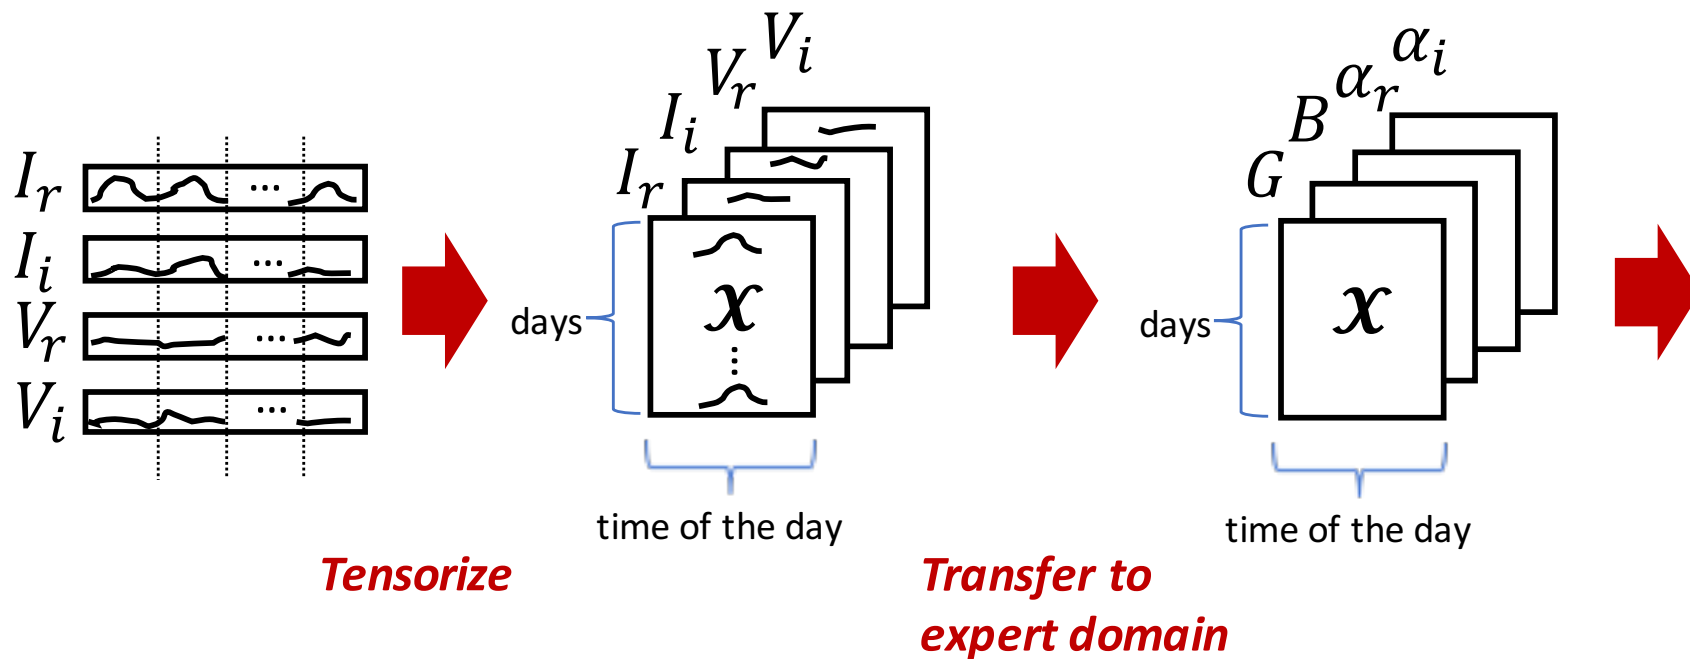
$$I_i(t) = B(t)V_r(t) + G(t)V_i(t) + \alpha_i(t)$$




PowerCast

$$I_r(t) = G(t)V_r(t) - B(t)V_i(t) + \alpha_r(t)$$

$$I_i(t) = B(t)V_r(t) + G(t)V_i(t) + \alpha_i(t)$$

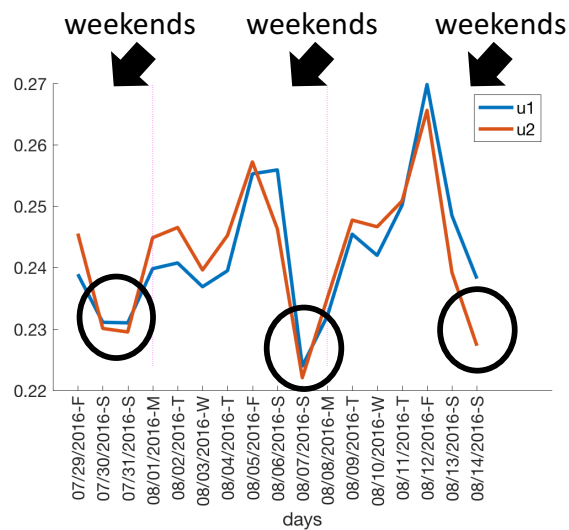


PowerCast




Tensor factors (concepts)

$$G \begin{matrix} B \\ \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} \begin{matrix} \alpha_i \\ \alpha_r \end{matrix} = \sum_{r=1}^R \lambda_r \times \begin{matrix} w_r \\ v_r \\ u_r \end{matrix}$$



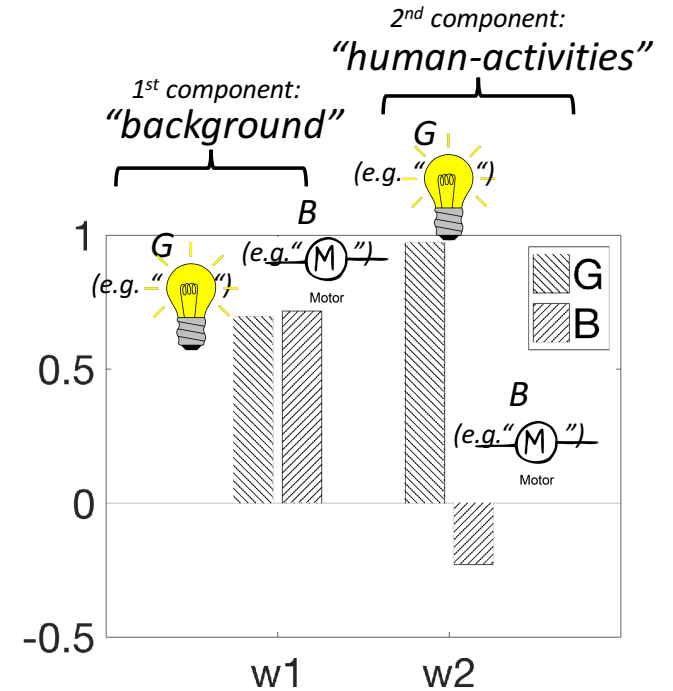
Long-term-concepts
(u_r)

SIAM, July 2017




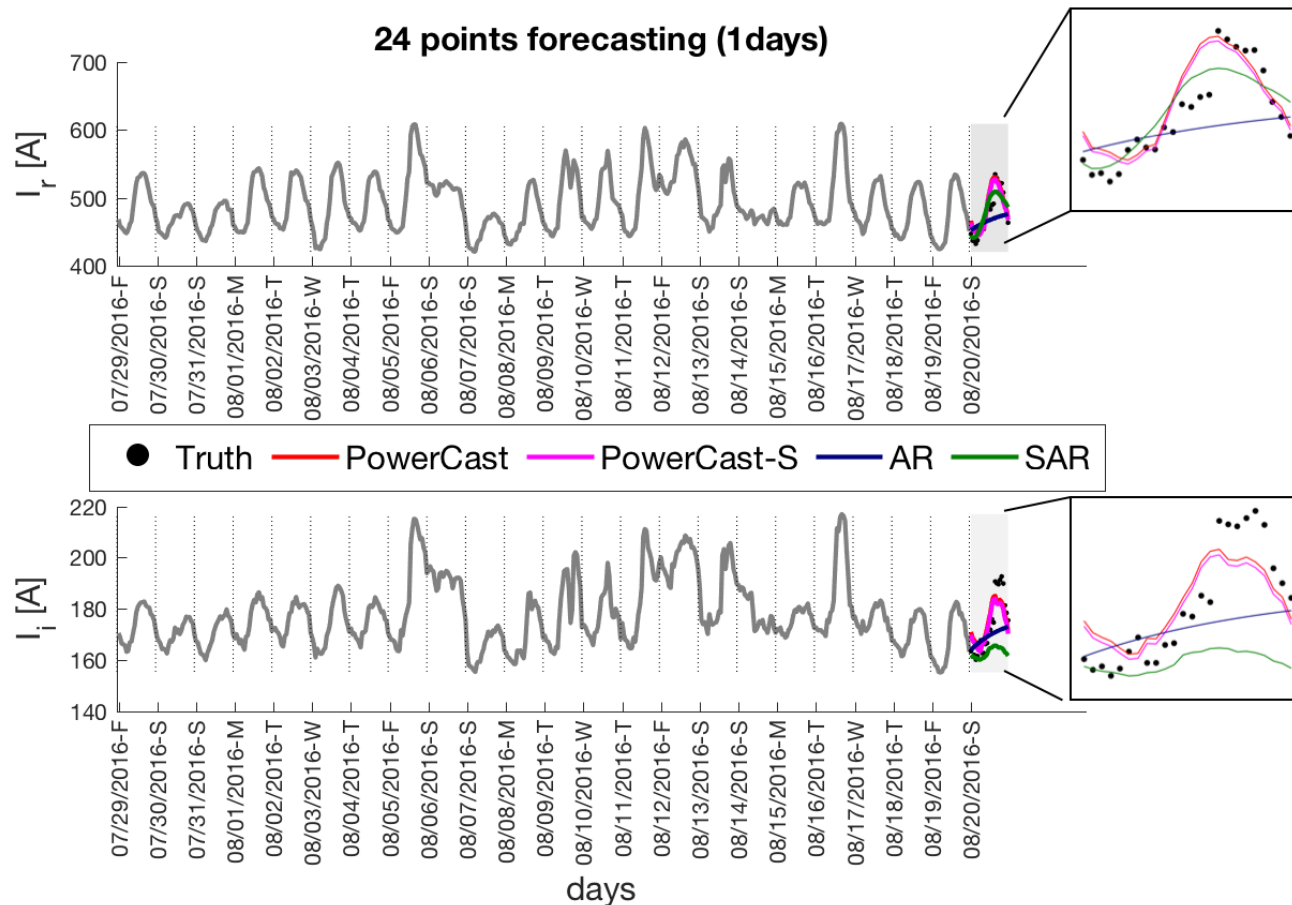
Daily-concepts
(v_r)

(c) C. Faloutsos, 2017



User-profile-concepts
(w_r)

Forecast




PowerCast forecasts 24-steps (1-day) ahead on CMU data more accurate than other competitors



Anomaly detection & explanation

Anomaly!

Q: What happened?




A:




‘background’ component increased!

?:

Case 1?



Case 2?




SIAM, July 2017

(c) C. Faloutsos, 2017

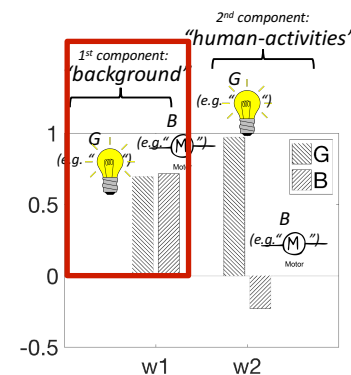
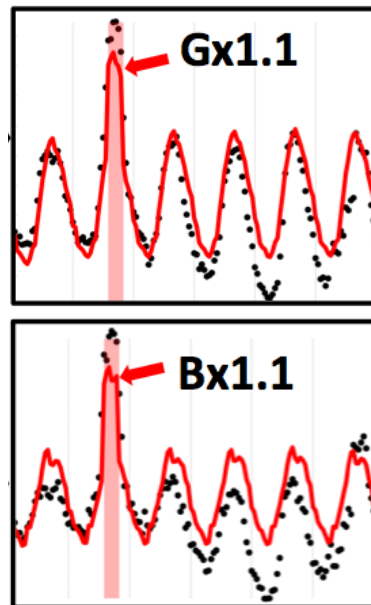
Anomaly detection & explanation

Anomaly!

Q: What happened?




A:




'background' component increased!

?:

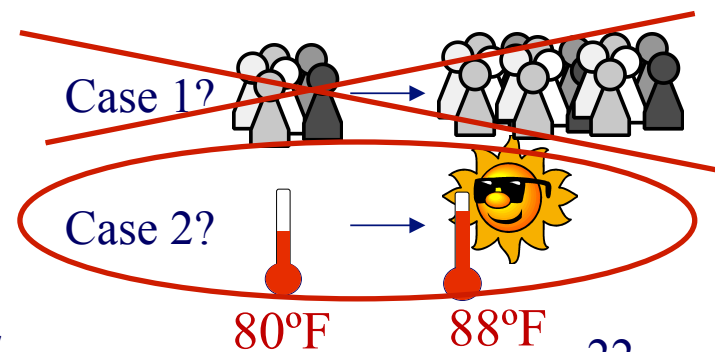
Case 1?



Case 2?




SIAM, July 2017




(c) C. Faloutsos, 2017

Roadmap



- Applications – pattern discovery

- Applications – anomaly detection



- Phone-call data


- Intrusion detection

- Algorithms


- Conclusions

Anomalies in phone-call data

- PARAFAC decomposition
- Results for who-calls-whom-when
 - 4M x 15 days




Anomaly detection in time-evolving graphs




- Anomalous communities in phone call data:
 - European country, 4M clients, data over 2 weeks


1 caller



5 receivers




4 days of activity




~200 calls to EACH receiver on EACH day!

Anomaly detection in time-evolving graphs




- Anomalous communities in phone call data:
 - European country, 4M clients, data over 2 weeks


1 caller



5 receivers




4 days of activity




~200 calls to EACH receiver on EACH day!


Anomaly detection in time-evolving graphs




- Anomalous communities in phone call data:
 - European country, 4M clients, data over 2 weeks




1 caller



5 receivers




4 days of activity



~200 calls to EACH receiver on EACH day!

Anomaly detection in time-evolving graphs




- Anomalous communities in phone call data:
 - European country, 4M clients, data over 2 weeks




Miguel Araujo, Spiros Papadimitriou, Stephan Günnemann, Christos Faloutsos, Prithwish Basu, Ananthram Swami, Evangelos Papalexakis, Danai Koutra. *Com2: Fast Automatic Discovery of Temporal (Comet) Communities*. PAKDD 2014, Tainan, Taiwan.

Roadmap





- Applications – pattern discovery
- Applications – anomaly detection
 - Phone-call data
 - Intrusion detection
- Algorithms
- Conclusions




ParCube at Work: Port-Scanning


LBNL Network Data




Roadmap



- Applications – pattern discovery
- Applications – anomaly detection
- Algorithms
 - ➔ – ParCube (and TurboSMT)
 - S-HOT for higher order Tucker
- Conclusions




ParCube/Turbo-SMT: Triple-sparse Parallel Tensor & Coupled Decomposition




Speedup

4 Intel Xeon E74850
512Gb RAM, Fedora 14


Data size
~500Mb



Roadmap



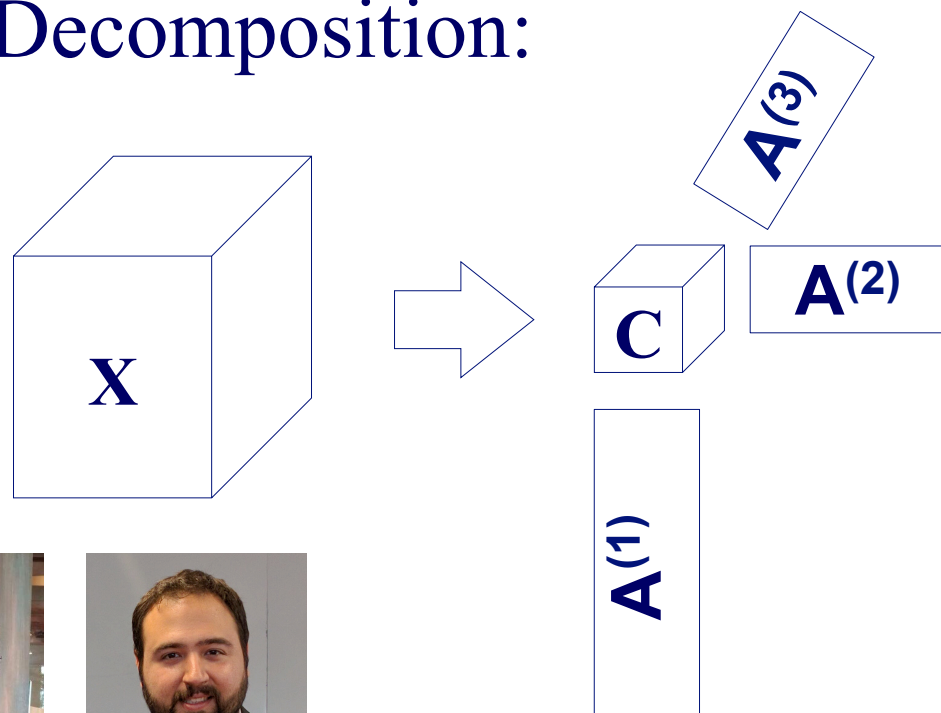


- Applications – pattern discovery
- Applications – anomaly detection
- Algorithms
 - ParCube (and TurboSMT)
 - S-HOT for higher order Tucker
- Conclusions



S-HOT: Scalable High-Order Tucker Decomposition

- Tucker Decomposition:



Jinhoh Oh, Kijung Shin, Evangelos E. Papalexakis, Christos Faloutsos, and Hwanjo Yu, *S-HOT: Scalable High-Order Tucker Decomposition*, WSDM 2017

High-Order Tucker Decomposition (Example)

- Input tensor
 - \mathbf{X} : a 5-way sparse tensor
(size: 1M ... 1M, #non-zeros: 100M)
- Output tensors:
 - \mathbf{C} : a 5-way core tensor
(size: 10 ... 10, #non-zeros : ~100K)
 - $\mathbf{A}, \dots, \mathbf{A}^{(5)}$: factor matrices
(size: 1M 10, #non-zeros : ~10M)

High-Order Tucker Decomposition (Main idea)

Algorithm 1: Tucker-ALS

Input : \mathcal{X} , a N -order tensor of $\mathbb{R}^{I_1 \times \dots \times I_N}$.
 J_1, \dots, J_N , the rank size of each mode.
 T , the number of iterations.

Output: $\{\mathbf{A}\}$, a set of factor matrices $\{\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}\}$ where $\mathbf{A}^{(n)} \in \mathbb{R}^{I_n \times J_n}$.
 \mathcal{G} , a N -order core tensor of $\mathbb{R}^{J_1 \times \dots \times J_N}$.

- 1 Initialize all $\mathbf{A}^{(n)}$
- 2 **for** $t \leftarrow 1..T$ **do**
- 3 **for** $n \leftarrow 1..N$ **do**
- 4 $\mathbf{Y}_{(n)} \leftarrow [\mathcal{X} \times_{-n} \{\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(n-1)}, \mathbf{A}^{(n+1)}, \dots, \mathbf{A}^{(N)}\}]_{(n)}$
- 5 $\mathbf{A}^{(n)} \leftarrow$ top- J_n left singular vectors of $\mathbf{Y}_{(n)}$
- 6 $\mathcal{G} \leftarrow \mathbf{Y}_{(N)} \times_N \mathbf{A}^{(N)T}$
- 7 **return** $\mathcal{G}, \{\mathbf{A}\}$

Huge, & dense

High-Order Tucker Decomposition (Main idea)

Algorithm 1: Tucker-ALS

Input : \mathcal{X} , a N -order tensor of $\mathbb{R}^{I_1 \times \dots \times I_N}$.
 J_1, \dots, J_N , the rank size of each mode.
 T , the number of iterations.

Output: $\{\mathbf{A}\}$, a set of factor matrices $\{\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}\}$ where $\mathbf{A}^{(n)} \in \mathbb{R}^{I_n \times J_n}$.
 \mathcal{G} , a N -order core tensor of $\mathbb{R}^{J_1 \times \dots \times J_N}$.

```

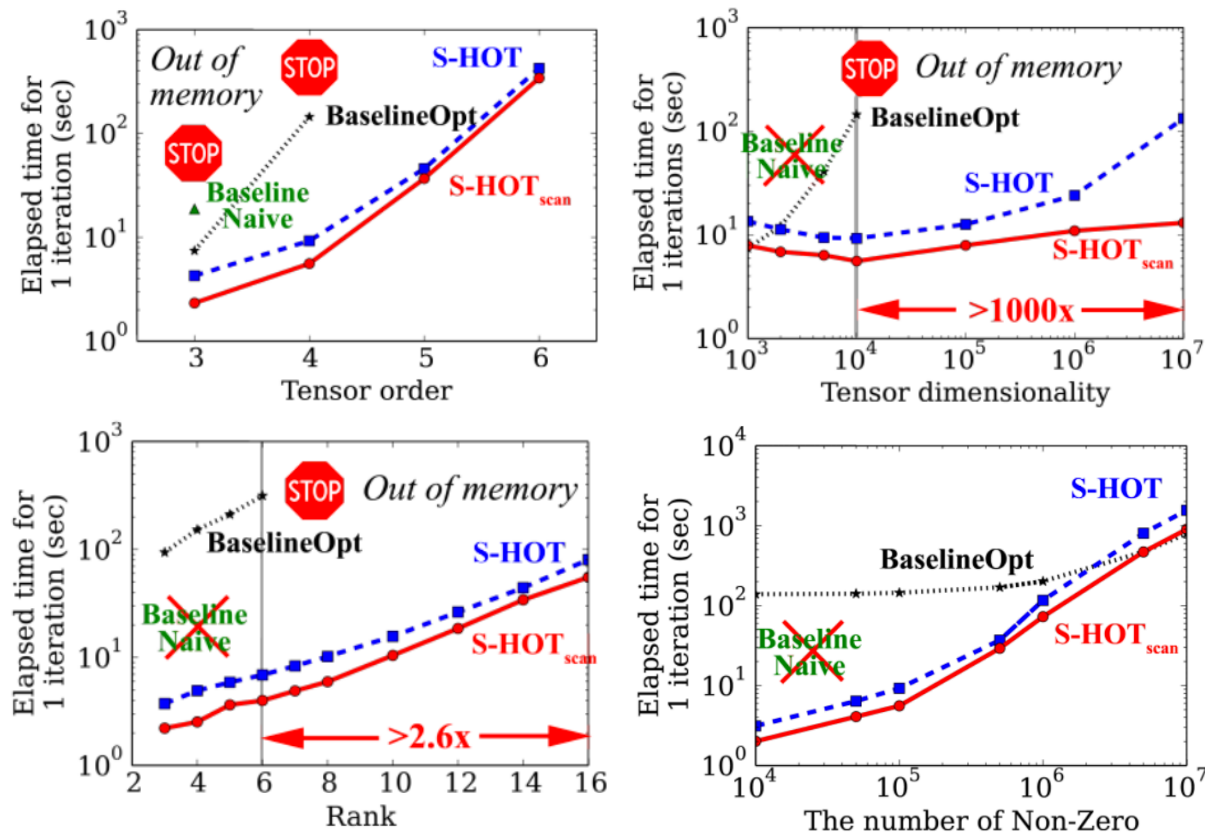
1 Initialize all  $\mathbf{A}^{(n)}$ 
2 for  $t \leftarrow 1..T$  do
3   for  $n \leftarrow 1..N$  do
4      $\mathbf{Y}_{(n)} \leftarrow [\mathcal{X} \times_{-n} \{\mathbf{A}^{(T)}\}]_{(n)}$ 
5      $\mathbf{A}^{(n)} \leftarrow$  top- $J_n$  left singular vectors of  $\mathbf{Y}_{(n)}$ 
6  $\mathcal{G} \leftarrow \mathbf{Y}_{(N)} \times_N \mathbf{A}^{(N)T}$ 
7 return  $\mathcal{G}, \{\mathbf{A}\}$ 

```

Huge, & dense
 -> DON'T
 materialize it

Scalability of S-HOT

- Baselines: Standard ALS (Naïve), & (opt)



Scalability of S-HOT

- Baselines: Standard ALS (Naïve), & (opt)
- $I = 1\text{M}$; $J=10$; $M=1\text{B}$; $N=5$ modes

Method	Space Requirements for Intermediate Data	
	(in Theory)	(in Example)
BaselineNaive	$M J^{N-1}$	$\sim 40\text{TB}$
BaselineOpt [12]	$I J^{N-1}$	$\sim 40\text{GB}$
HATEN2 [8]	$\max(I J^{N-1}, M(N-1)J)$	$\sim 160\text{GB}$
S-HOT	$\max(I, J^{N-1})$	$\sim 4\text{MB}$
S-HOT _{scan}	J^{N-1}	$\sim 40\text{KB}$

Discovery using S-HOT

- Microsoft Academic Graph

(42M papers × 25K venues × 115M authors × 54K keywords)

CS-related International Conference on Networking(ICN),
Wired/Wireless Internet Communications(WWIC),
Database and Expert Systems Applications(DEXA), Data Mining and Knowledge
Discovery, IEEE Transactions on Robotics, ...

Nanotech. Nature Nanotechnology, PLOS ONE, Journal of Ex-
perimental Nanoscience, Journal of Nanoscience and
Nanotechnology, Journal of Semiconductors, Trends
in Biotechnology, ...

Clinical European Journal of Cancer, PLOS Biology, Clin-
ical and Applied Thrombosis-Hemostasis, Journal
of Infection Prevention, RBMC Clinical Pharmacol-
ogy, Regional Anesthesia and Pain ...

Discovery using S-HOT

- Microsoft Academic Graph


(42M papers × 25K venues × 115M authors × 54K keywords)

CS-related International Conference on Networking(ICN),
Wired/Wireless Internet Communications(WWIC),
Database and Expert Systems Applications(DEXA), Data Mining and Knowledge
Discovery, IEEE Transactions on Robotics, ...

Roadmap



- Applications – pattern discovery
- Applications – anomaly detection
- Algorithms
 - ParCube (and TurboSMT)
 - S-HOT for higher order Tucker



- ➔ • Conclusions


Thanks



Disclaimer: All opinions are mine; not necessarily reflecting the opinions of the funding agencies

Thanks to: NSF IIS-0705359, IIS-0534205, CTA-INARC; Yahoo (M45), LLNL, IBM, SPRINT, Google, INTEL, HP, iLab

CONCLUSION#1





- **MANY** applications for tensors

Nouns

glass
tomato
bell

Questions

can you pick it up?
can you hold it in one hand?
is it smaller than a golfball?'





Group 3

CONCLUSION#2



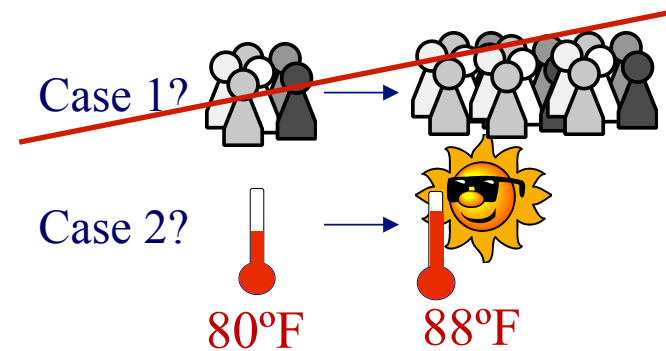
- Domain-experts: valuable

Q: What happened?



$$I_r(t) = G(t)V_r(t) - B(t)V_i(t) + \alpha_r(t)$$

$$I_i(t) = B(t)V_r(t) + G(t)V_i(t) + \alpha_i(t)$$





Thank you!



- Domain-experts: valuable

Q: What happened?



$$I_r(t) = G(t)V_r(t) - B(t)V_i(t) + \alpha_r(t)$$

$$I_i(t) = B(t)V_r(t) + G(t)V_i(t) + \alpha_i(t)$$

