



1. Subgraph Analysis

2. Propagation Methods

3. Latent Factor Models

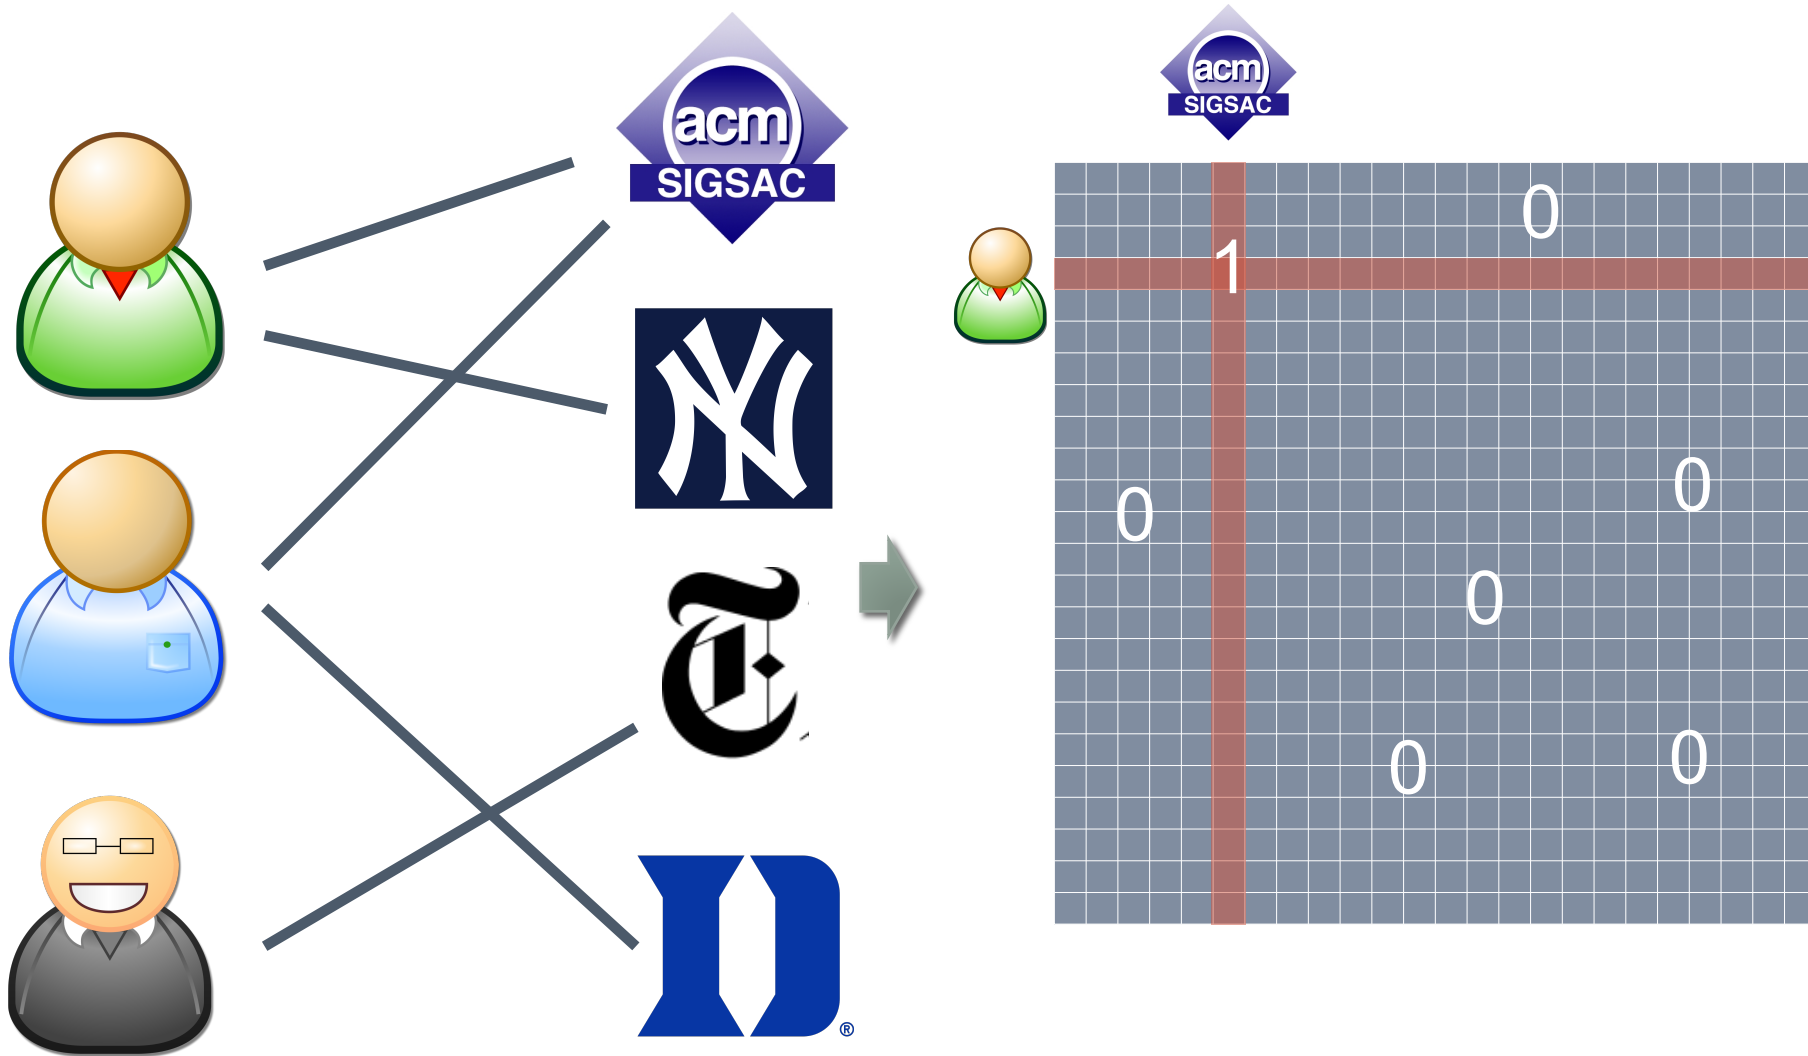
a) Background

b) Normal Behavior

c) Abnormal Behavior

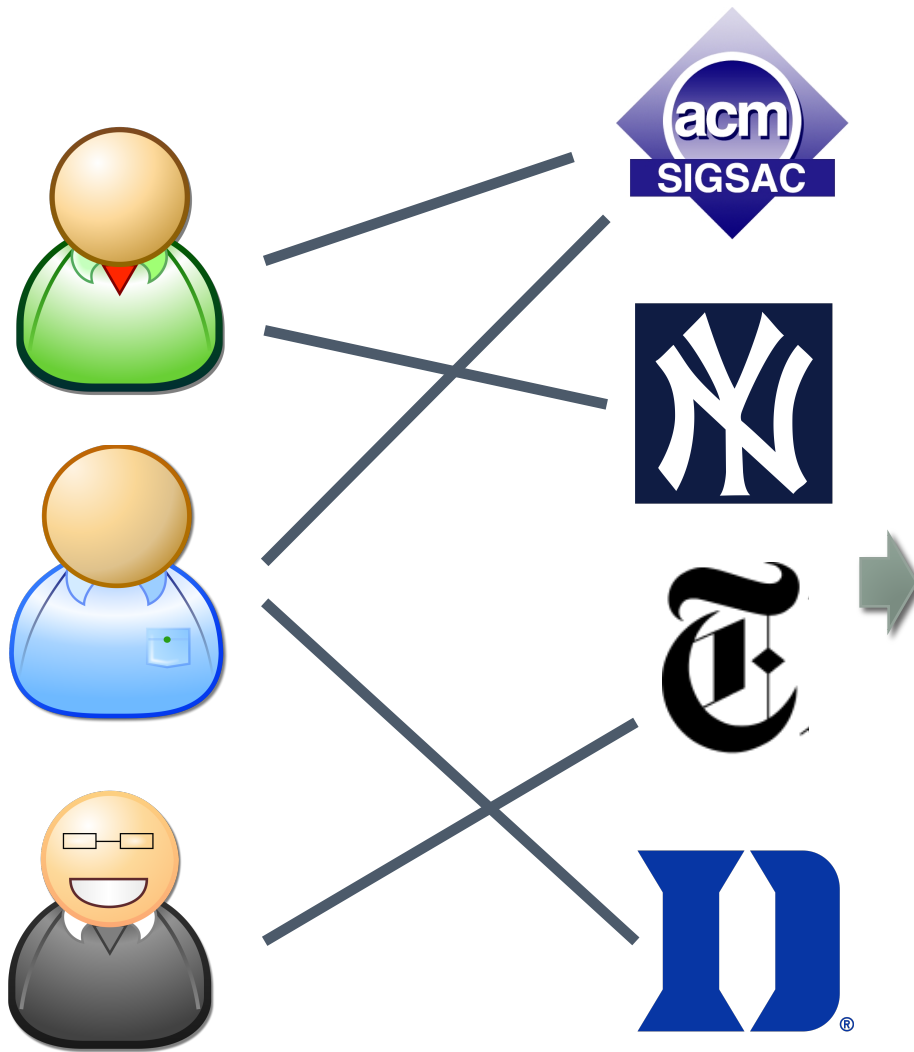


Matrix Modeling






Matrix Modeling



Matrix M





Matrix Modeling

HITS


Authoritativeness \vec{v} is first
eigenvector of $M^T M$

$$\vec{v} = c M^T M \vec{v}$$

Hubness \vec{u} is first
eigenvector of $M M^T$

$$\vec{u} = c M M^T \vec{u}$$

Matrix M





Matrix Modeling

HITS


Authoritativeness \vec{v} is first
eigenvector of $M^T M$

$$\vec{v} = c M^T M \vec{v}$$

Hubness \vec{u} is first
eigenvector of $M M^T$

$$\vec{u} = c M M^T \vec{u}$$

Matrix M



What about the other eigenvectors?



Matrix Modeling

Singular Value Decomposition

$U \Sigma V^T \approx M$


U	Σ	V
	1	0
	0	1
	0	0
	1	0

1	0	0	0
0	1	0	0
0	0	1	0
1	0	0	1



Matrix Modeling

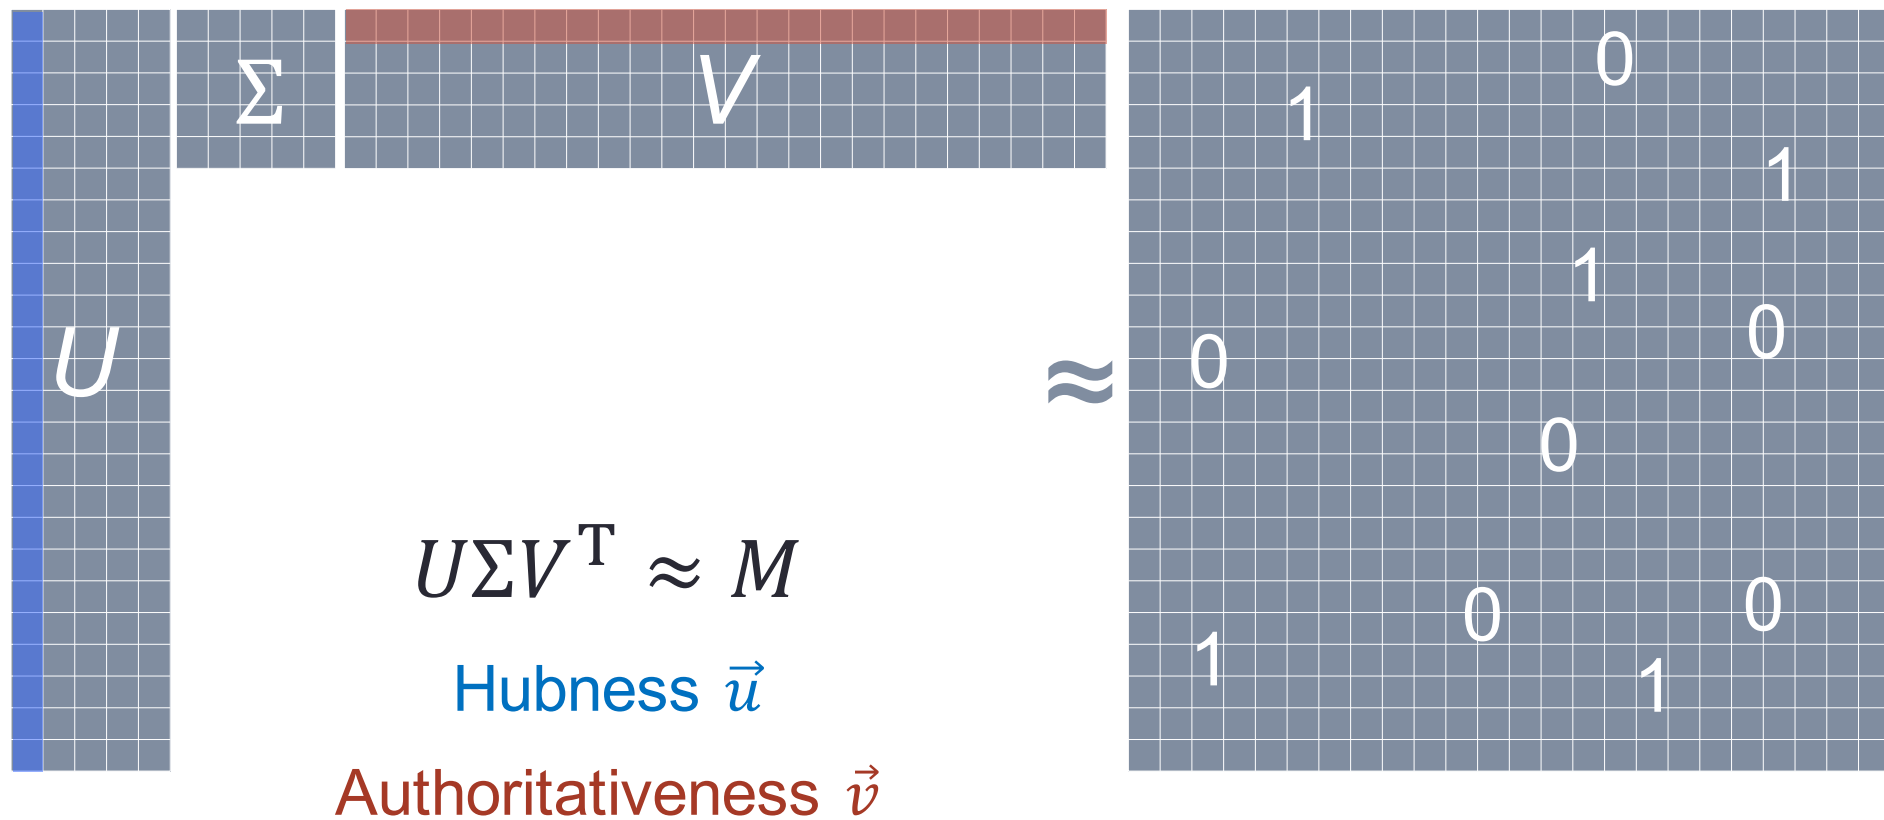
Singular Value Decomposition





Matrix Modeling

Singular Value Decomposition





Σ contains normalization for \vec{u} and \vec{v}



Matrix Factorization

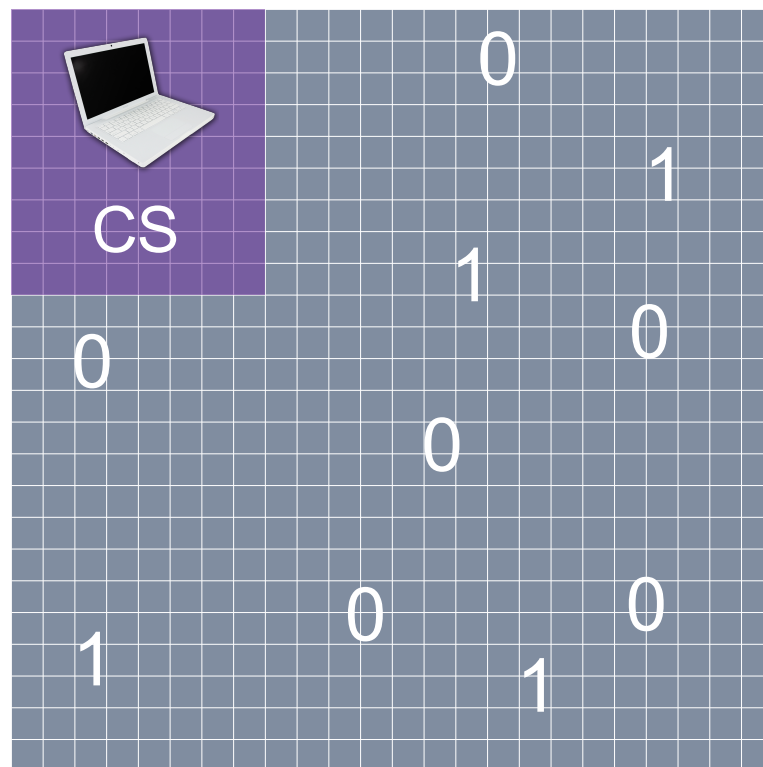
What does each eigenvector capture?



Topics

$$UV^T \approx M$$



Each factor captures a dense block in the matrix





Matrix Factorization


What does each eigenvector capture?



Topics


$$UV^T \approx M$$

Each factor captures a dense block in the matrix






Matrix Factorization





Matrix Factorization



Topics

.5	.1	-.1	.3	.9	•	.6	.2	.1	0	.8	≈	1
----	----	-----	----	----	---	----	----	----	---	----	---	---





1. Subgraph Analysis

2. Propagation Methods

3. Latent Factor Models

a) Background

b) Normal Behavior

c) Abnormal Behavior



Matrix Completion

Can't find singular vectors with missing entries. Instead,

$$\min_{U,V} \sum_{(i,j) \in M} (M_{i,j} - \vec{u}_i \cdot \vec{v}_j)^2$$



Matrix Completion

Can't find singular vectors with missing entries. Instead,

$$\min_{U,V} \sum_{(i,j) \in M} (M_{i,j} - \vec{u}_i \cdot \vec{v}_j)^2$$

Genres



Matrix Completion

Can't find singular vectors with missing entries. Instead,

$$\min_{U,V} \sum_{(i,j) \in M} (M_{i,j} - \vec{u}_i \cdot \vec{v}_j)^2$$

Genres

1.2 -1 .5 .8 -.5 • .6 .8 0 0 .1 ≈ 1

♥ ♪ 👻



Matrix Completion

Can't find singular vectors with missing entries. Instead,



$$\min_{U, V} \sum_{(i,j) \in M} (M_{i,j} - \hat{M}_{i,j})^2$$

$$\hat{M}_{i,j} = \vec{u}_i \cdot \vec{v}_j$$

Genres



Adding Latent Factors




Consider additional factors:

- Dataset mean μ
- Row (user) baseline $b_i \approx$
- Column (movie) baseline b_j

$$\hat{M}_{i,j} = \mu + b_i + b_j + \vec{u}_i \cdot \vec{v}_j$$

$$\min_{U,V} \sum_{(i,j) \in M} (M_{i,j} - \hat{M}_{i,j})^2$$





Adding Latent Factors

What if we know the **time** of the rating
(time of the edge being created)?

Collaborative Filtering with Temporal Dynamics

Yehuda Koren


KDD 2009





Adding Latent Factors


Mean Rating by Date (Netflix)





Adding Latent Factors

Mean Rating by Movie Age (Netflix)



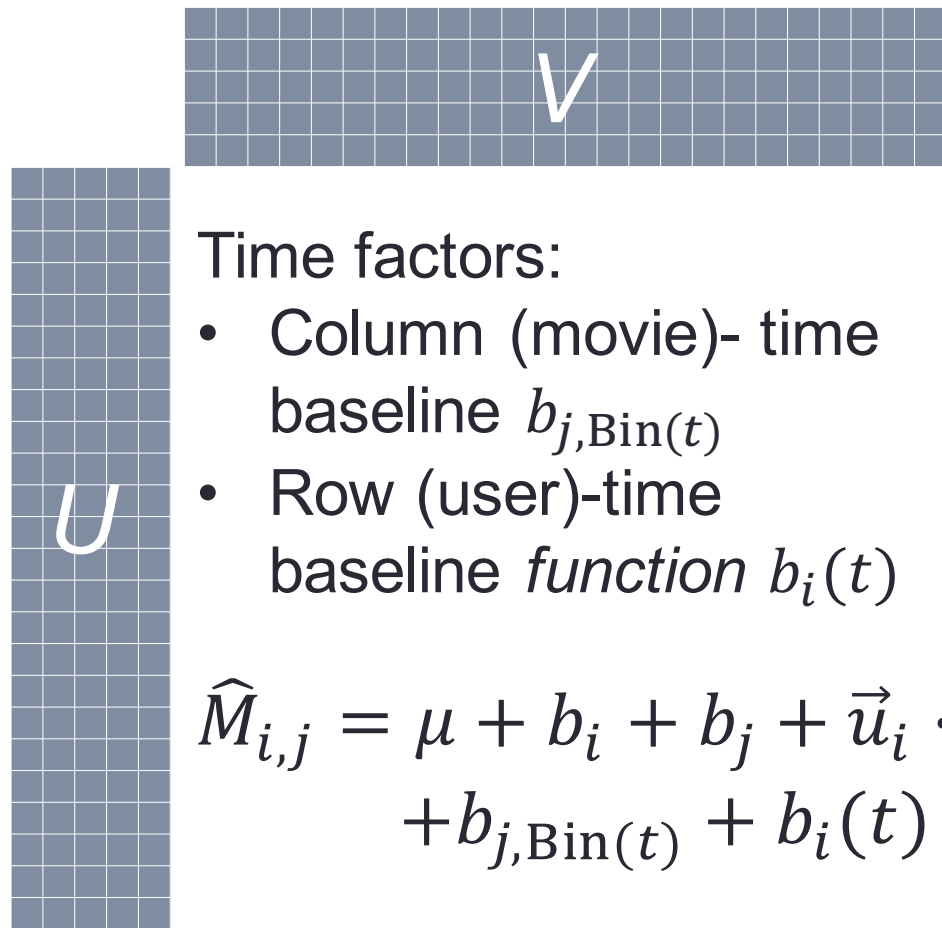
Collaborative Filtering with Temporal Dynamics

Yehuda Koren

KDD 2009



Adding Latent Factors




Time factors:

- Column (movie)- time baseline $b_{j, \text{Bin}(t)}$
- Row (user)-time baseline *function* $b_i(t)$

$$\hat{M}_{i,j} = \mu + b_i + b_j + \vec{u}_i \cdot \vec{v}_j + b_{j, \text{Bin}(t)} + b_i(t)$$








$$\min_{U,V} \sum_{(i,j) \in M} (M_{i,j} - \hat{M}_{i,j})^2$$

≈






Bayesian Modeling











Sample user factors from Normal distribution





Bayesian Modeling








Sample user factors from Normal distribution

Update mean based on user factors



Bayesian Modeling




~

Similarly sample movie factors



Bayesian Modeling



$$p(M_{i,j} | U, V) = \mathcal{N}(M_{i,j} | \vec{u}_i \cdot \vec{v}_j, \sigma^2)$$


Bayesian Probabilistic Matrix Factorization

Ruslan Salakhutdinov and Andriy Mnih

ICML 2008




Bayesian Modeling






Bayesian Modeling with Co-Clustering




Cluster users
with similar factors







Bayesian Modeling with Co-Clustering



Cluster 28	Cluster 30	Cluster 48
Simpsons	Scooby Doo	Star Trek
Family Guy	Spy Kids	Back to the Future
Monty Python	Stuart Little	Southpark
Curb your Enthusiasm	Dr. Dolittle	Lord of the Rings
The Twilight Zone	Lion King	Harry Potter
Arrested Development	Agent Cody Banks	The X-Files



Online Rating Models



Newegg (Hard Drive):

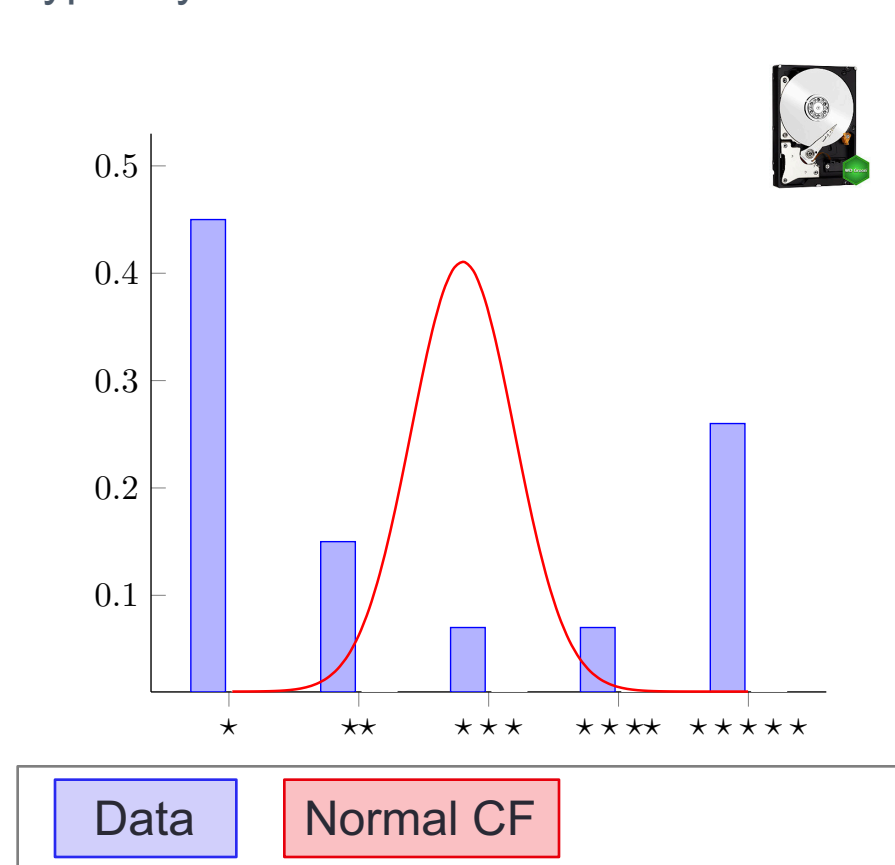
50%	(298)
14%	(83)
6%	(38)
5%	(31)
25%	(152)

App Store (5-star rating):

Yelp (Bulluck's Bar-B-Cue):

5 stars	22
4 stars	39
3 stars	15
2 stars	12
1 star	6

Typically fit a Gaussian - Minimize RMSE





Online Rating Models



Newegg rating distribution:

5 stars	50%	(298)
4 stars	14%	(83)
3 stars	6%	(38)
2 stars	5%	(31)
1 star	25%	(152)

Apple rating distribution:

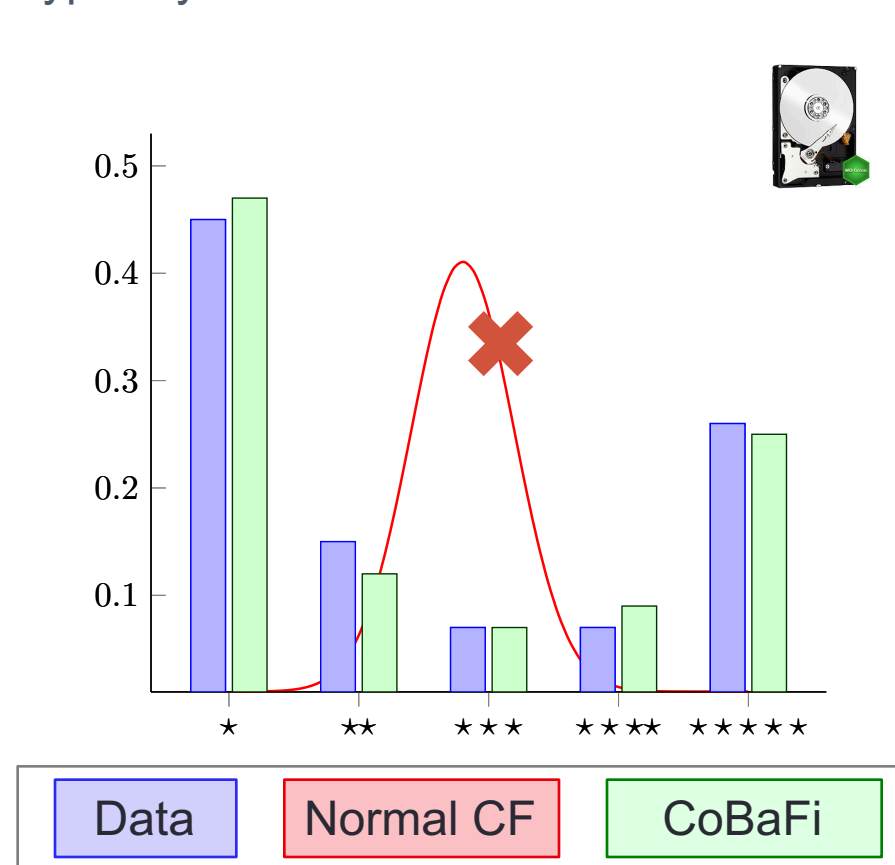
5 stars	~45%
4 stars	~35%
3 stars	~15%
2 stars	~5%
1 star	~0%

Yelp rating distribution:

5 stars	22
4 stars	39
3 stars	15
2 stars	12
1 star	6

Bulluck's Bar-B-Cue logo: SOUTHERN COOKING AND CATERING


Typically fit a Gaussian - Minimize RMSE




CoBaFi: Collaborative Bayesian Filtering
 Alex Beutel, Kenton Murray,
 Christos Faloutsos Alex Smola
 WWW 2014




Shape of Netflix reviews



Most Gaussian	Most skewed
The Rookie	The O.C. Season 2
The Fan	Samurai X: Trust and Betrayal
Cadet Kelly	Aqua Teen Hunger Force: Vol. 2
Money Train	Sealab 2001: Season 1
Alice Doesn't Live Here	Aqua Teen Hunger Force: Vol. 2
Sea of Love	Gilmore Girls: Season 3
Boiling Point	Felicity: Season 4




Stars



Stars

Movies



TV Shows

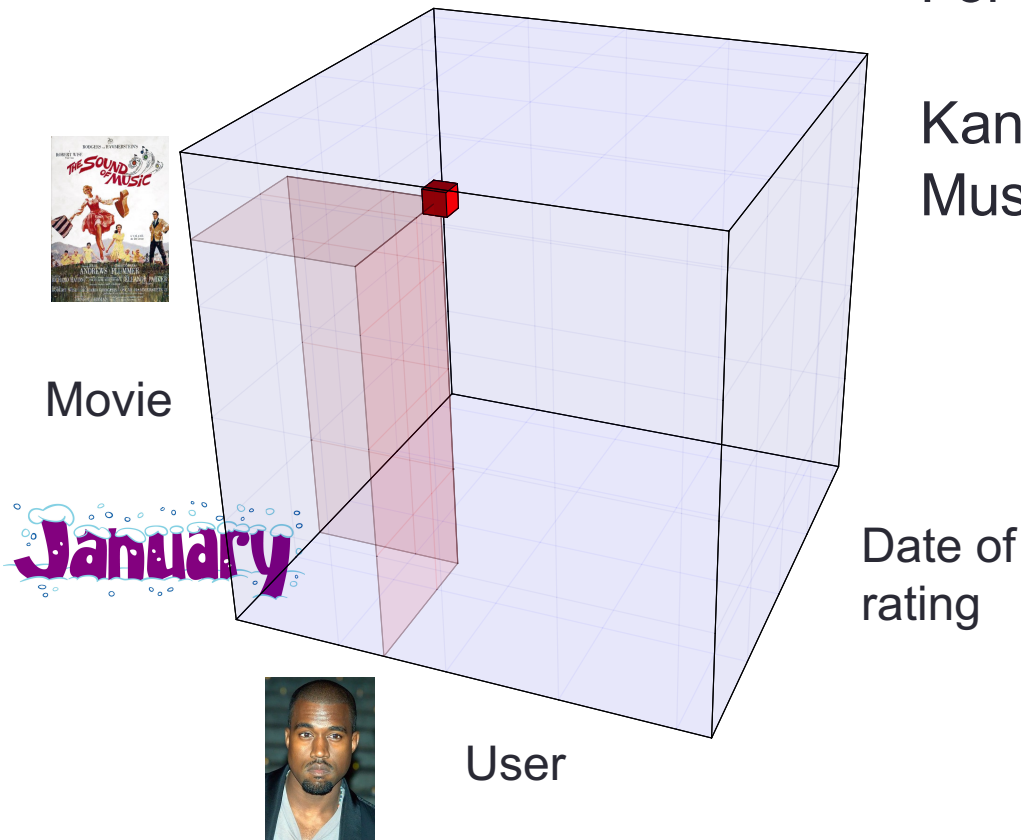


What is a tensor?

- Tensors are used for structured data > 2 dimensions
- Think of as a 3D-matrix

For example:


Kanye West rated The Sound of Music five stars last January.





Tensor Decomposition

Kanye West rated The Sound of Music five stars last January.



$$X \approx U \otimes V \otimes W$$

$$X_{i,j,k} \approx \sum_{r=1}^{\text{Rank}} U_{i,r} V_{j,r} W_{k,r}$$




Graph Clustering with Tensors

Multiple possible views
of the DBLP network:

1. Who-cites-whom
2. Co-authorship
3. Using same words in title

Graph View

Author



Author


Do more Views of a Graph help? Community Detection and Clustering in Multi-Graphs
Evangelos E. Papalexakis, Leman Akoglu, Dino Ienco

FUSION 2013





Graph Clustering with Tensors




Sparse Tensor Factorization




Graph Clustering with Tensors


DBLP-1



(a) citation




(b) co-auth.




(c) co-term


DBLP-2



(a) citation



(b) co-auth.



(c) co-term




Graph Clustering with Tensors

Dataset	Baseline	GraphFuse
DBLP-1	0.12	0.30
DBLP-2	0.08	0.12

Modeling Accuracy

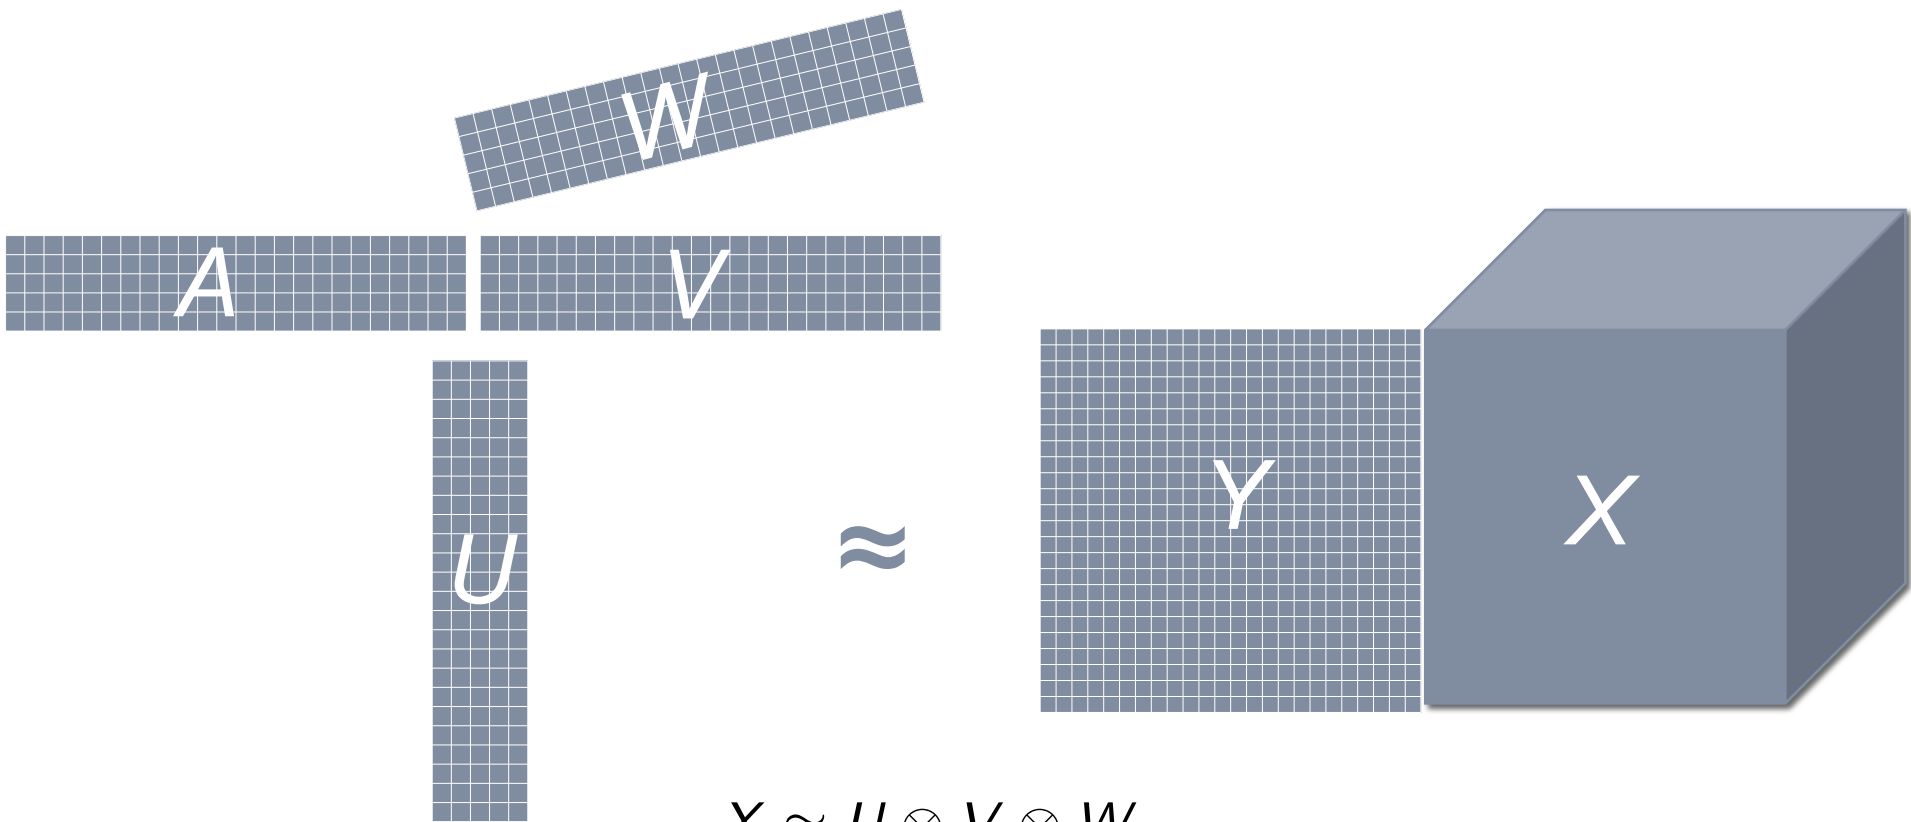


Coupled Matrix + Tensor Decomposition





Coupled Matrix + Tensor Decomposition



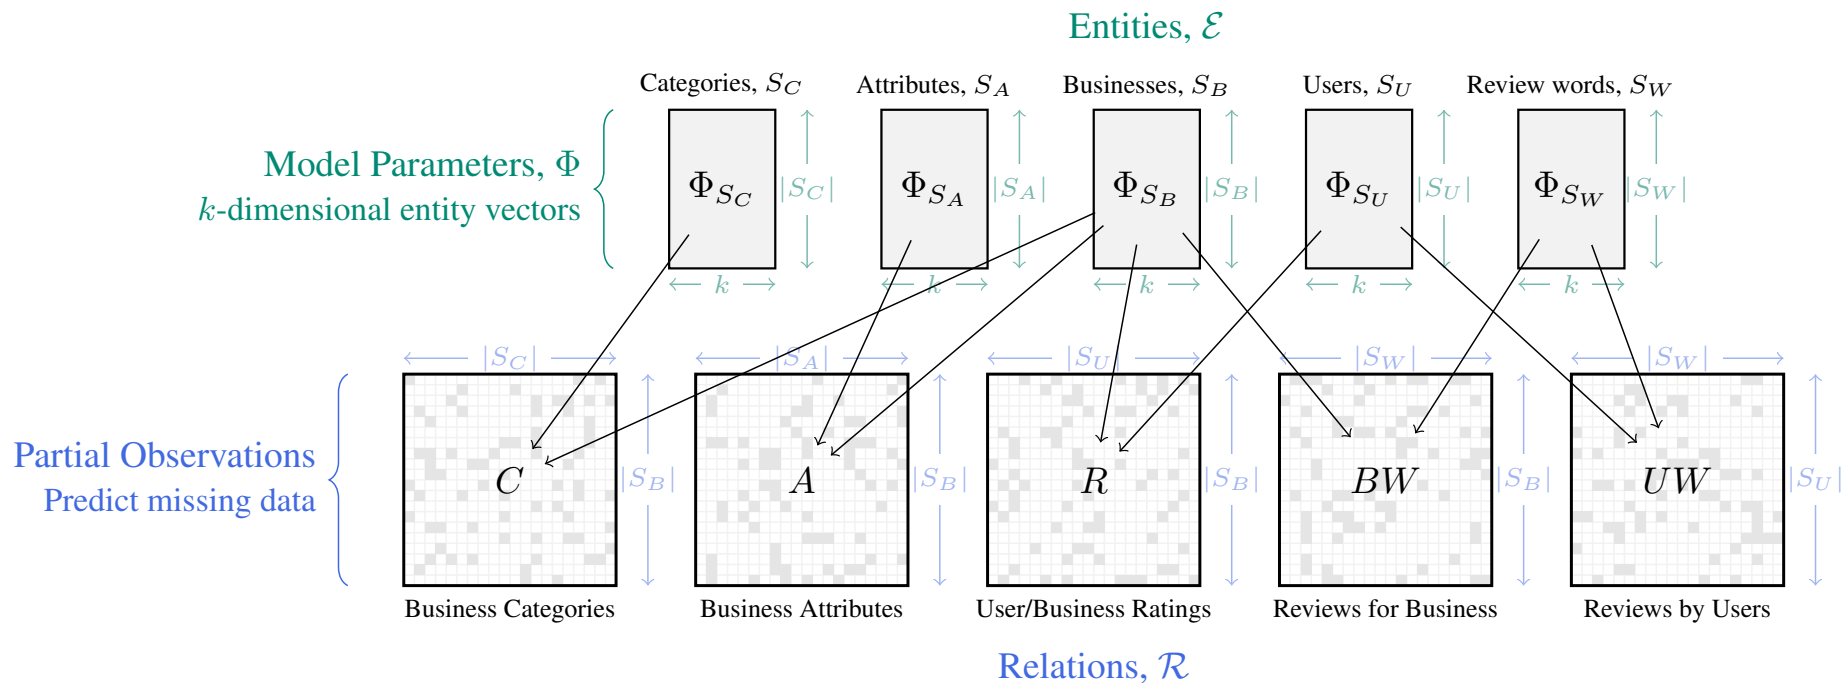
$$X \approx U \otimes V \otimes W$$

$$Y \approx UA^T$$

$$\min_{U, V, W, A} \|X - U \otimes V \otimes W\|_F^2 + \|Y - UV^T\|_F^2$$



Joint Factorization




Collective Factorization for Relational Data:
 An Evaluation on the Yelp Datasets
 Nitish Gupta, Sameer Singh





Joint Factorization

PR Curve (Ratings)



Most valuable:

1. Ratings
2. Review text
3. Business Categories



Collective Factorization for Relational Data:
 An Evaluation on the Yelp Datasets
 Nitish Gupta, Sameer Singh



1. Subgraph Analysis

2. Propagation Methods

3. Latent Factor Models

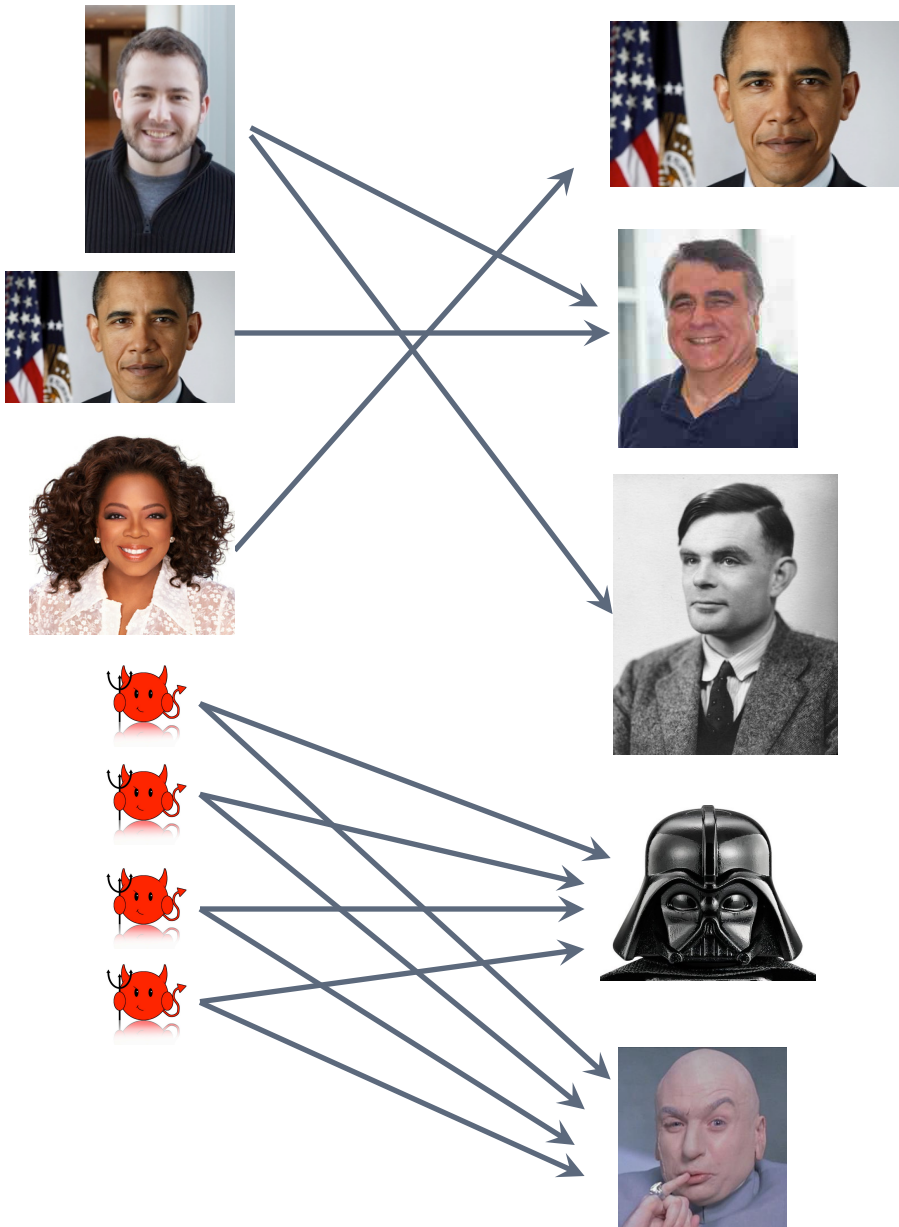
a) Background

b) Normal Behavior

c) Abnormal Behavior

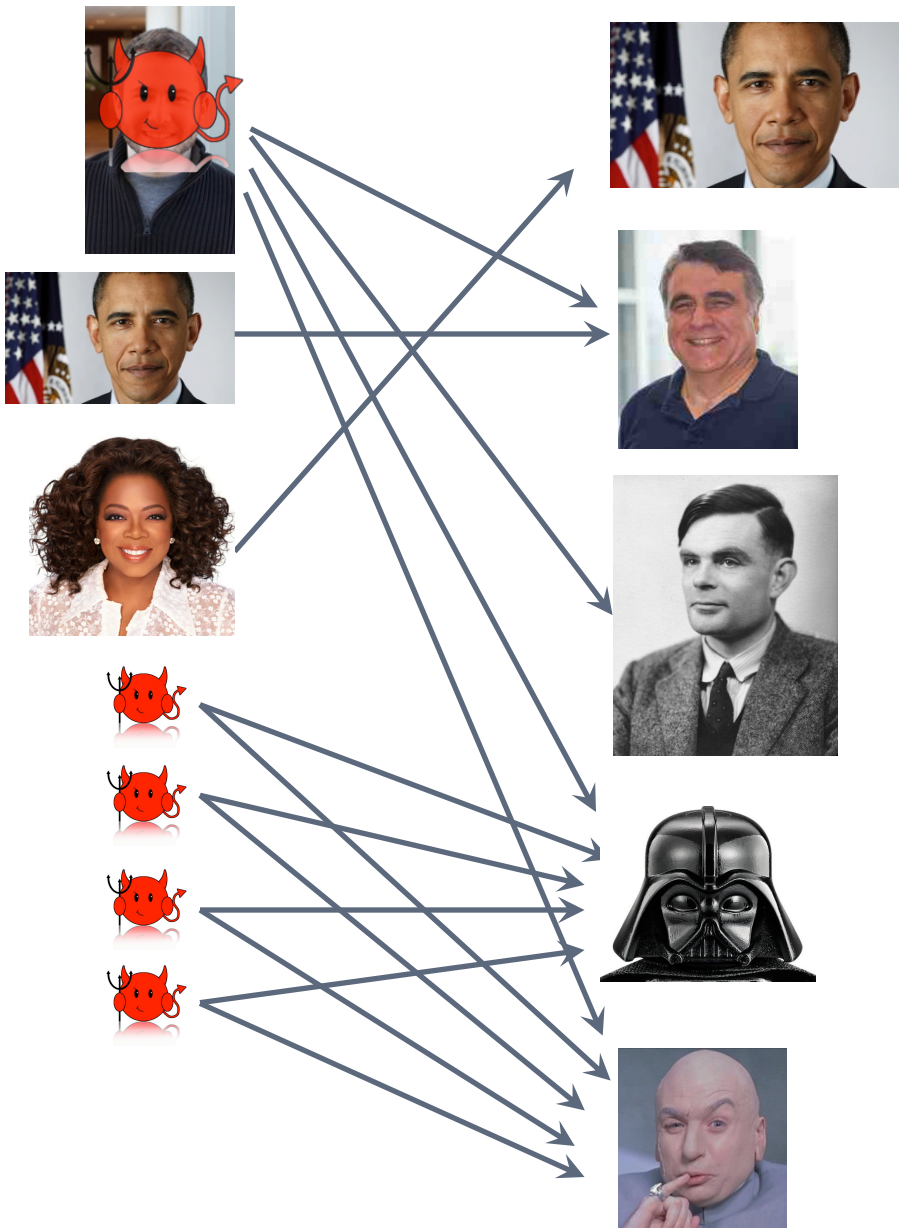


Fraud Detection



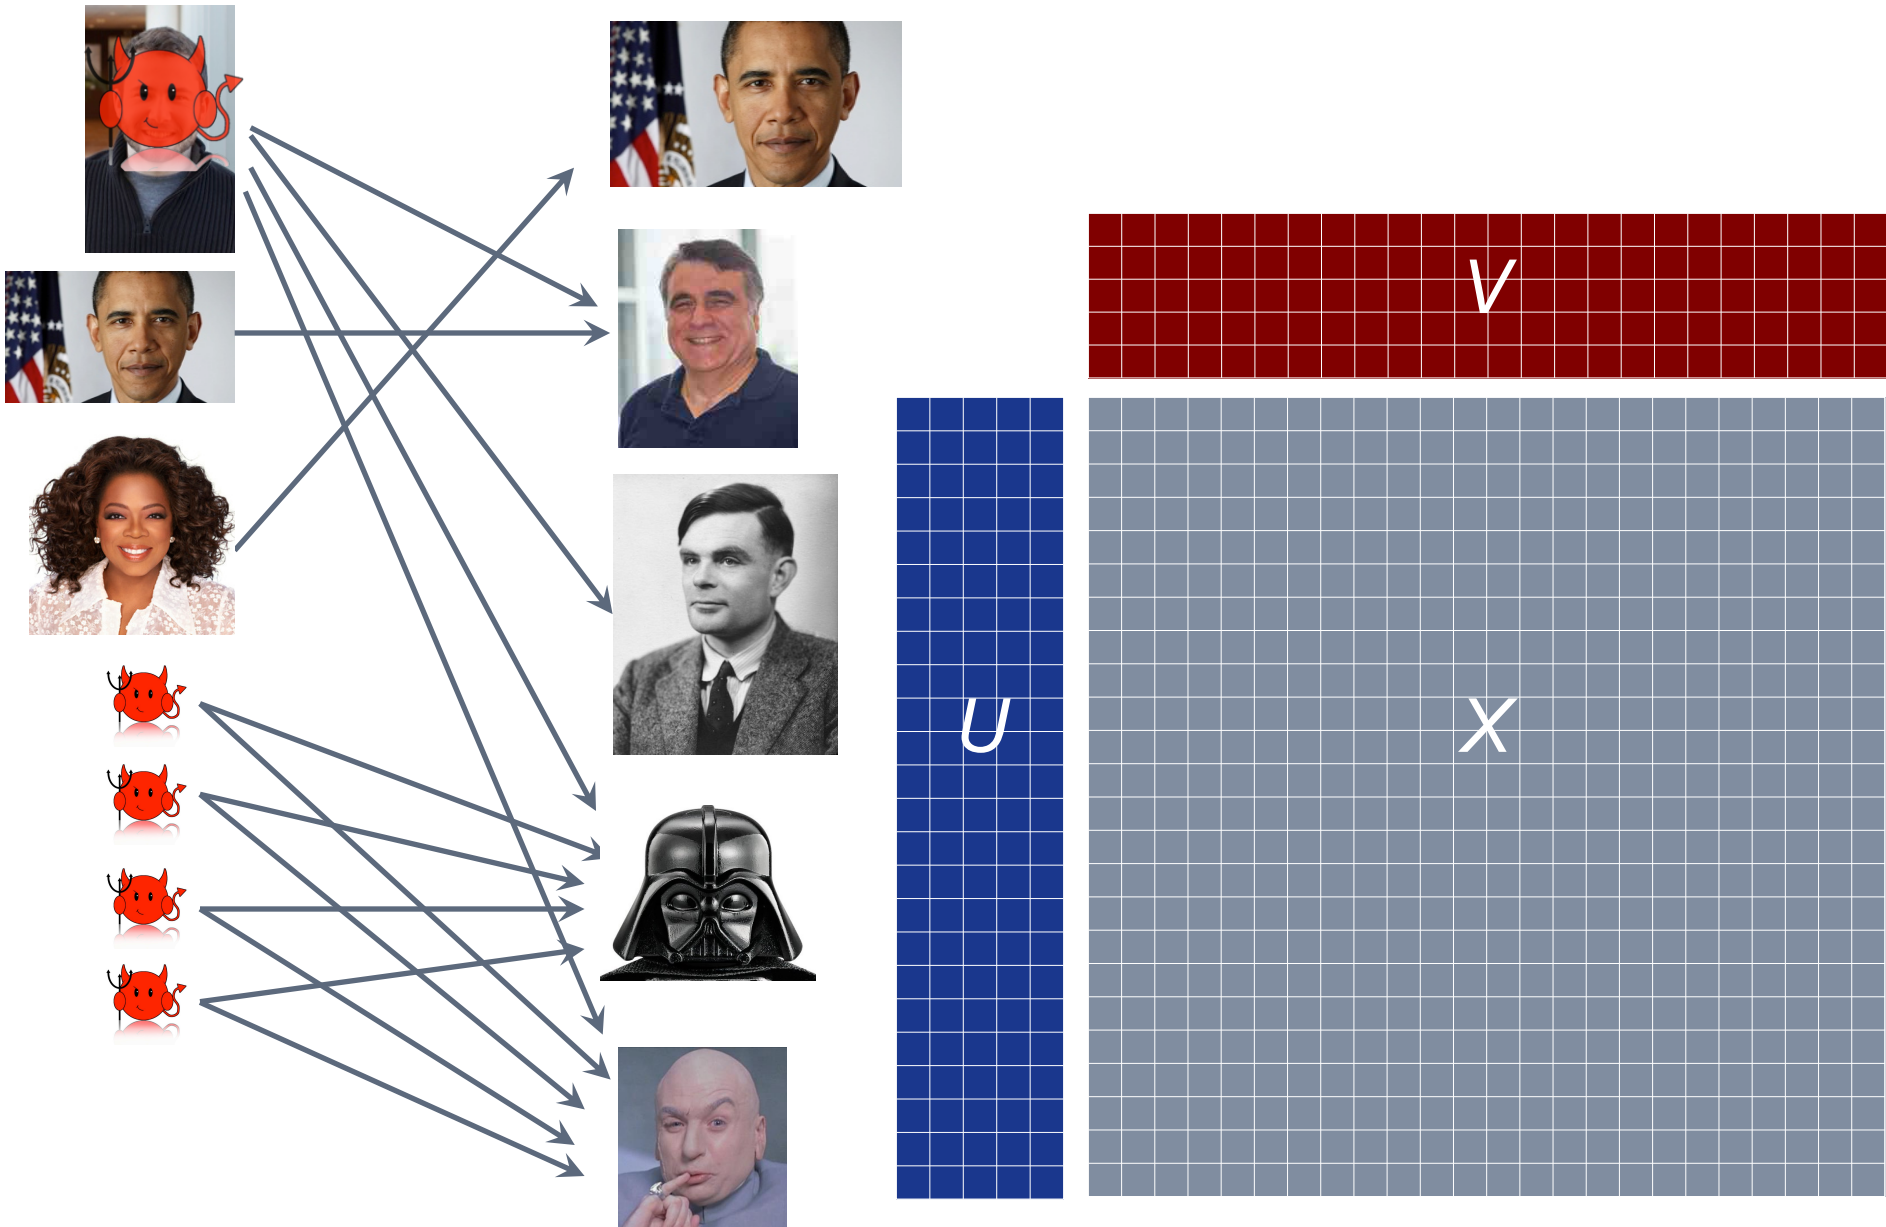


Fraud Detection






Fraud within a factorization





Fraud within a factorization





Fraud within a factorization



1.5	1	-0.5	-2	1
-----	---	------	----	---

Followers

U



?	?	?	?	?
---	---	---	---	---



?	?	?	?	?
---	---	---	---	---





?	?	?	?	?
---	---	---	---	---

Followees


V

X



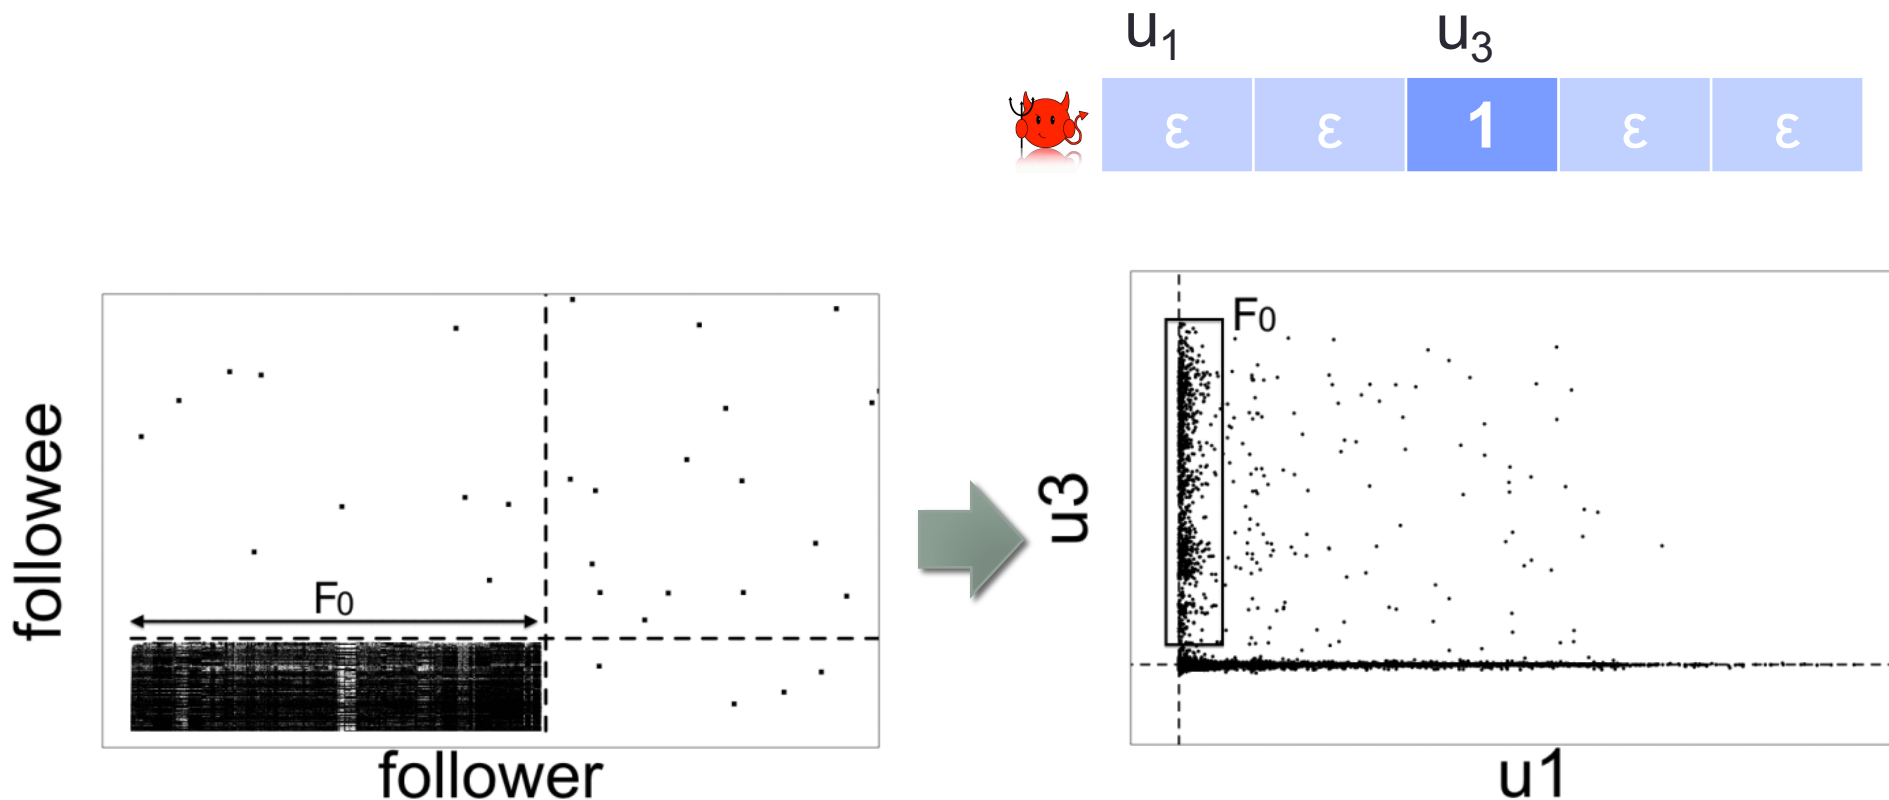


Fraud within a factorization







Fraud within a factorization





Fraud within a factorization



Inferring Strange Behavior from Connectivity Pattern in Social Networks

Meng Jiang, Peng Cui, Alex Beutel,
Christos Faloutsos, Shiqiang Yang.

PAKDD, 2014








Fraud within a factorization



0.1	ϵ	1	ϵ	ϵ
-----	------------	---	------------	------------





Inferring Strange Behavior from Connectivity Pattern in Social Networks

Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang.


PAKDD, 2014



Fraud within a factorization



Username: a#####
 Birthday: January 1st



Inferring Strange Behavior from Connectivity Pattern in Social Networks

Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, Shiqiang Yang.


PAKDD, 2014



Complementary Fraud Detection

Followees

Limited r



Followers

U

X




Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective

Neil Shah, Alex Beutel, Brian Gallagher,

Christos Faloutsos



ICDM, 2014.





Complementary Fraud Detection

960 fraudsters safely following 960 customers






Complementary Fraud Detection

Followees

Limited r



Followers



11111111

Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective


Neil Shah, Alex Beutel, Brian Gallagher,

Christos Faloutsos

ICDM, 2014.



Complementary Fraud Detection



Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective



Neil Shah, Alex Beutel, Brian Gallagher,

Christos Faloutsos

ICDM, 2014.



Complementary Fraud Detection



2 TWEETS | 1,054 FOLLOWING | 88 FOLLOWERS

[Follow Laurie House](#)

- Laurie House** @nothingfor14137 2 Jul 09
Cool! video: <http://twurl.nl/aakgla>
- Laurie House** @nothingfor14137 23 Jun 09
an Easy Way to Make Money- Fully Automated go <http://bit.ly/lxIM5X> :)hahaha i ment

- sungard55** @sungard55
- sungard54** @sungard54
- sungard53** @sungard53
- sungard52** @sungard52



Lekan Olawole Lowe @loweinc 26 Jul 09
Sign up free and Get 400 followers a day using <http://tweeteradder.com>




Lekan Olawole Lowe @loweinc 26 Jul 09
Get 400 followers a day using <http://www.tweeterfollow.com>



Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective
Neil Shah, Alex Beutel, Brian Gallagher, Christos Faloutsos
ICDM, 2014.



Complementary Fraud Detection



2 TWEETS | 1,054 FOLLOWING | 88 FOLLOWERS

[Follow Laurie House](#)

- Laurie House** @nothingfor14137 2 Jul 09
Cool! video: <http://twurl.nl/aakgla>
- Laurie House** @nothingfor14137 23 Jun 09
an Easy Way to Make Money- Fully Automated go <http://bit.ly/lxIM5X> :)hahaha i ment

93% Precision
70% of accounts missed by Twitter

- sungard55** @sungard54
Lekan Olawole Lowe @loweinc 26 Jul 09
Sign up free and Get 400 followers a day using <http://tweeteradder.com>
- Lekan Olawole Lowe** @loweinc 26 Jul 09
Get 400 followers a day using <http://www.tweeterfollow.com>



Spotting Suspicious Link Behavior with fBox: An Adversarial Perspective
 Neil Shah, Alex Beutel, Brian Gallagher,
 Christos Faloutsos
 ICDM, 2014.








Practitioner's Guide

Method	Graph Type	Node Attributes	Edge Attributes	Seed Labels
EigenSpokes	Directed+			
Get-the-Scoop	Directed+			
fBox	Directed+			
CoBaFi	Bipartite+		✓	
CDOutliers	Undirected	✓		



Detecting Fraud within Recommendation

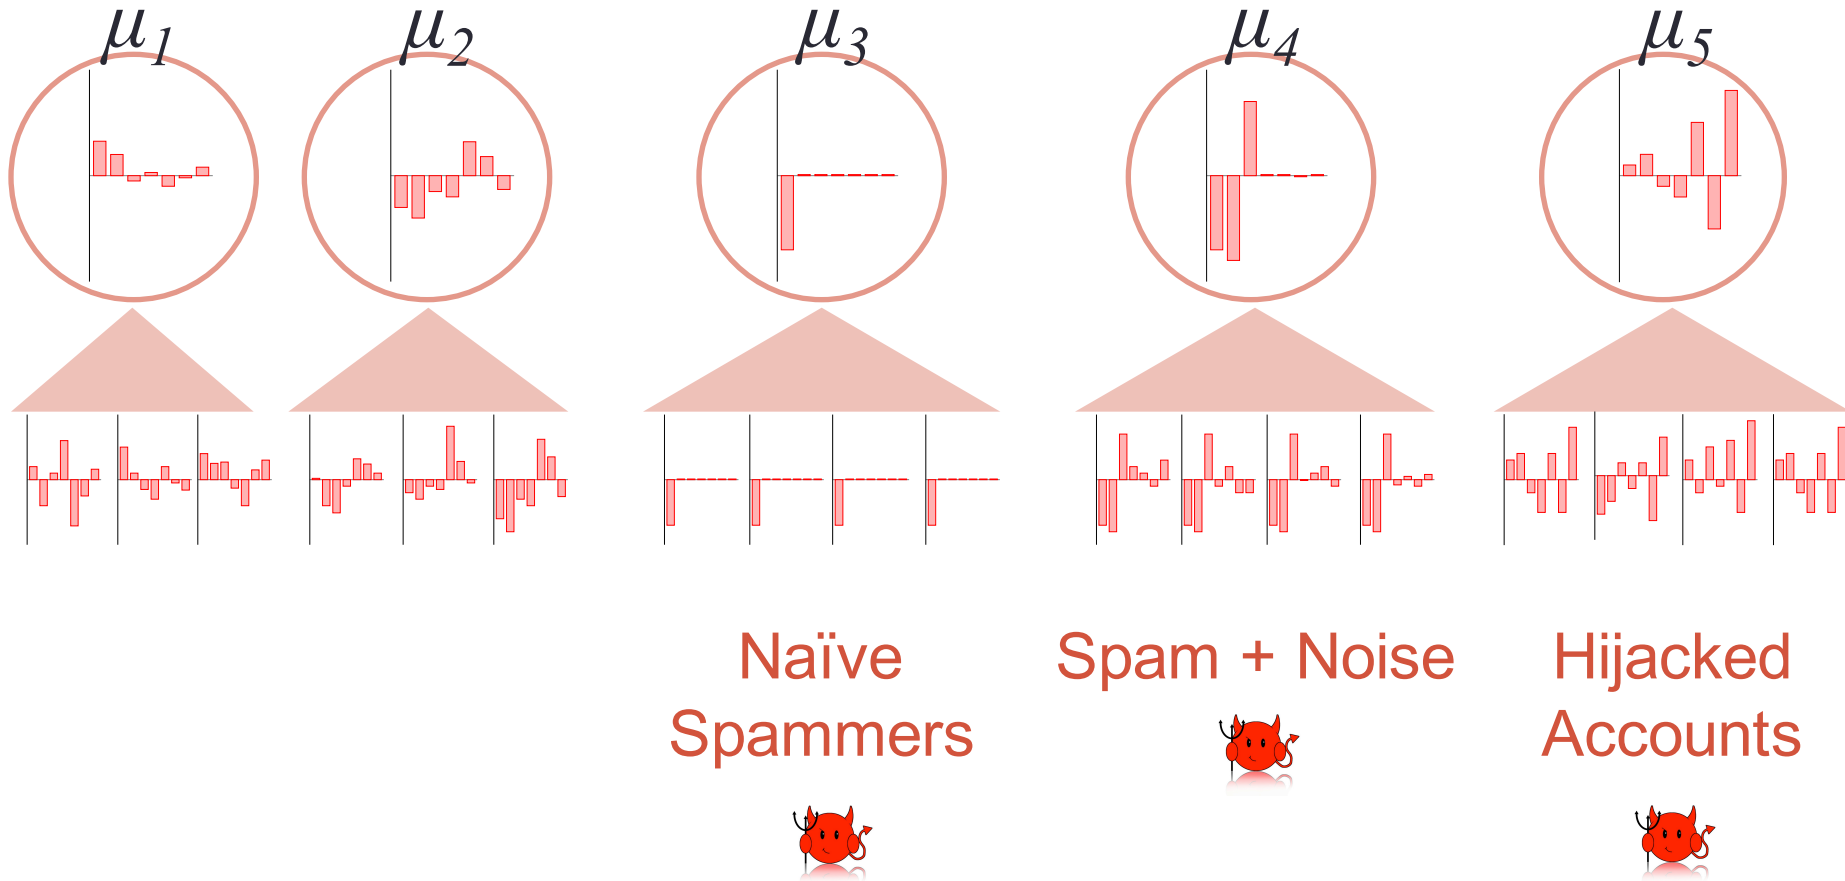
	1.4	0.8	-0.2	-1.5	0.6
	1.5	1	-0.5	-2	1
	?	?	?	?	?
	?	?	?	?	?
	?	?	?	?	?

U

<i>V</i>									
<i>X</i>									




Clustering Fraudsters






Clustered Fraudsters


Clustered naïve spammers



Clustered hijacked accounts




Clustered “attacked” movies



83% are clustered together



Outliers in Joint Factorization




Enforce $U_1 \approx U_2$ and $U_1, U_2, V_1, V_2 \geq 0$

Community Distribution Outlier Detection in Heterogeneous Information Networks

Manish Gupta, Jing Gao, and Jiawei Han

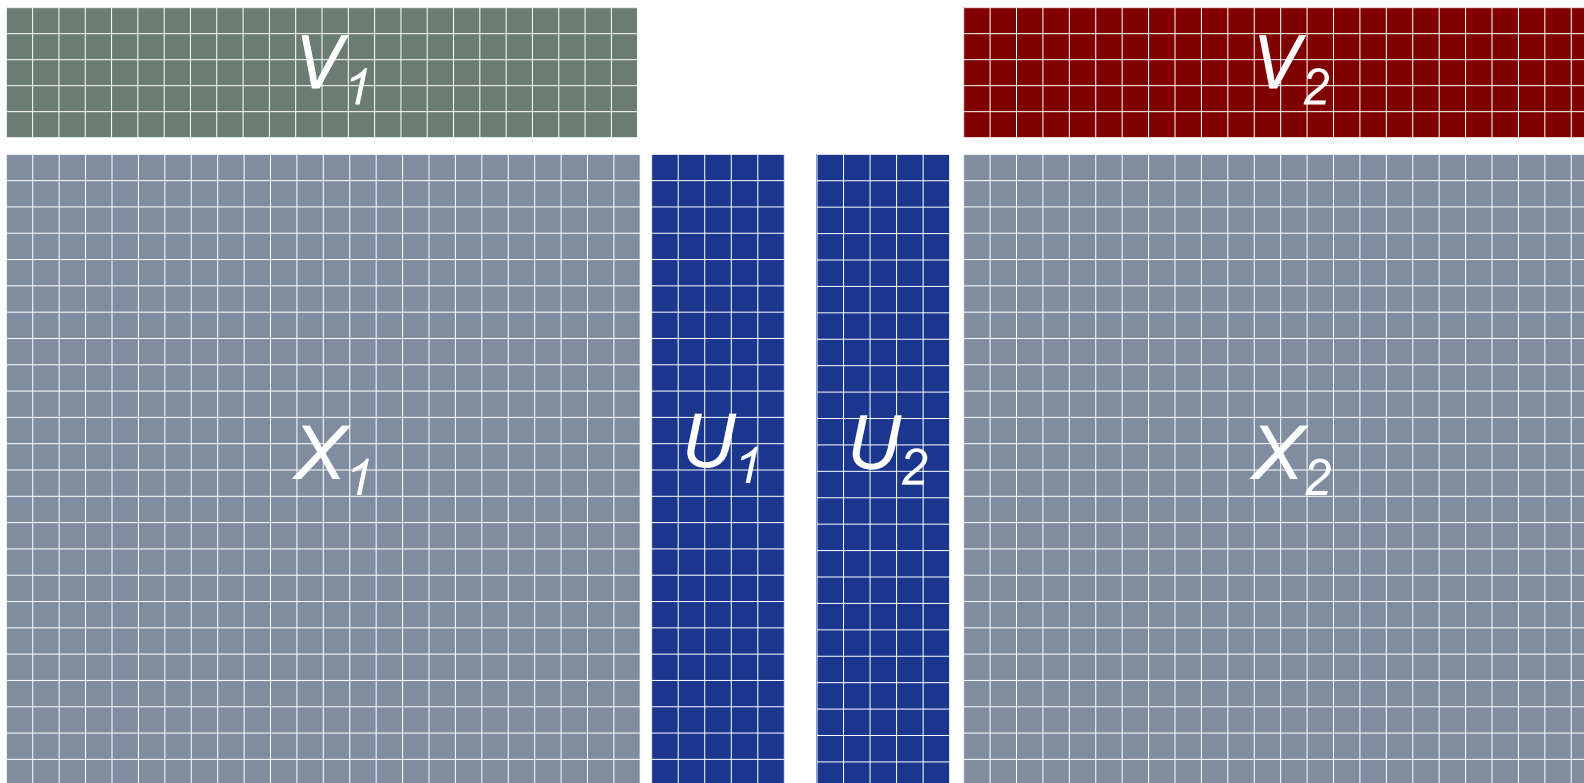
ECML/PKDD 2013





Outliers in Joint Factorization

Interesting design of X_1 and X_2 ; see paper for details




Enforce $U_1 \approx U_2$ and $U_1, U_2, V_1, V_2 \geq 0$



Outliers in Joint Factorization

Rows of V_2 represent common patterns in X_2 (cluster centroids)





Enforce $U_1 \approx U_2$ and $U_1, U_2, V_1, V_2 \geq 0$



Outliers in Joint Factorization

An anomaly is a row of X_i that is *not* similar to any row in V_i

Rows of V_2 represent common patterns in X_2 (cluster centroids)






Practitioner's Guide

Method	Graph Type	Node Attributes	Edge Attributes	Seed Labels
EigenSpokes	Directed+			
Get-the-Scoop	Directed+			
fBox	Directed+			
CoBaFi	Bipartite+		✓	
CDOutliers	Undirected	✓		



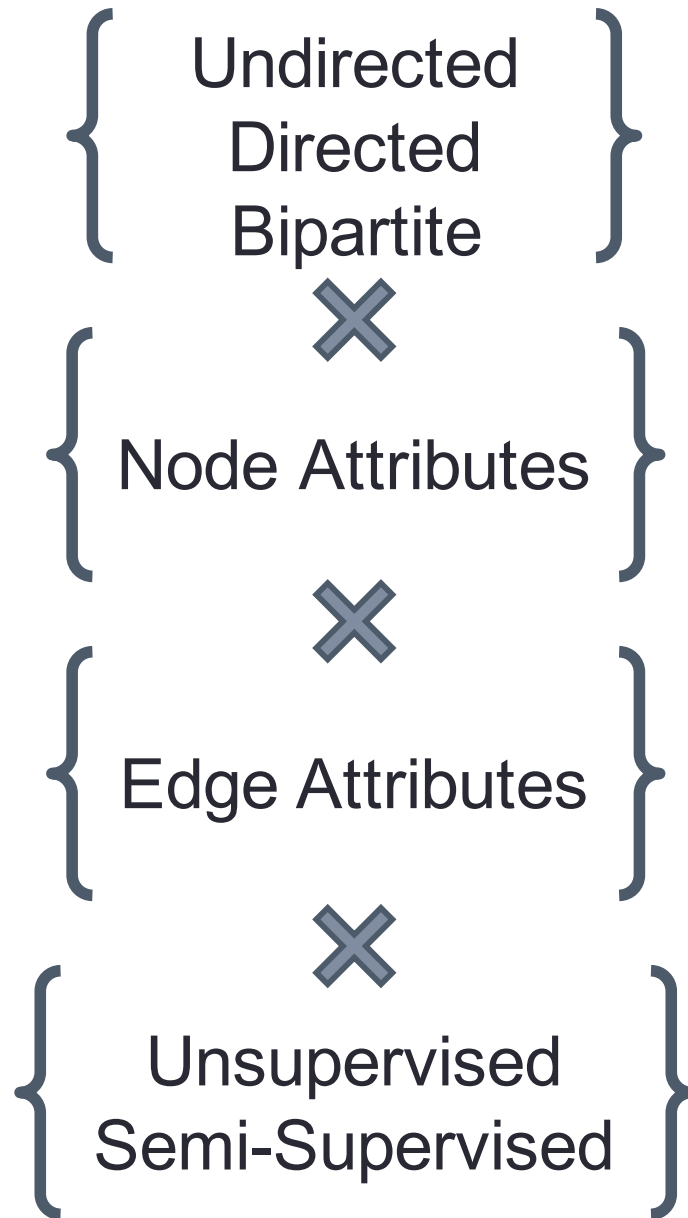
Recap

- SVD captures communities of interest
- Bayesian methods can:
 - Handle missing values
 - Give factorization models (-> patterns, & anomalies)
- Group-outliers: spotted by CoBaFi, Get-the-Scoop, etc.







CONCLUSION





Open Problems / Opportunities

 **P1. Complex data:** How should we integrate data from multiple data sources?






Open Problems / Opportunities



P2. Adversarial analysis: Can we offer provable guarantees on detecting fraud and spam?






Open Problems / Opportunities




P3. Early detection: Can we detect fraudsters before they cause significant damage?





Summary

Local Subgraph Analysis: Patterns and Features e.g. using ego-nets







Summary

Propagation Methods

“Guilt-by-association”

“Importance-by-association” = PageRank


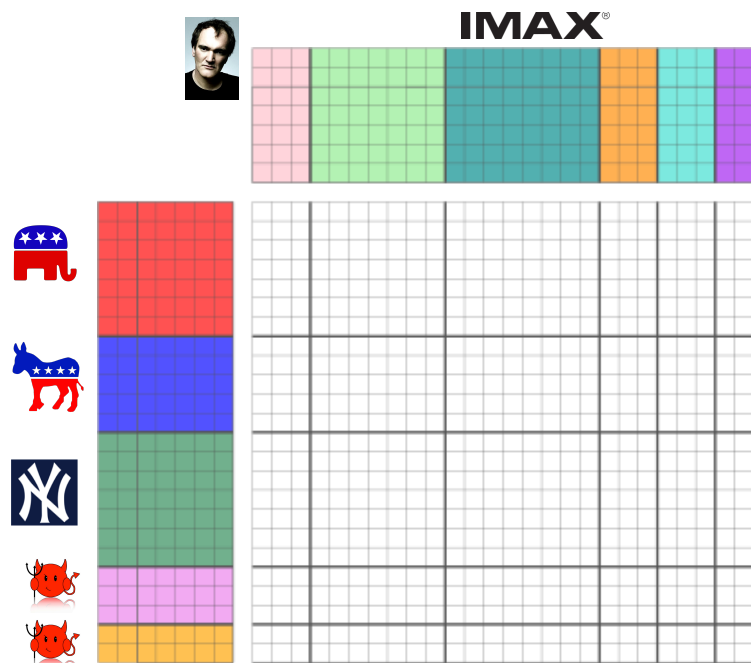




Summary

Latent Factor Models



Find multiple communities, patterns and anomalies.





Take Away

User Modeling and Fraud Detection
are two sides of the same coin.





Thanks again to



NSF Grant No. IIS-1408924, IIS-1408287,
CAREER 1452425, DGE-1252522, ...



Questions?



Carnegie Mellon University



Stony Brook University

References and resources available at cs.cmu.edu/~abeutel/ccs2015_tutorial

CCS 2015 A. Beutel, L. Akoglu, C. Faloutsos 46

Pattern: Ego-net Power Law Density

POSTNET

<http://www.size-more.co.uk/2005/08/i-feel-some-movies-coming-on.html>

<http://mstapundit.com/archives/025235.php>

— $1.1094x + (-0.21414) = y$
 - - - $1.1054x + (-0.21432) = y$
 — $2.1054x + (-0.51535) = y$

Oddball: Spotting anomalies in weighted graphs
 Leman Akoglu, Mary McGlohon, Christos Faloutsos
 PAKDD 2010

CCS 2015 A. Beutel, L. Akoglu, C. Faloutsos 25

Semi-supervised Classification

Given a graph and labels for some nodes, can we learn the labels for the other nodes?

CCS 2015 A. Beutel, L. Akoglu, C. Faloutsos 9

Matrix Factorization

What does each eigenvector capture?

$UV^T \approx M$

Each factor captures a dense block in the matrix