

## Case Study 4: Collaborative Filtering

### Graph-Parallel Problems

### Synchronous v. Asynchronous Computation

Machine Learning for Big Data  
CSE547/STAT548, University of Washington

Emily Fox

February 25<sup>th</sup>, 2014

©Emily Fox 2014

1

## Map-Reduce Abstraction

Map: *Transforms a data element*

- Data-parallel over elements, e.g., documents
- Generate (key,value) pairs
  - "value" can be any data type

*In this example:*  
(UW, 17)  
(Mary, 1)  
(UW, 1)  
(Mary, 1)

*Example: word count*  
map(document)  
for word in doc  
emit(word, 1)

Reduce: *Take all values associated w/ a key and aggregate*

- Aggregate values for each key
- Must be commutative-associate operation
- Data-parallel over keys
- Generate (key,value) pairs

*reduce('UW', [1, 17, 0, 0, 12])*  
*emit('UW', 30)*

*Reduce(word, count+list(int))*  
*c = 0*  
*for i in count*  
*c += count[i]*  
*emit(word, c)*

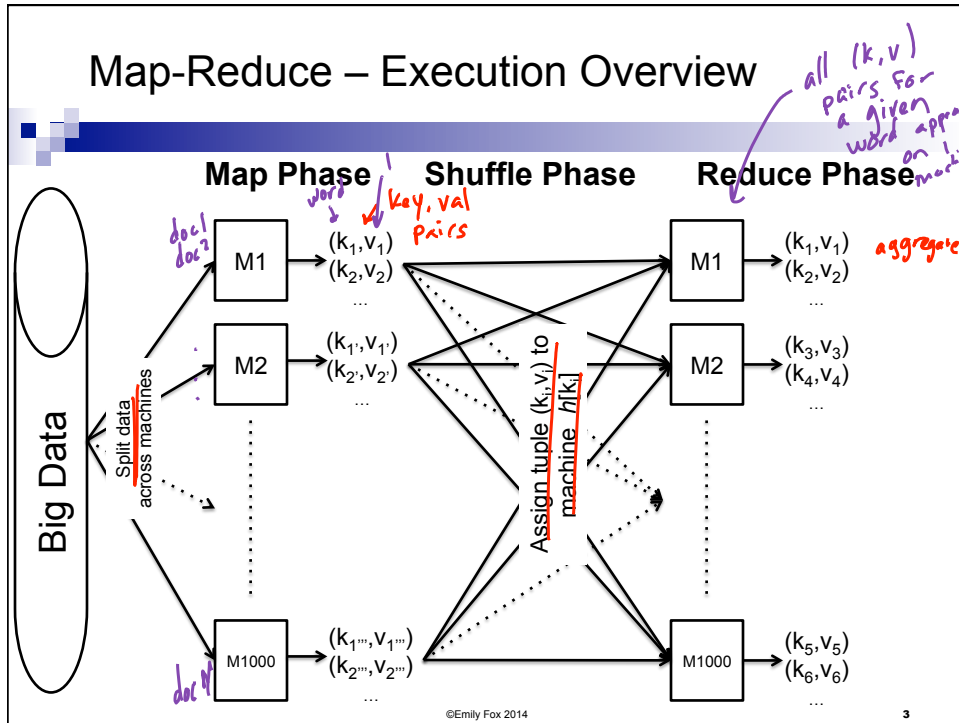
Map-Reduce has long history in functional programming

- But popularized by Google, and subsequently by open-source Hadoop implementation from Yahoo!

©Emily Fox 2014

2

## Map-Reduce – Execution Overview



## Issues with Map-Reduce Abstraction

- Often all data gets moved around cluster
  - Very bad for iterative settings ← in shuffle phase
- Definition of Map & Reduce functions can be unintuitive in many apps
  - Graphs are challenging ←
- Computation is synchronous ← MapReduce waits for all mappers to finish

# SGD for Matrix Factorization in Map-Reduce?

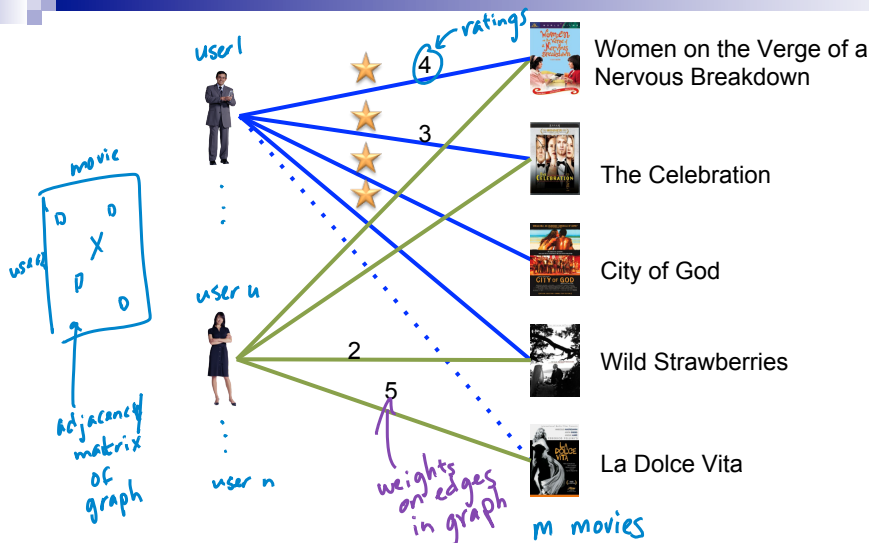
$$\begin{bmatrix} L_u^{(t+1)} \\ R_v^{(t+1)} \end{bmatrix} \leftarrow \begin{bmatrix} (1 - \eta_t \lambda_u) L_u^{(t)} - \eta_t \epsilon_t R_v^{(t)} \\ (1 - \eta_t \lambda_v) R_v^{(t)} - \eta_t \epsilon_t L_u^{(t)} \end{bmatrix}$$

$$\epsilon_t = L_u^{(t)} \cdot R_v^{(t)} - r_{uv}$$


- Map and Reduce functions???
- Map-Reduce:
  - Data-parallel over all mappers
  - Data-parallel over reducers with same key
- Here, one update at a time!

does not nicely fit into data parallel setting

# Matrix Factorization as a Graph



# Flashback to 1998



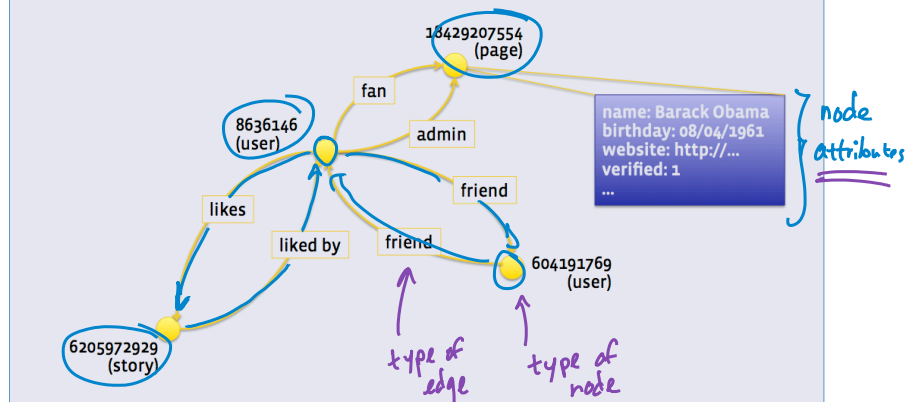
First Google advantage:  
a **Graph Algorithm** & a **System to Support it!**

*^ "page rank" ← see this later*

## Facebook Graph

**Data model**  
Objects & Associations

*more than just friend/friend interactions*



Slide from Facebook Engineering presentation 8

©Emily Fox 2014


## Label a Face and Propagate



©Emily Fox 2014

9

## Pairwise similarity not enough...



©Emily Fox 2014

10

# Propagate Similarities & Co-occurrences for Accurate Predictions

grandma

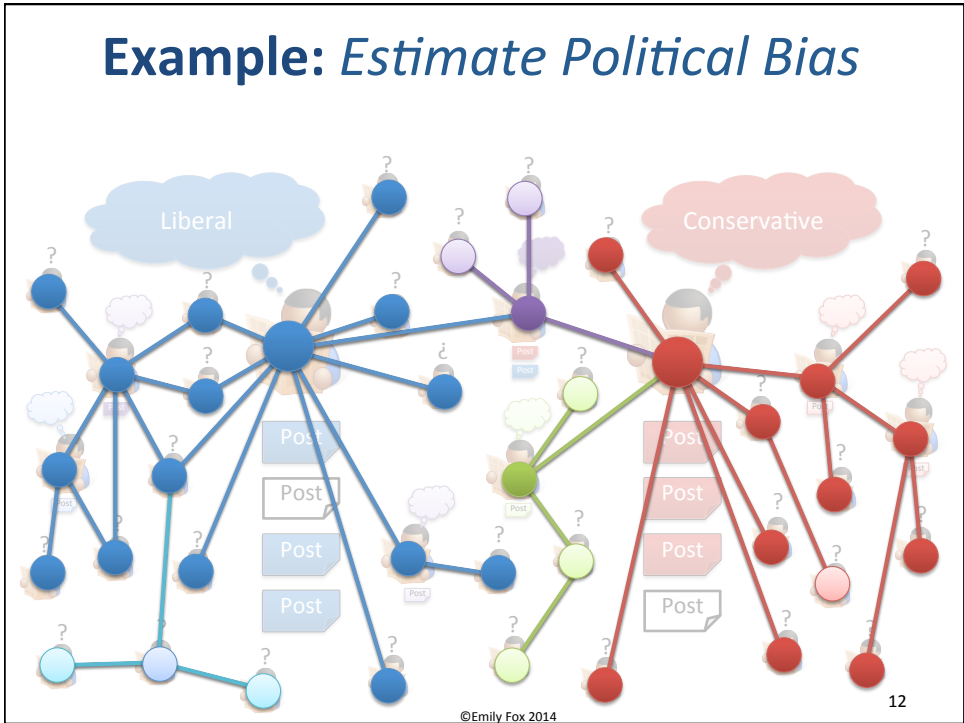
similarity edges

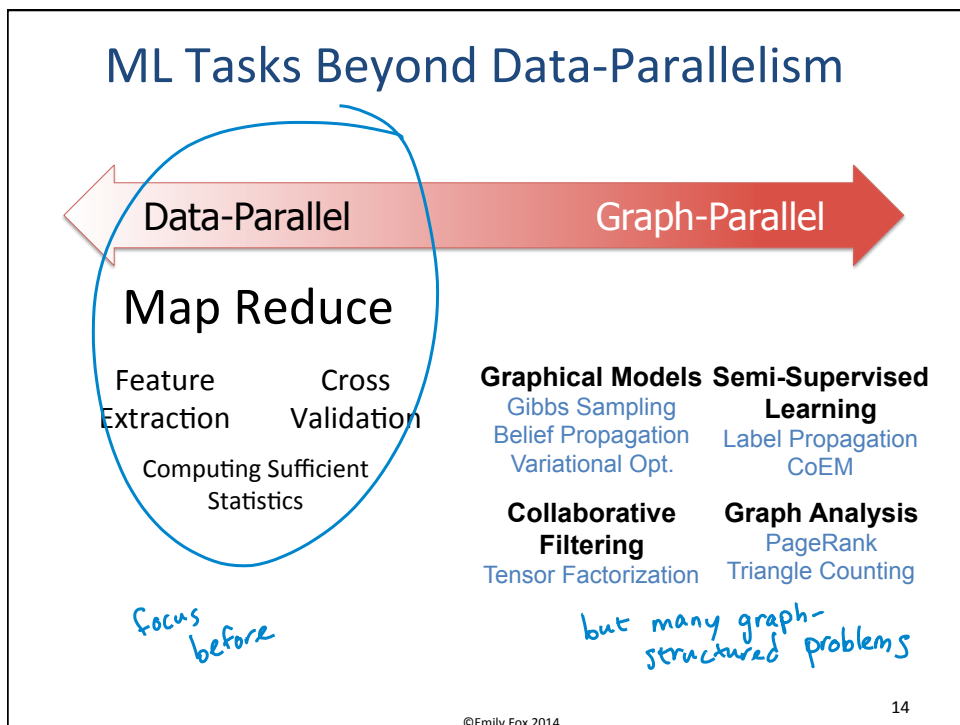
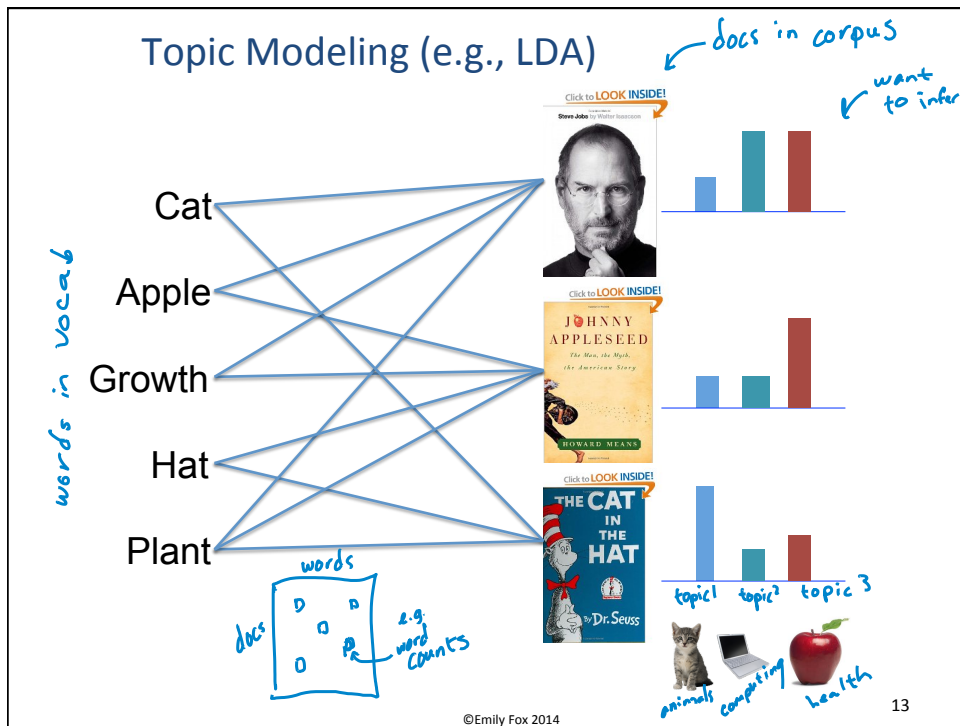
grandma!!!

co-occurring faces further evidence

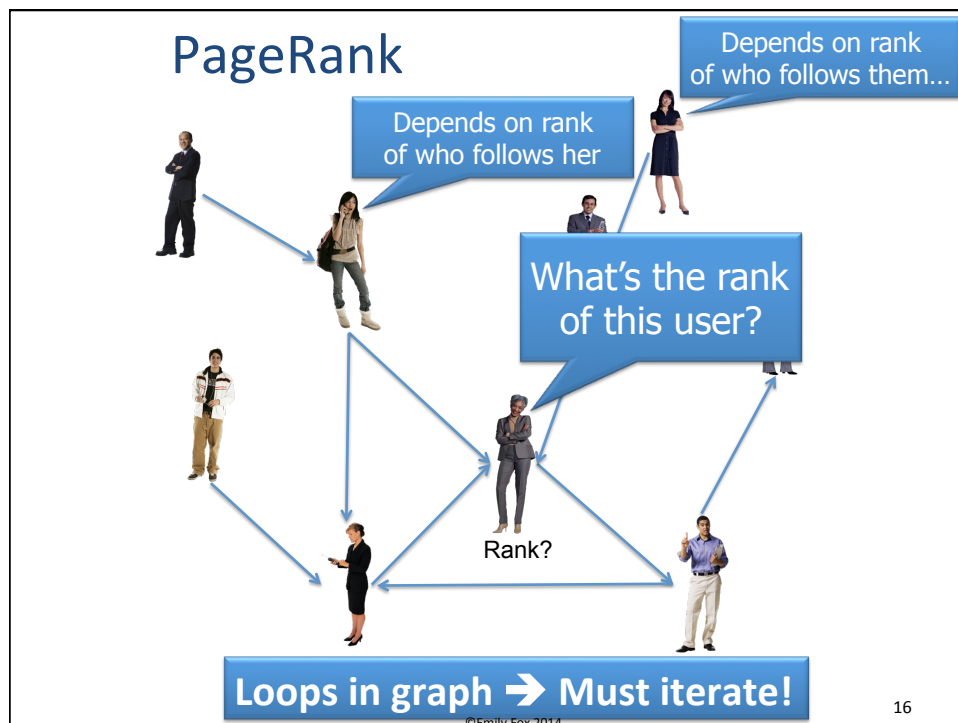
11

## Example: Estimate Political Bias





# Example of a Graph-Parallel Algorithm





## PageRank Iteration

*rank of user  $i$*

$$R[i] = \alpha + (1 - \alpha) \sum_{(j,i) \in E} w_{ji} R[j]$$

*edges in graph*

$$R[0] = 0.15 + 0.85(0.2R[1] + 0.5R[2] + 0.3R[3])$$

*0.15*

*prob. of landing on 0 directly*

- $\alpha$  is the random reset probability
- $w_{ji}$  is the prob. transitioning (similarity) from  $j$  to  $i$

17  
©Emily Fox 2014

## Properties of Graph Parallel Algorithms

### Dependency Graph

### Local Updates

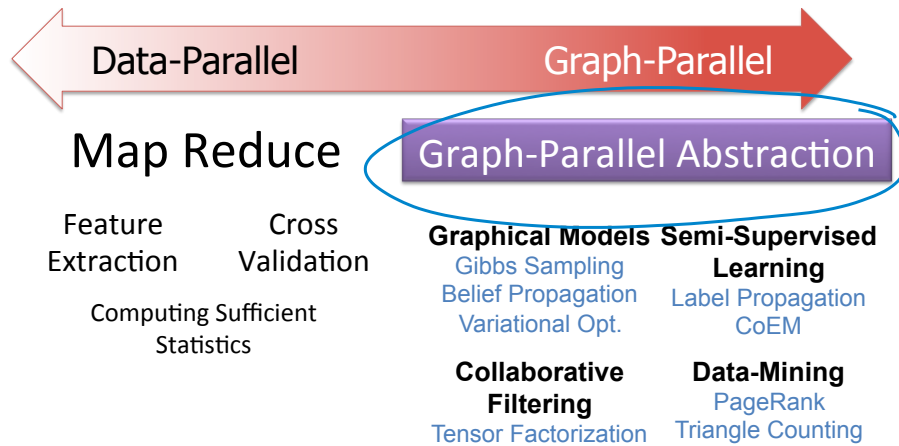
*algorithm shown on prev. slide*

### Iterative Computation

*guaranteed to conv. for page rank*

18  
©Emily Fox 2014

## Addressing Graph-Parallel ML



©Emily Fox 2014

19

## Graph Computation:

*Synchronous*

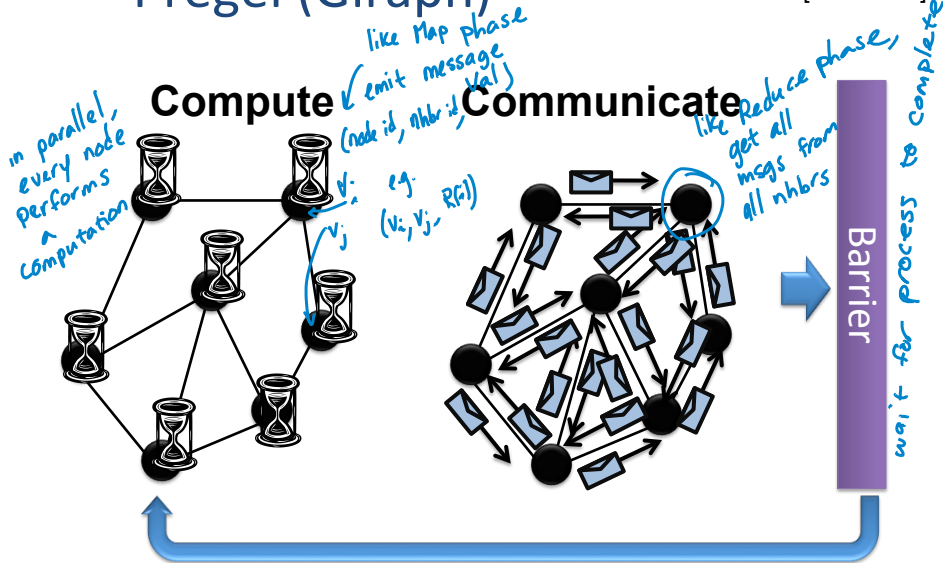
v.

*Asynchronous*

*key questions*

# Bulk Synchronous Parallel Model: Pregel (Giraph)

[Valiant '90]

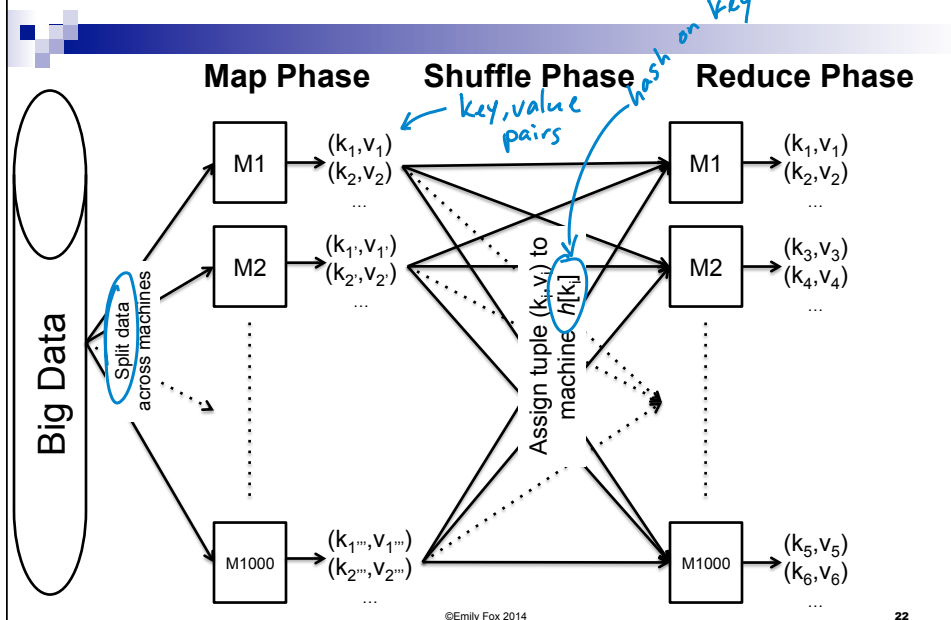


©Emily Fox 2014

21

Recall:

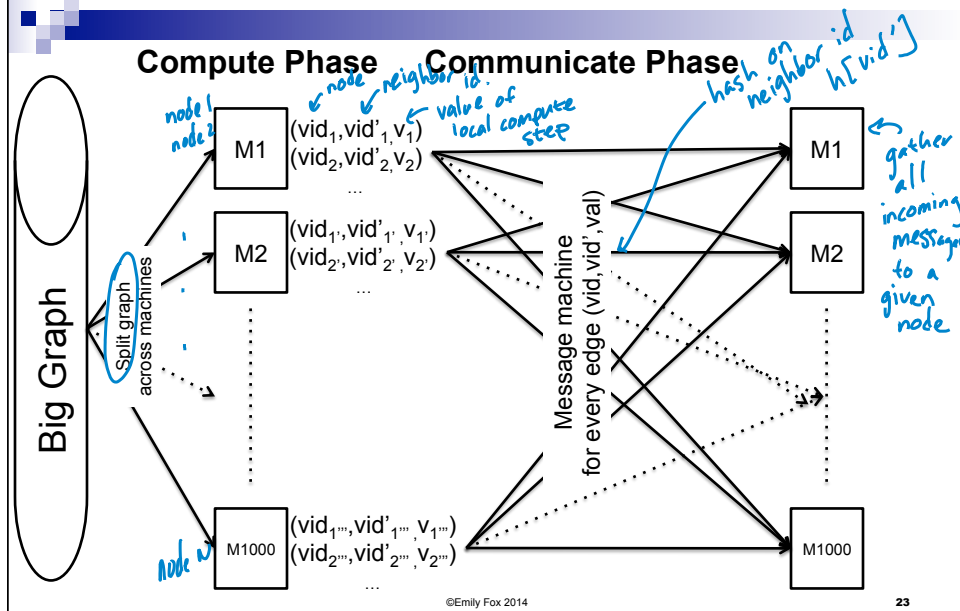
## Map-Reduce – Execution Overview



©Emily Fox 2014

22

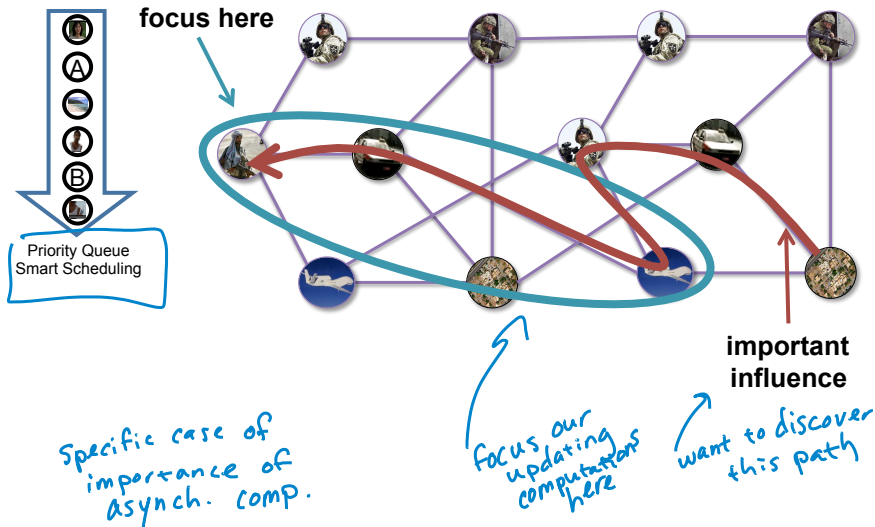
## BSP – Execution Overview



*Bulk synchronous  
parallel model  
provably inefficient  
for some ML tasks*

# Analyzing Belief Propagation

[Gonzalez, Low, G. '09]




©Emily Fox 2014


25

# Asynchronous Belief Propagation


Challenge = Boundaries



Synthetic Noisy Image



Cumulative Vertex Updates

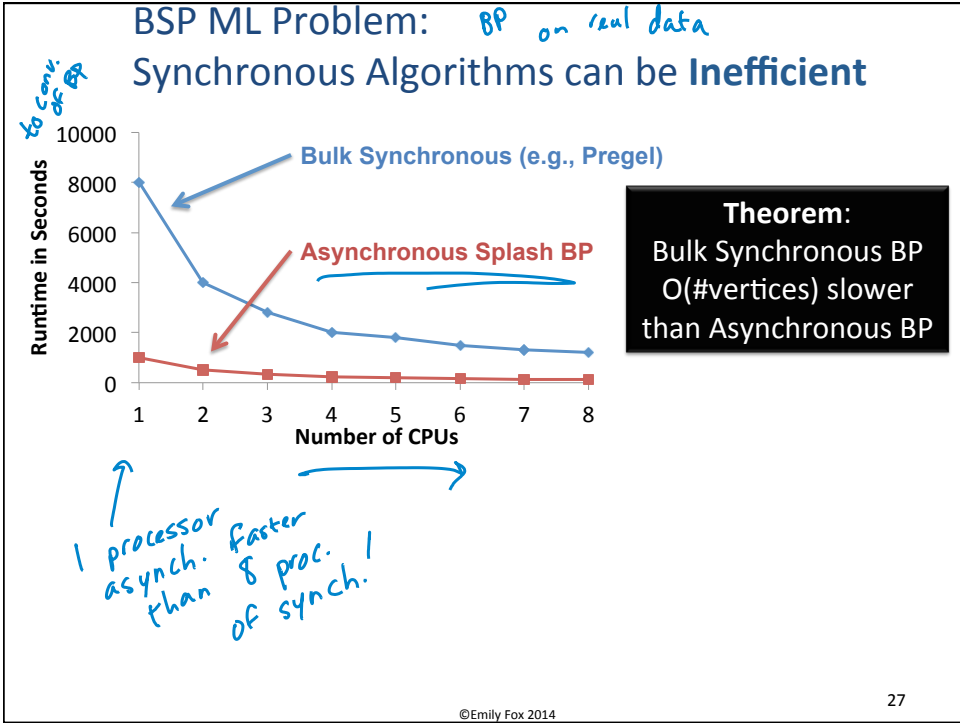


Graphical Model

Algorithm identifies and focuses on hidden sequential structure

©Emily Fox 2014

26



## Synchronous v. Asynchronous

<ul style="list-style-type: none"> <li>■ Bulk synchronous processing:           <ul style="list-style-type: none"> <li>□ Computation in phases               <ul style="list-style-type: none"> <li>■ All vertices participate in a phase                   <ul style="list-style-type: none"> <li>□ Though OK to say no-op</li> </ul> </li> <li>■ All messages are sent</li> </ul> </li> <li>□ Simpler to build, like Map-Reduce               <ul style="list-style-type: none"> <li>■ No worries about <u>race conditions</u>, → barrier guarantees data consistency</li> <li>■ Simpler to make fault-tolerant, save data on barrier</li> </ul> </li> <li>★ □ Slower convergence for many ML problems               <ul style="list-style-type: none"> <li>□ In matrix-land, called Jacobi Iteration</li> <li>□ Implemented by Google Pregel 2010</li> </ul> </li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>■ Asynchronous processing:           <ul style="list-style-type: none"> <li>□ Vertices see <u>latest information from neighbors</u> <ul style="list-style-type: none"> <li>■ Most closely related to sequential execution</li> </ul> </li> <li>□ Harder to build:               <ul style="list-style-type: none"> <li>■ Race conditions can happen all the time                   <ul style="list-style-type: none"> <li>□ Must protect against this issue</li> </ul> </li> <li>■ More complex fault tolerance</li> <li>■ When are you done?</li> <li>■ Must implement scheduler over vertices</li> </ul> </li> <li>★ □ Faster convergence for many ML problems               <ul style="list-style-type: none"> <li>□ In matrix-land, called Gauss-Seidel Iteration</li> <li>□ Implemented by GraphLab 2010, 2012</li> </ul> </li> </ul> </li> </ul>
---	---

©Emily Fox 2014 28

## Case Study 4: Collaborative Filtering

# GraphLab

Machine Learning for Big Data  
CSE547/STAT548, University of Washington

Emily Fox


February 25<sup>th</sup>, 2014

©Emily Fox 2014

29

## The GraphLab Goals

Know how to solve ML problem on 1 machine



- 4 components:
- Data graph
  - Update fn
  - Scheduler
  - Consistency model

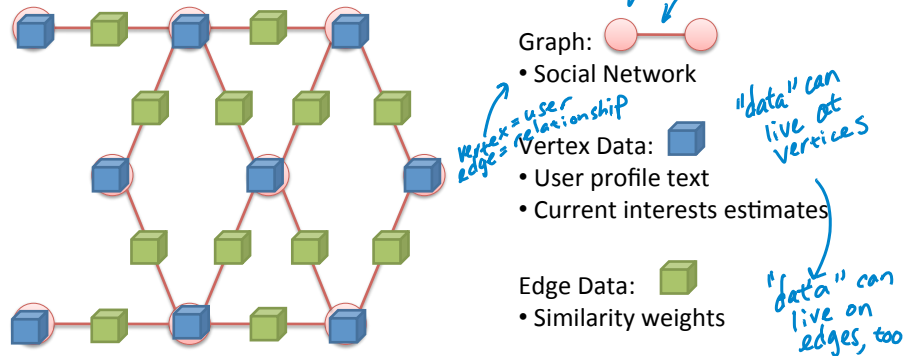
Efficient parallel predictions

©Emily Fox 2014

30

## ① Data Graph

Data associated with vertices and edges



©Emily Fox 2014

31

How do we *program*  
**graph** computation?

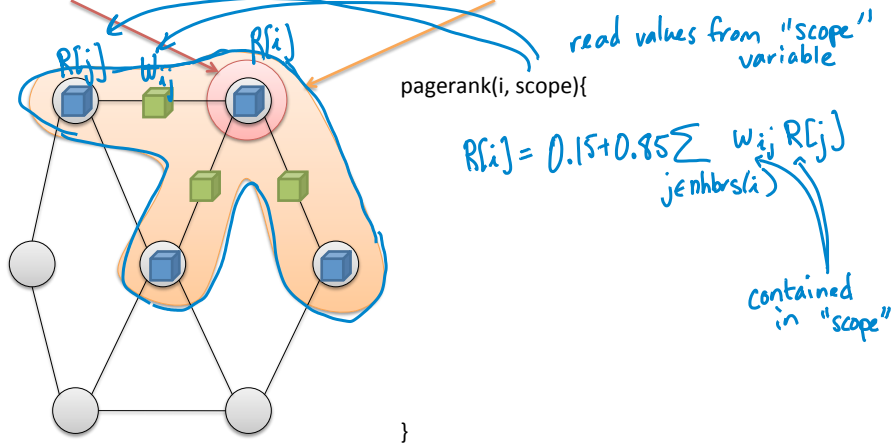
“Think like a Vertex.”

-Malewicz et al. [SIGMOD'10]



## ② Update Functions


User-defined program: applied to **vertex** transforms data in **scope** of vertex



©Emily Fox 2014

33

## Update Function Example: Connected Components




1. Initialize all nodes w/ unique labels
2. Pick a node:  
component = min(self, neighbors)
3. Return to step 2

©Emily Fox 2014

34

## Update Function Example: Connected Components




$\swarrow$  label of node  $i$   
 $\text{init } \text{component}[i] = i$   
 $\text{update}(i, \text{scope}) \}$   
 $\text{component}[i] =$   
 $\min[\text{component}[i],$   
 $\min_{j \in \text{nbr}(i)} \text{component}[j]]$   
 $\}$  all in "scope"

©Emily Fox 2014

35

## ③ The Scheduler

The **scheduler** determines order vertices are updated



©Emily Fox 2014

36

## Example Schedulers

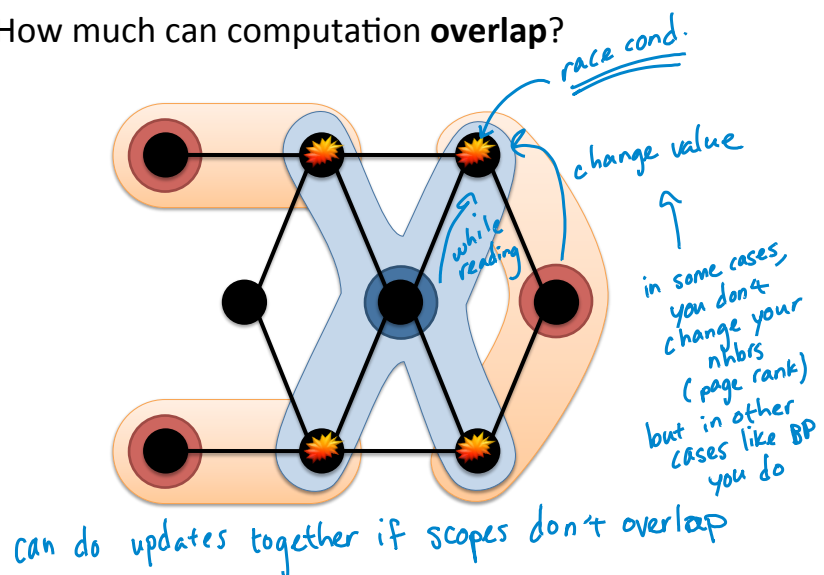
- Round-robin
- Selective scheduling (skipping):
  - round robin but jump over un-scheduled vertice
- FIFO
- Prioritize scheduling (e.g. splash BP)
  - Hard to implement in a distributed fashion
    - Approximations used (each machine has its own priority queue)

©Emily Fox 2014

37

## ④ Ensuring Race-Free Code

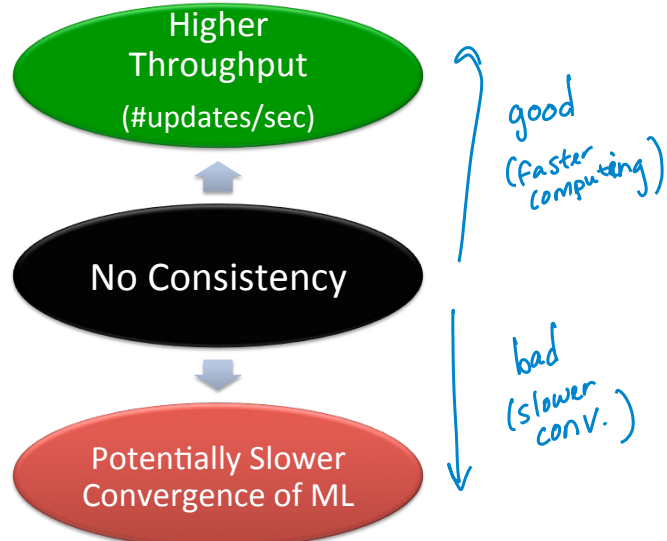
How much can computation **overlap**?



©Emily Fox 2014

38

## Need for Consistency?




©Emily Fox 2014

39

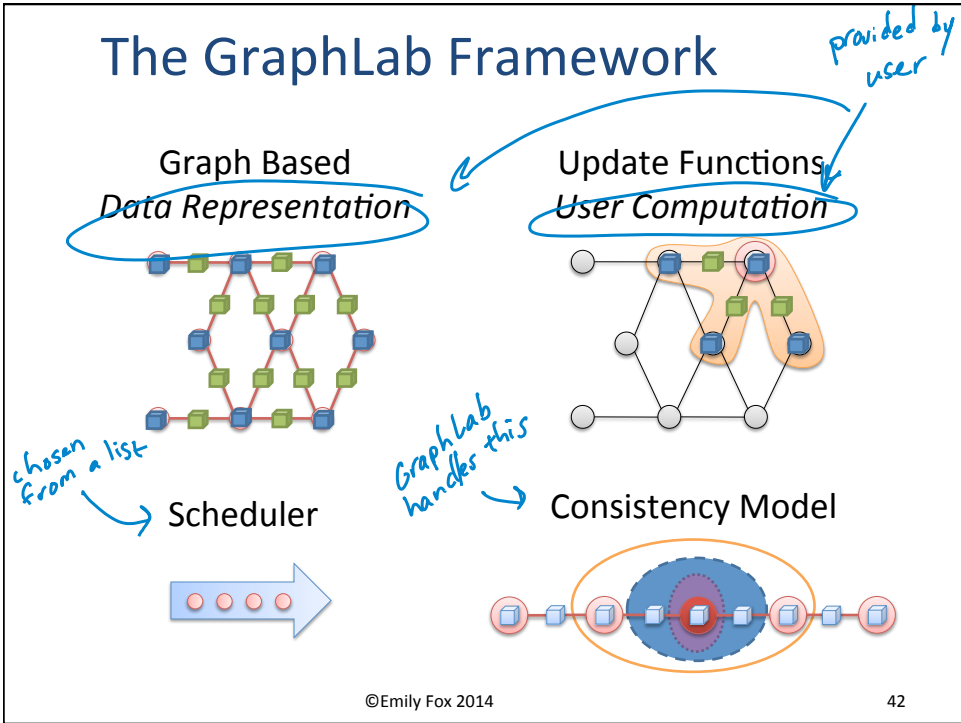

## GraphLab Ensures Sequential Consistency

For **each parallel execution**, there exists a **sequential execution** of update functions which produces the same result



©Emily Fox 2014

40



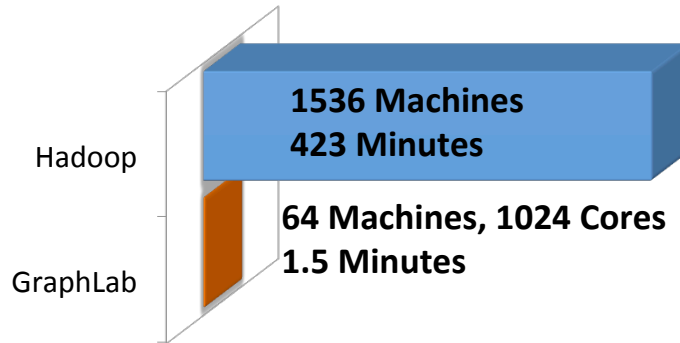
## Triangle Counting in Twitter Graph



40M Users  
1.2B Edges

**Total:**

**34.8 Billion Triangles**



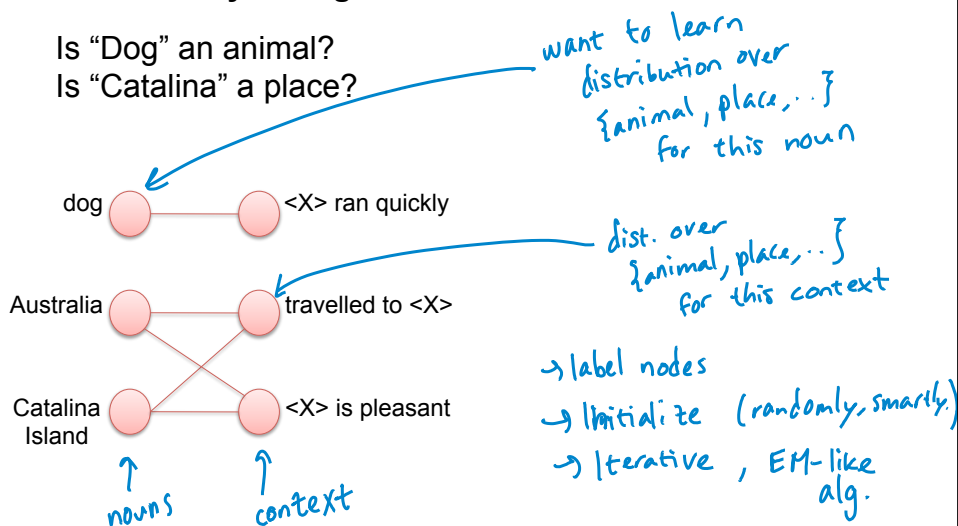
Hadoop results from [Suri & Vassilvitskii '11] ©Emily Fox 2014

43

## CoEM (Jones et al., 2005)

### Named Entity Recognition Task

Is "Dog" an animal?  
Is "Catalina" a place?



©Emily Fox 2014

44

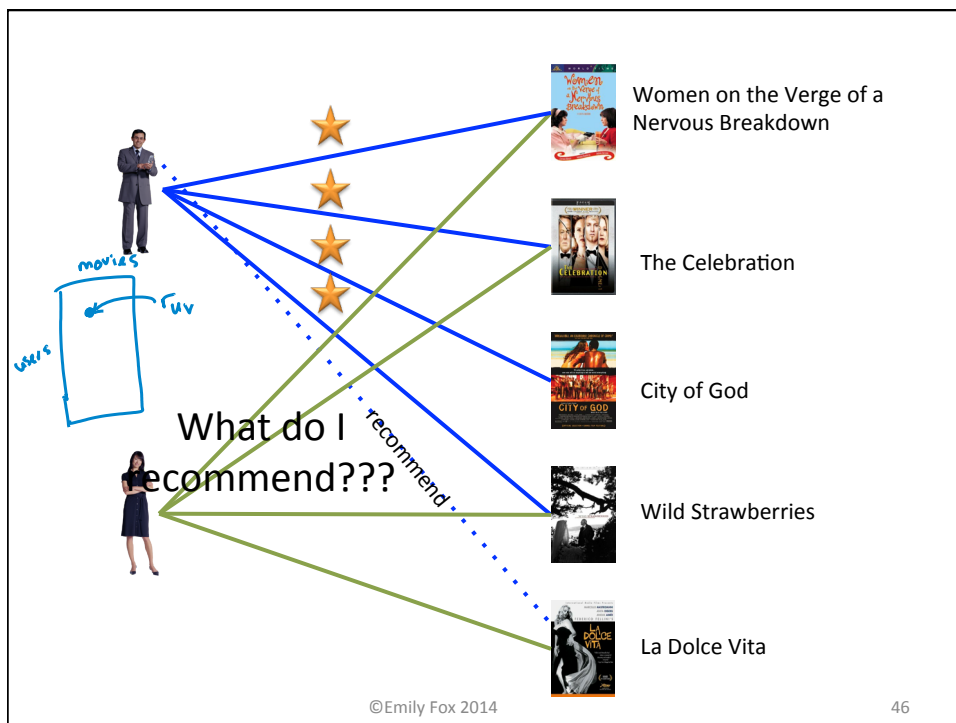
## Never Ending Learner Project (CoEM)

**Vertices: 2 Million**  
**Edges: 200 Million**

Hadoop	95 Cores	7.5 hrs
<b>Distributed GraphLab</b>	<b>32 EC2 machines</b>	<b>80 secs</b>

©Emily Fox 2014

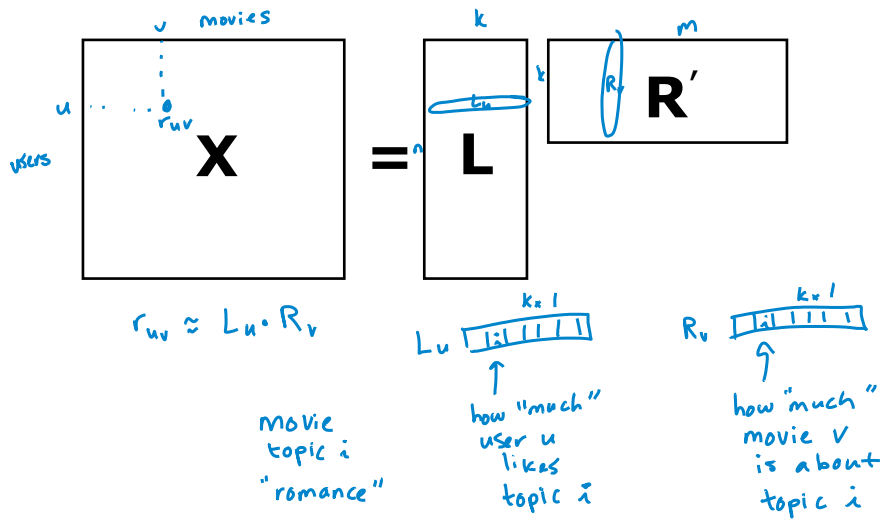
45



©Emily Fox 2014

46


## Interpreting Low-Rank Matrix Completion (aka Matrix Factorization)



©Emily Fox 2014

47

## Matrix Completion as a Graph



$X_{ij}$  known for black cells  
 $X_{ij}$  unknown for white cells  
 Rows index users  
 Columns index movies

©Emily Fox 2014

48



# Coordinate Descent for Matrix Factorization: Alternating Least-Squares

$$\min_{L,R} \sum_{(u,v):r_{uv} \neq ?} (L_u \cdot R_v - r_{uv})^2 + \lambda_u \|L\| + \lambda_v \|R\|$$

- Fix movie factors, optimize for user factors

★ □ Independent least-squares over users  $\min_{L_u} \sum_{v \in V_u} (L_u \cdot R_v - r_{uv})^2 + \lambda_u \|L\|$

- Fix user factors, optimize for movie factors

★ □ Independent least-squares over movies  $\min_{R_v} \sum_{u \in U_v} (L_u \cdot R_v - r_{uv})^2 + \lambda_v \|R\|$

- System may be underdetermined: use regularization

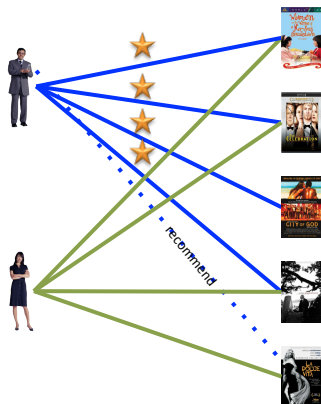
- Converges to local optima

©Emily Fox 2014

49

# Alternating Least Squares Update Function

$$\min_{L_u} \sum_{v \in V_u} (L_u \cdot R_v - r_{uv})^2 \quad \min_{R_v} \sum_{u \in U_v} (L_u \cdot R_v - r_{uv})^2$$



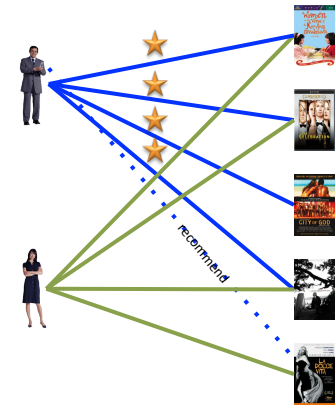
©Emily Fox 2014

50

# SGD for Matrix Factorization in GraphLab

$$\epsilon_t = L_u^{(t)} \cdot R_v^{(t)} - r_{uv}$$

$$\begin{bmatrix} L_u^{(t+1)} \\ R_v^{(t+1)} \end{bmatrix} \leftarrow \begin{bmatrix} (1 - \eta_t \lambda_u) L_u^{(t)} - \eta_t \epsilon_t R_v^{(t)} \\ (1 - \eta_t \lambda_v) R_v^{(t)} - \eta_t \epsilon_t L_u^{(t)} \end{bmatrix}$$



©Emily Fox 2014

51

# Bayesian PMF Example

*\* Full Bayesian approach  
place priors on  $\phi$   
as well!*

- Latent user and movie factors:

$$L_u \sim N(\mu_u, \Sigma_u) \quad u=1, \dots, n$$

$$R_v \sim N(\mu_v, \Sigma_v) \quad v=1, \dots, m$$

- Observations  $r_{uv} \sim N(L_u^T R_v, \sigma_r^2)$


- Hyperparameters:

$$\phi = \{ \underbrace{\mu_u, \Sigma_u}_{\phi_u}, \underbrace{\mu_v, \Sigma_v}_{\phi_v}, \underbrace{\sigma_r^2}_{\phi_r} \}$$

- Want to predict new movie rating:

$$p(r_{uv}^* | X, \phi) = \int p(r_{uv}^* | L_u, R_v) p(L, R | X, \phi) dL dR$$

*↑ new rating*  
*↑ obs. ratings*



*(more general than before)*

*new user/movie combo*

*posterior given obs. so far*

©Emily Fox 2014

52

# Bayesian PMF Gibbs Sampler

- Outline of Bayesian PMF sampler

1. Init  $L^{(1)}, R^{(1)}$

2. For  $k=1, \dots, \text{Niter}$

(i) Sample hyperparams  $\phi^{(k)} = \{\phi_u^{(k)}, \phi_v^{(k)}, \phi_r^{(k)}\}$

(ii) For each user  $u=1, \dots, n$  sample in parallel

$$L_u^{(k+1)} \sim P(L_u | X, R^{(k)}, \phi^{(k)})$$

(iii) For each movie  $v=1, \dots, m$  sample in parallel

$$R_v^{(k+1)} \sim P(R_v | X, L^{(k+1)}, \phi^{(k)})$$

Very similar to ideas of ALS (systematically)

©Emily Fox 2014

53

# Bayesian PMF Example

- For user  $u$ :

$$p(L_u | X, R, \phi_u) \propto p(L_u | \phi_u) \prod_{v \in V_u} p(r_{uv} | L_u, R_v, \phi_r)$$

$$\propto N(L_u | \mu_u, \Sigma_u) \prod_{v \in V_u} N(r_{uv} | L_u R_v, \sigma_r^2)$$

$$= N(L_u | \tilde{\mu}_u, \tilde{\Sigma}_u) \leftarrow \text{via conjugacy}$$

$$\text{where } \tilde{\Sigma}_u^{-1} = \Sigma_u^{-1} + \sigma_r^{-2} \sum_{v \in V_u} R_v R_v^T$$

$$\tilde{\mu}_u = \tilde{\Sigma}_u (\sigma_r^{-2} \sum_{v \in V_u} r_{uv} R_v + \Sigma_u \mu_u)$$

posterior is in the same family as prior


- Symmetrically for  $R_v$  conditioned on  $L$  (breaks down over movies)
- Luckily, we can use this to get our desired posterior samples

©Emily Fox 2014

54

## PMF Gibbs Sampling in GraphLab

$$p(L_u | X, R, \phi_u) = N(\tilde{\mu}_u, \tilde{\Sigma}_u) \quad \tilde{\Sigma}_u = \Sigma_u^{-1} + \sigma_r^{-2} \sum_{v \in V_u} R_v R_v^T \quad \tilde{\mu}_u = \tilde{\Sigma}_u \left( \sigma_r^{-2} \sum_{v \in V_u} r_{uv} R_v + \Sigma_u \mu_u \right)$$



©Emily Fox 2014

55

GraphLab

Release 2.2 available now

<http://graphlab.org>

Documentation... Code... Tutorials... (more on the way)

GraphChi 0.1 available now

<http://graphchi.org>

## What you need to know...

- Data-parallel versus graph-parallel computation
- Bulk synchronous processing versus asynchronous processing
- GraphLab system for graph-parallel computation
  - Data representation
  - Update functions
  - Scheduling
  - Consistency model
- ALS, SGD and Gibbs for matrix factorization/PMF in GraphLab

©Emily Fox 2014

64

## Reading

- Papers under “Case Study IV: **Parallel Learning with GraphLab**”
- Optional:
  - Parallel Splash BP  
<http://www.ml.cmu.edu/research/dap-papers/dap-gonzalez.pdf>

©Emily Fox 2014

65

# Acknowledgements

- Slides based on Carlos Guestrin's GraphLab talk