Machine Learning Techniques in Social Computing (社会计算)

Irwin King

king@cse.cuhk.edu.hk
http://www.cse.cuhk.edu.hk/~king

Department of Computer Science & Engineering
The Chinese University of Hong Kong

Alexa as of Nov. 2008	USA	CHINA	Global
1	Google	Baidu	Yahoo
2	Yahoo	QQ	Google
3	Myspace	Sina	YouTube
4	YouTube	Google.cn	Windows Live
5	Facebook	Taobao	Facebook
6	Windows Live	163	MSN
7	MSN	Yahoo	Myspace
8	Wikipedia	Google	Wikipedia
9	EBay	Sohu	Blogger
10	AOL	Youku	Yahoo.jp

What's On the Menu?

- Web 2.0 and Social X
- Social Computing (社会计算)
- Some Interesting Problems
  - Collaborative Filtering (協同過濾)
  - Query Suggestion (查询建议)

What's On the Menu?

- Web 2.0 and Social X
- Social Computing (社会计算)
- Some Interesting Problems
  - Collaborative Filtering (協同過濾)
  - Query Suggestion (查询建议)

Web 2.0

- Web as a medium vs. **Web as a platform**
- Read-Only Web vs. **Read-and-Write Web**
- Static vs. **Dynamic**
- Restrictive vs. **Freedom & Empowerment**
- Technology-centric vs. **User-centric**
- Limited vs. **Rich User Experience**
- Individualistic vs. **Group/Collective Behavior**
- Consumer vs. **Producer**
- Transactional vs. **Relational**
- Top-down vs. **Bottom-up**
- People-to-Machine vs. **People-to-People**
- Search & browse vs. **Publish & Subscribe**
- Closed application vs. **Service-oriented Services**
- Functionality vs. **Utility**
- Data vs. **Value**

Machine Learning Techniques for Social Computing, Irwin King, MLA'08, Nanjing, China, November 8-9, 2008

Social Networking

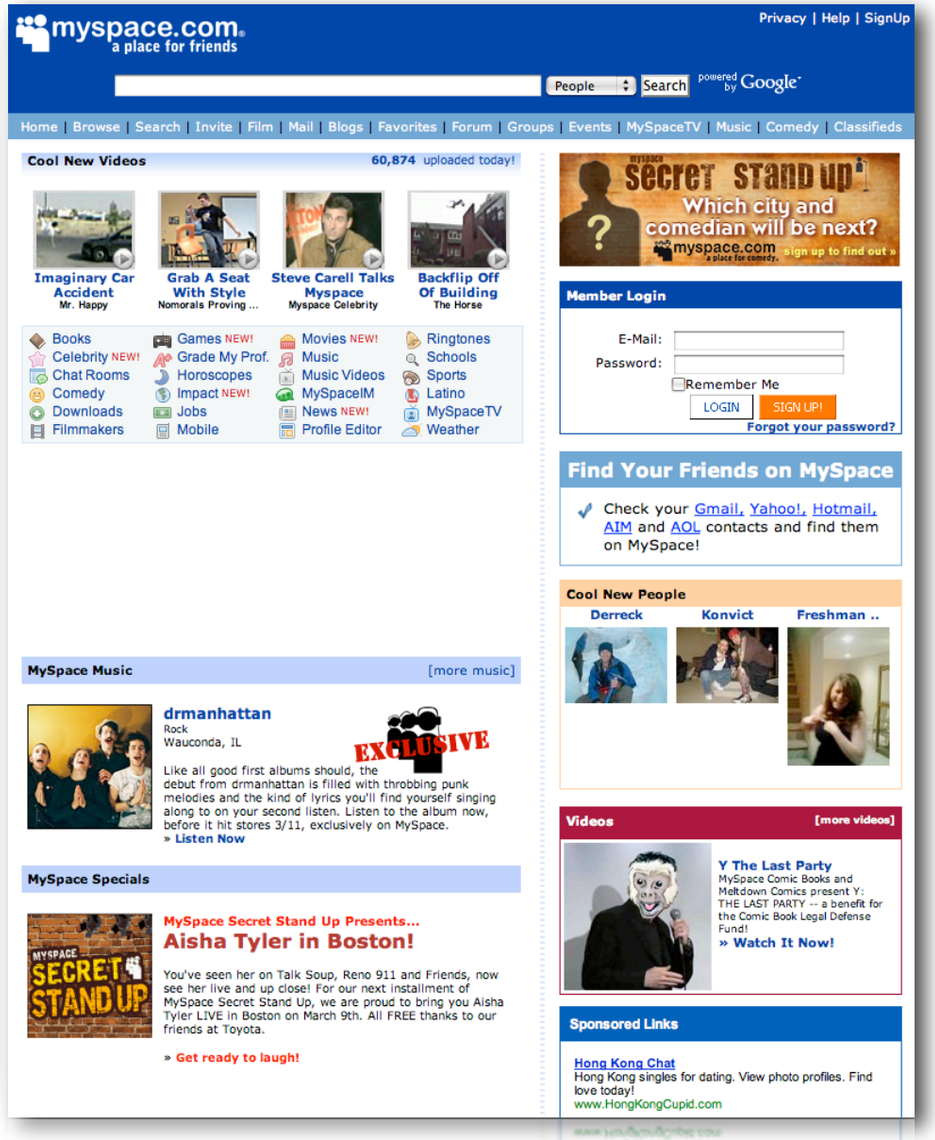

Social Search

- Social Search Engine
- Leveraging your social networks for searching

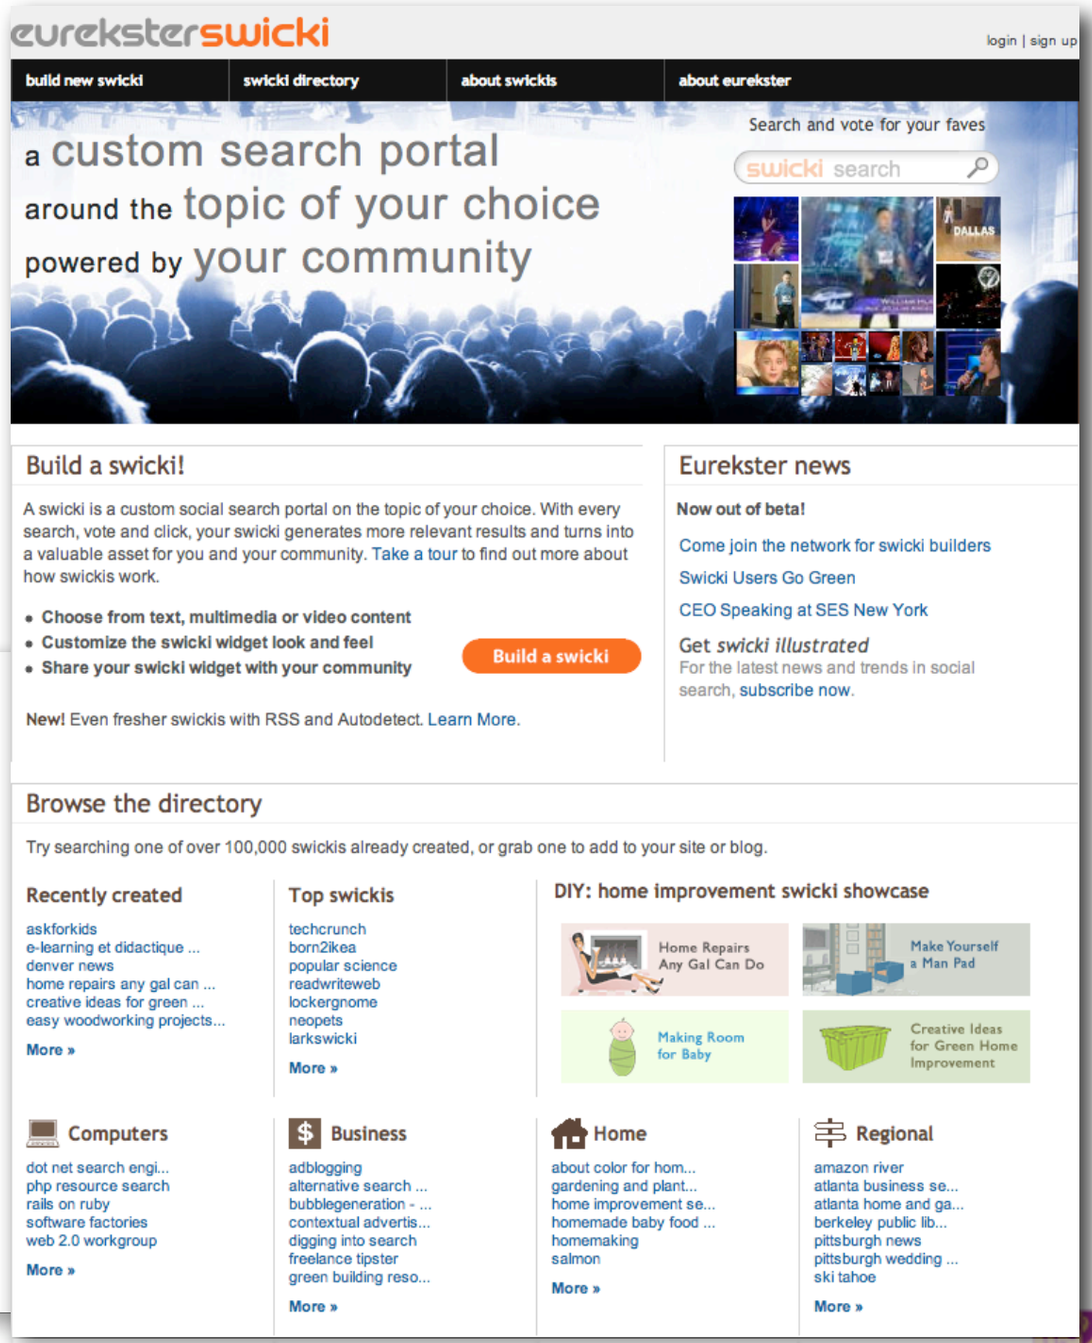

Social Bookmarking

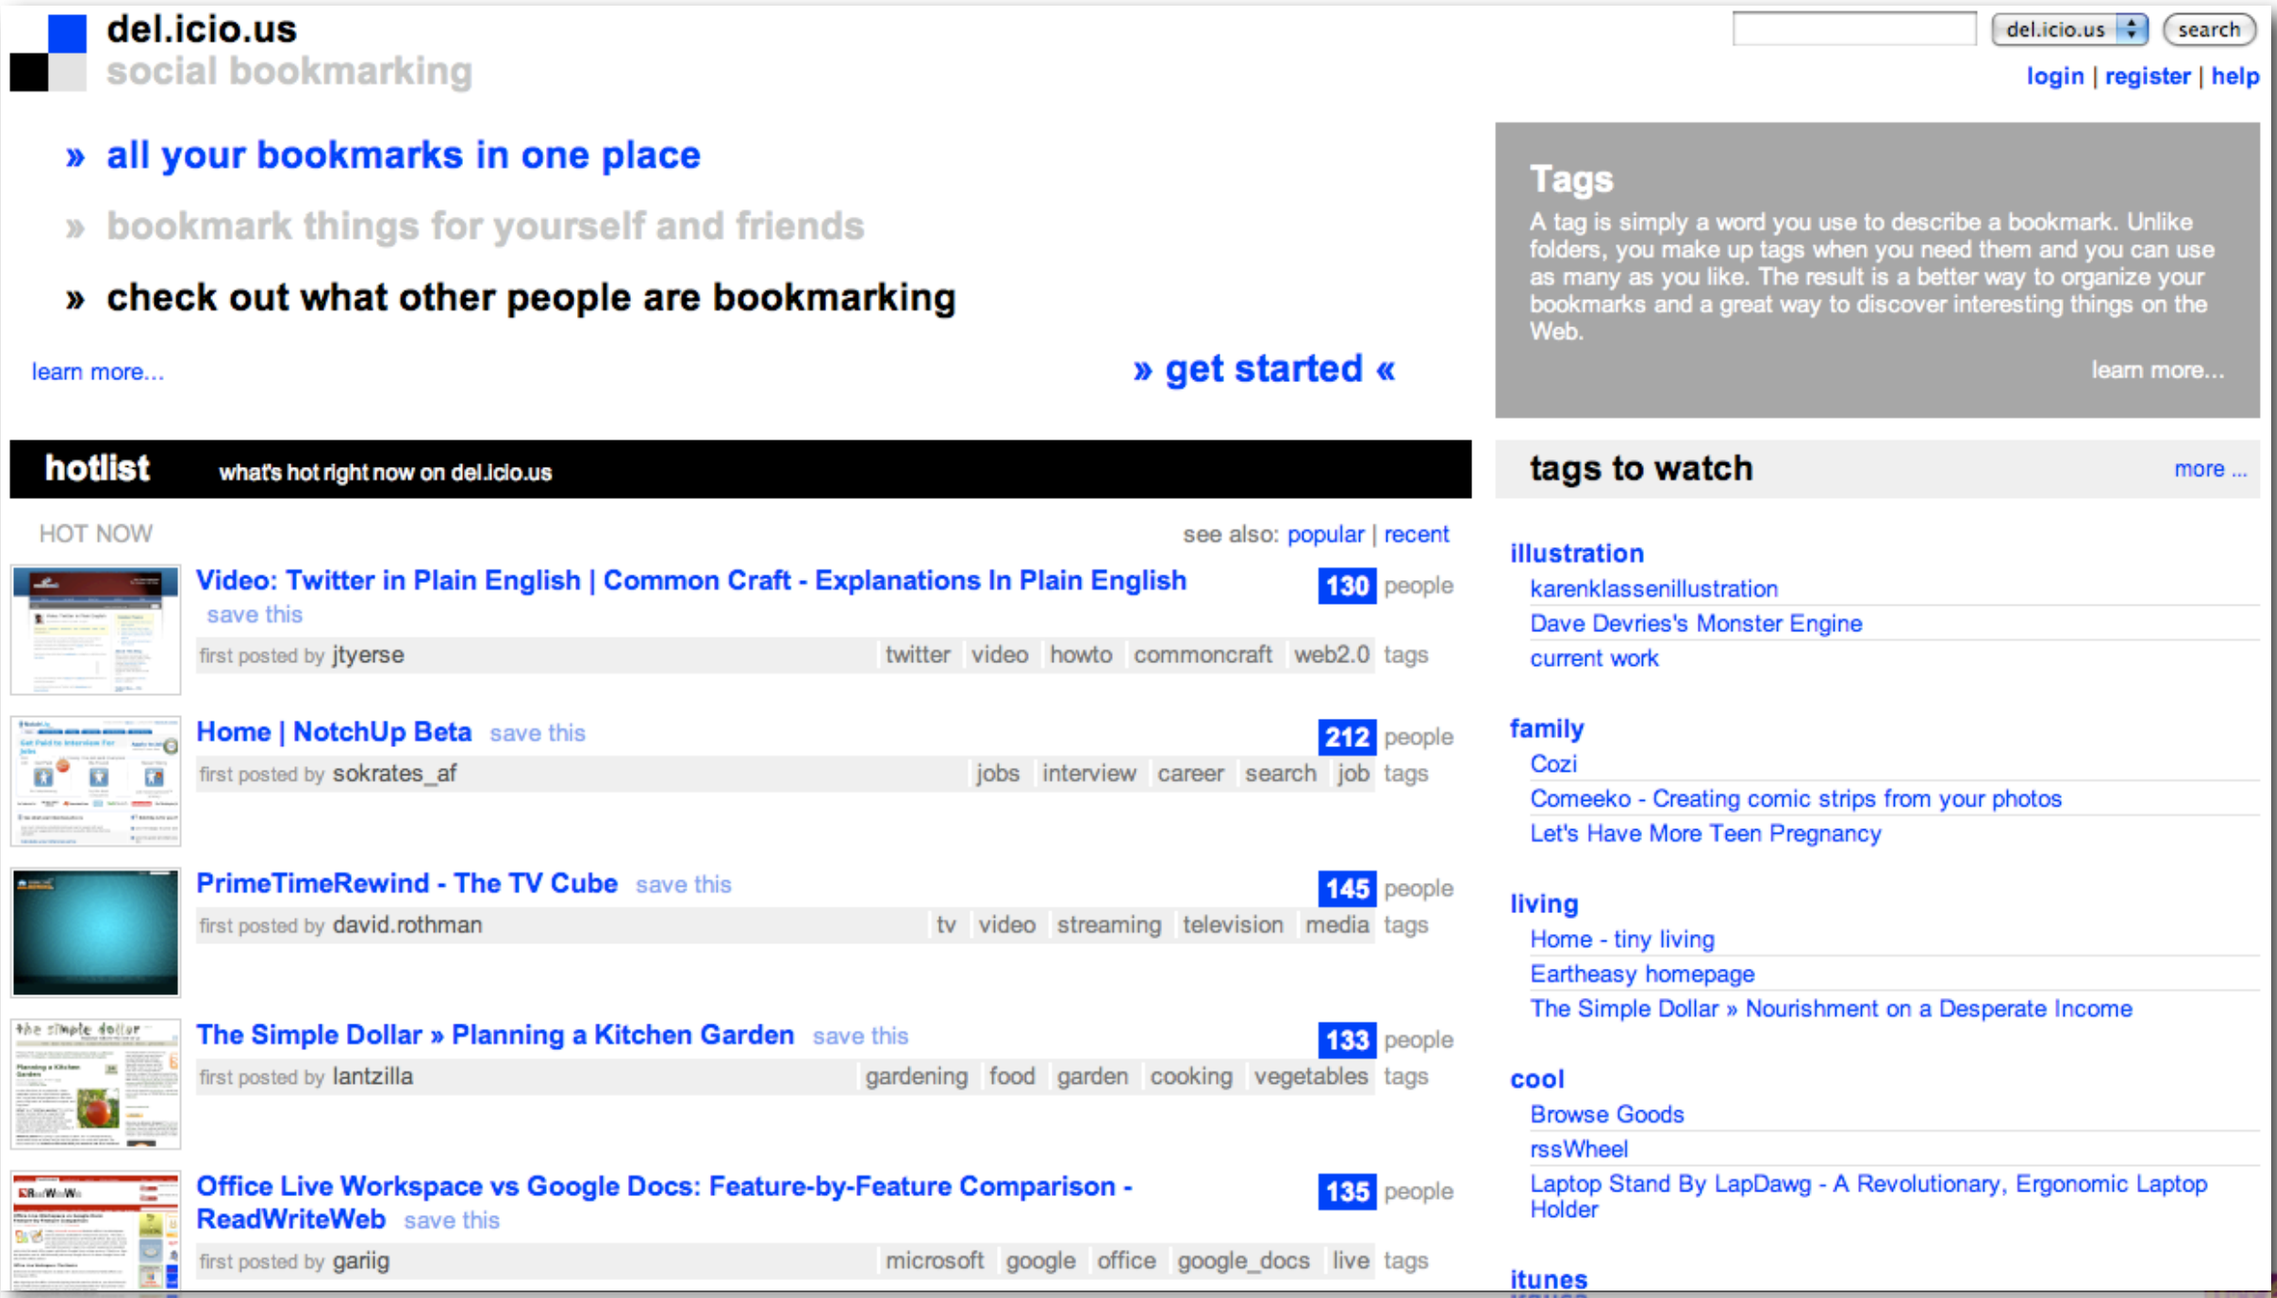

Social Media

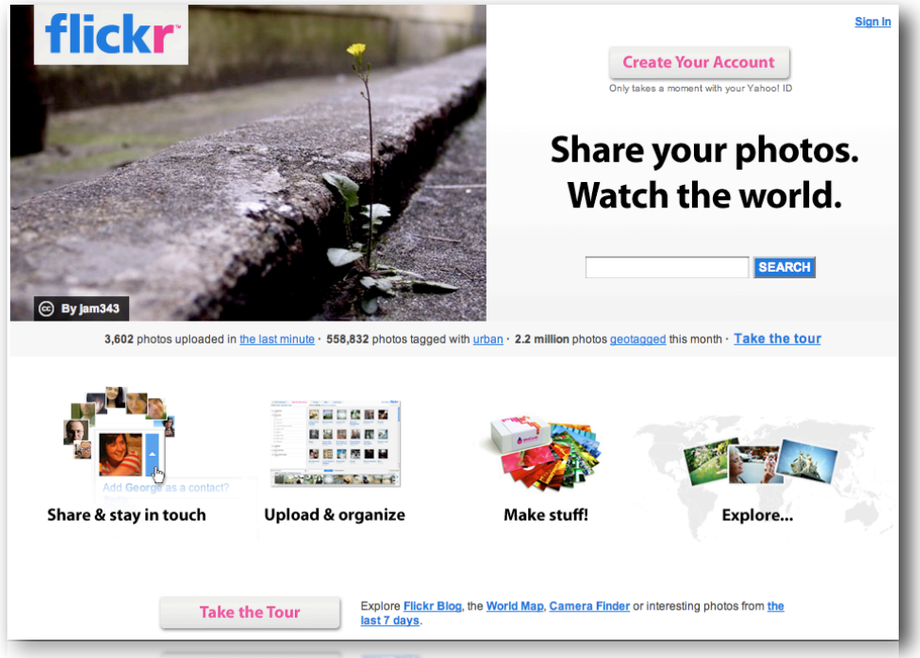

Social News/Mash Up

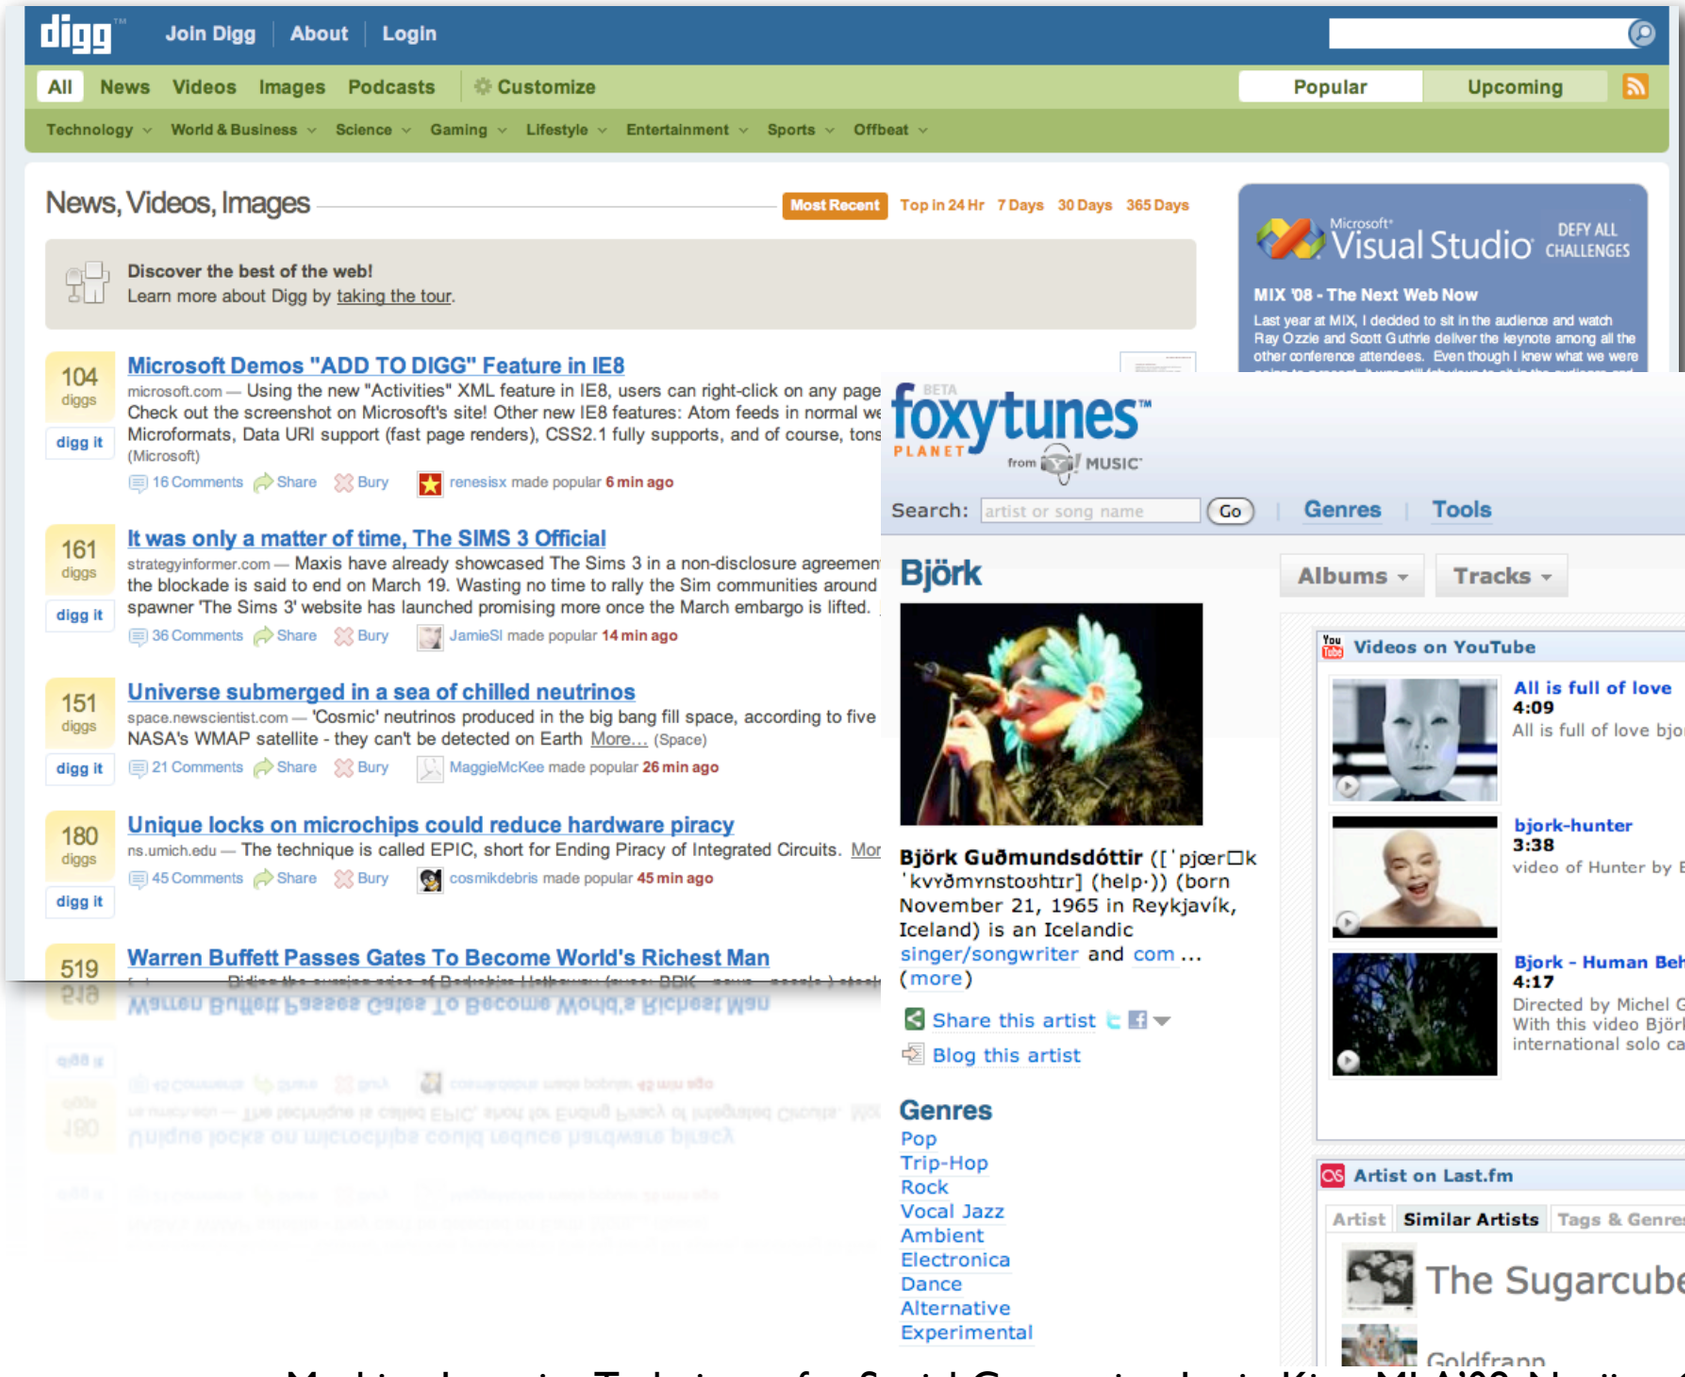

Social Knowledge Sharing

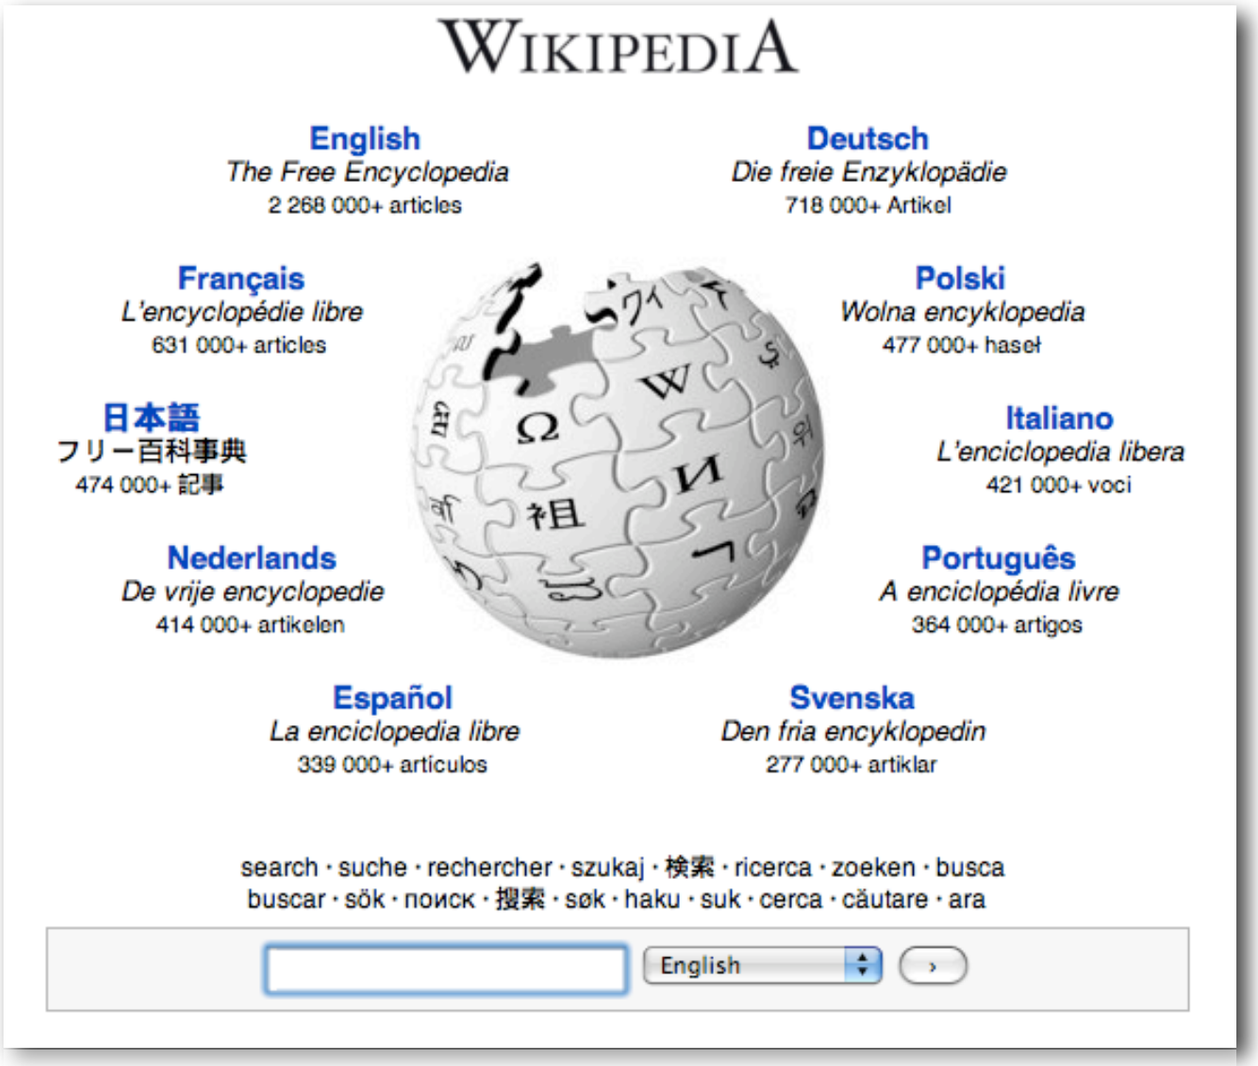

Social Marketing

- Viral marketing
- Who are the brokers?
- Who can exert the most influence on buying/ selling?
- How much should one advertise?

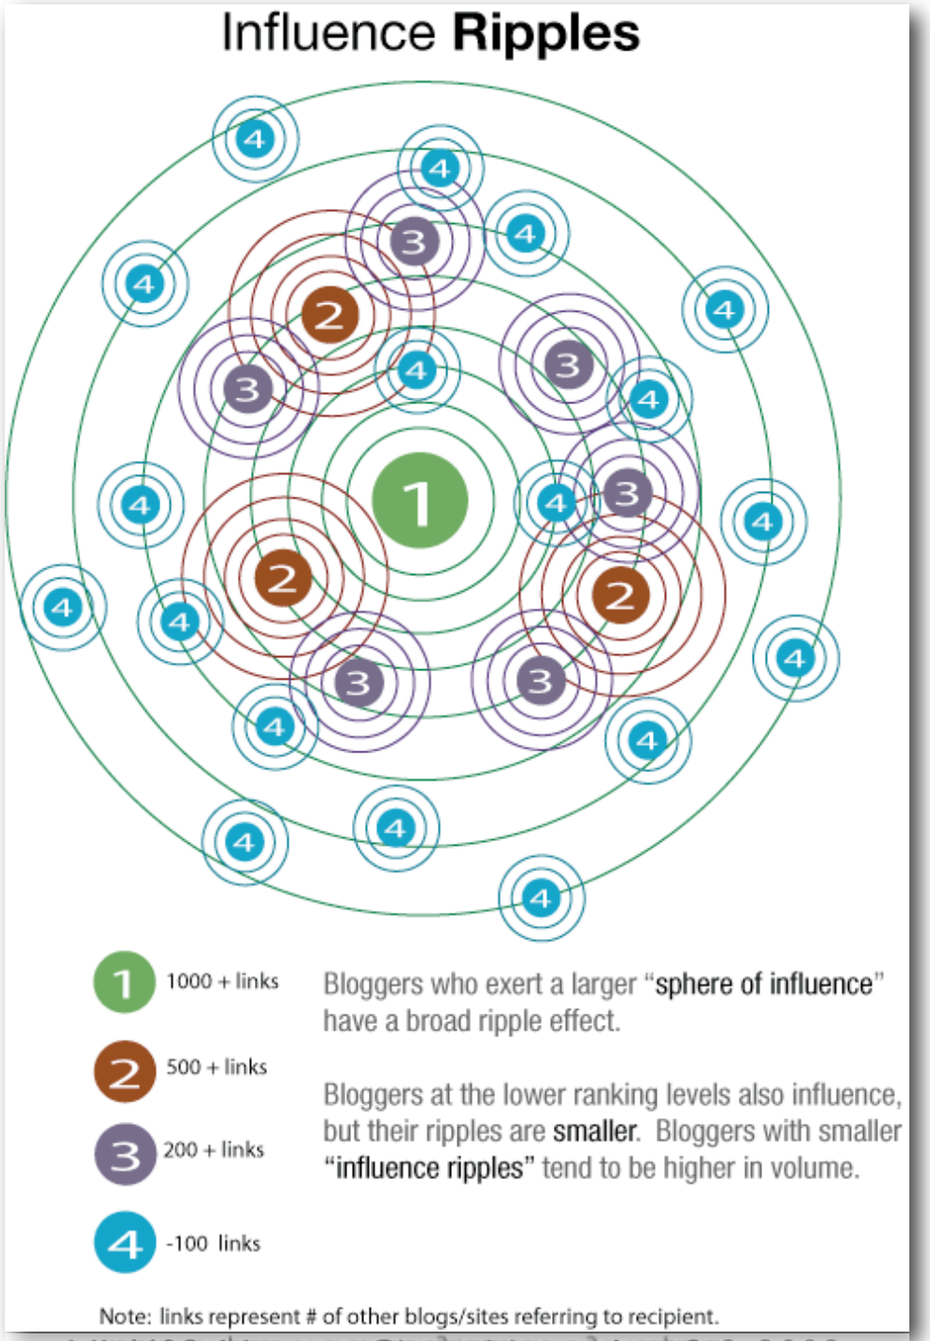

Social/Human Computation

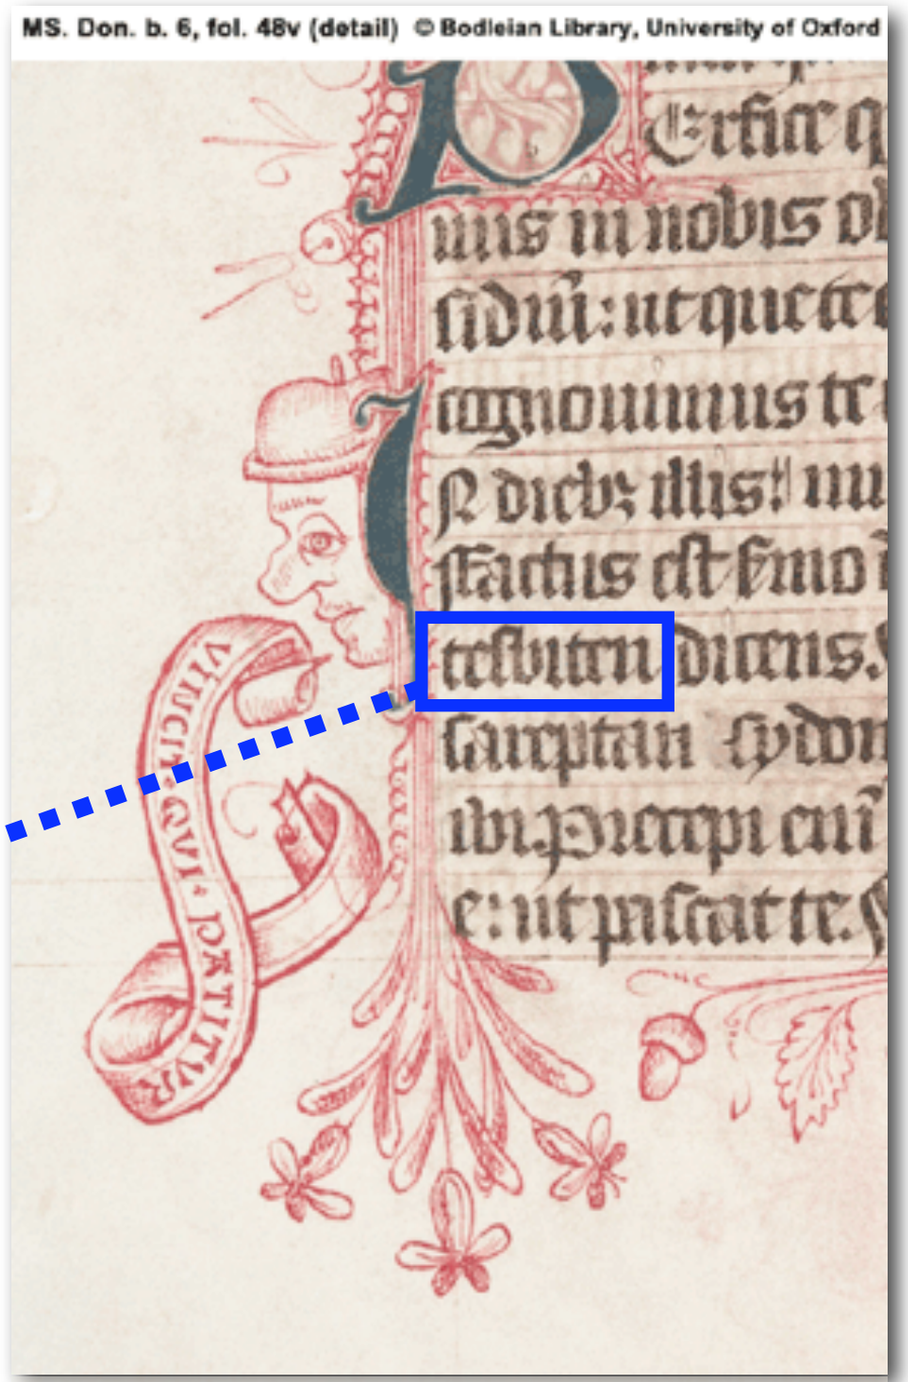

Web 2.0 Revolution

革命 vs. 演变

The Three C's

Connectivity (人等之间的关系)

Collaboration (协作)

Communities (社会团体)

What's On the Menu?

- Web 2.0 and Social X
- Social Computing (社会计算)
- Some Interesting Problems
  - Collaborative Filtering (協同過濾)
  - Query Suggestion (查询建议)

Social Relations

同伴;支持者;共谋人;市场;社区;团体;组织

crew

族,队;组;班;群;窝

teams

populations

binary

identity (身份)

presence (存在)

social role (社会角色)

squad

cardinal

reputation (信誉)

organizations

expertise (专长)

integer

trust (信任)

cohorts

markets

ownership (所有权)

real

communities

accountability (问责制)

partners

knowledge (知识)

groups

Social Computing

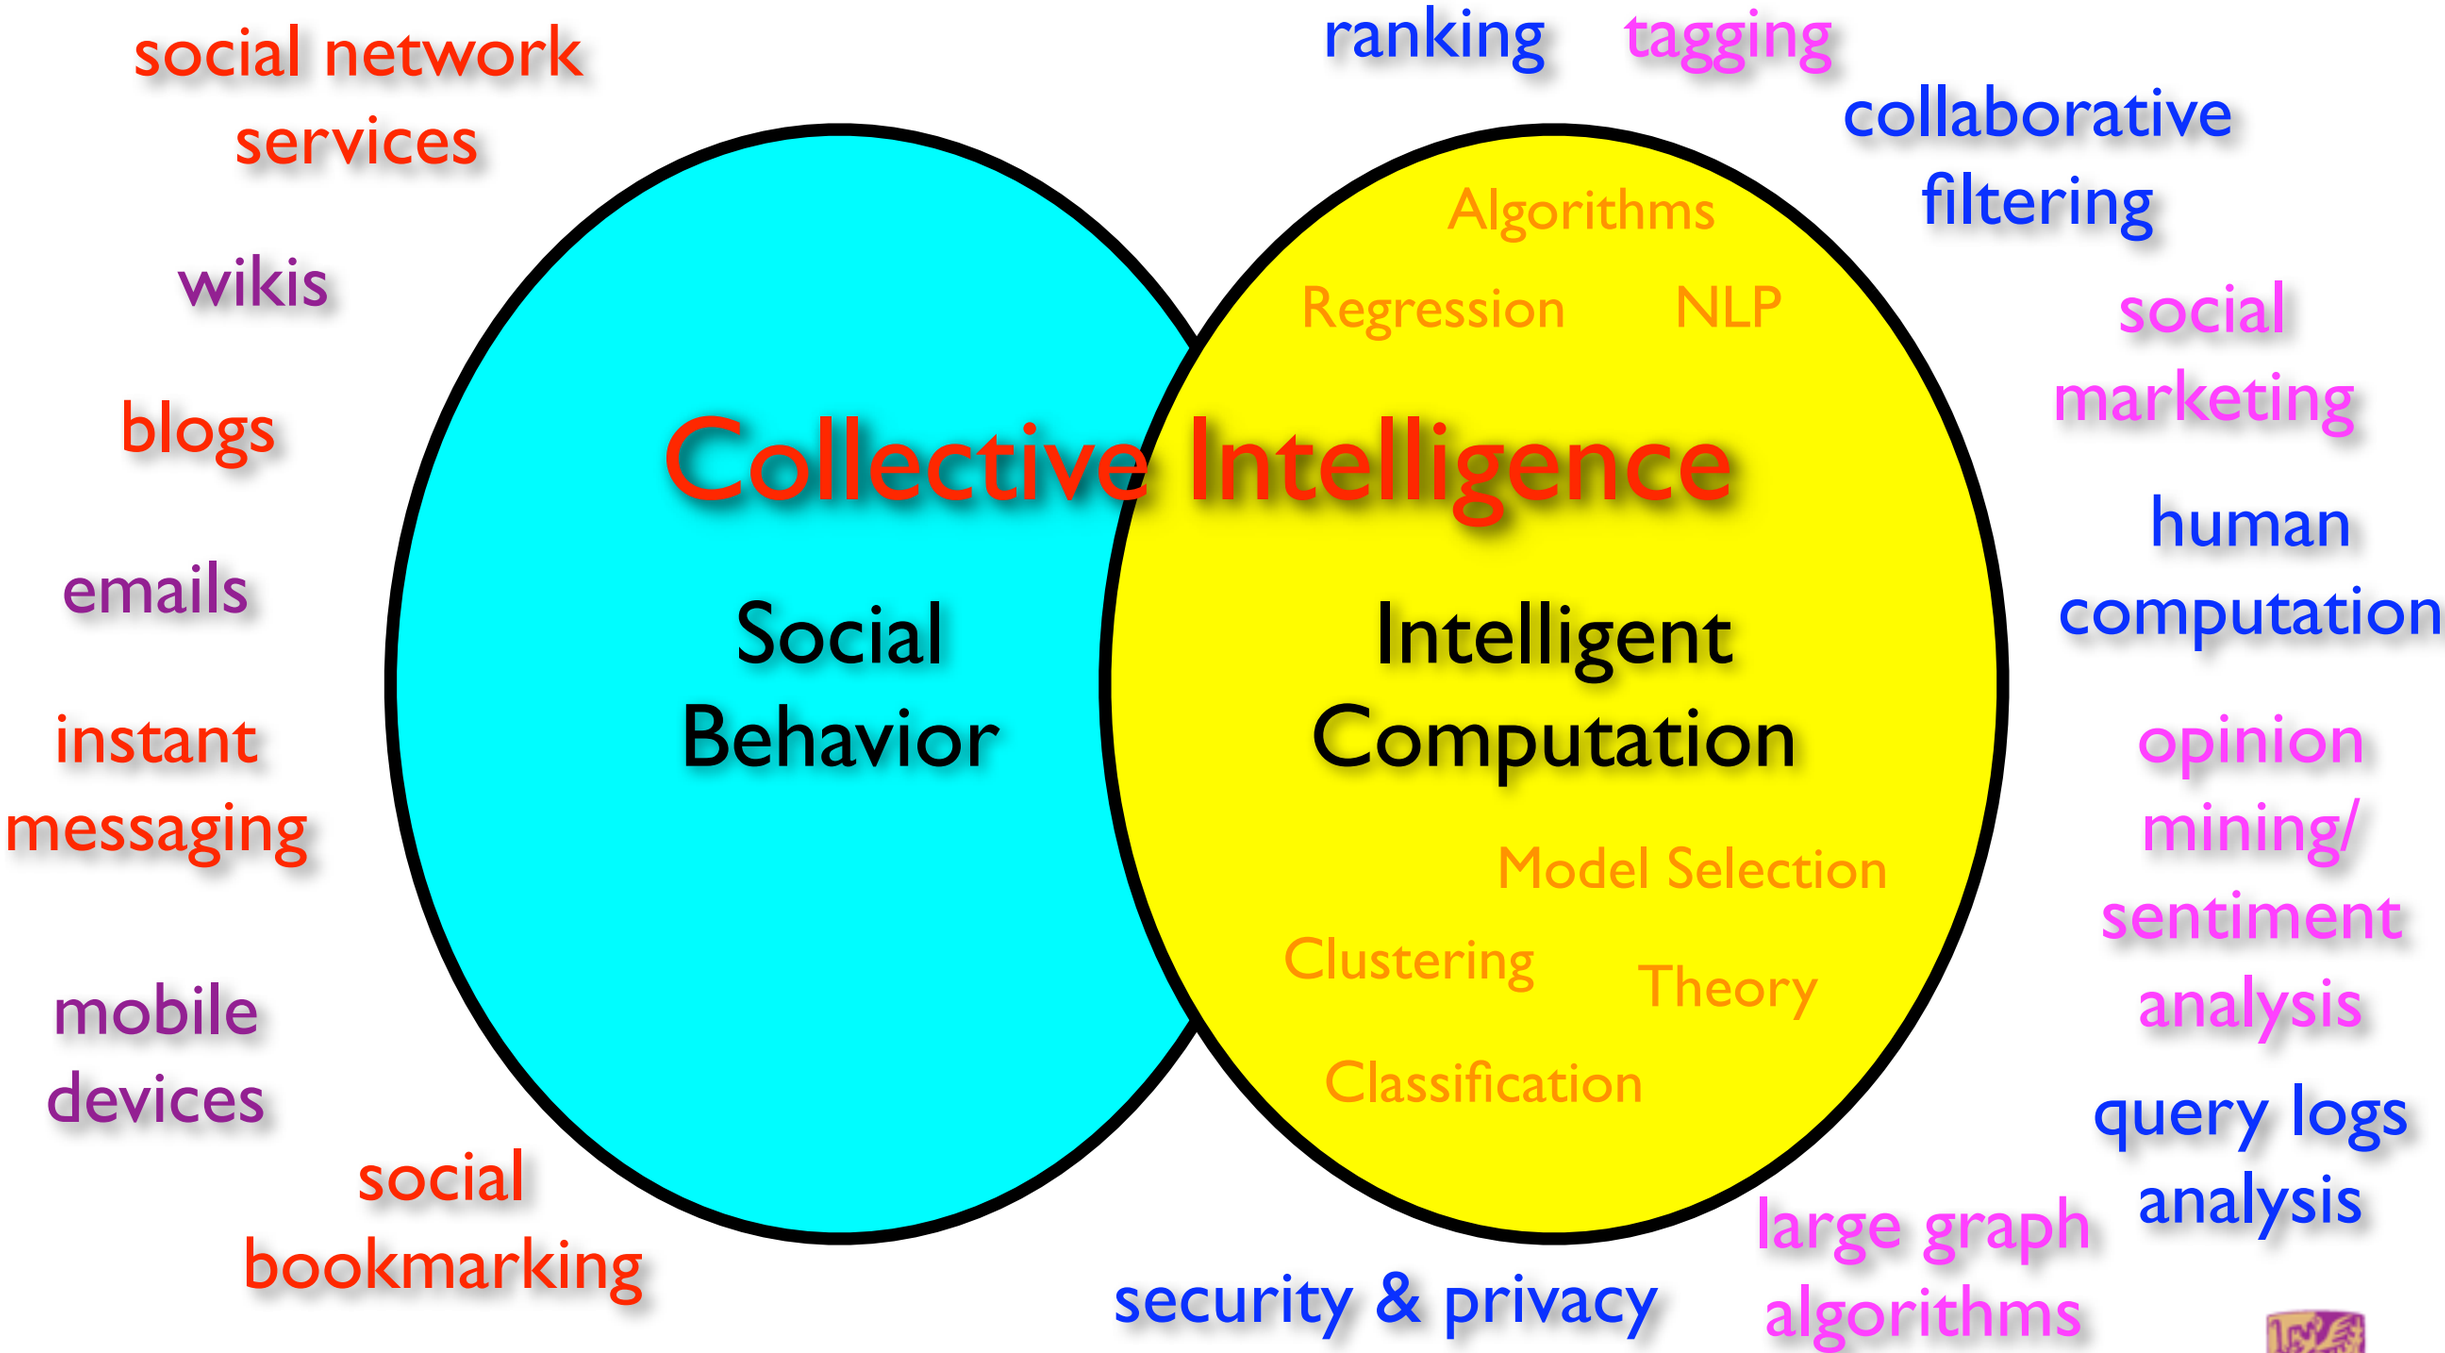

Social Computing (SC)

- Social computing is a general term for an area of computer science that is concerned with the intersection of *social behavior* (社会行为) and *computational systems*. *Wikipedia*

- *A social structure in which technology puts power in communities, not institutions.* *Forrester*

- *Forms of web services where the value is created by the collective contributions (集体贡献) of a user population.*

Issues

- Theory and models (理论与模型)
- Mining (数据挖掘) of existing information, e.g., spatial (relations) and temporal (time) domains
  - Dealing with partial and incomplete (局部和不完整) information, e.g., collaborative filtering, ranking, tagging, etc.
- Scalability and algorithmic (可扩展性和算法) issues
- Security and privacy (安全和隐私) issues
- Monetization (货币) of social interactions

Machine Learning in SC

- Classification, clustering, regression, etc.
- New insights on the data
  - Social relations are often hidden (latent)
  - Change data from (x, y) to $(x, c_1(x), c_2(x), \cdots, y)$
- $c(x)$ = context in *tags*, *relations*, *ratings*, etc.
- data type = *binary*, *integer*, *real*, *cardinal*, etc.

What's On the Menu?

- Web 2.0 and Social X
- Social Computing (社会计算)
- Some Interesting Problems
  - Collaborative Filtering (協同過濾)
  - Query Suggestion (查询建议)

What's On the Menu?

- Web 2.0 and Social X
- Social Computing (社会计算)
- Some Interesting Problems
  - Collaborative Filtering (協同過濾)
  - Query Suggestion (查询建议)

A Better Mousetrap?

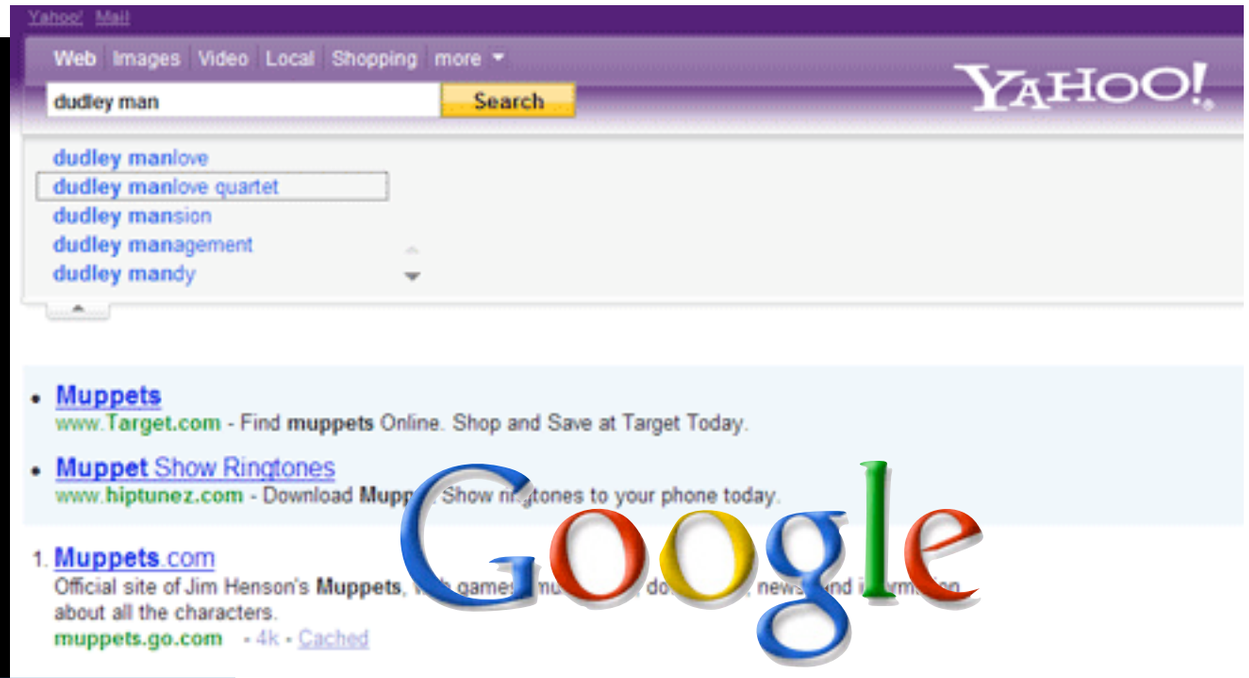

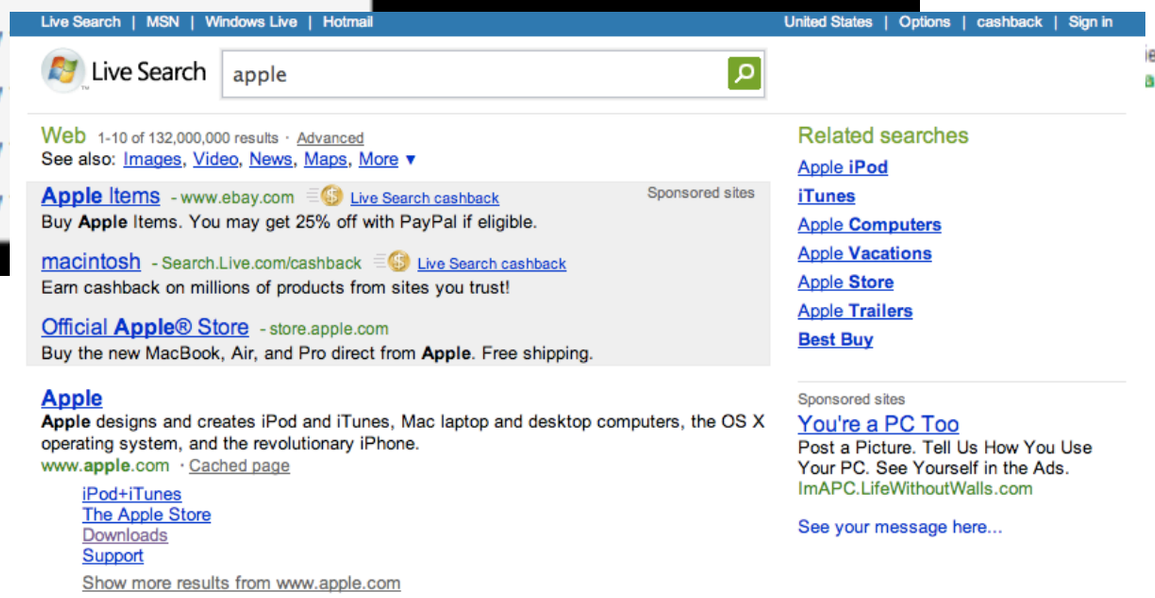

Challenges

- Queries contain ambiguous and new terms
  - apple: "apple computer" or "apple pie"?
  - NDCG:?
- Users tend to submit short queries consisting of only one or two words
  - almost 20% one-word queries
  - almost 30% two-word queries

- Users may have little or even no knowledge about the topic they are searching for!

What is Clickthrough Data

- Query logs recorded by search engines

$$\langle u, q, l, r, t \rangle$$

Table 1: Samples of search engine clickthrough data

ID	Query	URL	Rank	Time
358	facebook	http://www.facebook.com	1	2008-01-01 07:17:12
358	facebook	http://en.wikipedia.org/wiki/Facebook	3	2008-01-01 07:19:18
3968	apple iphone	http://www.apple.com/iphone/	1	2008-01-01 07:20:36
...	...	...	...	...

- Users' relevance feedback to indicate desired/preferred/target results

Joint Bipartite Graph

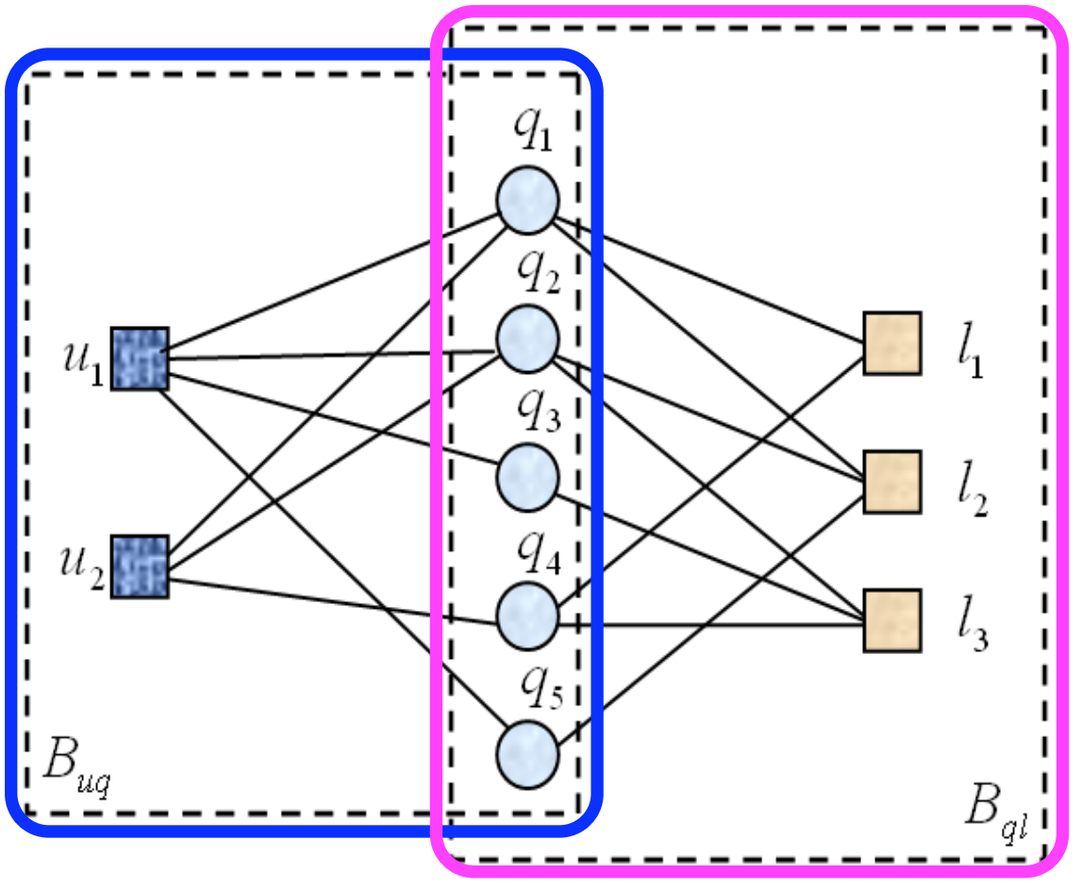

$B_{uq} = (V_{uq}, E_{uq})$
$V_{uq} = U \cup Q$
$U = \{u_1, u_2, ..., u_m\}$
$Q = \{q_1, q_2, ..., q_n\}$
$E_{uq} = \{(u_i, q_j) |$ there is an edge from u_i to $q_j\}$ is the set of all edges.
The edge (u_i, q_j) exists in this bipartite graph if and only if a user u_i issued a query q_j.

$B_{ql} = (V_{ql}, E_{ql})$
$V_{ql} = Q \cup L$
$Q = \{q_1, q_2, ..., q_n\}$
$L = \{l_1, l_2, ..., l_p\}$
$E_{ql} = \{(q_i, l_j) |$ there is an edge from q_i to $l_j\}$ is the set of all edges.
The edge (q_j, l_k) exists if and only if a user u_i clicked a URL l_k after issuing an query q_j.

Key Points

- Two-level latent semantic analysis

Level 1

- Consider the use of a joint user-query and query-URL bipartite graphs for query suggestion
- Use matrix factorization for learning query features in constructing the Query Similarity Graph

Level 2

- Use heat diffusion for similarity propagation for query suggestions

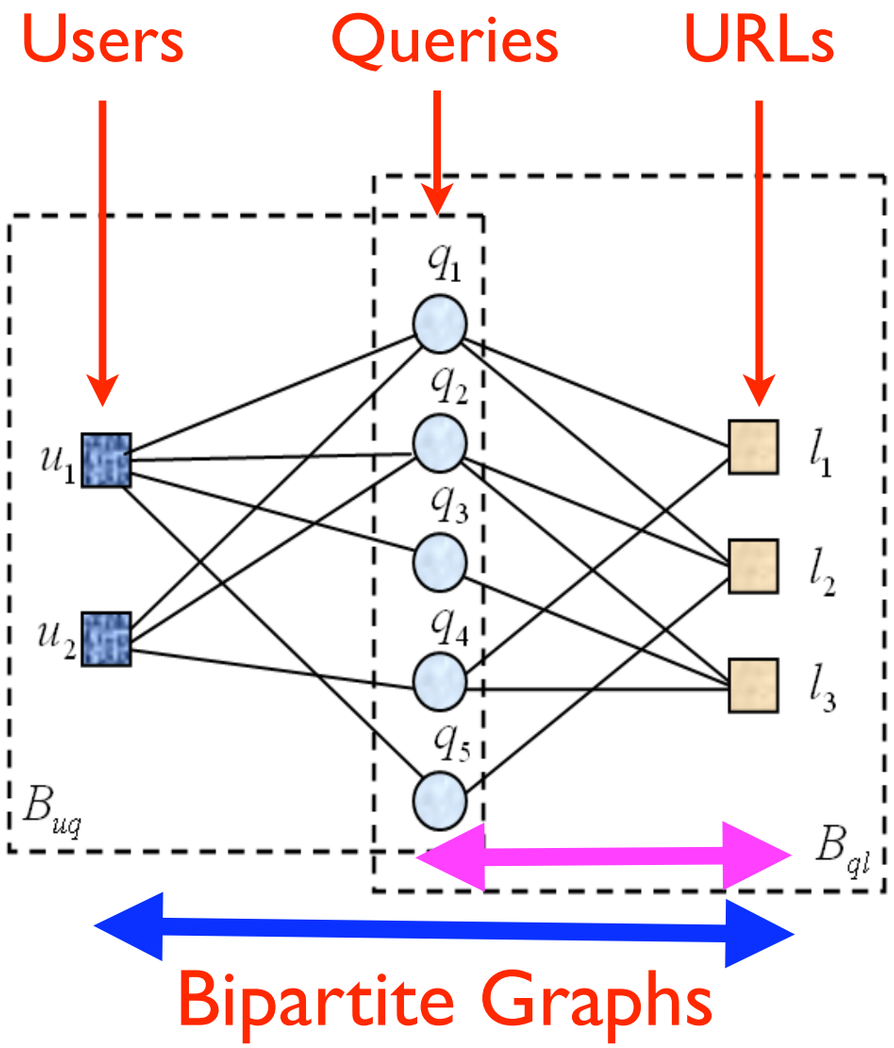

Bipartite Graphs

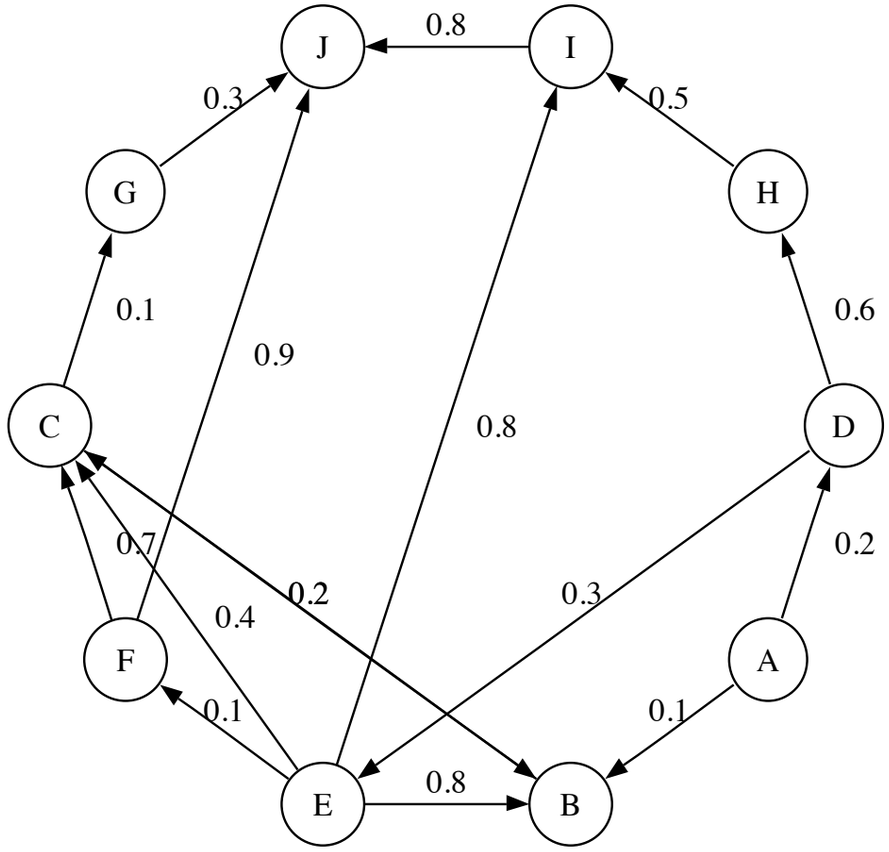

Query Similarity Graph

- Queries are issued by the users, and which URLs to click are also decided by the users
- Two distinct users are similar if they issued similar queries
- Two queries are similar if they are issued by similar users

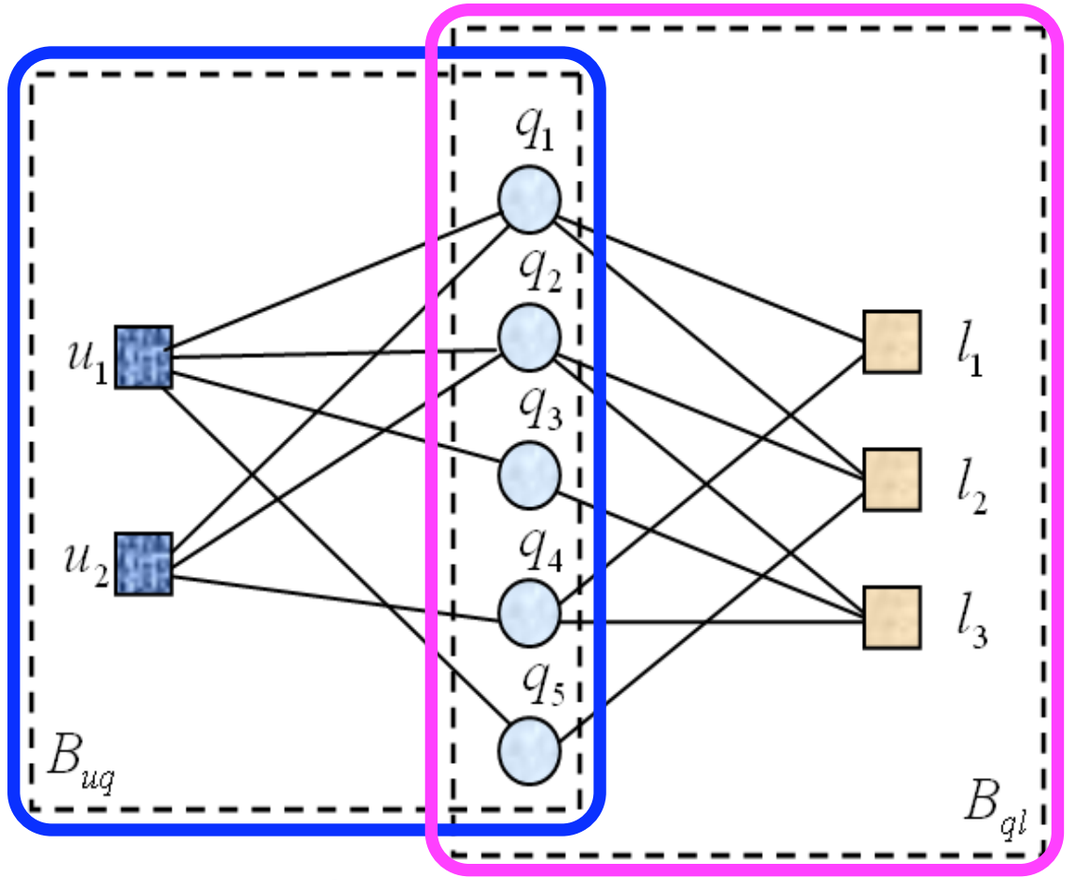

r^*_{ij} Normalized weight, how many times u_i issued q_j

s^*_{jk} Normalized weight, how many times q_j is linked to l_k

U_i L-dimensional vector of user u_i

Q_j L-dimensional vector of query q_j

L_k L-dimensional vector of URL l_k

$$\mathcal{H}(R,U,Q) = \min_{U,Q} \frac{1}{2} \sum_{i=1}^{m} \sum_{j=1}^{n} I^R_{ij} (r^*_{ij} - g(U_i^T Q_j))^2 + \frac{\alpha_u}{2} \|U\|_F^2 + \frac{\alpha_q}{2} \|Q\|_F^2$$

$$\mathcal{H}(S,Q,L) = \min_{Q,L} \frac{1}{2} \sum_{j=1}^{n} \sum_{k=1}^{p} I^S_{jk} (s^*_{jk} - g(Q_j^T L_k))^2 + \frac{\alpha_q}{2} \|Q\|_F^2 + \frac{\alpha_l}{2} \|L\|_F^2$$

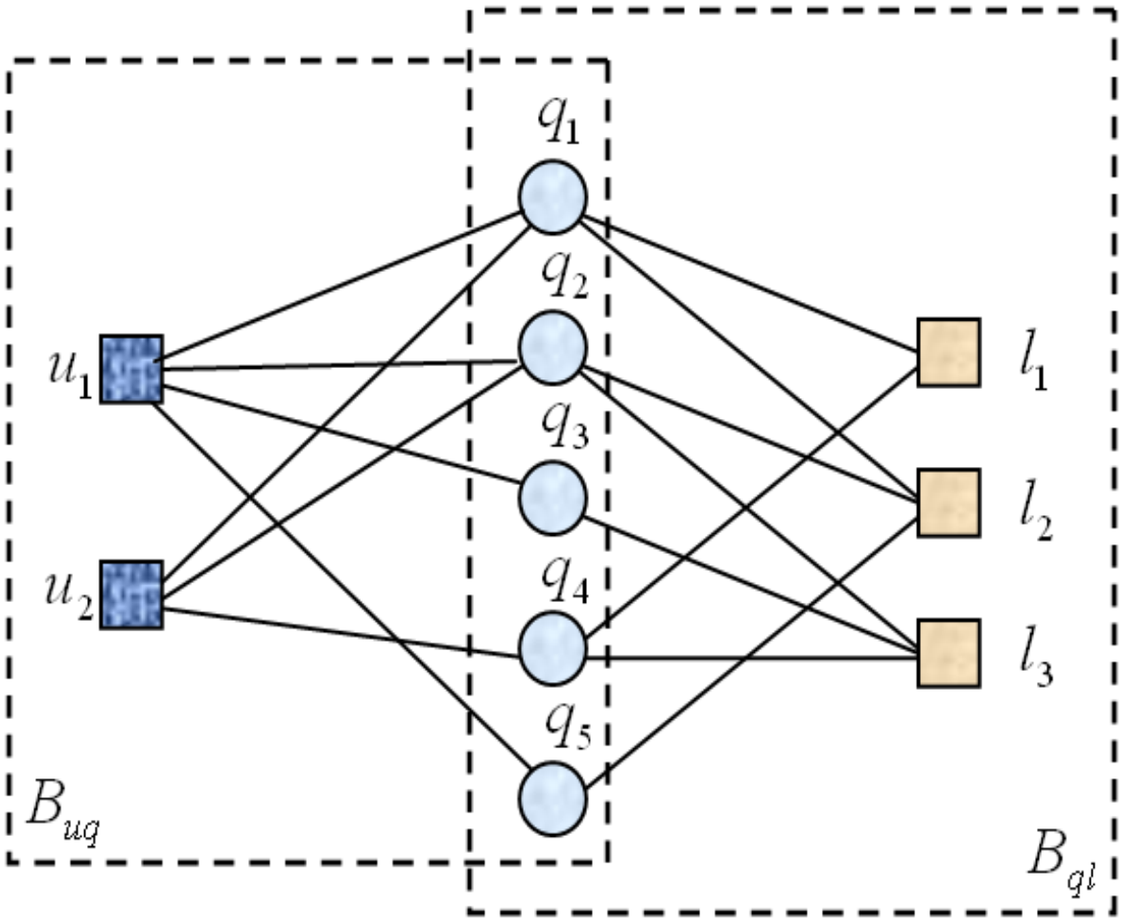

$$\mathcal{H}(S,R,U,Q,L) =$$
$$\frac{1}{2}\sum_{j=1}^{n}\sum_{k=1}^{p} I_{jk}^{S}(s_{jk}^{*}-g(Q_j^T L_k))^2 + \frac{\alpha_r}{2}\sum_{i=1}^{m}\sum_{j=1}^{n} I_{ij}^{R}(r_{ij}^{*}-g(U_i^T Q_j))^2$$
$$+\frac{\alpha_u}{2}\|U\|_F^2 + \frac{\alpha_q}{2}\|Q\|_F^2 + \frac{\alpha_l}{2}\|L\|_F^2,$$

- A local minimum can be found by performing gradient descent in U_i, Q_j and L_k

Gradient Descent Equations

$$\frac{\partial \mathcal{H}}{\partial U_i} = \alpha_r \sum_{j=1}^{n} I_{ij}^R g'(U_i^T Q_j)(g(U_i^T Q_j) - r_{ij}^*)Q_j + \alpha_u U_i,$$

$$\frac{\partial \mathcal{H}}{\partial Q_j} = \sum_{k=1}^{p} I_{jk}^S g'(Q_j^T L_k)(g(Q_j^T L_k) - s_{jk}^*)L_k + \alpha_r \sum_{i=1}^{m} I_{ij}^R g'(U_i^T Q_j)(g(U_i^T Q_j) - r_{ij}^*)U_i + \alpha_q Q_j,$$

$$\frac{\partial \mathcal{H}}{\partial L_k} = \sum_{j=1}^{n} I_{jk}^S g'(Q_j^T L_k)(g(Q_j^T L_k) - s_{jk}^*)Q_j + \alpha_l L_k,$$

Only the Q matrix, the queries' latent features, is being used to generate the query similarity graph!

Query Similarity Graph

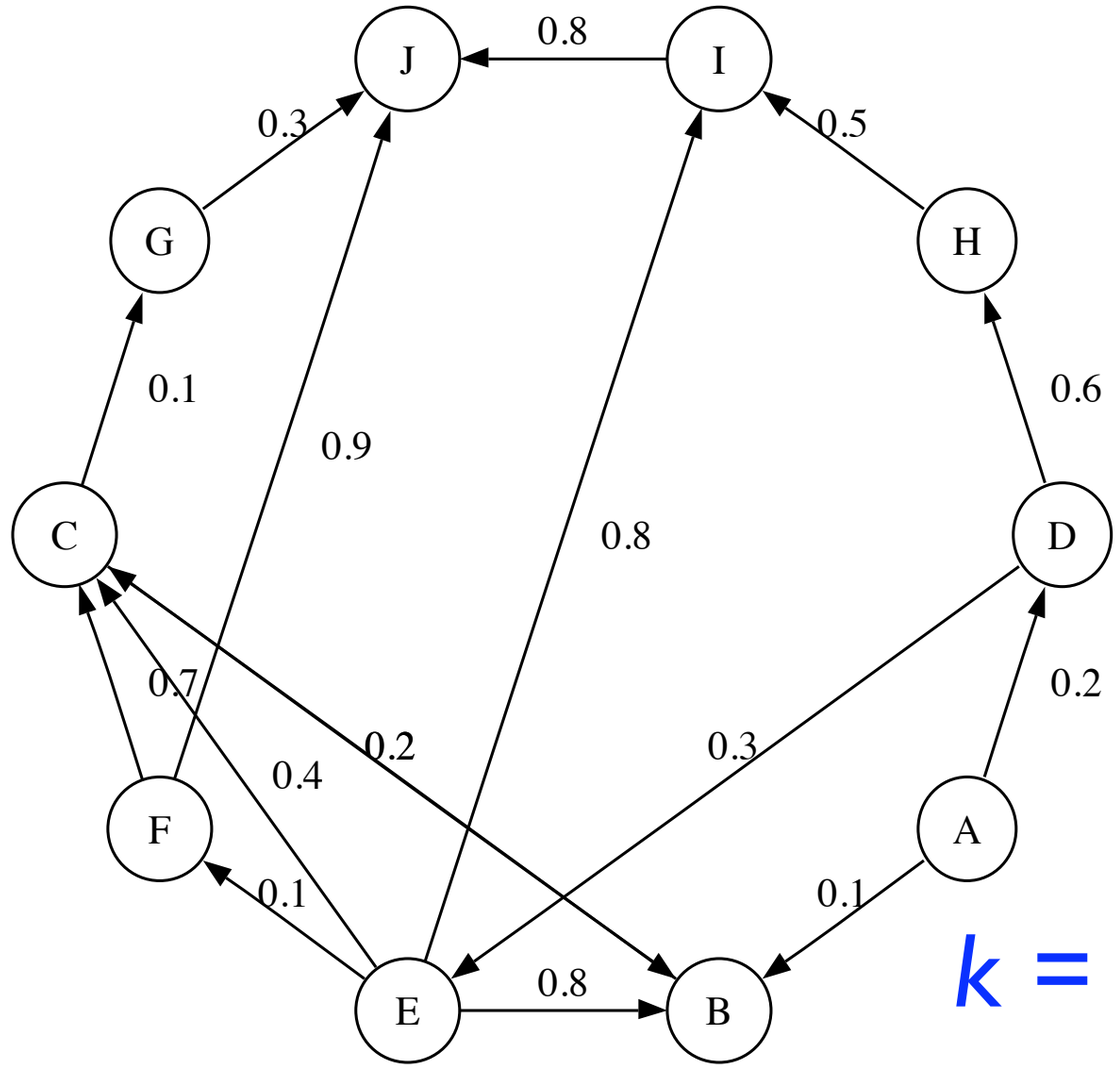

- Similarities are calculated using queries' latent features
- Only the top-*k* similar neighbors (terms) are kept

Similarity Propagation

- Based on the Heat Diffusion Model
- In the query graph, given the heat sources and the initial heat values, start the heat diffusion process and perform *P* steps
- Return the Top-*N* queries in terms of highest heat values for query suggestions

Heat Diffusion Model

- Heat diffusion is a physical phenomena
- Heat flows from high temperature to low temperature in a medium
- Heat kernel is used to describe the amount of heat that one point receives from another point
- The way that heat diffuse varies when the underlying geometry

$$\rho C_P \frac{\partial T}{\partial t} = Q + \nabla \cdot (k \nabla T)$$

ρ	Density
C_P	Heat capacity and constant pressure
$\frac{\partial T}{\partial t}$	Change in temperature over time
Q	Heat added
k	Thermal conductivity
∇T	Temperature gradient
$\nabla \cdot \mathbf{v}$	Divergence

Heat Diffusion Process

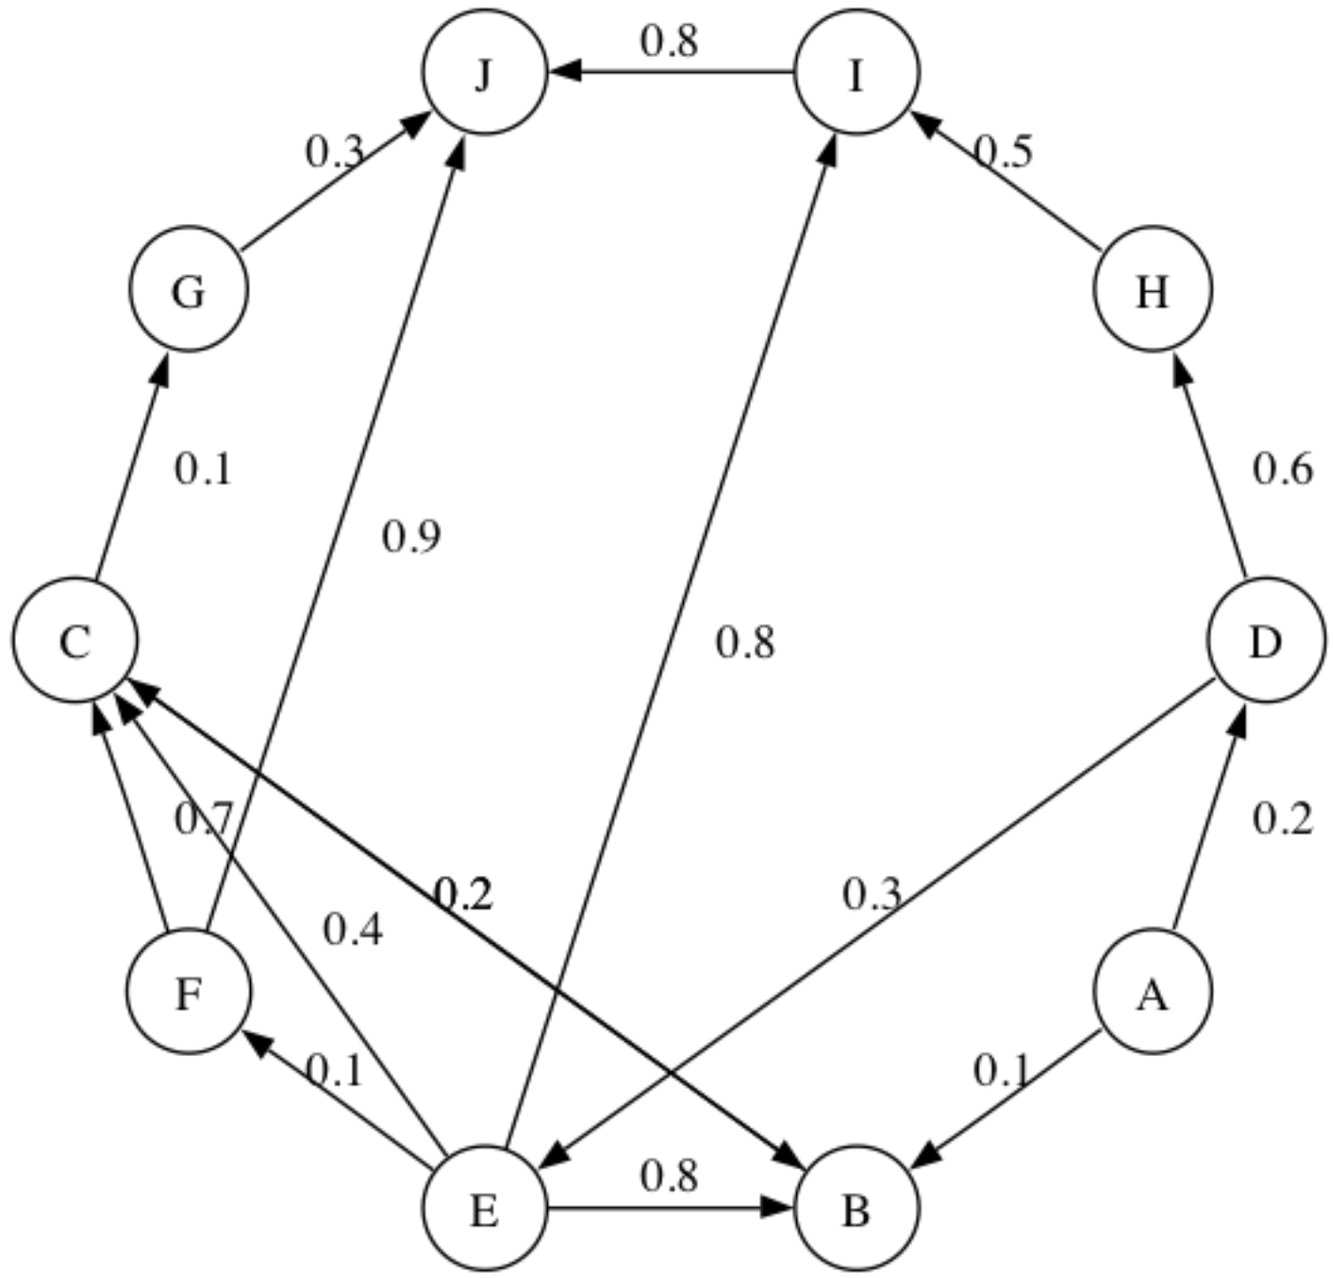

Similarity Propagation Model

$$\frac{f_i(t+\Delta t) - f_i(t)}{\Delta t} = \alpha \left(-\frac{\tau_i}{d_i} f_i(t) \sum_{k:(q_i,q_k)\in E} w_{ik} + \sum_{j:(q_j,q_i)\in E} \frac{w_{ji}}{d_j} f_j(t) \right) \quad (1)$$

$$\mathbf{f}(1) = e^{\alpha \mathbf{H}} \mathbf{f}(0) \quad (2)$$

$$H_{ij} = \begin{cases} w_{ji}/d_j, & (q_j, q_i) \in E, \\ -(\tau_i/d_i) \sum_{k:(i,k)\in E} w_{ik}, & i = j, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

$$\mathbf{f}(1) = e^{\alpha \mathbf{R}} \mathbf{f}(0), \quad \mathbf{R} = \gamma \mathbf{H} + (1-\gamma)\mathbf{g}\mathbf{1}^T \quad (4)$$

α	Thermal conductivity
d_i	Heat value of node i at time t
$f_i(t)$	Heat value of node i at time t
w_{ik}	Weight between node i and node k
$\mathbf{f}(0)$	Vector of the initial heat distribution
$\mathbf{f}(1)$	Vector of the heat distribution at time 1
τ_i	Equal to 1 if node i has outlinks, else equal to 0
γ	Random jump parameter, and set to 0.85
$\mathbf{g}$	Uniform stochastic distribution vector

Discrete Approximation

- Compute $e^{\alpha \mathbf{R}}$ is time consuming
- We use the discrete approximation to substitute

$$\mathbf{f}(1) = \left(\mathbf{I} + \frac{\alpha}{P}\mathbf{R}\right)^P \mathbf{f}(0)$$

- For every heat source, only diffuse heat to its neighbors within *P* steps
- In our experiments, *P* = 3 already generates fairly good results

Query Suggestion Procedure

- For a given query *q*

1. Select a set of *n* queries, each of which contains at least one word in common with *q*, as heat sources
2. Calculate the initial heat values by

$$f_{\hat{q}_i}(0) = \frac{|\mathcal{W}(q) \cap \mathcal{W}(\hat{q}_i)|}{|\mathcal{W}(q) \cup \mathcal{W}(\hat{q}_i)|}$$

q = "Sony"
"Sony" = 1
"Sony Electronics" = 1/2
"Sony Vaio Laptop" = 1/3

3. Use $\mathbf{f}(1) = e^{\alpha \mathbf{R}} \mathbf{f}(0)$ to diffuse the heat in graph
4. Obtain the Top-*N* queries from $\mathbf{f}(1)$

Physical Meaning of α

- If set α to a large value
  - The results depend more on the query graph, and more semantically related to original queries, e.g., travel => lowest air fare
- If set α to a small value
  - The results depend more on the initial heat distributions, and more literally similar to original queries, e.g., travel => travel insurance

Experimental Dataset

Data Source	Clickthrough data from AOL search	After Pre-Processing
Collection Period	March 2006 to May 2006 (3 months)	
Lines of Logs	19,442,629	
Unique user IDS	657,426	192,371
Unique queries	4,802,520	224,165
Unique URLs	1,606,326	343,302
Unique words		69,937

Query Suggestions

Table 2: Examples of LSQS Query Suggestion Results ($k = 50$)

Testing Queries	Suggestions				
	$\alpha = 10$			$\alpha = 1000$	
	Top 1	Top 2	Top 3	Top 4	Top 5
michael jordan	michael jordan shoes	michael jordan bio	pictures of michael jordan	nba playoff	nba standings
travel	travel insurance	abc travel	travel companions	hotel tickets	lowest air fare
java	sun java	java script	java search	sun microsystems inc	virtual machine
global services	ibm global services	global technical services	staffing services	temporary agency	manpower professional
walt disney land	world of disney	disney world orlando	disney world theme park	disneyland grand hotel	disneyland in california
intel	intel vs amd	amd vs intel	pentium d	pentium	centrino
job hunt	jobs in maryland	monster job	jobs in mississippi	work from home online	monster board
photography	photography classes	portrait photography	wedding photography	adobe elements	canon lens
internet explorer	ms internet explorer	internet explorer repair	internet explorer upgrade	microsoft com	security update
fitness	fitness magazine	lifestyles family fitness	fitness connection	womens health magazine	family fitness
m schumacher	schumacher	red bull racing	formula one racing	ferrari cars	formula one
solar system	solar system project	solar system facts	solar system planets	planet jupiter	mars facts
sunglasses	replica sunglasses	cheap sunglasses	discount sunglasses	safilo	marhon
search engine	audio search engine	best search engine	search engine optimization	song lyrics search	search by google
disease	grovers disease	liver disease	morgellons disease	colic in babies	oklahoma vital records
pizzahut	pizza hut menu	pizza coupons	pizza hut coupons	papa johns pizza coupon	papa johns
health care	health care proxy	universal health care	free health care	great west healthcare	uhc
flower delivery	global flower delivery	online florist	flowers online	send flowers	virtual flower
wedding	wedding guide	wedding reception ideas	wedding decoration	unity candle	centerpiece ideas
astronomy	astronomy magazine	astronomy pic of the day	star charts	space pictures	comet

Comparisons

Table 3: Comparisons between LSQS and SimRank

	Top 1	Top 2	Top 3	Top 4	Top 5
jaguar					
LSQS	jaguar cat	jaguar commercial	jaguar parts	jaguarundi	leopard
SimRank	american black bear	bottlenose dolphin	leopard	margay	jaguarundi
apple					
LSQS	apple computers	apple ipod	apple diet	apple vacations	apple bottom
SimRank	ipod troubleshooting	apple quicktime	apple ipods	apple computers	apple software

Table 4: Accuracy Comparisons

Accuracy	**LSQS**	**SimRank**
By Experts	0.8413	0.7101
By ODP	0.6823	0.5789

ODP, Open Directory Project, see http://dmoz.org

Summary

- Propose an offline novel joint matrix factorization method using user-query and query-URL bipartite graphs for learning query features
- Propose an online diffusion-based similarity propagation and ranking method for query suggestion

Conclusion

- Social Computing is a paradigm shift (範式轉變)!
- Novel views on the spatial and temporal relationship among social entities!
- Great opportunities in a new research direction!

On-Going Research

Machine Learning

- Direct Zero-norm Optimization for Feature Selection (ICDM'08)
- Semi-supervised Learning from General Unlabeled Data (ICDM'08)
- Learning with Consistency between Inductive Functions and Kernels (NIPS'08)
- An Extended Level Method for Efficient Multiple Kernel Learning (NIPS'08)
- Semi-supervised Text Categorization by Active Search (CIKM'08)
- Transductive Support Vector Machine (NIPS'07)
- Global and local learning (ICML'04, JMLR'04)

Web Intelligence

- Effective Latent Space Graph-based Re-ranking Model with Global Consistency (WSDM'09)
- Formal Models for Expert Finding on DBLP Bibliography Data (ICDM'08)
- Learning Latent Semantic Relations from Query Logs for Query Suggestion (CIKM'08)
- RATE: a Review of Reviewers in a Manuscript Review Process (WI'08)
- MatchSim: link-based web page similarity measurements (WI'07)
- Diffusion rank: Ranking web pages based on heat diffusion equations (SIGIR'07)
- Web text classification (WWW'07)

Collaborative Filtering

- Recommender system: accurate recommendation based on sparse matrix (SIGIR'07)
- SoRec: Social Recommendation Using Probabilistic Matrix Factorization (CIKM'08)

Human Computation

- An Analytical Study of Puzzle Selection Strategies for the ESP Game (WI'08)
- An Analytical Approach to Optimizing The Utility of ESP Games (WI'08)

Acknowledgments

- Prof. Michael R. Lyu
- Prof. Jimmy Lee
- Dr. Kaizhu Huang
- Dr. Haixuan Yang
- Thomas Chan (M.Phil)
- Hongbo Deng (Ph.D.)
- Zhenjiang Lin (Ph.D.)
- Hao Ma (Ph.D.)
- Haiqin Yang (Ph.D.)
- Xin Xin (Ph.D.)
- Zenglin Xu (Ph.D.)
- Chao Zhou (Ph.D.)

Q & A

http://www.cse.cuhk.edu.hk/~king