



Ranking linked data

Web graph, PageRank, Topic-specific PageRank and HITS

Web Search


Graph Data: Social Networks



Facebook social graph

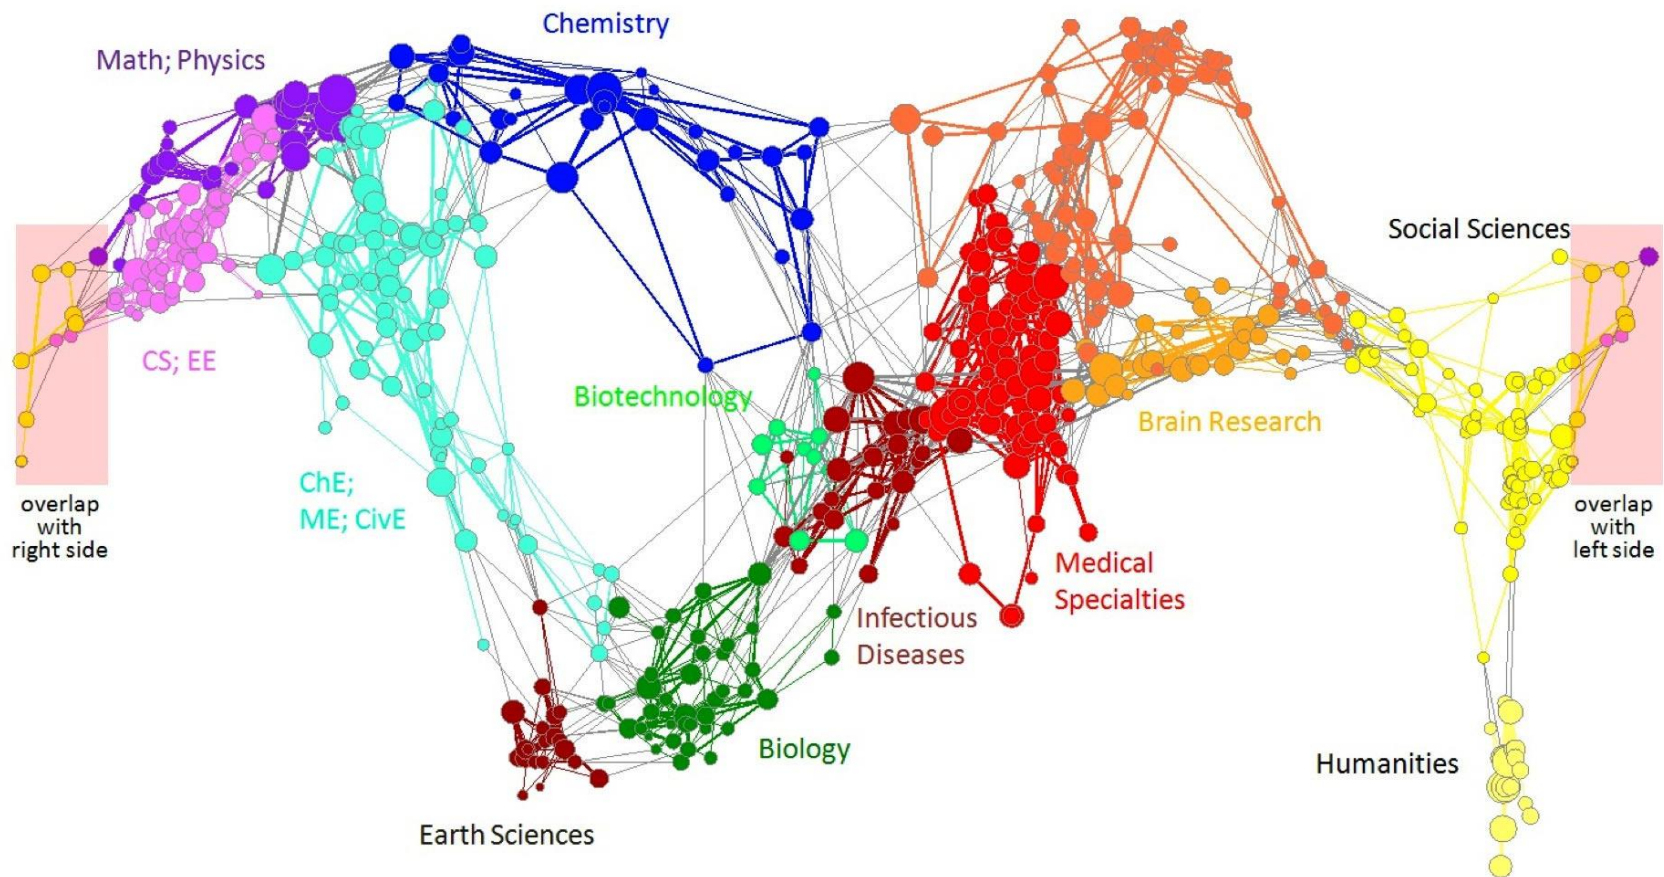
4-degrees of separation [Backstrom-Boldi-Rosa-Ugander-Vigna, 2011]

Graph Data: Media Networks



Connections between political blogs
Polarization of the network [Adamic-Glance, 2005]


Graph Data: Information Nets




Citation networks and Maps of science
[Börner et al., 2012]

Ranking linked data

- Links are inserted by humans.
- **They are one of the most valuable judgments of a page's importance.**
- A link is inserted to denote an association. The anchor text describes the type of association.



The Web as a directed graph

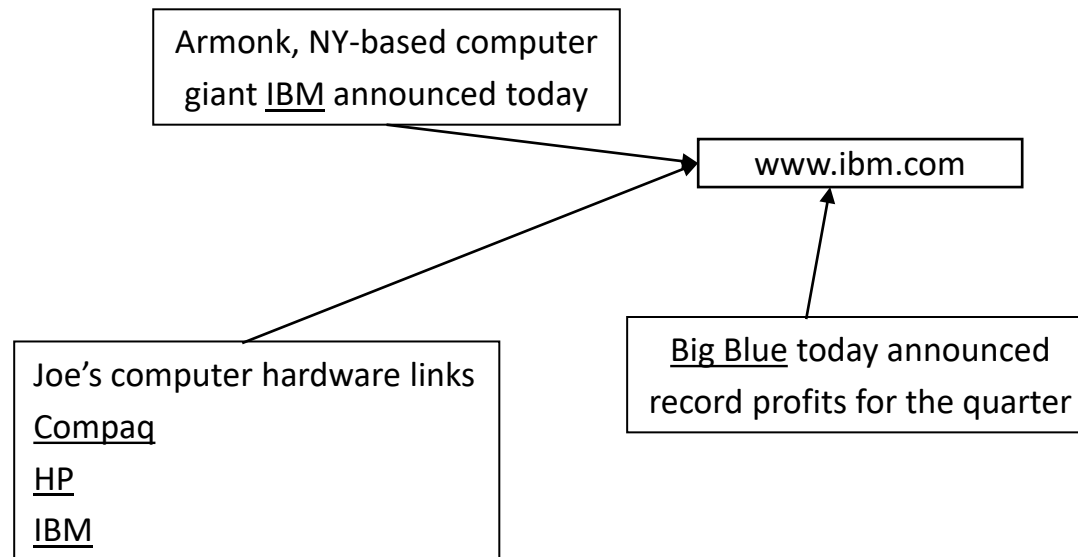


Assumption 1: A hyperlink between pages denotes author perceived relevance (quality signal)

Assumption 2: The anchor of the hyperlink describes the target page (textual context)

Anchor text

- When indexing a document D , include anchor text from links pointing to D .



Indexing anchor text

Sec 21.1.1 Page 8


- Can sometimes have unexpected side effects - *e.g., evil empire.*
- Can boost anchor text with weight depending on the authority of the anchor page's website
 - E.g., if we were to assume that content from cnn.com or yahoo.com is authoritative, then trust the anchor text from them

Citation analysis

- Citation frequency
- Co-citation coupling frequency
 - Co-citations with a given author measures “impact”
 - Co-citation analysis [Mcca90]
- Bibliographic coupling frequency
 - Articles that co-cite the same articles are related
- Citation indexing
 - Who is author cited by? [Garf72]
- PageRank preview: Pinski and Narin '60s


Query-independent ordering

- First generation: using link counts as simple measures of popularity.
- Two basic suggestions:
 - Undirected popularity:
 - Each page gets a score = the number of in-links plus the number of out-links ($3+2=5$).
 - Directed popularity:
 - Score of a page = number of its in-links (3).




Incoming and outgoing links

- The popularity of a page is related to the number of incoming links
 - Positively popular
 - Negatively popular
- The popularity of a page is related to the popularity of pages pointing to them



PageRank scoring


- Imagine a browser doing a random walk on web pages:
 - Start at a random page
 - At each step, go out of the current page along one of the links on that page, equiprobably



- “In the steady state” each page has a long-term visit rate - use this as the page’s score.

Not quite enough

- The web is full of dead-ends.
 - Random walk can get stuck in dead-ends.
 - Makes no sense to talk about long-term visit rates.




Teleporting

- At a dead end, jump to a random web page.
- At any non-dead end, with probability 10%, jump to a random web page.
 - With remaining probability (90%), go out on a random link.
 - 10% - a parameter.
- Result of teleporting:
 - Now cannot get stuck locally.
 - There is a long-term rate at which any page is visited.
 - How do we compute this visit rate?

The random surfer


- The PageRank of a page is the probability that a given random “Web surfer” is currently visiting that page.




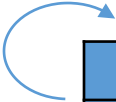
- This probability is related to the incoming links and to a certain degree of browsing randomness (e.g. reaching a page through a search engine).

Markov chains

- A Markov chain consists of n states, plus an $n \times n$ transition probability matrix \mathbf{P} .
- At each step, we are in exactly one of the states.
- For $1 \leq i, j \leq n$, the matrix entry P_{ij} tells us the probability of j being the next state, given we are currently in state i .



Transitions probability matrix





	A	B	C	D
A	0	1	1	1
B	1	0	0	0
C	0	1	0	1
D	0	1	0	0

	A	B	C	D
A	0	P_{ab}	P_{ac}	P_{ad}
B	P_{ba}	0	0	0
C	0	P_{cb}	0	P_{cd}
D	0	P_{db}	0	0

Ergodic Markov chains

- A Markov chain is ergodic if
 - you have a path from any state to any other
 - For any start state, after a finite transient time T_0 , the probability of being in any state at a fixed time $T > T_0$ is nonzero.




Ergodic Markov chains

- For any ergodic Markov chain, there is a unique long-term visit rate for each state.
 - Steady-state probability distribution.
- Over a long time-period, we visit each state in proportion to this rate.
- It doesn't matter where we start.

The PageRank of Web page i corresponds to the probability of being at page i after an infinite random walk across all pages (i.e., the stationary distribution).

PageRank


- The rank of a page is related to the number of incoming links of that page and the rank of the pages linking to it.



$$PR(A) = \frac{(1 - d)}{N - 1} + d \cdot \left[\frac{PR(B)}{OL(B)} + \frac{PR(C)}{OL(C)} \right]$$

PageRank: formalization

- The RandomSurfer model assumes that the pages with more inlinks are visited more often




- The rank of a page j is computed as:


$$\Pr(p_j) = \frac{1-d}{N-1} + d \sum_{i \in \text{Parents}(j)} \frac{\Pr(p_i)}{OL_i}$$

where OL_i is the number of outgoing links of page i and p_i is the PageRank of that page

Transitions probability matrix




	A	B	C	D
A	0	1	1	1
B	1	0	0	0
C	0	1	0	1
D	0	1	0	0



	A	B	C	D
A	0	P_{ab}	P_{ac}	P_{ad}
B	P_{ba}	0	0	0
C	0	P_{cb}	0	P_{cd}
D	0	P_{db}	0	0

Example


- Consider three Web pages:



- Compute the PageRank of each page on the following Web graph. Consider $d = 0.9$.

$$\Pr(p_j) = \frac{1-d}{N-1} + d \sum_{i \in \text{Parents}(j)} \frac{\Pr(p_i)}{OL_i}$$

Power iteration method



$$\Pr(p_j) = \frac{1-d}{N-1} + d \sum_{i \in \text{Parents}(j)} \frac{\Pr(p_i)}{OL_i}$$

Page	Iteration 0	Iteration 1	Iteration 2
A	$\frac{1}{3}$	$\frac{0.1}{2}$...
B	$\frac{1}{3}$	$\frac{0.1}{2} + 0.9 \cdot \left(\frac{1}{3}\right)$...
C	$\frac{1}{3}$	$\frac{0.1}{2} + 0.9 \cdot \left(\frac{1}{3} + \frac{1}{1}\right)$...

Power iteration method

```
n = X.shape[0]
X = X.copy()
incoming_counts = np.asarray(X.sum(axis=1)).ravel()

print("Normalizing the graph")
for i in incoming_counts.nonzero()[0]:
    X.data[X.indptr[i]:X.indptr[i + 1]] *= 1.0 / incoming_counts[i]
dangle = np.asarray(np.where(np.isclose(X.sum(axis=1), 0),
                                1.0 / n, 0)).ravel()

scores = np.full(n, 1. / n, dtype=np.float32) # initial guess
for i in range(max_iter):
    print("power iteration #%d" % i)
    prev_scores = scores
    scores = (alpha * (scores * X + np.dot(dangle, prev_scores))
              + (1 - alpha) * prev_scores.sum() / n)
    # check convergence: normalized L_inf norm
    scores_max = np.abs(scores).max()
    if scores_max == 0.0:
        scores_max = 1.0
    err = np.abs(scores - prev_scores).max() / scores_max
    print("error: %.6f" % err)
    if err < n * tol:
        return scores

return scores
```

PageRank-Iterate(G)

$P_0 \leftarrow e/n$

$k = 1$

repeat

$P_{k+1} \leftarrow (1 - d)e + dA^T P_k$;

$k = k + 1$;

until $\|P_{k+1} - P_k\|_1 < \varepsilon$

return P_{k+1}

PageRank: issues and variants

- How realistic is the random surfer model?
 - What if we modeled the back button? [Fagi00]
 - Surfer behavior sharply skewed towards short paths [Hube98]
 - Search engines, bookmarks & directories make jumps non-random.
- Biased Surfer Models
 - Weight edge traversal probabilities based on match with topic/query (non-uniform edge selection)
 - Bias jumps to pages on topic (e.g., based on personal bookmarks & categories of interest)

Topic Specific PageRank [Have02]

- The idea is that links between pages of the same category are more likely to be followed.
- Conceptually, we use a random surfer who teleports, with ~10% probability, using the following rule:
 - Selects a category (say, one of the 16 top level categories) based on a query & user -specific distribution over the categories
 - Teleport to a page uniformly at random within the chosen category.
 - Do not teleport to a pages outside the chosen category.

Topic Specific PageRank - Implementation

- **offline:** Compute PageRank distributions wrt individual categories
 - Query independent model as before
 - Each page has multiple PageRank scores – one for each category, with teleportation only to that category
- **online:** Distribution of weights over categories computed by query context classification
 - Generate a dynamic PageRank score for each page - weighted sum of category-specific PageRanks

Offline: Web page topic classifier


- Web pages have specific topics that can be detected by some classifier.
- Links are more likely between topics of the same topic.
- Links between pages of the same topic are more likely to be followed.

Offline: Topic Specific PageRank

- Compute the PageRank of each page for each topic.
- Key difference: The teleporting is canceled between pages of different topics.

Page	PR_Health	PR_Sports	PR_Travel	PR_Business
1	0.54	0.23	0.78	0.0190
2 (sports)	0.65	0.18	0.25	0.0182
3	0.78	0.09	0.72	0.0178
...				


Online: Query topic classification



Query category = 80% sports + 20% health

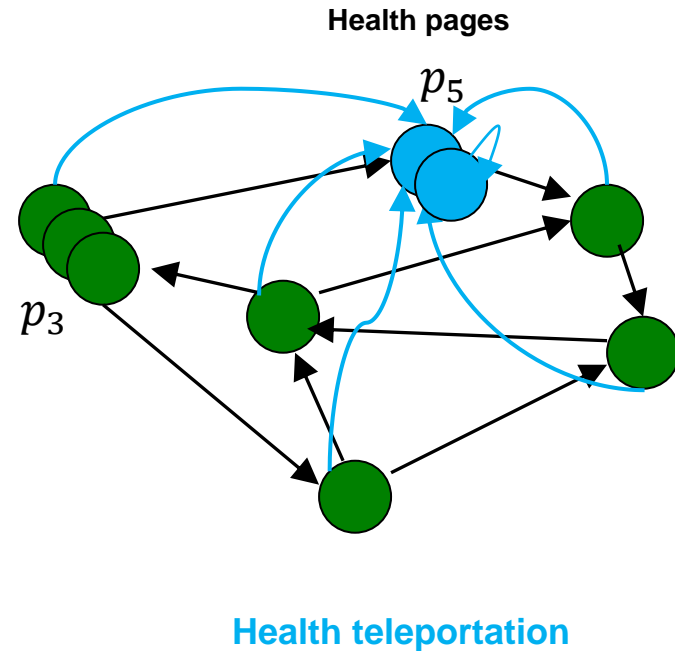
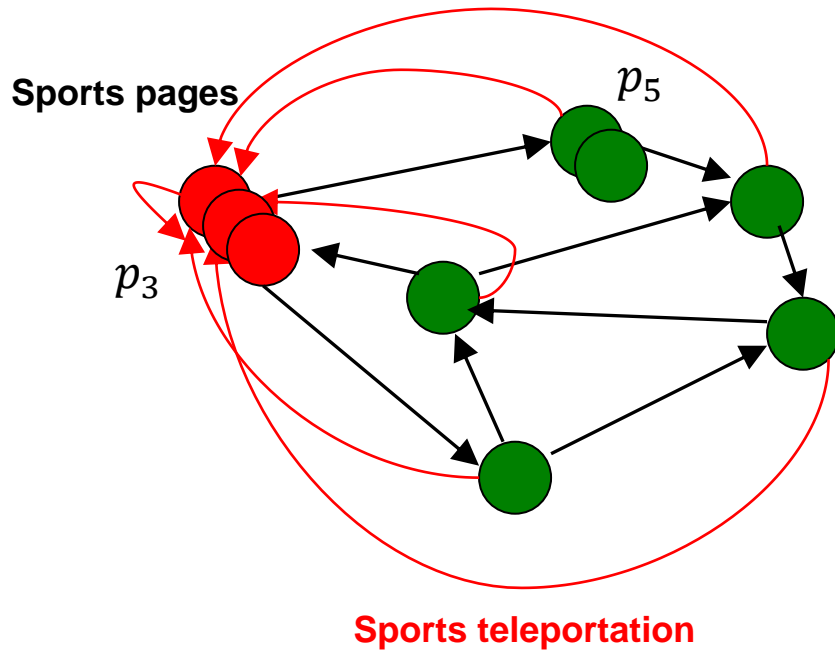
Example

- Consider a query on a given set of Web pages with the following graph:



- The query has 80% probability of being about Sports.
- The query has 20% probability of being about Health.

Non-uniform Teleportation



$$PR_{\text{Sports}}(p_3) = \frac{1-d}{N-1} + d \sum_{i \in \text{Parents}(p_3)} \frac{\Pr(p_i)}{OL_i}$$


$$PR_{\text{Sports}}(p_5) = \sum_{i \in \text{Parents}(p_5)} \frac{\Pr(p_i)}{OL_i}$$

$$PR_{\text{Health}}(p_3) = \sum_{i \in \text{Parents}(p_3)} \frac{\Pr(p_i)}{OL_i}$$

$$PR_{\text{Health}}(p_5) = \frac{1-d}{N-1} + d \sum_{i \in \text{Parents}(p_5)} \frac{\Pr(p_i)}{OL_i}$$

Interpretation

Query has 80% sports probability and 20% health probability




$$pr = (0.8 PR_{\text{sports}} + 0.2 PR_{\text{health}})$$

80% Sports PageRank, 20% Health PageRank

Hyperlink-Induced Topic Search (HITS) - Klei98

- In response to a query, instead of an ordered list of pages each meeting the query, find two sets of inter-related pages:
 - *Hub pages* are good lists of links on a subject.
 - e.g., “Bob’s list of cancer-related links.”
 - *Authority pages* occur recurrently on good hubs for the subject.
- Best suited for “broad topic” queries rather than for page-finding queries.
- Gets at a broader slice of common *opinion*.

The hope




Best Mobile+Net+TV bundles.

High-level scheme

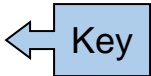
- Extract from the web a base set of pages that could be good hubs or authorities.
- From these, identify a small set of top hub and authority pages;
 - iterative algorithm.

Base set and root set

- Given text query (say *browser*), use a text index to get all pages containing *browser*.
 - Call this the root set of pages.
- Add in any page that either
 - points to a page in the root set, or
 - is pointed to by a page in the root set.
- Call this the base set.




Distilling hubs and authorities


- Compute, for each page x in the base set, a hub score $h(x)$ and an authority score $a(x)$.
- Initialize: for all x , $h(x) \leftarrow 1$; $a(x) \leftarrow 1$;
- Iteratively update all $h(x)$, $a(x)$; 
- After iterations
 - output pages with highest $h()$ scores as top hubs
 - highest $a()$ scores as top authorities.

Iterative update

- Repeat the following updates, for all x :



$$a(x) \leftarrow \sum_{y \mapsto x} h(y)$$



$$h(x) \leftarrow \sum_{x \mapsto y} a(y)$$

How many iterations?

- Claim: relative values of scores will converge after a few iterations:
 - in fact, suitably scaled, $h()$ and $a()$ scores settle into a steady state!
- We only require the relative orders of the $h()$ and $a()$ scores - not their absolute values.
- In practice, ~ 5 iterations get you close to stability.

PageRank in Social Media

- Can be applied to rank people, entities, news sources, etc.
- The Who to Follow services is a:
 - personalized version of PageRank
 - combines PageRank with recommender systems

WTF: The Who to Follow Service at Twitter

Pankaj Gupta, Ashish Goel, Jimmy Lin, Aneesh Sharma, Dong Wang, Reza Zadeh

Twitter, Inc.
@pankaj @ashishgoel @lntool @aneeshs @dongwang218 @reza_zadeh

ABSTRACT

WTF ("Who to Follow") is Twitter's user recommendation service, which is responsible for creating millions of connections daily between users based on shared interests, common connections, and other related factors. This paper provides an architectural overview and shares lessons we learned in building and running the service over the past few years. Particularly noteworthy was our design decision to process the entire Twitter graph in memory on a single server, which significantly reduced architectural complexity and allowed us to develop and deploy the service in only a few months. At the core of our architecture is Casowary, an open-source in-memory graph processing engine we built from scratch for WTF. Besides powering Twitter's user recommendations, Casowary is also used for search, discovery, promoted products, and other services as well. We describe and evaluate a few graph recommendation algorithms implemented in Casowary, including a novel approach based on a combination of random walks and SALSA. Looking into the future, we revisit the design of our architecture and comment on its limitations, which are presently being addressed in a second-generation system under development.

Categories and Subject Descriptors: H.2.8 [Database Management]: Database applications *Data mining*

General Terms: Algorithms, Design

Keywords: graph processing, link prediction, Hadoop

1. INTRODUCTION

The lifeblood of a vibrant and successful social media service is an active and engaged user base. Therefore, maintaining and expanding the active user population is a top priority, and for Twitter, this is no exception. At the core, Twitter is a platform remarkably simple in concept: a user "follows" other users to subscribe to their 140-character tweets, which may be received on a variety of clients (e.g., the twitter.com website, mobile clients on iPhones and Android devices, etc.). The vibrancy of the service, whether in informing users of relevant breaking news or connecting them to communities of interest, derives from its users—all 200 million of them, collectively posting over 400 million tweets every day (as of early 2013).

One important way to sustain and grow Twitter is to help users, existing and new, discover connections. This is the

goal of WTF ("Who to Follow").¹ In the current interface, the WTF box is prominently featured in the left rail of the web client as well as in many other contexts across multiple platforms. WTF suggests Twitter accounts that a user may be interested in following, based on shared interests, common connections, and a number of other factors. Social networking sites such as Facebook and LinkedIn have comparable offerings as well. We identify two distinct but complementary facets to the problem, which we informally call "interested in" and "similar to". For example, a user interested in sports might follow *Benigno*, but we probably wouldn't consider that user similar to *Benigno*. On the other hand, two users might be considered similar based on their shared interest in, say, basketball, or if they follow many of the same users. Twitter also exposes profile similarity as a product feature, visible when visiting a user's profile page. Throughout this paper, our discussion of user recommendations covers both these aspects. Based on the homophily principle, similar users also make good suggestions. Besides powering user recommendations, WTF is also used for search relevance, discovery, promoted products, and other services as well.

This paper provides an overview of Twitter's WTF service, a project that began in spring 2010 and went into production the same summer.² Quite explicitly, our goal here is not to present novel research contributions, but to share the overall design of a system that is responsible for creating millions of connections daily and lessons that we have learned over the past few years.

We view this paper as having the following contributions:

- First, we explain and justify our decision to build the service around the assumption that the entire graph will fit in memory on a single machine. This might strike the reader as an odd assumption, especially in the era of "big data", but our approach has worked well and in retrospect we believe the decision was correct in the broader context of being able to launch the product quickly.

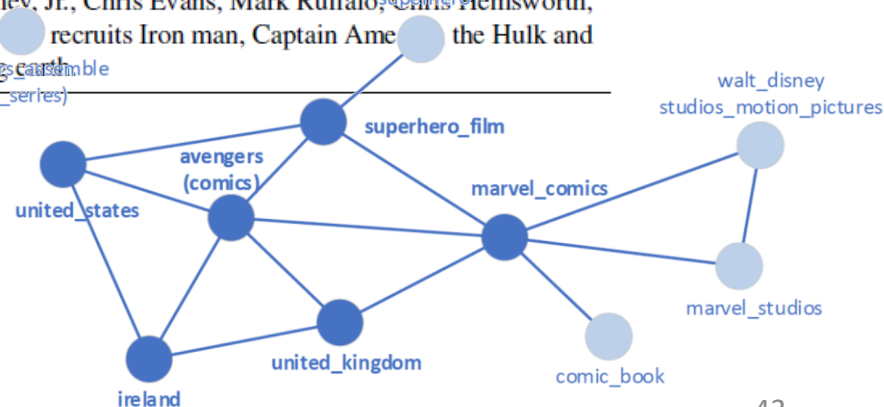
- Second, we describe the complete end-to-end architecture of the WTF service. At the core is Casowary, our open-source in-memory graph processing engine.

¹The confusion with the more conventional expansion of the acronym is intentional and the butt of many internal jokes. Also, it has not escaped our attention that the name of the service is actually ungrammatical; the persons should properly be in the objective case, as in "whom to follow".

²<http://blog.twitter.com/2010/07/discovering-who-to-follow.html>


PageRank in Conversational Agents

Query Turn:	Who are The Avengers?
Passage 1	The Avengers (2012 film) Marvel's The Avengers (classified under the name Marvel Avengers Assemble in the United Kingdom and Ireland), or simply The Avengers , is a 2012 American superhero film based on the Marvel Comics superhero team of the same name, produced by Marvel Studios and distributed by Walt Disney Studios Motion Pictures.
Passage 2	Marvel's The Avengers (Marvel Avengers Assemble in the UK and Ireland) more commonly known as The Avengers , is a 2012 American superhero film , scripted and directed by Joss Whedon, based on the Marvel Comics superhero team of the same name. The film stars an ensemble cast consisting of Robert Downey, Jr., Chris Evans, Mark Ruffalo, Chris Hemsworth, Scarlett Johansson, Jeremy Renner, Tom Hiddleston, Clark Gregg, Cobie Smulders, Stellan Skarsgård and Samuel L. Jackson. In The Avengers , Nick Fury (Jackson), director of the peacekeeping organization S.H.I.E.L.D., recruits Iron Man (Downey), Captain America (Evans), the Hulk (Ruffalo), and Thor (Hemsworth) to form a team that must stop Thor's adoptive brother Loki (Hiddleston) from subjugating Earth.
Passage 3	The Avengers (also known as Marvel's The Avengers and classified in the UK and Ireland under the title Marvel Avengers Assemble) is a 2012 American superhero film produced by Marvel Studios and distributed by Walt Disney Studios Motion Pictures, based on the Marvel Comics superhero team of the same name.
T5-BASE Answer	The Avengers is a 2012 american superhero film based on the Marvel comics superhero team of the same name. the film stars an ensemble cast consisting of Robert Downey, Jr., Chris Evans, Mark Ruffalo, Chris Hemsworth, Scarlett Johansson and Jeremy Renner. In the film, Nick Fury recruits Iron man, Captain Ame the Hulk and Thor to form a team that must stop Loki from subjugating earth.




Summary

- Web graphs denote a relation of relevance between edges
- Introduced a new way of modeling the value of Web links.
- Key algorithms: PageRank, Topic Specific PageRank, HITS
- References (and slides based on):



[Chapter 5](#) of Jure Leskovec, Anand Rajaraman, Jeff Ullman, “Mining of Massive Datasets”, Cambridge University Press, 2011.



[Chapter 21](#) of Christopher D. Manning, Prabhakar Raghavan and Hinrich Schütze, “Introduction to Information Retrieval”, Cambridge University Press. 2008.