INFORMATION RETRIEVAL USING MULTI-ARMED BANDITS

Paul Lagrée

January 20, 2015

Université Paris-Sud

Information Retrieval

Much work has been done to improve keyword based web search

- Popularity of web pages using **Web links** (PageRank)
- **Personalization**, contextual models
- Query **intent detection**
- **Learning to rank**

AS-YOU-TYPE SEARCH IN SOCIAL MEDIA

Search And Recommendation

Different perspectives are possible for moving search closer to recommendation

- **relevance-oriented** query suggestion (recommend queries which are the most relevant to the input query)
- **diversification**: consider several aspects of the input query
- **personalization**: adapt suggestions to the user (user-centric approach)

See *Personalized Query Suggestion With Diversity Awareness* for an approach combining the two opposite settings.

Motivation

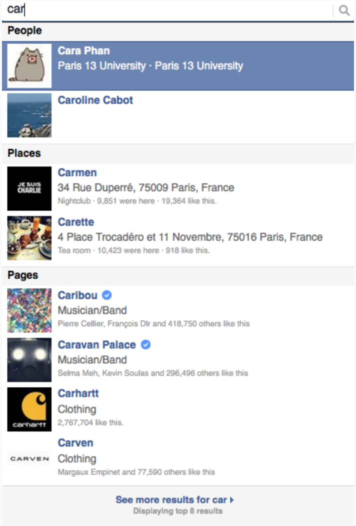

- **Social-aware** method (network-based)
- As-you-type search, handling **prefixes**
- **Real-time** approach (200ms maximum)
- **Incremental** computation: exploits what has already been computed

Problem Definition

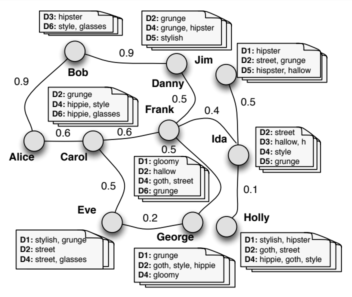

- **Triples** (*user*, *item*, *tag*): items can be movies, documents, photos
- **Weighted similarity network** between users (reflects proximity, friendship, similarity and can be computed using tags, item-tags, social links)

Given a **query** $Q = (k_1, ..., k_l, prefix)$ and a **seeker** s, we aim at finding the top-k social-aware answer.

Score Model

For a given **tag** t and **seeker** s, the **score** is

$$score(item|s,t) = \alpha \times textual(t, item) + (1-\alpha) \times social(item|s,t)$$

where $\alpha \in [0,1]$ gives how much we want the answer to be social.

If $\alpha = 1$, we come back to the classical web search.

The **social score** is defined as

$$social(item \mid s,t) = \sum_{v \text{ tagged } item \text{ with tag } t} \sigma(s,v)$$

Intuition of the algorithm

At each round:

- Visit of
  - next neighbour in the graph *u* (**social branch**)
  - **or** next element in inverted lists (**textual branch**)
- If social branch, exploration of user space using **trie** structures for matching tag prefixes
- Updates of computed scores, of bounderies, top-*k* checking
- Addition of neighbours of *u* as in Dijkstra algorithm

Experimental Framework

Two datasets: **Amazon** and **Twitter** (**Tumblr** incoming)

	Amazon	Twitter
Users	886,814	458,117
Items	252,891	1,662,464
Tags	91,352	550,157

Experimental Framework

Experiments:

- given one triple (u, i, t), do the search with keyword t with user u as the seeker
- **metric**: ranking of i in the result
- **precision recall** (test dataset D of N triples)

$$P(k) = \frac{\#\{triple \mid ranking < k, triple \in \mathsf{D}\}}{N}$$

Results - α impact

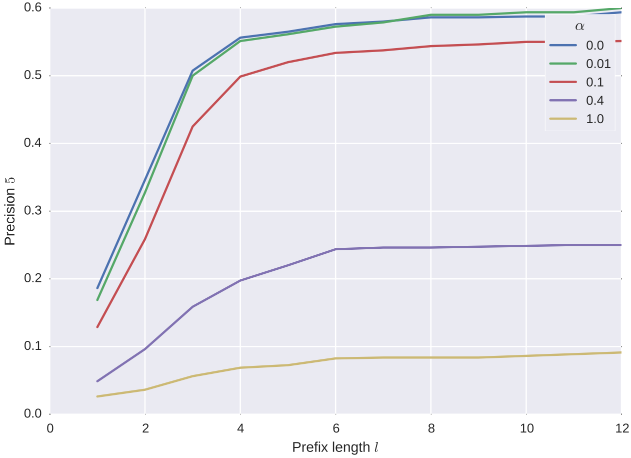

Figure: Impact of α, the social weight

Results - Connectivity impact

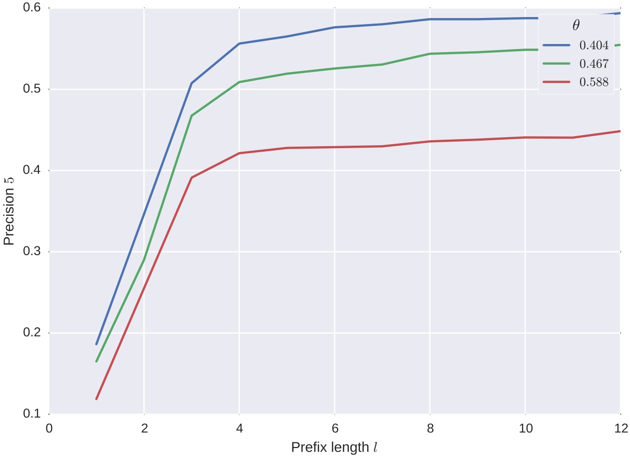

Figure: Impact of connectivity for θ, the similarity threshold

How can MABs be used?

Several aspects could be improved:

- Setting a **seeker-specific** value for α based on user historic
- Learning about queries
  - **Global** information needs
  - **Social** perspectives
  - **Mixed** information needs

SOCIAL NETWORK SEARCH WITH MABS

Problem

- Based on *Social Network Search as a Volatile Multi-armed Bandit Problem* by Z. Bnaya, R. Puzis, R. Stern and A. Felner.
- **Objective:** Given a **target** (user profile), finding as many information about it while crawling the social graph
- **Main idea: Equivalence classes** of nodes used as arms of a *Volatile* Multi-Armed Bandit framework.

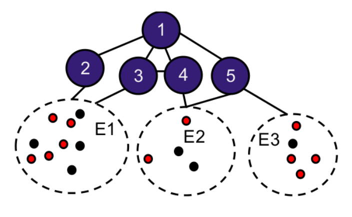

Problem: **Structural Equivalence Classes** (SEC) are volatile

- A new SEC can appear
- A SEC can become the best arm at any time
- Each arm is associated to a **lifespan** (z_i, t_i)

Two kinds of nodes:

- **Internal** nodes (already acquired)
- **Frontier** nodes (added when acquiring a previous node)

Question: which node to acquire next?

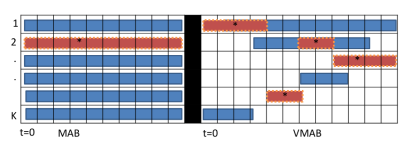

Figure: Best arm along time (red)

Evolution of SECs

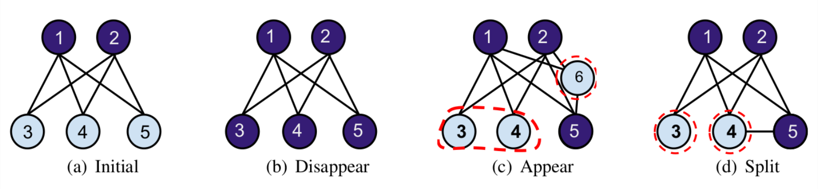

Figure: SEC possible events

The MAB framework needs to handle such volatility.

VMAB Policy

At each round, VMAB chooses the arm a_i with a valid lifespan maximizing

$$\overline{R_i} + \sqrt{\frac{2 \cdot \ln(n - z_i)}{N_i(n)}}$$

- $\overline{R_i}$ is the average reward observed on arm i
- $N_i(n)$ corresponds to the number of time i has been pulled
- **UCB-like** policy
- Usually, a parameter M is added and corresponds to **virtual acquisitions** (lower exploration of new arms)

VMAB Policy

Definition

The expected regret R_n^v of a VMAB policy is

$$R_n^v = \sum_{t=1}^{n} \mu^*(t) - \mu(I(t))$$

where $\mu^*(t)$ is the expected reward of the optimal arm at turn t and $I(t)$ is the arm selected at turn t.

Theorem (VMAB regret)

Let B be the number of epochs *up to time n. The* ***accumulated regret*** *R_n of VUCB policy is $\mathcal{O}(B \cdot \ln(n))$.*

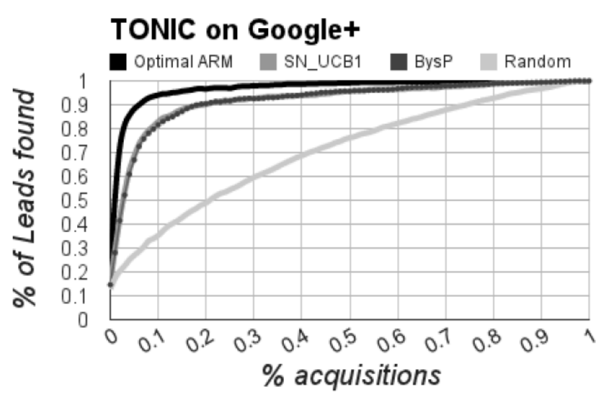

Figure: Recall of TONIC search

LEARNING DIVERSE RANKINGS

Motivation

Given previous work, we have a top-k answer for a query Q, which can be a mix of **social** and **global** relevance.

Question: Can we improve the answer by changing the ranking?

Problem

Let D be a set of n items.

- We aim at finding a **ranking** of k items B^*.
- The ranking needs to be **diverse**. Documents are potentially relevant to the query for different reasons (e.g. *Jaguar*)
- Online learning relying on Multi-Armed Bandits to **minimize** the total number of poor rankings.

Maximal Marginal Relevance

One basic solution to the problem is to use the **Maximal Marginal Relevance** (MMR).
Given a similarity measure between documents and queries sim_1 and a similarity measure between pairs of documents sim_2, MMR iteratively selects the documents

$$d_i = \arg\max_{d \in D} \left(\lambda \cdot sim_1(d, q) - (1 - \lambda) \cdot \max_{d_j \in S} sim_2(d, d_j)\right)$$

Ranked Bandits Algorithm

- The algorithm runs a MAB instance MAB_i on each rank i.
- Each document d is an arm of the MAB_i instance
- MAB_1 is responsible for choosing which document to show at rank 1
- MAB_2 determines which document is shown at rank 2, unless it is the same as rank 1. Then we choose a random document.
- If a document chosen by a MAB_i insance is clicked by the user, a reward 1 is given to the arm of the corresponding MAB_i

Experiments

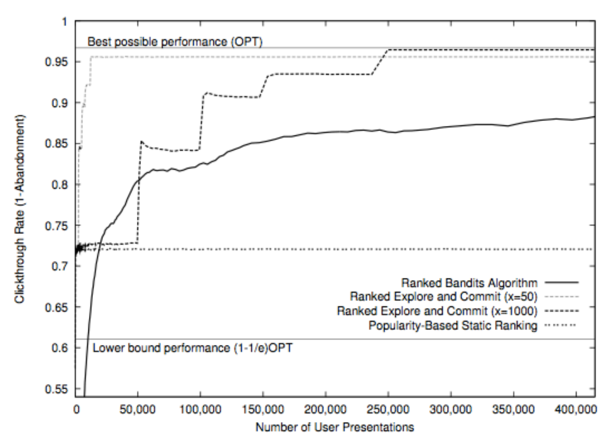

Figure: Clickthrough rate of the learned ranking function of the number of user presentations

Questions?

References I

D. Jiang, K. W.-T. Leung, J. Vosecky and W. Ng
Personalized Query Suggestion With Diversity Awareness.
In ICDE, pages 400-411, 2014.

S. Maniu and B. Cautis
Network-aware Search in Social Tagging Applications: Instance Optimality Versus Efficiency.
Proceedings of the 22nd ACM International Conference on Conference on Information & Knowledge Management, CIKM 2013.

B. Cautis, P. Lagrée and H. Vahabi
A Network-Aware Approach for Searching As-You-Type in Social Media.
To be published, 2015.

F. Radlinski, R. Kleinberg and T. Joachims
Learning diverse rankings with multi-armed bandits.
In Proceedings of the 25 th ICML, 2008.

Z. Bnaya, R. Puzis, R. Stern and A. Felner
Social Network Search as a Volatile Multi-armed Bandit Problem.

HUMAN, 2(2):pp–84, 2013.