York University at TREC 2013: Contextual Suggestion Track

Sushma Yalamarti, Mariam Daoud, Jimmy Huang
Information Retrieval and Knowledge Management Research Lab
York University, Toronto, Canada
{vsushma, daoud, jhuang}@yorku.ca

Abstract

This paper presents our participation in the Contextual Suggestion Track of TREC 2013. The goal of this track is to investigate search techniques for complex information needs that are highly dependent on context and user interests. To achieve this goal, we propose a semantic user profile modeling for personalized place recommendation. For the semantic user profile model construction, we construct a place ontology based on the Open Directory Project (ODP), a hierarchical ontology scheme for organizing websites. Previously rated attractions by the user are mapped to the place ontology where we represent positive and negative preferences under each place type. For a new user context represented by geographic coordinates, we rerank the top 50 suggestions returned by Google places API using the user profile.

General Terms

Measurement, Performance, Experimentation

Keywords

Context Based Google Suggestions, Semantic User Profile Model, Place Ontology, Re-ranking

1 Introduction and Motivation

With technological advancements in mobile technologies like mobile data networks (GPRS and WCDMA), positioning systems (GPS), mobile phones and personal digital assistants (PDAs) , there has been an increase in the use of internet-accessing devices. It is now possible to offer up-to-date information and online services to people whenever and wherever they are. With the increasing ubiquity of internet-accessing smart phones, it is now possible to offer personalized, context-sensitive recommendations. These online services are mostly useful for people in places they have never visited before. Often, tourists do not know their way, nor which restaurants, museums, shops, public services, etc are available to them. The number of potential places to visit can be quite overwhelming, especially in touristic regions . Personalized Context-Aware Recommender systems can help a tourist to find places matching his interests and his current location.

For example, In a mobile context this contextual suggestion system might take the form of an app that recommends interesting places and activities based on the users location and personal preferences. Imagine a user with a young family travelling for a summer vacation to Philadelphia, Pennsylvania. A contextual suggestion system might recommend to spend time at a kids friendly museum (http://www.pleasetouchmuseum.org/), or a walk in the beautiful historic Penn Treaty

Park (http://penntreatypark.org/), or to spend the day splashing in the waters at the Clementon Park Splash World (http://www.clementonpark.com).

In first year of TREC Contextual Suggestion track 2012 [DHCK$^+$12], a single task was introduced to evaluate contextual suggestion from the openWeb. As input to the task participants were given a set of example suggestions, a set of user preference profiles, and a set of geotemporal contexts. The task was to take the profiles and contexts and to produce up to 50 ranked suggestions for each combination of profile and context. Participants gathered suggestions from the open Web.

Each profile corresponds to a single user, and indicates that user's preference with respect to each sample suggestion. For example, one suggestion might be to have a beer at the Dogfish Head Alehouse, and the profile might include a negative preference with respect to this suggestion. Each suggestion includes a title, description, and an associated URL. Each context corresponds to a particular geotemporal location, including city, day of the week, time of day, and season. For example, the context might be Gaithersburg, Maryland, on a weekday evening in the fall. The geographical contexts are very coarse-grained (i.e., an entire city) to help simplify the task.

The goal of the TREC Contextual Retrieval Track 2012 was to encourage research in information retrieval techniques on providing useful place recommendations based on the user profile and geotemporal data. Chief part of the work involved collecting place suggestions form open web sources like Google Places, Yelp, Foursquare etc based on location and time and re-ranking these suggestions based on a ranking model that takes into account the user preferences.

Most of the previous work, model the user profile in terms of bag of words, or a category list from Yelp. Our intention in this work is to model semantically the user profile where multiple user interests represent place types and are modelled by place type concepts of the ODP ontology. We consider specific and general concepts when modeling the user interests. We construct a Place ontology by selecting concepts from the ODP ontology related with place types such as Museums, Parks, Restaurants, Shopping, Travel, Arts and Sports. The main distinctive feature that characterizes our user profile construction is to mix semantic representation with both granularity and user feedback. For each user, we construct a profile that consists of an instance of the Place ontology. Each concept of the Place ontology is associated with a positive and negative vector of terms. The latter are derived from the positively and negatively rated sample place suggestions respectively. Our personalised recommendation is then based on re-ranking place suggestions by exploiting the user profile depending on how well they match a positive/negative user interest.

This paper is organized as follows: Section 2 presents related work. Section 3 details our semantic user profile modeling and personalized recommendation. Section 4 presents experiments and results. In Section 5, we conclude the paper and present our future work.

2 Related Work

In the past, there has been a lot of work done on context based information retrieval, where the chief portion of work extensively discusses the correlations among query, context and document in the contextual retrieval setting. Focusing on various context approaches, some are query related, some approaches are based on document context and some are user based context approaches. [HH05][Hua05] uses the related text contextual information for query expansion in Okapi retrieval system". Approaches based on using the document context has been studied in detail in [YHL11] which utilizes linkage information from the citation graph has been shown to be effective. In [DH13] geographic and temporal data has been used as 'context' for effective mining of information that target what, when and where components in news articles search. Patients related contexts are utilized to build a bayesian-based personalized recommendation model for the application of health care [ZHH$^+$12]. Additionally, many studies have been performed for personalized information retrieval for place suggestion using context information in the form of geo-temporal data.

In [MTP12], a place recommendation approach has been proposed based on the intersection of candidate venues provided by Google Places and Foursquare. Suggestions are ranked according to a linear combination of two scores: one that captures the venue's appropriateness at the given time, by noting the number of people checked in at regular intervals, and another score that captures its appropriateness given the user's prior likes and dislikes. To match the suggestions with profiles, the description of each suggestion was first reduced to a BM25-weighted vector. Vectors were summed, with fixed multipliers of 0.75 for positive examples and -0.25 for negative examples. The resulting term weights were used to score the description for each candidate suggestion.

Fasilkom UI from Universitas Indonesia derives a user model based on Yelp's category list. A combination of user model and geolocation is used to generate the place suggestions. Search was expanded to similar categories if no results can be found in a particular category for a current location. As for scoring, they use the review rating for the places to produce the ranked results.They also try to apply diversity to results to make suggestions more interesting to the users..

HP Labs China uses a context-aware recommendation approach to produce the ranked list of suggestions. In their approach they employed Matrix Factorization for collaborative filtering to learn the latent factor of each user in profiles, and also used its category information to learn user's general preference. They use SVD++ to predict the scores of all suggestions and later use pairwise ranking model to rank a list of suggestions for a user. They use the contextual post-filtering approach to adjust the resulting set of suggestions with the help of the category feature and the information extracted from Yelp. The description for each suggestion is generated by a human defined template which includes location, category and many other features about the suggestion.

[HC12] defined a retrieval framework that combines two modules. First, a context processing module consists of using Google Places API that takes geographic coordinates and a set of place types as inputs to retrieve a list of places. Second, a preference processing consists of result personalization according to user interests. Profiles are represented by vectors relying on the Vector Space Model. For each profile, the positive user preferences were based on the positively rated examples and vice versa. The fine-grained approach consists in defining positive and negative preferences as sets of positive and negative preference examples (one vector per example). A similarity score between each place vector (from Google Places) and each preference vector based on the cosine measure was then computed.

In [MS12], recommendations were collected by using the location context as search query in Google Places and were ranked by their textual similarity to the user profiles, based on a TF-IDF measure. Initially, the cosine similarity of an initial recommendation to the positive profile determined the ranking. Further, the textual similarity was based on a point-wise Kullback Leibler divergence score which is based on the probability of observing a term in the set of examples that were rated positively, compared to the probability of observing it in the set as a whole. The sum of Kullback Leibler scores for the terms that occur in the initial recommendations determined the rank, which was later combined with a number of other rankings based on Google Search, popularity and categories in order to improve the ranking. As a final step, items that did not match the temporal context were filtered out.

[AY12] first identifies context-independent queries from the combined with new location information are sent to multiple Web search engines, such as Google, Yelp, Yellow Pages, and Bing, to crawl and build a large pool of potential contextual suggestions. A learning to rank model is then used to merge and re-rank those potential suggestions from the pool for each query, context, and user profile combination. Learning to rank model utilizes three types of profiles: a general profile based on the training suggestions given in the Toronto examples provided by TREC 2012, a city profile containing well-known attractions for each city, and a personal profile based on a user's personal interests. Moreover, the learning model makes distinctions among major personal interests, minor personal interests, and negative personal interests for all personal profiles. The detections of major, minor and negative personal interests are done by statistical analysis across users, examples, and context-independent query types.

[AM12] approaches the TREC task by devising an algorithm based on the regular expression for extracting time and address information from different websites for each place and then using this extracted information to rank suggestions along with user's context and preference.

[KKH12] uses Wikitravel as a source for travel suggestions. Wikitravel is a well structured community-based travel guide for destinations all over the world. From pages dedicated to cities in US, suggestions are extracted for sightseeing, shopping, eating and drinking. Descriptions from positive examples in the user profiles are used as queries to rank all suggestions. The ranked suggestions are then filtered based on the location of the user ignoring temporal aspects of the context

[PY12] collects candidate suggestions by web crawl from multiple online sources such as Yelp and Foursquare based on the geographical information from the 50 contexts. Later, ranked candidate suggestions based on their similarity to the personal profile and that to the contexts (i.e., geographic and temporal information). The ranking function is computed based on similarity between a suggestion and the places that the user like and the dis-similarity between the suggestion and the places disliked by the user. The similarities are computed based on the either the category or description of the suggestions.

[LQ12] designed a spider framework to crawl websites from tripadvisor, in order to collect candidate pages related to attractions, restaurants etc. Pre-processing involves extracting useful information, including name, homepage, rate and comment. Next, context filtering based on season, time, and weather are applied. User modeling based on td-idf and cosine similarity function is used for re-ranking using the example suggestions.

3 Our approach for Personalized Recommendation

Personalization is a significant component in recommender systems. Capturing information about user search interests is a chief component in personalization. One of the various definitions attributed to the concept of personalization [ADOM] is given as follows: “Personalization is the ability to provide content and services that are tailored to individuals based on knowledge about their preferences and behavior”. A user profile can be used to understand the characteristics of a person, and also the user interests in specific places he/she has visited/rated before. This information can be exploited by the place recommender system by taking into account the persons' characteristics and preferences.

Personalized place recommender systems aim at helping a jetsetter with the crucial travel planning decisions that he will face before travel or while on-the-move. These recommender systems necessitate the need to acquire the knowledge of interests and wants, either explicitly (by asking) or implicitly (by mining the user online activity), and suggest destinations to visit, points of interest, events or activities. The main objective of a recommender system is to ease the information search process of the traveller and additionally provide personalised search results on-the-move.

We tackle these requirements by proposing a new type of semantic-based recommender system. Our approach involves constructing a place ontology, modeling a semantic user profile and devising a personalized recommendation.

3.1 A place ontology construction using the ODP

3.1.1 Overview of the ODP ontology

Semantic knowledge structures, such as ontology, can provide valuable domain knowledge and user information. An ontology is specifications of concepts and relations between them. It defines “content specific agreements” on vocabulary usage, sharing, and reuse of knowledge [ZGG+99]. When initially learning user interests, systems usually perform poorly until they collect enough relevant information. Since initial behaviour information is matched with existing concepts and relations between them, using ontologies as the basis of the profile helps to avoid or ease this problem. In our case, we use the Open Directory Project concept hierarchy [ODP02] (and associated, manually classified Web pages) as our reference ontology.

Ontologies are the structural frameworks for organizing information for domain knowledge representation. The Open Directory Project (ODP), also known as DMOZ (from directory.mozilla.org, its original domain name), is a multilingual open content directory of World Wide Web links. It is owned by AOL but it is constructed and maintained by a community of volunteer editors.

Open Directory Project (ODP) is a human-edited Index of Web sites, also known as DMOZ, an acronym for "Directory Mozilla". ODP is hosted and administered by Netscape Communications and Weblogs, Inc. and is associated with the Mozilla browser. According to the official Web site, ODP hosts the largest and most comprehensive Web site directory in the world. The purpose of the ODP is to list and categorize sites, not to rank or promote them. ODP uses a hierarchial ontology scheme for organizing site listings. Listings on a similar topic are grouped into categories which can then include more specific categories.

From ODP, we extracted a sub-directory relevant to place types. For example, as most recommendations we retrieved from Google Places are related to places to visit, we extracted the concepts associated with travel categories of the ODP.

3.1.2 Building a place ontology

To achieve personalization we build a place ontology based on the ODP ontology. We only extract concepts in ontology that are related to place types and that could represent the user interests in places. We consider both generic and specific concepts of the ODP ontology. Generic concepts represent general user interests among different place types. For example, a user is more interested in museums than restaurants or parks. A specific concept represents specific user interest under a generic place type. For instance, a user is more interested to visit a cultural museum than a science museum.

The place ontology takes into account several main concepts such as Museums, Travel, Food, Shopping etc. Specific concepts under each generic concept are also considered in order to capture specific user interests and also to differentiate between different user's tastes. Concepts such as science museum, history museum, arts and entertainment museum and cultural museum are different kinds of museums that allows capturing specific user interests.

There are different levels in this place ontology, the depth can vary between two to seven levels. The depth of the ontology construction signifies that a user's interest is captured at a granular level.

For instance, a user looking for a place to snack specifically a sandwich, such an interest can be mapped to one or more ontology sub-tree Business/Hospitality/Restaurant Chains/Sandwich and Deli and Recreation/Food.

Figure 1 presents a portion of the place ontology that we constructed based on the ODP ontology.

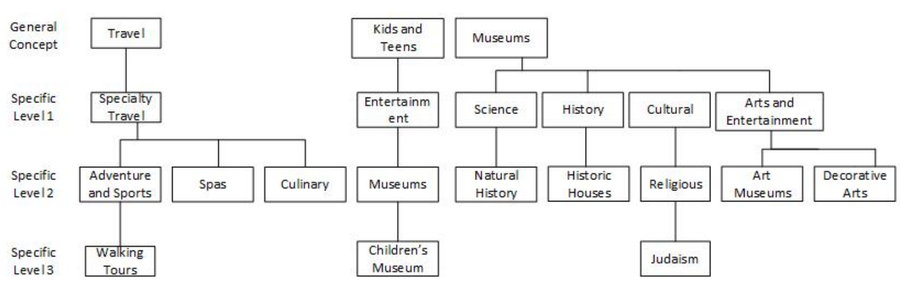

Figure 1: A portion of our place ontology

3.2 Semantic user profile modeling

We propose a new user profiling that refers to the construction of a profile via semantic means using the ODP place ontology. Each user profile is represented as an instance of the place ontology that gathers positive and negative user interests. More precisely, each concept (place type) of the ontology is represented with positive and negative term vectors.

For each user, we build the user profile by mapping the previously rated places on the Place ontology. This mapping results in classifying the rated web pages into one or more concepts of the place ontology. We map the user-rated example suggestions on the Place ontology. Each will be mapped to one or more ontology concepts. Each concept of the ontology is represented by a pair of contexts: a negative and a positive context. The positive context is represented by a vector of terms issued from the positively rated examples mapped to it. The negative context is also represented by a vector of terms issued from the negatively rated examples mapped to it. For instance, the web page http://clementonpark.com/ is classified under : Recreation/Theme Parks/Water Parks. The web page http://studio34yoga.com/ is mapped to Society/Religion and Spirituality/Yoga/Teachers and Centers on the Place Ontology.

Consider concept C_i of the place ontology, and a set S of previously rated pages classified under concept C_i, we partition S into positive set S^+ and negative set S^- where S^+ contains the web pages rated above 3 and S^- contains the set of web pages rated below 3 by the user. Each page in S^+ and S^- is represented with a term vector. Based on S+ and S- we derive positive context Vector C_i^+ and a negative context vector C_i^-. C_i^+ and C_i^- are derived as being the centroid vector in the sets S^+ and S^- respectively. The following matrix represents how to derive the positive and negative contexts for each concept. Each row represents one page, and each column represent a term.

Positive Concept Matrix: M^+ Negative Concept Matrix: M^-

$$
\begin{array}{c} \\ p_1^+ \\ p_2^+ \\ \vdots \\ p_n^+ \end{array}
\begin{pmatrix} t_1 & t_2 & . & t_n \\ w_{11} & w_{12} & . & w_{1n} \\ w_{21} & w_{22} & \vdots & w_{2n} \\ \vdots & \vdots & & \vdots \\ w_{n1} & w_{n2} & . & w_{nn} \end{pmatrix}
\qquad
\begin{array}{c} \\ p_1^- \\ p_2^- \\ \vdots \\ p_n^- \end{array}
\begin{pmatrix} t_1 & t_2 & . & t_n \\ w_{11} & w_{12} & . & w_{1n} \\ w_{21} & w_{22} & \vdots & w_{2n} \\ \vdots & \vdots & & \vdots \\ w_{n1} & w_{n2} & . & w_{nn} \end{pmatrix}
$$

The positive concept vector is calculated as the centroid vector of positively rated pages in M+ and the negative concept vector is calculated as the centroid vector of negatively rated pages in M-.The weight of term t in C_i^+ is calculated as the average weight of the term in the positive pages. The weight of term t in C_i^- is calculated as the average weight of the term in the negative pages. Hence the weight Wtk^+ of term t_k in the centroid C_i^+ is calculated as follows:

$$
Wt_k^+ = \sum_{i=1}^{n} w_{ik} \tag{1}
$$

The same formula applies for constructing the negative centroid vector C_i^-.

The user profile is then represented as an instance of the place ontology where each concept is associated with a positive vector and a negative vector. For instance, an example of a user profile with a Museum category (Reference/Museums/History) is represented by positive and negative weighted vectors as follows:
Negative concept vector: community=1, evolution=1, history=2, nation=3, time=1.
Positive concept vector: eastern=2, history=3, interest=1, penitentiary=1, philadelphia=4, state=2, visit=1.

3.3 Personalized recommendation

Our approach for a personalized recommendation is based on re-ranking the suggestions returned by Google places API using the user profile, thus by taking into account positive and negative user interests. Our main motivation for considering positive and negative user interests is to privilege the positive suggestions in the top ranks of the result list presented to the user. The re-ranking approach helps in differentiating between positive and negative suggestions that belong to the same place type.

Our approach is represented in algorithm 1.

Context Ctx, suggestion set S and profile P are inputs to our reranking approach. Ctx is a city location represented by its geographical coordinates (latitude, longitude). S a set of top 50 suggestions returned by a basic place recommender system for corresponding Ctx. P is the user

Algorithm 1 Re-ranking the suggestion set for personalized recommendation

INPUT: Context Ctx, Suggestions $S = \{S_1, S_2, ...S_{50}\}$, Profiles $P = \{C_1\ (C_1^+, C_1^-),\ C_2\ (C_2^+, C_2^-) ..., C_N\ (C_N^+, C_N^-)\}$
OUTPUT: Result Suggestions R = $\{R_1, R_2 ... R_{50}\}$,
for each $S_i \in S$ **do**
  for each $C_i \in P$ **do**
    //do the following steps
    $Score^+(S_i) = \cos(C_i^+, S_i)$
    $Score^-(S_i) = \cos(C_i^-, S_i)$
    $Score(S_i) = Score^+(S_i)$ -$Score^-(S_i)$
  end for
end for
// sort suggestion set according to score SORT S
// re-ranked result set R = S

profile represented in terms of place concepts C_1, C_2, ..., C_n of the place ontology. Each concept C_i is represented by two vectors C_i^+ and C_i^-. For each suggestion in S we calculate a positive score and a negative matching score with each concept C_i in P. Positive scores are calculated using the cosine similarity between the positive concept vector and the suggestion vector. Similarly, we build a negative score. The final score of a suggestion is calculated by subtracting the positive and negative score. This is repeated for the remaining 49 suggestions. We re-rank the suggestions in the descending order of their scores. This process is repeated for every context and profile pair.

4 Experiments and Results

4.1 Datasets:

The TREC 2013 contextual suggestion track have provided contexts, profiles, and examples with the seed city as Philadelphia, PA. The profiles contain list of attractions. Each attraction has two ratings. One rating for the attraction's title and description and other rating for the attraction's website. The ratings are given on a five-point scale based on how interested the user would be in going to the attraction if they were visiting the city it was in. Rating 4-5 is given for attractions the user strongly likes, rating 2 for being neutral, 1-2 rating for user attractions that strongly disinterests him/her and -1 or 0 for attraction websites that didn't load. Profiles are distributed in two files. The examples file contains the example suggestions that were rated by users and the profiles file contains the ratings given for the example suggestions.

Table 1: Example of place suggestion

id	title	description	url
51	Elfreths Alley Museum	Elfreths Alley Museum is a reputable museum. A lovely little piece of history. Definitely a must while visiting Philadelphia... To walk down the oldest residential street in the country is just something I think everyone should do at least once if in the area! I really enjoyed it.	http://www.elfrethsalley.org

Table 2: Example of profile

id	attraction_id	rating description	rating website
35	51	0	4
35	52	1	4
35	53	3	3

Contexts are provided by TREC 2013 and contain the cities suggestions need to be generated for. The cities in the contexts file are the primary cities of 50 randomly selected metropolitan areas (which are not part of a larger metropolitan area) excluding Philadelphia, PA.*
Each context has a unique identifier id (context_id), city, state, lat, and long.

Table 3: Example of contexts

id	city	state	lat	long
1	New York City	NY	40.71427	-74.006
2	Chicago	IL	41.85003	-87.6501
3	Los Angeles	CA	34.05223	-118.244

We use Google Places API to build our baseline suggestion set. For a given location i.e. context we use Google Radar API to get the top 200 suggestions. We take the top 50, and get description from Google Custom Search API, more precisely, we get website(url), type, location, etc from Google Place Detail search API.

4.2 Results

We submitted two runs to TREC contextual retrieval 2013. Descriptions of our runs are presented as follows:

- Run1: this is the baseline run. For each context, profile pair we query Google places and get the list of top 50 suggestions ranked by google.
- Run2: this is our personalized place recommendation run. It is based on re-ranking the baseline for each context, profile according to algorithm 1.

Table 4 presents the results according to p@5, MRR and TBG.

Table 4: Baseline and personalized recommendation results in terms of p@5, MRR, TBG. * and ** refers to significance test according to Wilcoxon statistical test.

Results						
Measures	P@5		MRR		TBG	
	Baseline	Ours	Baseline	Ours	Baseline	Ours
Average	0.3273	0.3309	0.4742	0.4636	1.2969	1.3483
%Improvement	1.10%*		-2.24%**		3.96%*	

The results show that our approach for personalized recommendation outperforms significantly the baseline results returned by Google places API at P@5 and TBG measures. The improvement has been shown statistically significant according to the Wilcoxon test at P@5 and TBG. However, our personalized approach has decreased the recommendation performance at MRR measure. This means that our approach has shifted more relevant suggestions to the top 5 results presented to the user. In terms of TBG, the gain obtained using our personalized approach has exceeded the gain obtained in the baseline. We believe that the profile accuracy and level of granularity has a great impact on the personalization effect. Also, there is a variability in the performance of the personalized recommendation for different profile/context pairs. We plan to analyse the result performance of the personalized search with respect to the quality of the user profiles and the diversity of the suggestion list.

5 Conclusions and Future Work

The TREC Contextual Suggestion Track presents a place recommendation task where the challenges are to model the user interests and exploit them for recommendation. In this paper, we present our participation in the Contextual Suggestion Track of TREC 2012. We submit two runs: one run based on google places API and another one based on reranking google places API results using a semantic user profile. The experimental results show that when we utilize semantic user profile modeling techniques, the performance improves in terms of P@5. This indicates that we need to focus in the same direction and elaborate ways to make full use of user profiles. Future work will focus on exploring strategies incorporating place type along with user profile modeling for identifying the most relevant recommended places to the user. Also, we plan to use some external resources to help build better user profiles and achieve diversity in the suggestion list.

Acknowledgements

This research is supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada and Mathematics of Information Technology and Complex Systems (MITACS).

References

[AM12] Kripabandhu Ghosh Abhishek Mallik, Mandar Mitra. Contextual suggestion 2012. 2012.

[AY12] Hui Yang Nazli Goharian Steve Kunath Ophir Frieder Andrew Yates, Dave DeBoer. (not too) personalized learning to rank for contextual suggestion (2012). In *The Twenty-First Text REtrieval Conference (TREC 2012) Notebook*. National Institute for Standards and Technology, 2012.

[DH13] Mariam Daoud and Jimmy Xiangji Huang. Mining query-driven contexts for geographic and temporal search. *International Journal of Geographical Information Science*, 27(8):1530–1549, 2013.

[DHCK+12] Adriel Dean-Hall, Charles L. A. Clarke, Jaap Kamps, Paul Thomas, and Ellen Voorhees. Overview of the trec 2012 contextual suggestion track. In *Proceedings of the Twenty-first Text Retrieval Conference (TREC)*, 2012.

[HC12] Gilles Hubert and Guillaume Cabanac. IRIT at TREC 2012 contextual suggestion track (regular paper). In Ellen M. Voorhees and Lori P. Buckland, editors, *Text REtrieval Conference (TREC), Gaithersburg, USA, 07/11/2012-09/11/2012*, page (on line), http://www-nlpir.nist.gov/, novembre 2012. National Institute of Standards and Technology (NIST).

[HH05] Xiangji Huang and Yan Rui Huang. Using contextual information to improve retrieval performance. In *IEEE GrC*, pages 474–481, 2005.

[Hua05] Xiangji Huang. Incorporating contextual retrieval into okapi. In *Proceedings of ACM SIGIR Workshop on IR in Context (IRiX'05),Salvador, Brazil*, 2005.

[KKH12] Marijn Koolen, Jaap Kamps, and Hugo Huurdeman. University of Amsterdam at the TREC 2012 contextual suggestion track: Exploiting community-based suggestions from wikitravel. In *The Twenty-First Text REtrieval Conference (TREC 2012) Notebook*. National Institute for Standards and Technology, 2012.

[LQ12] QianQian Wang Yue Liu ZhiHua Zhou Weiran Xu Guang Chen Jun Guo Lin Qiu, JunRui Peng. PRIS at TREC 2012 contextual suggestion track). In , page (on line). National Institute of Standards and Technology (NIST), 2012.

[MS12] Wessel Kraaij Maya Sappelli, Suzan Verberne. Tno and run at the trec 2012 contextual suggestion track: Recommending personalized touristic sights using google places. In *In 21st Text REtrieval Conference Notebook Proceedings (TREC 2012)*. National Institute for Standards and Technology, 2012.

[MTP12] David Milne, Paul Thomas, and Cecile Paris. Finding, weighting and describing venues: Csiro at the 2012 trec contextual suggestion track. In *Proceedings of the Twenty-first Text Retrieval Conference (TREC)*, 2012. `http://es.csiro.au/pubs/milne_trec12.pdf`.

[ODP02] Open directory project. http://dmoz.org, 2002.

[PY12] Hui Fang Peilin Yang. An Exploration of Ranking-based Strategy for Contextual Suggestion). In *In Proceedings of 21st Text REtrieval Conference (TREC 2012)*. National Institute of Standards and Technology (NIST), 2012.

[YHL11] Xiaoshi Yin, Jimmy Xiangji Huang, and Zhoujun Li. Mining and modeling linkage information from citation context for improving biomedical literature retrieval. *Inf. Process. Manage.*, 47(1):53–67, 2011.

[ZGG+99] Xiaolan Zhu, Susan Gauch, Lutz Gerhard, Nicholas Kral, and Alexander Pretschner. Ontology-based web site mapping for information exploration. In *CIKM*, pages 188–194. ACM, 1999.

[ZHH+12] Jiashu Zhao, Jimmy Xiangji Huang, Xiaohua Hu, C. Joseph Kurian, and William Melek. A bayesian-based prediction model for personalized medical health care. In *BIBM*, pages 1–4, 2012.