# TUTA1 at the NTCIR-10 Intent Task

Haitao Yu
Faculty of Engineering, The University of Tokushima
2-1 Minamijyousanjima-cho, Tokushima
+81-88-656-7304.JAPAN.770-8506
yu-haitao@iss.tokushima-u.ac.jp

Fuji Ren
Faculty of Engineering, The University of Tokushima
2-1 Minamijyousanjima-cho, Tokushima
+81-88-656-7304.JAPAN.770-8506
ren@is.tokushima-u.ac.jp

## ABSTRACT
In NTCIR-10, we participated in the subtask of Subtopic Mining. We classify test topics into two types: role-explicit topic and role-implicit topic. According to the topic type, we devise different approaches to perform subtopic mining. Specifically, for role-explicit topics, we propose an approach of modifier graph based subtopic mining. The key idea is that: The modifier graph corresponding to a role-explicit topic is decomposable into clusters with strong intra-cluster interaction and relatively weak inter-cluster interaction. Each modifier cluster intuitively reveals a possible subtopic. For role-implicit topics that generally express single information needs, we directly generate the ranked list through semantic similarities leveraging on lexical ontologies. The evaluation results show that our best Chinese subtopic mining run gets the first position among all the runs in terms of $D\#-nDCG$. However, our English subtopic mining runs show a poor performance, which is planned to be further improved in our future work.

## Team Name
TUTA1

## Subtasks
Subtopic Mining (Chinese, English)

## Keywords
Subtopic Mining, Modifier Graph, Graph Clustering

## 1. INTRODUCTION
To raise satisfaction of the average users, the technique of diversification seeks for a trade-off between relevance and diversity, namely ranking documents that are highly relevant to major information needs above those that are marginally relevant to minor information needs and satisfying users with different information needs at the same time [1, 2, 3, 4]. Toward this goal, the NTCIR community [5, 6] proposes a two-step strategy consisting of subtopic mining and document ranking, which correspond to the subtasks of *Subtopic Mining* and *Document Ranking* respectively. Subtopic mining aims at capturing the possible information needs encoded within a query, where an information need is denoted as a subtopic. Document ranking aims at generating a diversified result list, in order to satisfy users with different information needs.

In this paper, we describe our proposed framework for the subtask of subtopic mining. The remainder of this paper is organized as follows: The related studies are reviewed in section 2. Section 3 details the proposed framework for subtopic mining. In section 4, we analyze the results we obtained. Finally we conclude our work in section 5.

## 2. RELATED WORK
To understand the underlying intents of a query, Hu [7] proposes an approach leveraging on the structure of Wikipedia. A number of intent seeds are iteratively propagated through Markov random walk. Radlinski [8] introduces an undirected graph model for inferring query intents. For an input query, the click and reformulation information are combined to identify a set of related queries to construct the undirected graph. The underlying intents are represented with clusters of related queries. Sadikov [9] integrates session co-occurrence information to cluster the refinements of a query. A cluster of refinements is used to represent a unique intent. The possibility of a drift among intents is also considered. However, most of the current studies rely on the click information in a query log and treat the whole query as the minimum granularity of analysis, which will not work well when processing queries with extremely sparse click information or previously unseen queries [10].

By learning lessons from previous studies, we perform subtopic mining at a fine-grained word-level, which is robust enough against the sparseness problem. Moreover, we encapsulate intent roles to facilitate subtopic mining, which is detailed in the following sections.

## 3. PROPOSED FRAMEWORK
To get a better query representation and understanding, Yu and Ren [11] introduce two intent roles: *kernel-object* and *modifier*. Based on the intent roles, they further classify user queries into two classes: role-explicit query and role-implicit query. Ignoring the order information, a role-explicit query can be given as $ko+\{mo\}$, where $ko$ denotes the kernel-object and $mo$ denotes a co-appearing modifier. On the contrary, queries that can't be well represented with kernel-object and modifier are defined as role-implicit queries. As subtopic strings and official topics are either real queries extracted from a query log or query-like strings formulated by evaluators, they can be analogously classified as role-explicit ones and role-implicit ones. Across our work, the model proposed by Yu and Ren [11] is used to identify the kernel-object and modifier. For convenience, we let function $ko()$ mean to get the kernel-object, and function $mo()$ mean to get the modifiers.

For a given topic, we manage to collect a set of sufficient subtopic strings, which are regarded as all possible expressions of subtopics. The following resources are used to generate the subtopic strings: (1) Query log; (2) Query suggestions; (3) Search snippets returned by a major search engine.

### 3.1 Subtopic Mining for Role-explicit Topics

Before describing the detailed procedures, we introduce the following necessary definitions:

**Word-level Co-session**: For two distinct words, if their parent subtopic strings appear in the same user session, we say the two words are co-session words.

**Word-level Co-click**: For two distinct words, if their parent subtopic strings share the same clicked documents, we say the two words are co-click words.

**Co-parent**: For two distinct words, if there exists a subtopic string that includes the two words as child words, we say the two words are co-parent words.

**Co-kernel-object Elements**: For a specific kernel-object, there exist a set of subtopic strings sharing the same kernel-object. E.g., the two subtopic strings 哈利波特游戏 (Harry Potter game) and 哈利波特小说 (Harry Potter fiction) share the same kernel-object 哈利波特 (Harry Potter). We call this kind of subtopic strings as co-kernel-object elements.

As studied by Yu and Ren [11], kernel-object abstracts the core object or topic indicated by a query, the differences among information needs represented by co-kernel-object elements are mainly determined by the co-appearing modifiers. Therefore, the co-kernel-object elements can be regarded as different expressions of kernel-object oriented subtopics.

For a role-explicit topic, we firstly identify its kernel-object $ko$. Then we scan the query log, all the queries that contain $ko$ as a substring are extracted out as subtopic strings. All the query suggestions are directly used as subtopic strings. We extract noun phrases (NP) and verb phrases (VP) from the search snippets, the ones that include $ko$ a substring are viewed as subtopic strings. Based on the set of subtopic strings, we distill a set of co-kernel-object subtopic strings, which are used to generate the target ranked list. Instead of disregarding the role-implicit subtopic strings that include $ko$ as a substring, we directly view $ko$ as its kernel-object, the remaining part are directly segmented into modifiers. This naïve intent role annotation guarantees the coverage of co-kernel-object elements expressing all possible subtopics.

Based on the set of co-kernel-object subtopic strings, we construct a modifier graph, which is an undirected, weighted graph. It consists of:

(1) Node: each node corresponds to a distinct modifier;

(2) Edge: For each pair of modifiers, the edge function $f(mo_i, mo_j, u) \rightarrow R$ determines an edge, where $u$ is a scenario parameter corresponding to co-parent, co-session, co-click and semantic similarity respectively.

Specifically, for $u_1 = CoParent$, the corresponding quantity is calculated as:

$$w(mo_i, mo_j, u_1) = \frac{|CoParent(mo_i, mo_j)|}{\max_{i,j} |CoParent(mo_i, mo_j)| + 1} \quad (1)$$

where $|CoParent(mo_i, mo_j)|$ represents co-parent times of the two modifiers.

For $u_2 = CoSession$, it is analogous quantified as:

$$w(mo_r, mo_w, u_2) = \frac{|CoSession(mo_r, mo_w)|}{\max_{r,w} |CoSession(mo_r, mo_w)| + 1} \quad (2)$$

where $|CoSession(mo_r, mo_w)|$ represents the co-session times.

For $u_3 = CoClick$, it is analogously quantified as:

$$w(mo_r, mo_w, u_3) = \frac{|CoClick(mo_r, mo_w, \mu_3)|}{\max_{r,w} |CoClick(mo_r, mo_w, \mu_3)| + 1} \quad (3)$$

where $|CoClick(mo_r, mo_w, \mu_3)|$ represents the co-click times.

For $u_4 = SemanticSimilarity$, it is quantified as:

$$w(mo_i, mo_j, u_4) = \frac{SSimilarity(mo_i, mo_j)}{\max_{i,j} SSimilarity(mo_i, mo_j) + 1} \quad (4)$$

where $SSimilarity(mo_i, mo_j)$ represents the semantic similarity. For Chinese topics, the semantic similarity is calculated based on HowNet[1] using the way proposed by Liu [12]. For English topics, it is calculated based on WordNet[2] using the way proposed by [13].

Finally, the weight of an edge is calculated through linearly combining the different scenario parameters in an equal manner as:

$$w(mo_i, mo_j) = \sum_{k=1}^{|u|} \frac{1}{|u|} w(mo_i, mo_j, u_k) \quad (5)$$

The previous studies show shat: (1) Queries among the same session are generally alternative queries that imply similar information needs [14, 15]. (2) The document clicks positively reflects result preference and relevance, queries for which users have clicked on the same documents usually imply a similar information need [16, 17]. (3) Words with a high semantic similarity are commonly observed in query reformulations that express the similar information need [18, 19]. Obviously, the modifier edges are biased towards modifiers reflecting a similar subtopic and the modifiers reflecting different subtopics are weakly connected. Once we decompose the modifier graph into clusters of modifiers based on the interaction of modifiers, it is reasonable to believe that each cluster of modifiers represents a different subtopic. Towards this direction, for a modifier graph, we get a group of modifier clusters by performing the LinLog method [20], which requires no pre-defined parameters. Each cluster of modifiers is used to represent a different subtopic. For each modifier cluster, we generate a corresponding cluster of subtopic strings. Namely, for each modifier, we get its parent subtopic string and add it into the parent subtopic string cluster. Specifically, for subtopic strings that their child modifiers locate in different modifier clusters, the following rules are considered:

[1] http://www.keenage.com/

[2] http://wordnet.princeton.edu/

(1) the parent subtopic string belongs to the cluster, of which the corresponding modifier cluster encloses the most child modifiers. (2) For the case of two modifier clusters enclosing an equal number of modifiers, the parent subtopic string belongs to the cluster, of which the corresponding modifier cluster encloses the most frequent modifier. Our intuition is that: the number of modifiers is proportional to the expression power of the underlying subtopic. Less frequent modifier is viewed to specify the more frequent one.

To facilitate generating the ranked list, we introduce the following definitions:

**Popularity**: popularity quantifies the relative importance of a possible subtopic, which is given as:

$$Popularity(sc) = Fre(sc) / Fre(SC) \tag{6}$$

where $sc = \{tStr\}$ denotes a cluster of subtopic strings, $SC = \{sc\}$ denotes the total clusters, $Fre(tStr)$ denotes the frequency of a subtopic string, $Fre(sc) = \sum_{tStr \in sc} Fre(tStr)$ denotes the cluster frequency, $Fre(SC) = \sum_{sc \in SC} Fre(sc)$ denotes the sum of cluster frequency. Moreover, we assume that there exists a probability distribution $p(mo)$ for modifiers, which describes how users formulate mutually-independent modifiers to express a subtopic. Using the Laplace Smoothing method, it is given as:

$$p(mo) = \frac{Fre(mo) + 1}{\sum Fre(mo) + V(mo)} \tag{7}$$

where $V(mo)$ denotes the vocabulary of distinct modifiers, $Fre(mo)$ denotes the times that $mo$ appears.

**Effectiveness**: Based on $p(mo)$, we define the effectiveness of a subtopic string expressing a subtopic as:

$$Eff(tStr) = p(mo(tStr)) = \prod_{mo \in mo(tStr)} p(mo) \tag{8}$$

To generate the target ranked list $L$ of subtopic strings, we define the gain value of ranking a subtopic string at the $i-th$ position as:

$$g(tStr, i) = Popularity(sc) / \log(i+1) \tag{9}$$

In equation 5, $sc$ is the cluster to which subtopic string $tStr$ belongs.

Going further, we quantify the quality of the list with the top-k subtopic strings as:

$$quality(L) = \beta N_{1+}(L) / |SC| + (1-\beta)\sum_{j=1}^{k} g(tStr, j) \tag{10}$$

where $N_{1+}(L)$ denotes the number of distinct clusters to which the current subtopic strings belong. $N_{1+}(L)/|SC|$ is used to approximate the subtopic recall. $\sum_{j=1}^{k} g(tStr, j)$ aims at ranking subtopic strings indicating major subtopics in higher positions. $\beta$ is a tuning parameter, which we set as 0.5 in our work. Essentially, the above definitions are inspired by the metric $D\#-nDCG$ [21] designed for diversified document ranking. By iteratively selecting the $k-th$ subtopic string that maximizes the quality of $L$, we can generate the target ranked list with a pre-defined size.

## 3.2 Subtopic Mining for Role-implicit Topics

As role-implicit topics generally express single information needs, we directly use the extracted subtopic strings. For the target ranked list, we sort the subtopic strings by their edit distance again the given topic and select the top-n (n is the required size of list L) to form the target ranked list. In our work, a variant edit distance proposed by Xia [22] is adopted.

# 4. EVALUATION RESULTS

## 4.1 Dataset

As for the subtask of "Subtopic Mining", the official resources include query suggestions for each topic and the Chinese query log SogouQ[3]. To ensure the coverage of expressions for all possible subtopics, we also utilize the top-100 (if exist) search snippets returned by Google to extract NP&VP segments, which are used for generating subtopic strings. For Chinese topics, SogouQ is used to calculate the word-level co-session and word-level co-click and all the proposed scenario parameters are used. Due to no matching query log, we only use the scenario parameters $u_1 = CoParent$ and $u_4 = SemanticSimilarity$ for English topics.

## 4.2 Results and Discussion

$D\#-nDCG$ is used as the primary evaluation metric for subtopic mining, which is a linear combination of $I-rec$ and $D-nDCG$.

$$D\#-nDCG@l = \gamma I-rec@l + (1-\gamma)D-nDCG@l \tag{11}$$

where $I-rec$ [23] measures diversity, $D-nDCG$ [24] measures overall relevance across subtopics, $\gamma$ is a tuning parameter that balances subtopic diversity and subtopic relevance, $l$ denotes the cutoff value.

For English subtopic mining and Chinese Subtopic Mining, Table 1 and Table 2 summarize the mean $I-rec$, $D-nDCG$ and $D\#-nDCG$ values with a cutoff value of 10 respectively. The runs of all participants are sorted by $D\#-nDCG$. The highest value in each column is underlined. Our runs are shown in bold.

From the obtained results in Table 1 and Table 2, we can draw the following conclusion: For the runs of Chinese subtopic mining, our proposed framework achieves the best performance in terms of $I-rec$ and $D\#-nDCG$. As most of the Chinese topics are role-explicit (92 out of 98 official topics), our modifier-graph based framework can incorporate a series of word-level knowledge, e.g., click information and semantic knowledge derived from the lexical ontology HowNet. Moreover, the proposed framework is robust to the sparseness problem commonly occurs when utilizing click information at a whole query level. However, our runs for English subtopic mining show a poor performance. We currently say that it is because of nonuse

[3] http://www.sogou.com/labs/dl/q.html

of a corresponding query log. We will further study this problem in our future work.

**Table 1. English subtopic mining runs ranked by D#-nDCG@10**

| Run Name | I-rec @10 | D-nDCG @10 | D#-nDCG @10 |
|---|---|---|---|
| THUIR-S-E-1A | 0.4107 | 0.3498 | 0.3803 |
| THUIR-S-E-3A | 0.3971 | 0.3492 | 0.3732 |
| THUIR-S-E-2A | 0.3908 | 0.3506 | 0.3707 |
| THUIR-S-E-4A | 0.3842 | 0.3517 | 0.3680 |
| THUIR-S-E-5A | 0.3748 | 0.3550 | 0.3649 |
| THCIB-S-E-2A | 0.3797 | 0.3499 | 0.3648 |
| KLE-S-E-4A | 0.3951 | 0.3282 | 0.3617 |
| THCIB-S-E-1A | 0.3785 | 0.3384 | 0.3584 |
| hultech-S-E-1A | 0.3099 | 0.3991 | 0.3545 |
| THCIB-S-E-3A | 0.3681 | 0.3383 | 0.3532 |
| THCIB-S-E-5A | 0.3662 | 0.3215 | 0.3438 |
| THCIB-S-E-4A | 0.3502 | 0.3323 | 0.3413 |
| KLE-S-E-2A | 0.3772 | 0.3028 | 0.3400 |
| hultech-S-E-4A | 0.3141 | 0.3566 | 0.3353 |
| ORG-S-E-4A | 0.3350 | 0.3156 | 0.3253 |
| SEM12-S-E-1A | 0.3318 | 0.3094 | 0.3206 |
| SEM12-S-E -2A | 0.3380 | 0.3020 | 0.3200 |
| SEM12-S-E-4A | 0.3328 | 0.2994 | 0.3161 |
| SEM12-S-E-5A | 0.3259 | 0.2977 | 0.3118 |
| ORG-S-E-3A | 0.3366 | 0.2842 | 0.3104 |
| KLE-S-E-3A | 0.3140 | 0.2895 | 0.3018 |
| KLE-S-E-1A | 0.2954 | 0.2719 | 0.2836 |
| ORG-S-E-2A | 0.2789 | 0.2564 | 0.2677 |
| SEM12-S-E-3A | 0.2933 | 0.2258 | 0.2595 |
| hultech-S-E-3A | 0.2475 | 0.2498 | 0.2486 |
| ORG-S-E-1A | 0.2398 | 0.2203 | 0.2300 |
| ORG-S-E-5A | 0.2532 | 0.1976 | 0.2254 |
| hultech-S-E-2A | 0.2263 | 0.2180 | 0.2221 |
| **TUTA1-S-E-1A** | 0.1892 | 0.1756 | 0.1824 |
| LIA-S-E-4A | 0.1655 | 0.1740 | 0.1698 |
| **TUTA1-S-E-2A** | 0.1724 | 0.1569 | 0.1646 |
| LIA-S-E-2A | 0.0278 | 0.0380 | 0.0329 |
| LIA-S-E-3A | 0.0298 | 0.0261 | 0.0280 |
| LIA-S-E-1A | 0.0213 | 0.0296 | 0.0255 |

**Table 2. Chinese subtopic mining runs ranked by D#-nDCG@10**

| Run Name | I-rec @10 | D-nDCG @10 | D#-nDCG @10 |
|---|---|---|---|
| **TUTA1-S-C-1A** | 0.4184 | 0.4686 | 0.4435 |
| THUIS-S-C-1A | 0.3881 | 0.4923 | 0.4402 |
| THUIR-S-C-3A | 0.3786 | 0.4987 | 0.4386 |
| **TUTA1-S-C-2A** | 0.4030 | 0.4655 | 0.4343 |
| THUIS-S-C-4A | 0.4036 | 0.4620 | 0.4328 |
| THUIR-S-C-5A | 0.3892 | 0.4757 | 0.4324 |
| THUIR-S-C-1A | 0.3839 | 0.4802 | 0.4321 |
| THUIR-S-C-2A | 0.3839 | 0.4775 | 0.4307 |
| THUIR-S-C-4A | 0.3792 | 0.4698 | 0.4245 |
| ICRCS-S-C-3A | 0.4046 | 0.4413 | 0.4229 |
| THUIS-S-C-3A | 0.3953 | 0.4504 | 0.4228 |
| ICRCS-S-C-1A | 0.3821 | 0.4219 | 0.4020 |
| ORG-S-C-1A | 0.3644 | 0.4336 | 0.3990 |
| ORG-S-C-4A | 0.3334 | 0.4516 | 0.3925 |
| THUIS-S-C-2A | 0. 3622 | 0.4157 | 0.3890 |
| ORG-S-C-3A | 0.3366 | 0.4407 | 0.3886 |
| ICRCS-S-C-2A | 0.3704 | 0.4024 | 0.3864 |
| KECIR-S-C-2B | 0.3743 | 0.3941 | 0.3842 |
| ORG-S-C-5A | 0.3091 | 0.4175 | 0.3633 |
| ORG-S-C-2A | 0.3163 | 0.4098 | 0.3630 |
| KECIR-S-C-1B | 0.3341 | 0.3763 | 0.3552 |
| KECIR-S-C-3B | 0.3001 | 0.3227 | 0.3114 |
| KECIR-S-C-4B | 0.2917 | 0.3081 | 0.2999 |

# 5. CONCLUSIONS AND FUTURE WORK

In our work of NTCIR-10, we tested the framework for subtopic mining mainly based on modifier graph clustering. The experimental results demonstrate the effectiveness of the proposed framework. However, there is still much space to be further improved. E.g., (1) for English subtopic mining, the effectiveness of incorporating the word-level knowledge derived from a query log should be further tested; (2) When constructing the modifier graph, filtering the noisy subtopic strings by raw frequency is of value to be explored; (3) Instead of directly using cluster frequency to quantify the popularity of a subtopic, other approaches should also be explored and tested; (4) In this paper, we mainly focus on role-explicit topics, and treat role-implicit topics by simply using the edit distance to rank the subtopic strings. In the future, we will make further efforts to explore these problems and improve our proposed framework.

# 6. REFERENCES

[1] R. Agrawal, S. Gollapudi, A. Halverson and S. Ieong. Diversifying Search Results. In Proceedings of the 2nd WSDM, pp. 5-14, 2009.

[2] Rafiei D., Bharat K. and Shukla A. Diversifying Web Search Results. In Proceedings of the 19th WWW, pp. 781-790, 2010.

[3] G. Capannini, F. M. Nardini, R. Perego and F. Silvesti. Efficient Diversification of Web Search Results. In Proceedings of the VLDB Endowment, 4(7), pp. 451-459, 2011.

[4] R. L. T. Santos, C. Macdonald and I. Ounis. On the Role of Novelty for Search Result Diversification. Information Retrieval, 15(5), pp. 478-502, 2012.

[5] R. H. Song, M. Zhang, T. Sakai, M. P. Kato, Y. Q. Liu, M. Sugimoto, Q. L. Wang and N. Orii. Overview of the NTCIR-9 INTENT Task. In Proceedings of NTCIR-9, pp. 82-105, 2011.

[6] T. Sakai, Z. C. Dou, T. Yamamoto, Y. Q. Liu, M. Zhang and R. H. Song. Overview of the NTCIR-10 INTENT-2 Task, 2013.

[7] J. Hu, G. Wang, F. Lochovsky, J. T. Sun and Z. Chen. Understanding User's Query Intent with Wikipedia. In Proceedings of the 18th WWW, pp. 471-480, 2009.

[8] F. Radlinski, M. Szummer and N. Craswell. Inferring Query Intent from Reformulations and Clicks. In Proceedings of the 19th WWW, pp. 1171-1172, 2010.

[9] E. Sadikov, J. Madhavan, L. Wang and A. Halevy. Clustering Query Refinements by User Intent. In Proceedings of the 19th WWW, pp. 841-850, 2010.

[10] F. Bonchi, R. Perego, F. Silvestri, H. Vahabi and R. Venturini. Recommendations for the Long Tail by Term-graph Graph. In Proceedings of the 20th WWW, pp. 15-16, 2011.

[11] H. T. Yu and F. J. Ren. Role-explicit Query Identification and Intent Role Annotation. In Proceedings of the 21st CIKM, pp. 1163-1172, 2012.

[12] T. Liu and S. J. Li. Word Similarity Computing based on Hownet. Computational Linguistics and Chinese Language Processing, 7: 59-76 (in Chinese), 2002.

[13] J. J. Jiang and D. W. Conrath. Semantic Similarity Based on Corpus Statistics and Lexical Taxonomy. In Proceedings of ROCLING X, 1997.

[14] J. Huang and E. N. Efthimiadis. Analyzing and Evaluating Query Reformulation Strategies in Web Search Logs. In Proceedings of 18th CIKM, pp. 77-86, 2009.

[15] X. H. Wang, D. Chakrabarti and K. Punera. Mining Broad Latent Query Aspects from Search Sessions. In Proceedings of the 15th SIGKDD, pp. 867-876, 2009.

[16] J. F. Gao, W. Yuan, X. Li, K. F. Deng and J. Y. Nie. Smoothing Clickthrough Data for Web Search Ranking. In Proceedings of the 32nd SIGIR, pp. 355-362, 2009.

[17] J. R. Wen, J. Y. Nie and H. J. Zhang. Clustering User Queries of a Search Engine. In Proceedings of 10th WWW, pp. 162-168, 2001.

[18] D. Shen, M. Qin, W. Z. Chen, Q. Yang and Z. Chen. Mining Web Query Hierarchies from Click-through Data. In Proceedings of the 22nd AAAI, pp. 341-346, 2007.

[19] X. X. Yin and S. Shah. Building Taxonomy of Web Search Intents for Name Entity Queries. In Proceedings of 19th WWW, pp. 1001-1010, 2010.

[20] A. Noack. Energy Models for Graph Clustering. Journal of Graph Algorithms and Applications, 11(2), pp. 453-480, 2007.

[21] T. Sakai and R. H. Song. Evaluating Diversified Search Results Using Per-intent Graded Relevance. In Proceedings of the 34th SIGIR, pp. 1043-1052, 2011.

[22] T. Xia. The Theory and Algorithm for Chinese Information Similarity (book in Chinese), ISBN: 978-7-5349-4434-5, 2009.

[23] C. X. Zhai, W. W. Cohen and J. Lafferty. Beyond Independent Relevance: Methods and Evaluation Metrics for Subtopic Retrieval. In Proceedings of the 26th SIGIR, pp. 10-17, 2003.

[24] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton and G. Hullender. Learning to Rank Using Gradient Descent. In Proceedings of the 22nd ICML, pp. 89-96, 2005.