A Robust Retrieval Engine for Proximal and Structural Search

Katsuya Masuda[†] Takashi Ninomiya^{†‡} Yusuke Miyao[†] Tomoko Ohta^{†‡} Jun'ichi Tsujii^{†‡}

[†] Department of Computer Science, Graduate School of Information Science and Technology,
University of Tokyo, Hongo 7-3-1, Bunkyo-ku, Tokyo 113-0033, Japan

[‡] CREST, JST (Japan Science and Technology Corporation)

Honcho 4-1-8, Kawaguchi-shi, Saitama 332-0012, Japan

{kmasuda, ninomi, yusuke, okap, tsujii}@is.s.u-tokyo.ac.jp

1 Introduction

In the text retrieval area including XML and Region Algebra, many researchers pursued models for specifying what kinds of information should appear in specified structural positions and linear positions (Chinenyanga and Kushmerick, 2001; Wolff et al., 1999; Theobald and Weilkum, 2000; Clarke et al., 1995). The models attracted many researchers because they are considered to be basic frameworks for retrieving or extracting complex information like events. However, unlike IR by keyword-based search, their models are not robust, that is, they support only exact matching of queries, while we would like to know to what degree the contents in specified structural positions are relevant to those in the query even when the structure does not exactly match the query.

This paper describes a new ranked retrieval model that enables proximal and structural search for structured texts. We extend the model proposed in Region Algebra to be robust by i) incorporating the idea of rankedness in keyword-based search, and ii) expanding queries. While in ordinary ranked retrieval models relevance measures are computed in terms of words, our model assumes that they are defined in more general structural fragments, i.e., *extents* (continuous fragments in a text) proposed in Region Algebra. We decompose queries into subqueries to allow the system not only to retrieve exactly matched extents but also to retrieve partially matched ones. Our model is robust like keyword-based search, and also enables us to specify the structural and linear positions in texts as done by Region Algebra.

The significance of this work is not in the development of a new relevance measure nor in showing superiority of structure-based search over keyword-based search, but in the proposal of a framework for integrating proximal and structural ranking models. Since the model treats all types of structures in texts, not only ordinary text structures like “title,” “abstract,” “authors,” etc., but also semantic tags corresponding to recognized named entities

or events can also be used for indexing text fragments and contribute to the relevance measure. Since extents are treated similarly to keywords in traditional models, our model will be integrated with any ranking and scalability techniques used by keyword-based models.

We have implemented the ranking model in our retrieval engine, and had preliminary experiments to evaluate our model. Unfortunately, we used a rather small corpus for the experiments. This is mainly because there is no test collection of the structured query and tag-annotated text. Instead, we used the GENIA corpus (Ohta et al., 2002) as structured texts, which was an XML document annotated with semantics tags in the field of biomedical science. The experiments show that our model succeeded in retrieving the relevant answers that an exact-matching model fails to retrieve because of lack of robustness, and the relevant answers that a non-structured model fails because of lack of structural specification.

2 A Ranking Model for Structured Queries and Texts

This section describes the definition of the relevance between a document and a structured query represented by the region algebra. The key idea is that a structured query is decomposed into subqueries, and the relevance of the whole query is represented as a vector of relevance measures of subqueries.

The region algebra (Clarke et al., 1995) is a set of operators, which represent the relation between the *extents* (i.e. regions in texts). In this paper, we suppose the region algebra has seven operators; four containment operators ($\triangleright$, $\triangleleft$, ∇ , $\nwarrow$) representing the containment relation between the extents, two combination operators ($\triangle$, ∇) corresponding to “and” and “or” operator of the boolean model, and ordering operator ($\diamond$) representing the order of words or structures in the texts. For convenience of explanation, we represent a query as a tree structure as

1	'([cons]>([sem]>"G#DNA_domain_or_region"))△("in"◇([cons]>([sem]>("G#tissue"▽"G#body_part"))))'
2	'([event]>([obj]>"gene"))△("in"◇([cons]>([sem]>("G#tissue"▽"G#body_part"))))'
3	'([event]>([obj]◇([sem]>"G#DNA_domain_or_region"))△("in"◇([cons]>([sem]>("G#tissue"▽"G#body_part"))))'
4	'([event]>([dummy]>"G#DNA_domain_or_region"))△("in"◇([cons]>([sem]>("G#tissue"▽"G#body_part"))))'

Table 1: Queries submitted in the experiments

3 Experiments

In this section, we show the results of our preliminary experiments of text retrieval using our model. Because there is no test collection of the structured query and tag-annotated text, we used the GENIA corpus (Ohta et al., 2002) as a structured text, which was an XML document composed of paper abstracts in the field of biomedical science. The corpus consisted of 1,990 articles, 873,087 words (including tags), and 16,391 sentences.

We compared three retrieval models, i) our model, ii) exact matching of the region algebra (*exact*), and iii) not-structured *flat* model. In the *flat* model, the query was submitted as a query composed of the words in the queries in Table 1 connected by the “and” operator (Δ).

The queries submitted to our system are shown in Table 1, and the document was “sentence” represented by “<sentence>” tags. Query 1, 2, and 3 are real queries made by an expert in the field of biomedicine. Query 4 is a toy query made by us to see the robustness compared with the *exact* model easily. The system output the ten results that had the highest relevance for each model².

Table 2 shows the number of the results that were judged relevant in the top ten results when the ranking was done using ρ_{sum} . The results show that our model was superior to the *exact* and *flat* models for Query 1, 2, and 3. Compared to the *exact* model, our model output more relevant documents, since our model allows the partial matching of the query, which shows the robustness of our model. In addition, our model outperforms the *flat* model, which means that the structural specification of the query was effective for finding the relevant documents. For Query 4, our model succeeded in finding the relevant results although the *exact* model failed to find results because Query 4 includes the tag not contained in the text (“<dummy>” tag). This result shows the robustness of our model.

Although we omit the results of using ρ_{sc} and ρ_{ic} because of the limit of the space, here we summarize the results of them. The number of relevant results using ρ_{sc} was the same as that of ρ_{sum} , but the rank of irrelevant

²For the *exact* model, ten results were selected randomly from the exactly matched results if the total number of results was more than ten. After we had the results for each model, we shuffled these results randomly for each query, and the shuffled results were judged by an expert in the field of biomedicine whether they were relevant or not.

Query	<i>our model</i>	<i>exact</i>	<i>flat</i>
1	10/10	9/10	9/10
2	6/10	5/5	3/10
3	10/10	9/9	8/10
4	7/10	0/0	9/10

Table 2: (The number of relevant results) / (the number of all results) in top 10 results.

results using ρ_{sc} was lower than that of ρ_{sum} . The results using ρ_{ic} varied between the results of the *flat* model and the results of ρ_{sum} depending on the value of λ .

4 Conclusions

We proposed a ranked retrieval model for structured queries and texts by extending the region algebra to be ranked. Our model achieved robustness by extending the concept of words to extents and by matching with sub-queries decomposed from a given query instead of matching the entire query or words.

References

- T. Chinenyanga and N. Kushmerick. 2001. Expressive and efficient ranked querying of XML data. In *Proceedings of WebDB-2001*.
- C. L. A. Clarke, G. V. Cormack, and F. J. Burkowski. 1995. An algebra for structured text search and a framework for its implementation. *The computer Journal*, 38(1):43–56.
- T. Ohta, Y. Tateisi, H. Mima, and J. Tsujii. 2002. GENIA corpus: an annotated research abstract corpus in molecular biology domain. In *Proceedings of HLT 2002*.
- S. E. Robertson and S. Walker. 2000. Okapi/Keenbow at TREC-8. In *TREC-8*, pages 151–161.
- A. Theobald and G. Weilkum. 2000. Adding relevance to XML. In *Proceedings of WebDB'00*.
- J. Wolff, H. Flörke, and A. Cremers. 1999. XPRES: a Ranking Approach to Retrieval on Structured Documents. Technical Report IAI-TR-99-12, University of Bonn.