

OLLSCOIL NA hÉIREANN
THE NATIONAL UNIVERSITY OF IRELAND
COLÁISTE NA hOLLSCOILE, CORCAIGH
UNIVERSITY COLLEGE, CORK

Summer Examinations 2005

Computer Science
CS4040: Information Retrieval and Organisation

Prof. M. Calder
Prof. G. Provan
Mr. H. Sorensen

Answer *four* questions

Paper Total: 160 Marks

All questions carry equal marks

Time: 3 Hours

1. Search Engines

You intend to develop a fully-functional (text-only) search engine for deployment on the internet. It is to support a rich set of query types, with well-defined semantics. For a given query, it will rank results based on conditions being satisfied, the number of query terms present, their specificity, where they are within the page, and on the page popularity.

- (a) State what query types you intend to support and the default semantics for a multi-term query. (5 marks)
- (b) Outline the structure of an *inverted index* suitable for support of the query types of (a). (10 marks)
- (c) The search engine will retrieve all Web pages that exactly match a user's query words (or expression). However, it is also to reward – to a lesser extent – pages containing variants of the query terms. How exactly might this be achieved? (6 marks)
- (d) Specify an appropriate *ranking algorithm* that might be employed by this search engine. (10 marks)
- (e) Specify – and explain – an expression by which *page popularity* might be computed. (4 marks)
- (f) Is this search engine susceptible to spoofing by Web page writers? Explain. If so, explain how such spoofing might be challenged and outline any detrimental effects your solution might have. (5 marks)

2. IR Models

(a) In the *Boolean Model*, queries must be expressed as Boolean expressions and must be cast into disjunctive normal form (DNF) prior to evaluation.

(i) Cast the following query into DNF:

Term1 OR Term2 AND (Term3 OR Term4) (3 marks)

(ii) Specify how the DNF query is evaluated against the document-term representation. (3 marks)

(iii) Why is the Boolean Model considered unsuitable for modern-day deployment in information retrieval? (3 marks)

(b) The *Vector Space Model* (VSM) is a commonly used model for document and query representation. Describe this model, paying particular attention to the following features:

(i) Give both the geometric and the algebraic representation of a document vector. (4 marks)

(ii) Define and explain why *term frequency / inverse document frequency (tf/idf)* is a suitable weighting function for this retrieval model. (4 marks)

(iii) State why and how the document and query vectors are generally normalised to be unit length from the origin. (4 marks)

(iv) Outline the basic storage structure for a collection of document representations. (3 marks)

(v) Specify the computational model whereby the storage structure outlined in (iv) is used to compute the ranking of stored documents relative to a query. (4 marks)

(vi) State what steps might be taken to increase the efficiency of the ranking algorithm. (3 marks)

(vii) Would you consider the VSM to be universally suited to textual information retrieval? Explain. (4 marks)

(c) The *Latent Semantic Indexing (LSI) Model* is a variation of the VSM that attempts to address some perceived shortcomings of the latter. State the essential difference between these models and the benefits and drawbacks of using the LSI approach. (5 marks)

3. Query/Retrieval Enhancement

(a) One approach to query enhancement involves the use of *relevance feedback*, whereby the results of one query phase can be used to adjust the input into the next phase, thereby allowing the phased (re)focusing of a query. With respect to either the *Vector Space Model* **or** the *Probabilistic Model* for information retrieval, address the following aspects of the process:

(i) Explain the basic principles involved in relevance feedback. (6 marks)

- (ii) Derive and explain the means by which a query or similarity expression is altered from one query phase to the next. (8 marks)
 - (iii) Specify how one might objectively assess the improvement brought about by the use of relevance feedback. (4 marks)
- (b) A second form of query enhancement involves *query expansion* – usually requiring the use of a *statistical thesaurus* so as to add synonymous terms to a query. This thesaurus itself can be constructed through either *local* or *global clustering*.
- (i) State the main differences between *local* and *global* clustering. (3 marks)
 - (ii) Describe the synonym-finding model of any one local clustering method. (9 marks)
- (c) Clustering can also be applied to documents, rather than to vocabulary terms. When combined with a suitable *visualisation* interface, it can prove beneficial to users.
- Outline a suitable visualisation tool that might be employed and state its advantages from the perspective of those requiring information retrieval. (10 marks)

4. Retrieval Evaluation; Document & Query Processing

- (a) *Precision* and *recall* are the most common measures of retrieval accuracy. They are usually combined to produce a *Precision:Recall (P:R) Graph*.
- (i) Define the terms precision and recall. (3 marks)
 - (ii) Specify the steps involved in carrying out experimental evaluation of P & R. (4 marks)
 - (iii) Assume that an IR system ranked documents as listed overleaf in the left hand column. Assume that the marked documents are those deemed by the user to be relevant. Construct the P:R graph for this case and comment of the apparent accuracy of the IR system. Repeat the process for the right-hand column (depicting another IR system using the same query & documents). Assume that you have prior knowledge that the collection contains 15 relevant documents. Briefly compare and contrast the IR systems involved. (13 marks)
 - (iv) IR systems can have very different types of uses – from casual to professional and from short-query to full-text, for example. For which type(s) of system is the P:R measure most appropriate? Explain. For a general-purpose search engine, what do you think would be a suitable evaluation measure? (6 marks)
- (b) There are several reasons for – and several means by which - terms can be added to, or omitted from, document and/or query representations.
- (i) Why would an IR system add terms to queries? State how. (5 marks)
 - (ii) Why would an IR system remove terms from queries? State how. (5 marks)
 - (iii) What effect would you expect these operations to have on precision & recall measures? (4 marks)

Doc 9		Doc 125	*
Doc 1		Doc 68	*
Doc 31		Doc 90	*
Doc 125	*	Doc 11	*
Doc 12		Doc 86	*
Doc 68	*	Doc 82	
Doc 90	*	Doc 9	
Doc 88	*	Doc 1	
Doc 77	*	Doc 10	*
Doc 56	*	Doc 16	
Doc 54	*	Doc 17	
Doc 33		Doc 229	
Doc 11	*	Doc 301	
Doc 8	*	Doc 2	
Doc 66	*	Doc 33	
Doc 225	*	Doc 7	
Doc 16		Doc 29	

5. Multimedia IR

- (a) Explain how an efficient approach to *colour image* information retrieval might be implemented. Pay particular attention to the following issues:
- (i) Specify an appropriate distance function between two colour images. (5 marks)
 - (ii) State what GEMINI feature extraction might be appropriate. (5 marks)
 - (iii) Specify how the distance in feature space might be computed. (4 marks)
 - (iv) What experiment(s) might be carried out to estimate the computational saving achieved by use of the GEMINI approach? Indicate the expected results. (5 marks)
 - (v) Explain whether your approach alleviates the *dimensionality curse* or the *cross-talk* problems inherent in multimedia retrieval. (4 marks)
- (b) *Information visualisation* can prove useful in the case of either video or audio content analysis. With respect to any one of these media, specify how this visualisation is achieved and used:
- (i) State which media feature(s) are being visualised and what display structure is employed. (5 marks)
 - (ii) Outline the computation involved in producing the data from which the display structure is derived. (7 marks)
 - (iii) State what interpretations and conclusions can be inferred from the visualisation. (5 marks)