

Explicit Diversification of Image Search

Jonathon S. Hare
jsh2@ecs.soton.ac.uk

Paul H. Lewis
phl@ecs.soton.ac.uk

Electronics and Computer Science, University of Southampton, United Kingdom

ABSTRACT

Search result diversification can increase user satisfaction in answering a particular information need. There are many ways of diversify search results. In some cases the user has a clear idea of how they would like to see their results diversified. This work presents a system that is capable of diversifying search results along specific user-specified axes of diversity.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval

Keywords

Image search; Diversity; Result diversification

1. INTRODUCTION

The information need of a user is often better satisfied by a retrieval system when the result set for a particular query shows many different aspects of that query. This is especially important when the query is poorly specified or ambiguous [4, 2]. By presenting a diverse range of results covering many possible representations of a query the probability of finding relevant images is increased. In terms of image search there are many ways of diversifying results; often these are implicit and involve removing near duplicates [3]. Sometimes, however, users have an explicit idea about how they would like their results to be diversified. The demonstration described in this work shows an image search engine with explicit result diversification.

Our linked-data diversity search tool uses semantic web technologies to perform web-based image searches along specific axes of diversity. The initial vision for this tool was described briefly in [5]. Unlike a normal search engine the application requires you to be a little more specific with your query - at a minimum it requires a “subject” of the search, and an “axis” along which you would like to see the results presented. For example, you could ask the engine to find images about “David Beckham”, organised by the various clubs he has played for. The implementation builds on top of existing web and image search engines and semantic web

tools. Specifically, Microsoft’s Bing Search API¹ is used as the basis for the image and web search, and DBpedia [1] is used as the semantic knowledge store that helps provide query diversification.

2. DEMO

A live version of the tool is deployed at <http://degas.ecs.soton.ac.uk/~jsh2/diversity>. Figure 1 illustrates a search for images of the objects manufactured by the Japanese car manufacturer Mazda. The subject of the search was “mazda”, and the axis of diversity was “manufacturer”. The performance of the tool is very heavily query dependent, and is affected by how well the query terms are modelled in DBpedia as well as the performance of the underlying search engine. However, in a number of practical queries that have been tested, the results are often quite good and they can be maximised through iterative querying by the user.

3. DESCRIPTION

An overview of the tool can be seen in Figure 2. The search engine itself works by consuming machine-readable data from DBpedia and combining it with links to image documents from Bing’s search engine. A flowchart of the processes taken within the tool is shown in Figure 3.

The search process starts with the formation of the query by the user. The search specification consists of a mandatory subject and diversity axis. In addition, extra contextual search terms can be provided to narrow down the search. The user interface of the tool allows queries to be entered in two different ways; firstly, the user can enter the subject term(s), context term and a term describing the diversity axis in three separate free-text fields. Alternatively, there is a “wizard” interface in which the subject of the search can be entered, and on the following screen the user is able to select from a pull-down list of diversity axes directly relevant to that subject, in addition to entering additional context terms.

Once the query has been specified the following steps are taken within the search engine:

1. The Bing API is used to perform a web search, restricted to the Wikipedia domain, for the subject terms. The top-matching Wikipedia URL is returned.
2. The Wikipedia URL for the subject terms is transformed to a DBpedia URI that represents the subject.

¹<https://datamarket.azure.com/dataset/bing/search>



Figure 1: Search for images of objects *manufactured by Mazda*




Figure 2: Overview of the diversity search tool.

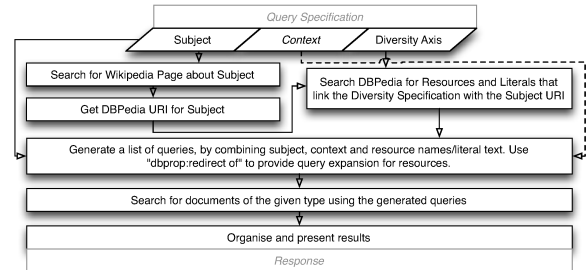


Figure 3: Flowchart illustrating the steps taken in producing diversified search results.

3. The diversity axis of a search is internally represented as a DBpedia URI. The URI is created as follows:
 - (a) If the user entered a URI directly, then this is used by the engine.
 - (b) If the user entered a free-text term for the axis, then the engine tries to guess the correct URI by performing a SPARQL query against DBpedia of the form `select distinct ?p where {<db_uri> ?p ?obj filter regex(?p, "axis_term")}`. If there are no results then the query is repeated with the `<db_uri>` and `?obj` swapped.
 - (c) If the user used the wizard to select the diversity axis, then the search engine already has the correct URI. The axis choices presented to the user are determined by querying the DBpedia SPARQL interface with two queries: `select distinct ?x ?y where {<db_url> ?x ?y}` and `select distinct ?x ?y where {?y ?x <db_url>}`. The set of returned `?x`'s that are of URI type are used to form the list of possible diversity axes.
4. The search engine now proceeds to look for resources related to the *subject* along the *diversity axis* using the following query: `select distinct ?x where {<db_url> <div_axis> ?x}`. If the result set of this query is empty, then the reverse query is performed: `select distinct ?x where {?x <div_axis> <db_url>}`. Each `?x` value is considered to be a point on the axis.
5. The search engine now iterates through the `?x` objects returned by the previous step and builds queries which are sent to the Bing image search API. The query sent to the API has the following form: `+"subject" +"context" +"axis_term"`. The actual construction of the `axis_term` part of the query depends on the type of the `?x` object in question; if the `?x` object is literal, then it is used directly. If it is a URI, then a lookup step takes place to determine a list of potential

terms that describe the URI, and each of these is used in turn for the Bing query.

6. Finally, the tool performs a rendering step to display the results of the search, organised by each of the `?x` points on the axis of diversity.

4. ACKNOWLEDGMENTS

The described work was funded by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement 231126 (LivingKnowledge).

5. REFERENCES

- [1] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives. DBpedia: A Nucleus for a Web of Open Data. *The Semantic Web*, pages 722–735, 2007.
- [2] H. Chen and D. R. Karger. Less is more: probabilistic models for retrieving fewer relevant documents. In *ACM SIGIR'06*, pages 429–436, New York, NY, USA, 2006. ACM.
- [3] M. Paramita, M. Sanderson, and P. Clough. Diversity in photo retrieval: overview of the ImageCLEFPhoto task 2009. In *CLEF working notes*, Corfu, Greece, 2009.
- [4] K. Song, Y. Tian, W. Gao, and T. Huang. Diversifying the image retrieval results. In *ACM MM'06*, pages 707–710, New York, NY, USA, 2006. ACM.
- [5] P. Zontone, G. Boato, F. G. B. D. Natale, A. D. Rosa, M. Barni, A. Piva, J. Hare, D. Dupplaw, and P. Lewis. Image diversity analysis: Context, opinion and bias. In *The First International Workshop on Living Web: Making Web Diversity a true asset*, volume 515. CEUR-WS, October 2009.