Faceted Browsing over Large Databases of Text-Annotated Objects

Wisam Dakka
Columbia University
wisam@cs.columbia.edu

Panagiotis G. Ipeirotis
New York University
panos@nyu.edu

Kenneth R. Wood
Microsoft Research
krw@microsoft.com

Abstract

We demonstrate a fully working system for multifaceted browsing over large collections of text-annotated data, such as annotated images, that are stored in relational databases. Typically, such databases can be browsed across multiple facets (by topic, genre, location, and so on) and previous user studies showed that multifaceted interfaces improve substantially the ability of users to identify items of interest in the database. We demonstrate a scalable system that automatically generates multifaceted browsing hierarchies on top of a relational database that stores the underlying text-annotated objects. Our system supports a wide range of ranking alternatives for selecting and displaying the best facets and the best portions of the generated hierarchies, to facilitate browsing. We combine our ranking schemes with Rapid Serial Visual Presentation (RSVP), an advanced visualization technique, which further enhances the browsing experience and demonstrate how to use prefetching techniques to overcome the latency issues that are inherent when browsing the contents of a relational database using multifaceted interfaces.

1. Introduction

Searching and browsing are the typical ways that users locate items of interest in collections of text and text-annotated objects. An example of such a database is the Corbis image database, where each image has a title, a free-text description, and a set of associated keywords. Recently, Yee et al. [3] showed that browsing interfaces that rely on multifaceted hierarchies represent a new powerful browsing paradigm, superior to interfaces that rely on single monolithic hierarchies.

Here, we demonstrate our techniques [1] that discover automatically the facets that can be used to browse an underlying database (Section 2). We also demonstrate how to enhance the ability of users to identify items of interest in the underlying database, by using ranking algorithms that take into consideration the available screen real estate and with the use of RVSP, an advanced visualization technique that exposes the contents of the underlying database, with minimal use of the screen real estate (Section 3). Finally, our system demonstrates how to enhance the browsing experience by using predictive prefetching techniques (Section 4).

2. Automatic Facet Discovery

With the increase of online databases, systems such as Netflix, Yahoo Shopping, or the Flamenco system at UC-Berkeley, have been utilizing multifaceted interfaces with *pre-defined* facets to help users locate items of interest. Our goal is to identify *automatically* facets from collections of text-annotated objects, which have no predefined facet structure. In [1], we presented a supervised technique that identifies the appropriate facet for each annotation keyword. For example, the keywords "*cat*" and "*dog*" are assigned to the "*Animals*" facet, while we put the words "*mountain*" and "*fields*" are assigned to the "*Topographic Features*" facet.

Demonstration Features: Our demonstration system takes as input an arbitrary database of textually annotated objects. Then, it identifies the potential facets for each object in the database and builds automatically the appropriate browsing hierarchies for these facets. The user can select one of the multiple image databases that we will have available and see how different facets emerge from databases. The user can observe how databases with different focal points generate different facets. For example, databases with images of desk objects do not contain the "*Animal*" or the "*Emotions*" facet, while databases with images of humans allow these facets to emerge.

3. Browsing through Multiple Hierarchies

Creating the multifaceted browsing structure is not the complete solution. The browsing hierarchies should be easy to navigate, and allow the users to explore quickly the underlying database. Our system will demonstrate the effectiveness of two of the techniques that we used to facilitate multifaceted browsing over collections of text-annotated objects.

3.1. Adaptive Category Ranking

For a large database, the size of the browsing hierarchy can be big, and it may not be possible to fit the contents of the hierarchy on the screen. The problem is even more pronounced when the user has only access to devices with

1-4244-0803-2/07/$20.00 ©2007 IEEE

a small screen window, such as a PDA or a smartphone. Therefore, it is important to make judicious choices about the portions of the hierarchy that are presented to the user. We have implemented a set of techniques for ranking the categories of a hierarchy, and we demonstrate our *merit-based ranking*, which ranks higher categories that expose the largest fraction of the databases with the lowest cost for the user. The cost is the time that the user spends going down the hierarchy, plus the time required to read the descriptions of the children categories and going through the retrieved objects. Hence, merit-based ranking takes into consideration the structural properties of the hierarchies and ranks higher branches that allow for easier and effective browsing.

Demonstration Features: The user will be able to modify the size of the visible part of the hierarchy, and the system will select the best categories to display, depending on the size of the window. Our system will adapt the hierarchies on-the-fly, using the available ranking alternatives.

3.2. Rapid Serial Visual Presentation

Rapid Serial Visual Presentation (RSVP) is a visualization technique that displays information (generally text) using a limited space [2]. RSVP techniques provide a trade-off between space and time, a trade-off of particular value when available space is limited. In our system, we use RSVP techniques for text, to supply the users with useful information about the categories they are intending to explore. For example, scrolling horizontally the titles of the subcategories under the selected category title allows the user to quickly capture the content of the category. As another example, "blinking" quickly the titles of the most visited images can also give the user an idea of why others have visited the selected category. Moreover, our system exploits RSVP to reveal structural characteristics about the categories that are presented in the user window. For instance, we vary the speed of the scrolling based on the size of the category or the number of subcategories, when the user just hovers on the category title, helps determine the structure of the underlying hierarchy.

Demonstration Features: In our demonstration, we illustrate the techniques described above, and present our ongoing research of utilizing RSVP techniques to enhance browsing. Users will be able to experience with the different RSVP features implemented in our windows user interface such as scrolling horizontally the titles of the subcategories under the selected category title, "blinking" subcategories in quick pace, and so on.

4. Prefetching for Interactive Browsing

A truly interactive browsing interface needs an interface that responds within milliseconds to the user actions. The most typical user action within our interface is the selection of a category from one of the available hierarchies. The selected category is not necessarily a child of the current category. It may belong to a different facet and it can be a parent or a child of the currently selected category for that facet. This seemingly simple action typically results in a sequence of several SQL statements. The system needs to update *all the object counts* for *all visible categories*, across all facets that are visible to the user. Such operations are expensive, especially when the user "rolls-up" across a facet and expands the selection to include more objects. ("Drilling-down" is comparatively simpler.) In this case, the object count updates may take several seconds to complete, and can make the browsing experience cumbersome and tiring.

Our system predicts the potential next click(s) of the user and precomputes the SQL statements required by the possible next mouse clicks. We prioritize the SQL precomputation, by giving higher priority to actions that can be generated by mouse clicks that are closest to the current mouse position. The time that the user spends browsing through the results is typically enough for our system to precompute all the SQL statements that can be generated from the next two clicks of the user. The resulting system gives the impression of real-time browsing even when we use a web browser and connect to the database from a remote location.

Demonstration Features: In our demonstration, users will be able to use our interface, and turn on and off the prefetching strategy to experience first-hand the difference in performance. We will demonstrate the feasibility of building such a browsing interface using a simple web browser, by using an AJAX interface.

5. Conclusions

We described our demonstration system that demonstrates multiple issues that arise when trying to build a *multifaceted browsing* interface over a database of text-annotated objects. The goal of this demo is to expose more database researchers to research problems that arise when trying to *browse* the contents of a text database, and let users understand the power of new interfaces for interacting with content stored in databases.

References

[1] W. Dakka, P. G. Ipeirotis, and K. R. Wood. Automatic construction of multifaceted browsing interfaces. In *CIKM 2005*.

[2] O. de Bruijn and R. Spence. Rapid serial visual presentation: A space-time trade-off in information presentation. In *AVI 2000*.

[3] K.-P. Yee, K. Swearingen, K. Li, and M. A. Hearst. Faceted metadata for image search and browsing. In *CHI 2003*.