

Preprint of

Chapter 4: User Interface Design (p.75)

from

Dynamic Taxonomies and Faceted Search

Theory, Practice, and Experience

Series: The Information Retrieval Series , Vol. 25

Sacco, Giovanni Maria; Tzitzikas, Yannis (Eds.)

2009, XVII, 340 p., Hardcover

ISBN: 978-3-642-02358-3

(c) Springer Science and Business Media, Inc.

The original publication is available at

<http://www.springer.com/978-3-642-02358-3>

Chapter 1

User Interface Design

Moritz Stefaner, Sébastien Ferré, Saverio Perugini, Jonathan Koren and Yi Zhang

As detailed in Chapter ??, system implementations for dynamic taxonomies and faceted search allow a wide range of query possibilities on the data. Only when these are made accessible by appropriate user interfaces, the resulting applications can support a variety of search, browsing and analysis tasks.

User interface design in this area is confronted with specific challenges. This chapter presents an overview of both established and novel principles and solutions. Based on a definition of core principles (see Section 1.1) and challenges (see Section 1.2), we define a taxonomy of navigation modes observed in existing applications (see Section 1.3). On that basis, design patterns for enabling these navigation modes in user interfaces (see Section 1.4) as well as extensions and related approaches (see Section 1.5) are discussed. The chapter closes with an approach to personalizing faceted search (see Section 1.6).

Moritz Stefaner

Interaction Design Lab, University of Applied Sciences Potsdam, Pappelallee 8–9, 14469 Potsdam, Germany, e-mail: moritz@stefaner.eu

Sébastien Ferré

Irisa, Université de Rennes 1, Campus de Beaulieu, F-35042 Rennes cedex, France, e-mail: ferre@irisa.fr

Saverio Perugini

Department of Computer Science, University of Dayton, 300 College Park, Dayton, OH 45469-2160 e-mail: saverio@udayton.edu

Jonathan Koren

University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA, USA 95064, e-mail: jonathan@soe.ucsc.edu

Yi Zhang

University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA, USA 95064, e-mail: yiz@soe.ucsc.edu

1.1 Principles

Extending traditional models of Information Retrieval, *search for digital resources* has lately been widely recognized as multi-step processes [66, 58, 9, 37]. To follow the terminology introduced in [37], a search usually involves an *initial constraint definition*, followed by an *orienteering and refinement* phase based on first inspections of the result, and finished with a closer examination of individual results in the so-called *endgame*.

In this context, the exploration of dynamic taxonomies [70] with *facet browsers* is often seen as a most promising candidates for "rich exploration of a domain across a variety of sources from a user-determined perspective" [49]. These make different aspects of the underlying data accessible in parallel. Selecting one of the values, and thus filtering the result set, restricts the available metadata values only to those occurring in the results. Consequently, the user is visually guided through an iterative process of query refinement and expansion, never encountering situations with zero results.

Applications for faceted search and dynamic taxonomies offer the following key features to support a wide range of search and browsing tasks:

- **Unrestricted query formulation over multi-dimensional classification**

Facet browsing applications impose no restrictions, in which order, or in which granularity filters are applied on a result set. Filters stem from various, orthogonal dimensions that can be combined by Boolean operators. This allows the formulation of complex queries, such as "All documents created before date A, related to topic B, and of file type C or D". The equal treatment of multiple dimensions differs from, e.g. typical web site structures or file systems, where a single taxonomy is the pre-dominant organization principle, and other metadata are only supplements for sorting or filtering.

- **Poka Yoke: No more Empty Result Sets**

One of the core principles of dynamic taxonomies to restrict the available filtering options in the given focus to only those, which will lead to a non-empty result set. Hence, the user can never run into a situation with zero results. This is opposed to the process in a typical *advanced search* situation, where first a complex boolean query is constructed, which is then evaluated on demand (see e.g. Figure 1.1). That, however, can result in empty result sets, often without further indication, which part of the query could be relaxed in order to retrieve some results. The exclusion of potentially frustrating situations by design is often referred to as Poka-Yoke principle¹.

- **Orienteering and Domain Understanding**


It is a common pattern to visualize the number of occurrences of a concept in the given focus. The simplest option is to provide it in brackets after the concept label (e.g. "Europe (5)"). Advanced techniques include the application of visual indicators, such as bar height or small bar charts (see Section 1.4.7).

¹ see e.g. <http://en.wikipedia.org/wiki/Poka-yoke>

Match **ALL** of the following rules:


Includes

Keyword Anywhere

 Add Another

but does not include

Keyword Anywhere

 Add Another

Limit my search results by:

Branch:

Any ☐ Only available items

Language: Any Language Format: Any Format Audience: Any

Date Published:

Between: and

Fig. 1.1 The advanced search interface for the Oakville Public library at <http://opl.bibliocommons.com/search>

This provides valuable *information scent* [64], i.e. “a user’s (imperfect) perception of the value, cost, or access path of information sources obtained from proximal cues” [89]). Orienteering, or “directed situated navigation” [82] is the process of reaching a goal through a series of small actions, supported by continuous evaluation of the respective focus. In this context, knowing beforehand, how many resources to expect after adding a concept as a filter, can be a valuable indicator of the utility of the filtering action. Additionally, this principle can be extended in order to foster domain understanding by learning about characteristic metadata distributions (see Section 1.4.7);

1.2 Challenges

The prototypical facet browsing application has at least two main interface areas: one for presenting facets and their values, one for displaying the result set. Additional components might include a detail view for selected resources and a breadcrumb strip for filter summary and selection history navigation (see Section 1.4.5).

Based on this basic setup, a number of dimensions can vary in the system and user interface design, and need to be carefully decided upon:

- Which is the data type of the different facets – nominal, hierarchical, ordinal, real valued?
- How are facet values presented to the user? Are all facets and values visible, or only a selection?
- Can the user select multiple values per facet? If so, does this result in conjunctive or disjunctive queries?

Based on these fundamental considerations in setting up a faceted navigation scheme, and designing an appropriate interface, the following recurring challenges in designing these systems will have to be tackled [34, 35, 54]:

- **Boolean query logic** A selection of single concepts from different facets is usually understood as conjunction (AND-query). If, however, multiple values within one facet are selectable (for instance, “red” and “green” from the “color” facet), depending on context and data set, either a conjunctive (“red” AND “green”) or a disjunctive (“red” OR “green”) interpretation are conceivable. If an application only uses one of these selection modes, this needs to be communicated to the user; if both are possible, separate controls for both modes will be needed (see Section 1.4.1).
- **Cluttered interfaces** The paradigm of making all filter options available in parallel naturally leads to the challenge of having to fit many controls and text fields on the user screen. Hence, clear visual structure and hierarchy as well as strategies to reduce visual clutter are vital. If a full exposure of all facets is not possible due to size constraints, strategies and user controls for showing and hiding, or expanding and collapsing facets will have to be integrated (see Sections 1.4.2 and 1.4.3).
- **Incorporating keyword search** A free-form keyword field in order to search for arbitrary terms in addition to the pre-defined classification scheme is a “key component to successful faceted search interfaces” [34]. One source of confusion can be the question, if the search field will act as a plain text filter (e.g. searching over titles and descriptions of the resources) or if it will also match classification terms. A third conceivable option is a “search within the results”, which just filters the result display, but does not act as a full-fledged facet. In either case, the relation of the free-form search to the rest of the filters has to be signaled clearly in order to avoid misconceptions (see Section 1.4.4).
- **Change blindness** Change blindness is a well-known psychological phenomenon [65]: a person viewing a visual scene apparently fails to detect large changes in the scene, if the change in the scene coincides with some visual disruption such as a saccade (eye movement) or a brief obscuration of the observed scene or image. This situation often occurs in web applications, where the web page briefly flashes after actions demanding a new server request. In this context, animated transitions can facilitate perception of changes in user interface design [39, 83, p. 84]. Perception of change is especially important for facet browsing, as the

sudden disappearance of list items after a click can be a source for misconceptions and confusion. Besides animation, clear marking of the current focus and the resulting effects are recommended (see Section 1.4.6).

1.3 Navigation Modes

As a basis for comparing user interface design patterns in the next section, this section defines and illustrates different navigation modes, that enable the user to navigate the available information space by consecutively applying operators on the query.

Given an infobase over a taxonomy (T, \leq) of concepts, a query is a Boolean combination of concepts. We recall the extension of such a query can be computed from the extensions of concepts by applying set operations: intersection for conjunction (*and*), union for disjunction (*or*), and complement for negation (*not*).

From this perspective, browsing an infobase consists in navigating from query to query. This is more general than defining browsing as navigating from sets of objects to sets of objects, because every query determines a set of objects, its extension, and not all navigation modes can be defined as a function from sets of objects to sets of objects. The queries are constructed by following navigation links or using interface controls. Most navigation links are provided by dynamic taxonomies, which also summarize the extension of the current query.

Based on an analysis of existing applications, we can distinguish the following navigation modes:

- *zoom-in* makes the query more specific,
- *zoom-out* makes it more general,
- *shift* replaces a part of the query by a related concept,
- *pivot* replaces the whole query by a related concept,
- *slice-and-dice* allows the disjunctive selection of multiple concepts within a facet,
- *range selection* offers the options to specify query intervals within ordinal or real value facets.

The change from a query to another query, and hence, from a focus to another focus, is defined as a *navigation link*. A navigation link is decomposed into a *selection* and a *navigation mode*. This means that a same selection can be used in different ways to reach different foci. In the simple case, a selection is a concept in the dynamic taxonomy of the current focus. In the general case, a selection is the disjunction of the concepts that are selected in the dynamic taxonomy (e.g., France or Germany or Italy). Controls in the interface can be activated to apply modifiers on such selections: adding negation (e.g., *not (Animal or Plant)* from the selection of *Animal* and *Plant*), replacing equalities by inequalities (e.g., *date >= 2002* from the selection of *date = 2002*). Given

those selections, the above navigation modes can be reduced to only two primitive navigation modes, *zoom* and *pivot*:

- *zoom-in* is a zoom on a selection whose extension contains some objects of the focus, but not all,
- *zoom-out* is a zoom on a selection whose extension contains all objects of the focus,
- *shift* is a combination of zoom-in and zoom-out,
- *pivot* is a basic mode,
- *slice-and-dice* is a zoom on a selection with disjunction,
- *range selection* is a zoom on a selection with inequalities.

An additional navigation mode is *querying-by-examples*, which defines the query from the selection of a set of objects, the examples.

The definitions of navigation modes rely on the fact that queries can be put in conjunctive normal form, i.e. conjunctive sets of simpler queries. For instance, `France and not date <= 2000 and (Building or Landscape)` is equivalent to `{France, not date <= 2000, Building or Landscape}`.

In the following, these interaction modes are illustrated with an example scenario using Camelis², a system for browsing a personal photo collection spanning the period 1999-2007. This collection contains 5,820 photos, which are described by date, location, event, type, visible persons and objects, and EXIF descriptors (e.g., time, flash, orientation).

Figure 1.2 shows a screenshot of Camelis. The current query is at the top. The extension of this query, i.e., the current focus, is at the right, where each object is represented by a thumbnail or a text snippet depending on its type. The dynamic taxonomy is at the left, in the form of concept trees whose nodes are expanded on demand. The number at the left of each concept represent its count in the current focus, and the font scale is logarithmically proportional to this count. At the top of the dynamic taxonomy, there are check and radio buttons to modify the current selection (insertion of negation and inequalities), and two buttons for applying the two primitive navigation modes (zoom and pivot).

1.3.1 Zoom-in

First, suppose some user, say Lisa, wants to find some photos from a trip in Australia at the conference ICFCA'04. She first expands the concept `Location`, and finds she has photos from Europe (5346), Africa (162), and Australia (148). After selecting the concept `Australia`³:

² The version used here is 1.4, and can be downloaded at <http://www.irisa.fr/LIS/ferre/camelis/>.

³ There are French words in the screenshot as it is a personal photo collection, but english translations are used in the text for better consistency.

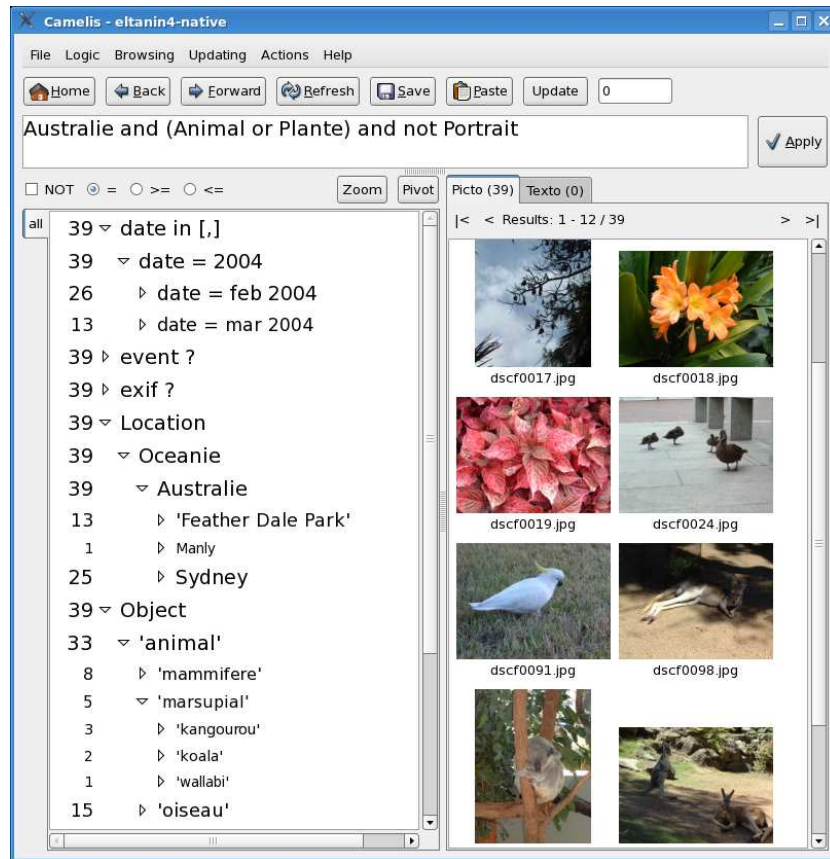


Fig. 1.2 The graphical interface of CAMELIS.

- the query becomes *Australia*,
- the extension displays 12 photos (out of 148),
- the concept *Australia* has now maximal font scale because all photos in the extension belong to it, and it is automatically expanded to show sub-locations of *Australia* (Lisa finds that she has been mainly in *Sydney* (105), and in the *Blue Mountains* (18)),
- the concepts *Europe* and *Africa* are no longer visible, because they are no longer relevant, i.e., count = 0.

Now she expands the concept *Type* and sees there are different types of photos: buildings (29), animals (34), and plants (6). She becomes interested in Australian organisms, so she selects both *Animal* and *Plant*, which leads her to the refined query *Australia and (Animal or Plant)*, whose extension contains 40 photos. One of these photos is a portrait, which Lisa does not want, so she selects the negation of *Portrait*. This leads her to the new query *Australia*

and (Animal or Plant) and not Portrait (39 photos). By expanding more concepts, she discovers that these photos were taken in February and March 2004, mainly in Sydney and at the Feather Dale Park, and that 5 photos of three different species of marsupials are present: kangaroo, koala, wallaby.

Figure 1.2 shows the interface obtained after the previous navigation operations. At this stage, Lisa can either browse the 39 photos in CAMELIS, or launch a slideshow in an external application.

These three navigation steps lead to local views with increasingly more precise queries, and hence increasingly smaller extensions. This is called *zoom-in* navigation, because it corresponds to moving towards smaller extensions. Its principle is to specialize the current query q by the selection x . A simple definition of the resulting query would be q and x , but this would entail redundancy in queries: e.g., *Australia* and *Sydney* which is equivalent to *Sydney* because *Sydney* is subsumed by *Australia* in the taxonomy. A better definition is to replace by x every part of the query that subsumes x :

$$q \rightsquigarrow (q \setminus \{y \in q \mid x \leq y\}) \cup \{x\} = \text{Min}_{\leq}(q \cup \{x\}).$$

The extension of the new query is $\text{extension}(q) \cap \text{extension}(x)$. Therefore, in case the extension of x contains the extension of the query, the new query has the same extension as the query q . So, we restrict zoom-in to selections x such that

$$\text{extension}(q) \cap \text{extension}(x) \neq \text{extension}(q),$$

i.e., to selections whose extension contains some of the objects of the focus, but not all.

1.3.2 Zoom-Out

During navigation, the user may want to remove or generalize concepts in the query so as to reach larger extensions: this is the *zoom-out* navigation mode. For instance, Lisa realises she needs more photos of animals and plants. The back button can be used to retract the previous refinement. Hence if she wants to remove the first refinement *Australia*, she needs to move three steps backwards, and then re-select the last two refinements. She could also edit the query by hand, but users usually prefer to navigate rather than to edit queries [33].

Besides, selections whose extension contains all objects in the focus, i.e. $\text{extension}(q) \subseteq \text{extension}(x)$, cannot be used for zoom-in. This makes them available for zoom-out. When such a selection is part of the query, it is removed from the query:

$$q \rightsquigarrow q \setminus \{x\}.$$

For instance, if Lisa selects *Australia*, the new query is (Animal or Plant) and not Portrait (282 photos from many locations). When such a selection

subsumes some query parts, it replaces those parts in the query:

$$q \rightsquigarrow (q \setminus \{y \in q \mid y \leq x\}) \cup \{x\} = \text{Max}_{\leq}(q \cup \{x\}).$$

For instance, if she selects `Pacific`, the new query is `(Animal or Plant)` and not `Portrait and Pacific`. It is now clear how zoom-in and zoom-out can be reduced to a single primitive navigation mode. When zooming on a selection, the relationship between this selection and the current query determines whether this is a zoom-in or a zoom-out.

Compared to existing approaches, i.e., a list of removable concepts, our approach has three advantages: (1) it is integrated into the dynamic taxonomy, (2) it allows the replacement of a concept by a more general one, and (3) it extends to selections with disjunction and negation by extending the ordering \leq to such selections.

1.3.3 Shift

Zoom-in and zoom-out can be combined in two forms of *shift* navigation modes. From the previous query `Australia and (Animal or Plant) and not Portrait`, Lisa first chooses to zoom-in on the concept `Plant`, resulting in the query `Australia and not Portrait and Plant` (6 photos). This is her starting point for shift navigation.

At this point, Lisa sees that 1 photo has also the type `Landscape`, which interests her. She selects this concept (zoom-in) and, since the result has only 1 photo, she generalizes it by removing the concept `Plant` from the query (zoom-out). Therefore, she has executed a shift from Australian plants (6 photos) to Australian landscapes (80 photos), replacing in the query the concept `Plant` by the concept `Landscape`. From there, she performs a new shift from the concept `Landscape` to the concept `Building`, resulting in 28 photos of Australian buildings. These navigation steps are suggested and supported by photos belonging to two concepts, i.e., by extensional relations [69]. This illustrates the relevance of assigning several types to photos, which is common in this photo infobase. The same would apply to persons visible on photos, as a photo can contain several people.

However, the same does not apply to locations, as a photo cannot be taken in two incomparable locations (e.g., in Australia and in Europe). Nonetheless, it is still possible to shift between locations, through the taxonomy of locations. Suppose Lisa wants to find building photos from other locations. She first generalizes `Australia` by `Location` in the query (zoom-out), and then browses suggested locations before selecting `Spain` (zoom-in). Thus, she has performed a shift from Australian buildings to Spanish buildings, and find 48 photos (mainly churches taken in the north-west of Spain in 2003).

The former form of shift is a zoom-in/zoom-out combination, and can be qualified as *extensional* because it relies on extensional relations in the infobase. The

latter form of shift is a zoom-out/zoom-in combination, and can be qualified as *conceptual* because it relies on conceptual relations in the taxonomy.

1.3.4 Pivot

The user may not remember a concept she wants to use to refine the query, but she can find it through another query. For instance, suppose Lisa wants to retrieve the photos of the building of some town. She does not remember which town it is, but she remembers that the ICFCA conference took place there in 2004. Therefore, she can first reach the query `event contains "ICFCA" and date = 2004` by zoom-in navigation. The resulting extension shows photos of ICFCA'04, and the dynamic taxonomy shows relevant information about these photos, such as precise dates, locations, and so on. By browsing the dynamic taxonomy, she discovers that Sydney, in Australia, is the location of ICFCA'04. Then, she can make the query become `Sydney`, and refine it to the desired query `Sydney and Building` by zoom-in. The concept `Sydney` plays the role of a *pivot* between the two queries.

Pivot navigation relies on the ability of DTs to answer queries not only by a set of objects (the extension), but also by a set of concepts (the dynamic taxonomy). In previous navigation modes, these concepts were added or removed from the query, whereas here they are used as new queries. Given a query q and a selection x , the query transformation is defined by

$$q \rightsquigarrow x.$$

Therefore, pivot navigation is a way to restart a search from the results of a first search. This kind of navigation has already been applied in collaborative web-sites [59, 98].

There is an interesting analogy with natural language. Indeed, the query above can be rephrased as “photos of buildings in the town, *where* the ICFCA conference took place in 2004”. The idea of pivot is reflected by the fact that Sydney occurs in the main sentence as “town”, and in the relative sentence as the relative pronoun “where”. The relative pronoun indicates which facet to browse for a pivot: e.g., “where” indicates a location, “when” indicates a date, and “who” indicates a person. Iterated pivot navigation then corresponds to nested relative sentences, such as “photos of buildings in the town, *where* the ICFCA conference took place in the year, *when* I also visited Hinterzarten”. The first pivot to be applied is the year 2004, and the second pivot is the town Sydney.

1.3.5 *Slice and Dice*

Section 1.3.1 shows how the query (France or Italy) and Building can be reached by performing a zoom-in successively on France or Italy and Building, thus selecting French and Italian buildings. The disjunction is introduced because both concepts France and Italy were selected when the zoom mode was activated. Now suppose Lisa wants to extend the selection to landscapes, while retaining the current selection of locations. She just has to extend the selection Building to the selection Building or Landscape, and activate the zoom mode. Because the new selection is more general than the old one, the zoom is interpreted as a zoom-out. According to the definition of zoom-out (Section 1.3.2), the new query is (France or Italy) and (Building or Landscape). Now, Lisa wants to refine the selection to Italy only. To this end, Lisa unselects France in the selection of locations, and applies the zoom mode. Because the new selection is more specific than the old one, the zoom is here interpreted as a zoom-in. According to the definition of zoom-in (Section 1.3.1), the new query is Italy and (Building or Landscape). This short navigation scenario demonstrates that in a query each facet can be refined and extended independently from other facets, simply by applying the zoom mode on the new selections. In fact, it is not necessary that the different concepts in a selection belong to the same facet, while this is the most common case.

1.3.6 *Range Selection*

Range selection is similar to slice-and-dice (Section 1.3.5), except disjunctions of concepts are replaced by inequalities as selections. This makes sense because an inequality $\text{date} \geq 2002$ is equivalent to the infinite disjunction $\text{date} = 2002$ or $\text{date} = 2003$ or \dots . Then, every range is the composition of two inequalities. For instance, the date range $\text{date in } [2002, 2007]$ is equivalent to $\text{date} \geq 2002$ and $\text{date} \leq 2007$. Therefore every range can be reached by two successive zoom-in steps on inequalities: one for the lower bound, and the other for the upper bound. It is assumed that the user interface allows the selection of inequalities even if the dynamic taxonomies contains only values (e.g., $\text{date} = 2003$).

Starting from $\text{date} \geq 2002$ and $\text{date} \leq 2007$ and France, the date range can be refined by zooming on $\text{date} \geq 2003$ or $\text{date} \leq 2006$ (zoom-in), or extended by zooming on $\text{date} \geq 2000$ or $\text{date} \leq 2008$ (zoom-out). The upper bound of the range can also be removed altogether by zooming on $\text{date} \leq 2007$ (zoom-out); and similarly for the lower bound. The above formulas are used to give the logic of the navigation, and to reduce range selection to basic navigation modes; but a user interface may render it in a more graphical way, e.g., with the help of a double slider on a scale covering the relevant values.

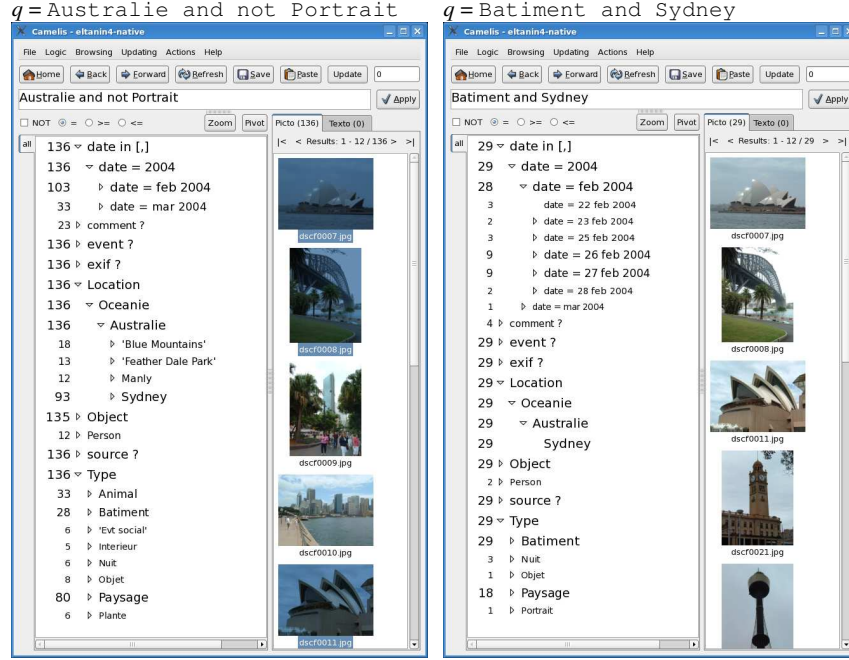


Fig. 1.3 A screenshot of CAMELIS before and after querying by examples.

1.3.7 Querying by Examples

A query can be determined by the selection of a subset of objects, thus supporting querying by examples. The idea is to construct the query as a conjunction of all most specific concepts which are shared by the selected objects O :

$$q \rightsquigarrow \text{Min}_{\leq} \{y \in T \mid O \subseteq \text{extension}(y)\}.$$

For instance, suppose Lisa starts with `Australie and not Portrait`. While browsing photos in the result, she sees interesting photos of buildings (e.g., 2 photos of the Opera, and 1 photo of the Harbour Bridge), and she would like to find more. By selecting them she moves to a new query that is the conjunction of the concepts shared by those 3 photos. As usual with this form of navigation, the resulting query is very specific and she receives no additional photos. At this stage, Lisa can use zoom-out navigation to generalize the query. Unlike approaches based on metrics, Lisa can choose which properties of the query should be generalized or removed [6]. By removing in the query concepts related to date and event, the query becomes `Sydney and Building`, and Lisa find 29 photos. Figure 1.3 shows the three selected photos in the initial query (left side), and the resulting view of the final query (right side). From there, she can further zoom-out, zoom-in to find

photos of modern buildings, or shift to find buildings in different countries. *Interactive query relaxation* [38] is similar, except that only one facet is retained in the generalized query. For instance, starting with the same photos, Lisa could reach the query *Sydney* or the query *Building*, but not *Sydney and Building*.

A special case of querying by examples is when selecting only one photo. Then there is only one object in the extent, because there are enough properties to uniquely characterize each photo, and the query contains all the object properties, which are more easily read in the dynamic taxonomy. So this is an easy way to access the full description of any object.

1.4 Design Patterns

This section gives an overview of solutions for solving the issues and challenges in the user interface design of applications for faceted browsing and dynamic taxonomies, with a special focus on how to enable the previously introduced navigation modes. Where applicable, these are referred to the respective user interface design patterns from established pattern libraries.

1.4.1 Selection Management

Filter selection and de-selection is of central importance in faceted search. The basic navigation modes of *zoom-in* and *zoom-out* are present in all examined user interfaces.

If only one concept should be selectable at a time within a facet (thus avoiding the possible confusion if multiple values are to be connected by AND or OR), traditional single-select controls such as *radio buttons*, *dropdown list controls* or *simple links* (in web applications) are advisable. The standard multi-select elements, on the other hand, are *check boxes*.

Interfaces that allow only one concept selection per facet support *shift* navigation in the easiest manner, since only one click is necessary to replace a selection with another one from the same facet. If multiple selections are allowed per facet (*slice-and-dice* navigation mode), a distinction has to be made between *zooming-in*, adding the clicked value to the active filters, and *shift*, i.e. replacing the previously selected concepts from the same facet.

For instance, the yelp⁴ web application provides check buttons for multi-select facets and simple links for facets with exclusive selection (see Figure 1.4). Alternatives for allowing both modes in a facet would be dedicated controls (e.g. a “jump to” button), or modifier keys (such as pressing “shift” while clicking).

⁴ <http://http://yelp.com>



Fig. 1.4 Mixing multi-select and single-select facets in the yelp (<http://yelp.com>) application.




Fig. 1.5 The ContentLandscape application (see 1.5.3) combines bar chart representations with slider controls for range selection.

For *range selection* navigation mode, *slider controls* can allow the specification of upper and lower bounds on the result set (see for Figures ?? and 1.5).

Additional UI functionality, however, is usually accompanied by additional complexity and visual clutter. Intelligently limiting users' options can help in allowing the user to focus on his core tasks without additional burden of rarely used functionality. For example, for a web shop application, it might be sufficient to split the "price facet" into 3–5 discrete regions from low- over mid-priced to expensive goods, instead of giving the more fine-grained option to filter from 37 to 82.

Either way, concept de-selection should be as easy as concept selection. Additionally, if breadcrumbs or a similar filter summary indicator are present, these should include the option to clear individual filters as well. Also, buttons for resetting single facets or all filter options can help to *zoom-out* quickly.

Pivoting is usually supported not directly in the facet panels, but from the detail views for single contents. First established in Web 2.0 applications [59], it has become a common practice that a metadata value clicked in the content presentation leads to a new view with the respective value as the only selected concept. The same holds for *Querying by examples*, as this action is intrinsically related to resource instances, and not to individual facet concepts. Consequently, querying by example is usually realized with context menus or buttons adjacent to the result list presentations or detail view.

If the data is only partially tagged, it is advisable to include a "no value assigned" concept, as for instance, demonstrated in the Exhibit prototype [46] (see Figure 1.6).

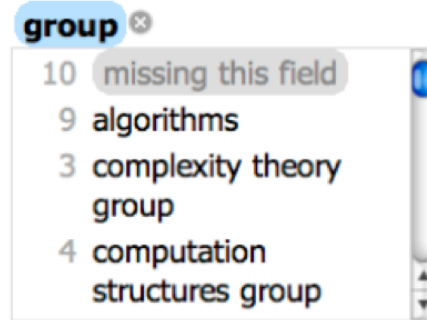


Fig. 1.6 The Exhibit (<http://simile.mit.edu/exhibit>) user interface signals missing concept assignments in a facet.

1.4.2 Revealing Hierarchy

For flat facets, i.e. not featuring a hierarchical relation between the concepts, simple list widgets are usually used. List sorting can either be alphabetical, or dynamically updated by the number of assigned items in the current result set. For navigating hierarchies, a number of different presentation and navigation options exist, which are discussed in the following.

1.4.2.1 Explorer Tree

The expandable explorer tree constitutes an established representation for hierarchical structures. This principle is, for example, used in the Camelis application (see Figure 1.2). Given the usually quite limited screen estate, however, the expanded lists often exceed the available facet widget space. This leads to the need for scrolling, which makes it difficult to orient in the hierarchical structure, especially if multiple levels are expanded.

1.4.2.2 Zoom and Replace

The Flamenco application⁵ [93] zooms into selected values, replacing the facet widget content with the level below the selected concept. The path from the root of the hierarchy to the current level is accessible via breadcrumbs in the header of each facet widget (see 1.7). This pattern works only for single-select facets, and optimizes for item presentation on one level. Consequently, navigation across the tree is facilitated.

⁵ Online demos available at <http://flamenco.berkeley.edu/>

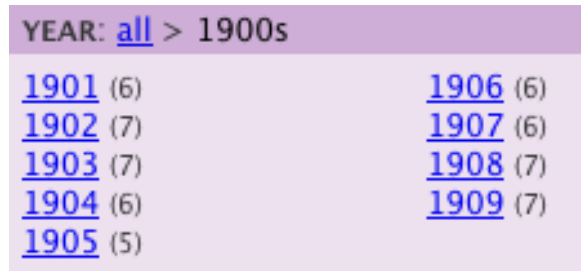


Fig. 1.7 Zoom-and-replace and breadcrumbs in the Flamenco application.

1.4.2.3 Collapsible Panels

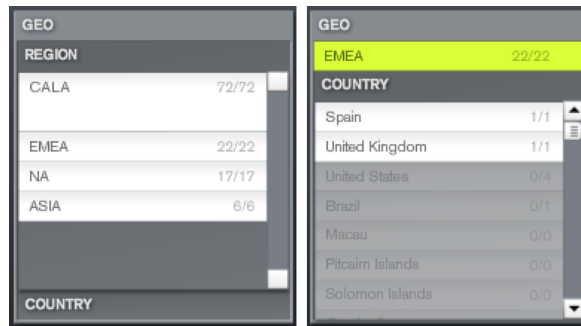


Fig. 1.8 The ContentLandscape application applies the collapsible panel pattern for zooming into concepts within a hierarchy.

The ContentLandscape application [79] features compact, hierarchical widgets based on the accordion pattern⁶, where each hierarchy level is represented as an individual accordion level. On concept selection, the respective level is collapsed, and the subsequent level opened, to allow further drill-down in the hierarchy. Moreover, opening a level is possible by simply clicking the respective accordion header (see Figure 1.8).

1.4.2.4 Continuous Zooming

In the FacetZoom prototype⁷ [23], hierarchical facets are displayed as space-filling widgets which allow a fast traversal across all levels while simultaneously main-

⁶ See e.g. <http://www.welie.com/patterns/showPattern.php?patternID=accordion>

⁷ Open source version available at <http://advancingusability.wordpress.com/2008/03/31/facetzoom-first-open-source-release/>



Fig. 1.9 The FacetZoom widget combines ideas from zoomable user interfaces (ZUIs) with faceted search

taining context. It supports both horizontal panning for exploring a whole hierarchy level, as well as *tap-and-center* navigation, allowing to dynamically zoom-in on tree nodes. For selected concepts, the child nodes are displayed on top of the widget. Navigating one level up the tree is supported by a bottom row presentation of the parent node.

1.4.3 Facet Management

A variety of options to overcome the problem that, often, more facets are available than can be put on screen at the same time, are discussed in [34] and [35]. The options range from *collapsible facet widgets* (such as used by, for instance, Getty images' faceted navigation interface⁸) over expandable filter areas ("More..." button) to dynamically selecting the shown facets based on the existing query (as demonstrated in the yelp application).

1.4.4 Keyword Search

As noted above, a free-form keyword field in order to search for arbitrary terms in addition to the defined classification scheme is a "key component to successful

⁸ <http://gettyimages.com>

faceted search interfaces” [34]. This task is especially challenging, since a search field can either act as a filter on the resources, e.g. searching over titles and descriptions, or if it can also match classification terms. The following sections discuss further variations within these two options.

1.4.4.1 Keyword Search as Additional Resource Filter

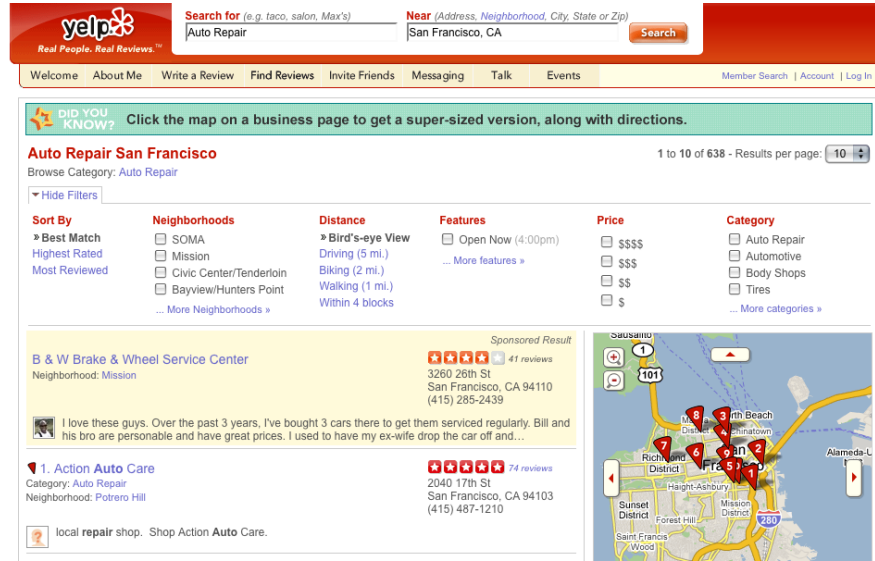


Fig. 1.10 The yelp application automatically selects the presented facets based on the search term.

When search engines are enhanced with faceted search, often, a keyword search is used to define the initial result set, which can be further refined by concept selections from facets. For instance, the yelp application (see Figure 1.10) asks for a topic (e.g. “auto repair”) and a location (e.g. “San Francisco”) to be entered, before entering the faceted search mode. In this application, displayed facets and concepts are selected dynamically depending on the type of query, i.e. a search for “auto repair” will yield different filtering options than one for “chinese restaurant”.

The Flamenco application [93] allows to choose between a full-text search over all results (overriding other filters) or within the current focus.

Obviously, when integrating keyword search with dynamic taxonomies, there might be zero hits, due to the unrestricted nature of the input. This violates the Poka-Yoke principle that we identified as one of the key features of applications in this domain. One solution could be to check for results after the user submits the keyword query, and leave the keyword filter in a “tentative” state if no results are found within the current focus. This would give the user the option to either zoom-

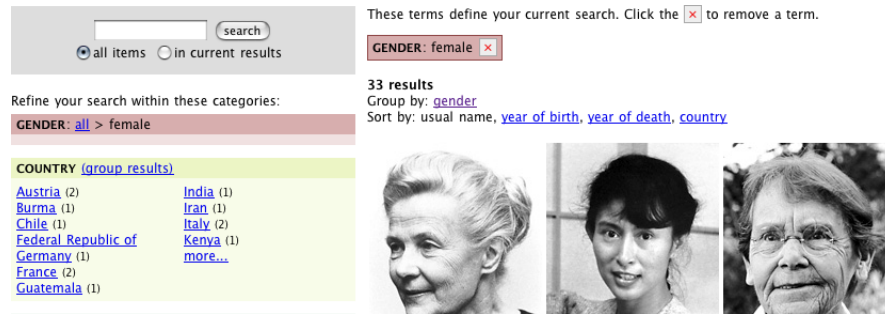


Fig. 1.11 The Flamenco application allows to choose between a full-text search over all results (overriding other filters) or within the current focus.

out some other filters, or re-tract the query. In this case, it would be helpful to have a preview, of how many results could be achieved, if the respective concept would be removed from the query.

1.4.4.2 Keyword search within facets

In order to avoid having to navigate large hierarchies, even though the target concept is already known by name, direct access to facet items can be achieved with a keyword search over the concept labels.

For instance, the /facet system [42], provides a keyword search box for each facet (see Figure 1.12). This interface dynamically suggests matching concept labels after the user has typed a few characters; only keywords that produce actual results are suggested. This makes the interaction often faster than manually navigating the tree. The results are presented in a collapsible tree structure. Additionally, if the target concept is known, but it is unclear, in which facet it is located, a global search box executes the described operation over all search boxes in parallel.

A similar approach is described in [8], which is even extended to finding facets by label, and can thus be applied to very large and heterogenous resource bases.



Fig. 1.12 The /facet system allows to quickly search within the concept labels of a facet

The ContentLandscape application [79] features a combo box component for quick access to concepts across all hierarchy levels (see Figure 1.13). It can be opened by clicking a search button in the facet box. In its initial state, the text field is empty, and a scrollable, alphabetical list presents all concept labels from that facet, regardless of depth level. When the user starts typing, this list is dynamically reduced to terms matching the query.

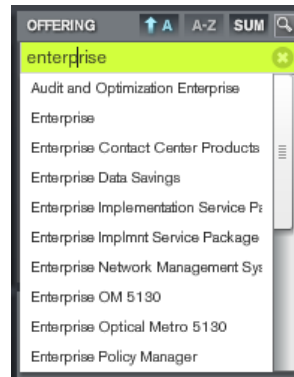


Fig. 1.13 Quick access to concepts with a combo box in the ContentLandscape application.

1.4.5 Filter Summary and History Navigation

Breadcrumbs can be used to summarize the current selection status in one central place in the user interface. These usually record the sequence of selection actions across all facets [34]. Breadcrumb entries should be clickable, leading to a zoom-out action on the respective concept. The footnote web site combines breadcrumbs with the option to refine with an additional keyword search (see Figure 1.14).




Fig. 1.14 The footnote web site combines a filter summary with the option to refine with an additional keyword search

1.4.6 Animated Transitions

Animated transitions can facilitate awareness of transformations and responses in user interface design [39, 88]. Perception of change is especially important for facet browsing, as the sudden disappearance of list items after click can be a source for misconceptions and confusion. In fact, studies have shown that so-called *change blindness* [65] is a common psychological phenomenon: changing details of visual scenes are often remain unnoticed, if the two states are separated by a short flash, as it is common, for example, in web applications. The Elastic Lists facet browser [78] demonstrates how smooth transitions can help in understanding filtering processes⁹.

1.4.7 Visualizing Proportions

As stated above, one common and useful technique is to exclude concepts with zero occurrences from the presented filter options, in order to avoid selections with zero results. For exploratory tasks, it can be useful to additionally see *how many* items match each of the respective concepts in the given focus.

Besides support for orienteering (see Section 1.1), analysing metadata distribution can constitute a valuable information source by itself, e.g. in order to understand what makes a data set special compared to the whole collection and to generate hypothesis about the underlying reasons.

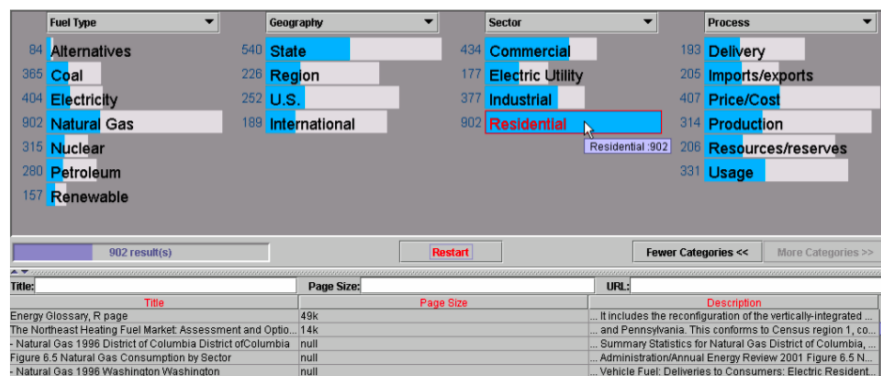


Fig. 1.15 The RAVE system visualizes metadata value proportions in horizontal bar charts.

The simplest option is to provide this information in brackets after the concept label (e.g. "Europe (5)"). While this presents an economic and easy solution, the user is left with the task of processing and understanding these numbers. Visualiza-

⁹ Interactive demo available at ...

tion can help to make the relevant information available pre-attentively and in an intuitive manner.

For instance, the RAVE system [96] visualizes metadata proportions in horizontal bar charts, below the concept label. While the graphic layout could be improved in order to introduce less visual clutter, the prototype shows how additional information about local and global weights can be integrated without loss of screen estate.

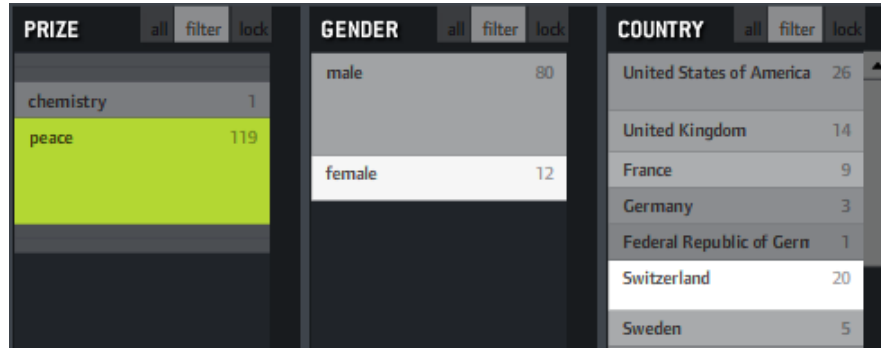


Fig. 1.16 Elastic lists indicate the number of matched resources in scaling list entry height. Additionally, unusually high proportions (compared to the global distribution) are indicated by brightness of the list entries.

In the Elastic Lists prototype [78], the height of a list item indicates the relative proportion of items associated with the respective metadata value in the given context. Additionally, a brighter color indicates that the current weight is significantly higher when compared to the global distribution. List entries with a weight of zero (i.e. not occurring in the current context) are collapsed to a minimal visible height.

The Visgets system [27] extends this principle by featuring a whole number of visualizations, with a weighted, coordinated brushing scheme (see Figure 1.17). The visualization elements include bar charts with range sliders, a map, and a tag cloud¹⁰. Visual representations for concepts and metadata values are scaled according to their global proportion. The coloring indicates, on the one hand, presence or absence of the respective value in the current result set. Additionally, on rollover on any concept or metadata value, more strongly associated items receive a higher opacity (*weighted brushing*).

1.5 Extensions and Related Approaches

Lately, the described principle of faceted search and dynamic taxonomies are being extended and translated to other types of search and browsing applications. This

¹⁰ see <http://www.welie.com/patterns/showPattern.php?patternID=tag-cloud>

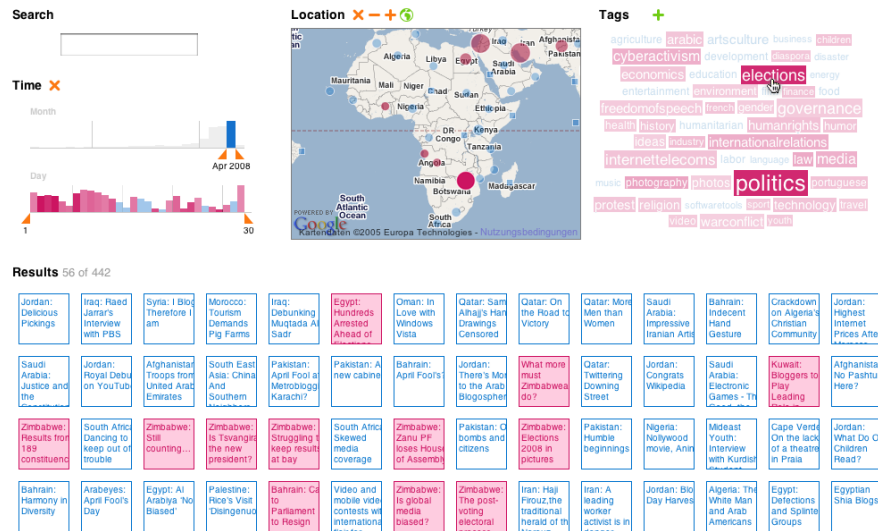


Fig. 1.17 Weighted, coordinated brushing in the visgets system.

section provides some brief pointers to current research in this area, and introduces the novel principle of out-of-turn interaction.

1.5.1 FaThumb

FaThumb [50] enables faceted search on mobile devices (see Figure 1.18). The filter area is grouped in nine zones, corresponding to the nine digit keys on mobile phones. The middle zone serves as a spatial overview during navigation. The surrounding eight zones allow the user to select hierarchy branches and repeatedly zoom in on subtrees. The left short shortcut key adds the currently selected concept to the query, the right one allows to quickly jump back to the top.

1.5.2 Browsing Related Entities

Usually, the type of resource entity to be browsed (e.g. book, car, web page...) remains fixed in a faceted browsing application. In [51], a conceptual prototype of a browsing application named *Humboldt* is described, which allows to switch the type of displayed entities based on relations to the current result set. In principle, the application allows to treat any facet value space as a search result list, and arrange the interface accordingly. To cite an example given in [51], "a user who filters films on

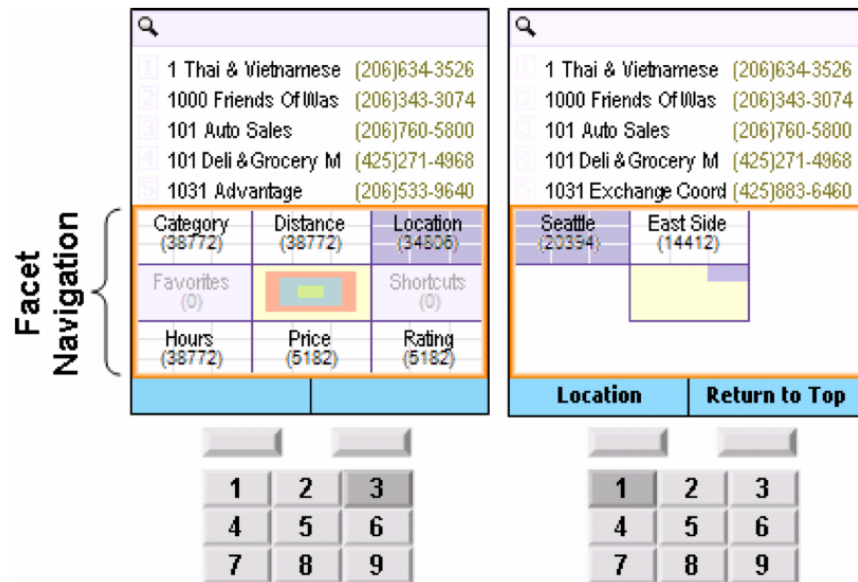


Fig. 1.18 Faceted search for small screens in the FaThumb prototype.

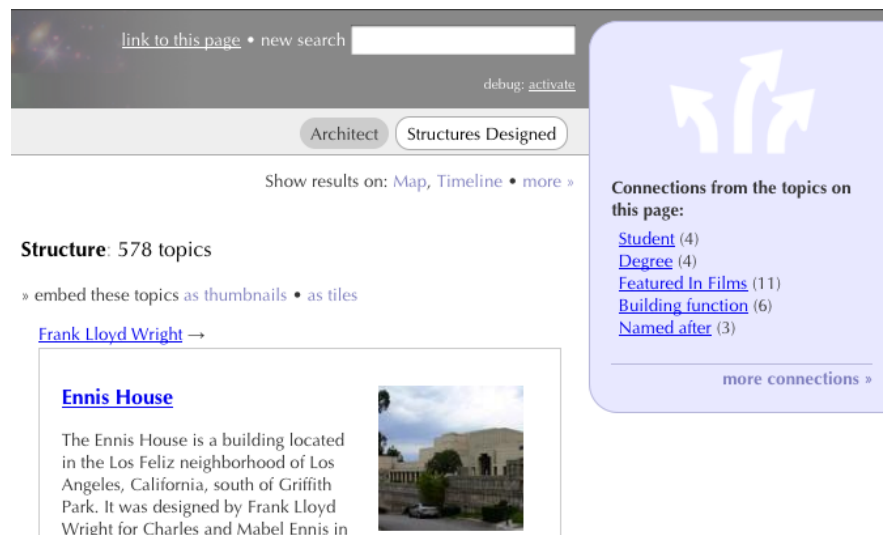


Fig. 1.19 The parallax application allows to jump to related sets of items from a faceted browsing situation.

certain directors and then pivots¹¹ on actors will see all actors in the result list, who are related to any film in the previous result list.”

This principle has also been demonstrated in the parallax application based on the freebase¹² public database [45] (see Figure 1.19). It allows, for example, to query the data set for architects, then filter down to all modern architects: a classical zoom-in. The novel principle, however, is that the user can explore related collections, like the buildings they designed, their birth places etc. in the same facet browsing space. The jump to these new results is offered in a “connections” box on the top right of the interface. History navigation for these “related set” browsing steps is provided by a breadcrumb control.

1.5.3 Resource Analytics

Understanding resource production, use and distribution across departments, regions, and product groups is one of the core challenges of knowledge management in the enterprise [68]. ‘What are the most downloaded contents?’, ‘do the presentation materials for a given product cover all important sales regions?’, ‘what parts of my resource collection are growing? and which are declining?’ are typical questions in this area.

The ContentLandscape application¹³ [79] is part of the BizSphere¹⁴ application suite and uses faceted browsing and search in order to facilitate the understanding of resource distributions. In addition to traditional result set views, a dashboard view presents statistical measures for the resource set in the current selection (see Figure 1.20). It features visualizations of trend measures such as the quarter to quarter growth, a detailed age histogram, and the rating distribution. Moreover, the coverage of the selected resource set with respect to the three main taxonomies ‘region’, ‘offering’ and ‘resource type’ is presented in squarified treemaps [16, 76]. At first glance, this visual display allows to see, for instance, if all product groups are represented by resources in the current selection, and in which specificity. This statistical analysis can be further decomposed into several matrix views. Inspired by OLAP approaches [20], these allow the user to split the result set according to up to three dimensions, and compare the statistical measures for the resulting sub-collections in parallel.

¹¹ Note that the semantics of *pivoting* in this case differs from the definition introduced in section 1.3.4.

¹² <http://freebase.com>

¹³ <http://moritz.stefaner.eu/projects/content-landscape>

¹⁴ <http://bizsphere.com>

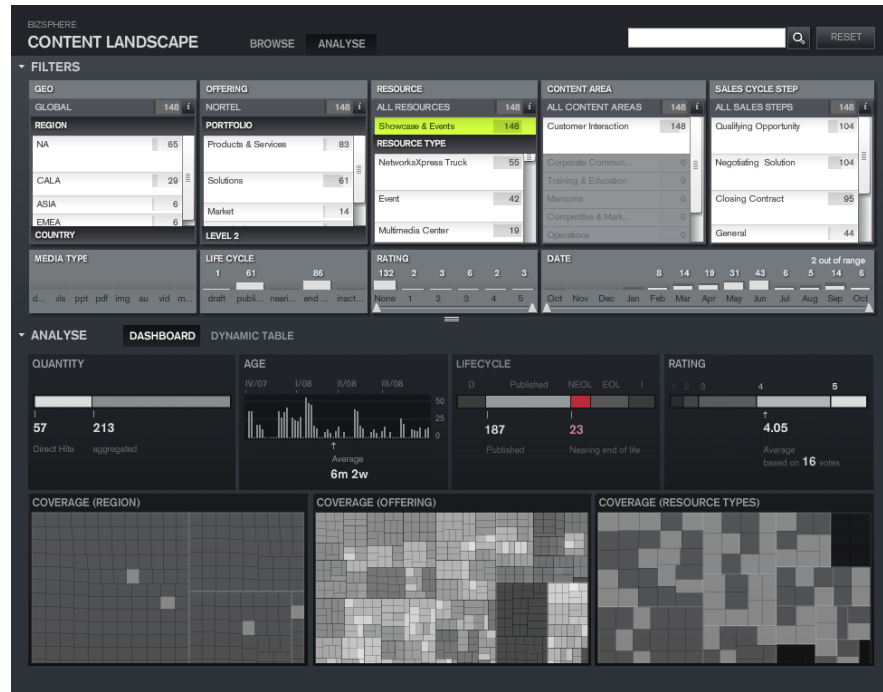


Fig. 1.20 The dashboard view of the ContentLandscape application

1.5.4 Out-of-turn Interaction

Out-of-turn interaction [61] is a technique for navigating hierarchical websites which augments traditional browsing by empowering the user to supply a hyperlink label which is presented beyond the current webpage (hence out-of-turn) to initiate a search over the site's hierarchical schema. When the system receives an out-of-turn input, it removes all paths through the site which do not contain a hyperlink labeled with the input and removes the hyperlink labeled with the input from the remaining paths. Fig. ?? illustrates how a sample hierarchy with a structure similar to that of PVS would be pruned based on supplying 'Republican' out-of-turn. Notice that all paths leading to the webpages of Democratic politicians (nodes 20, 21, 24, 25, 26, 27, 30, 31, 32, and 33) have been removed. In addition, the hyperlinks labeled 'Republican' in the remaining paths (those leading to nodes 22, 23, 28, 29, 34, 35, 36, and 37) have been removed.

Fig. 1.21 illustrates an out-of-turn interaction through a browser toolbar we call *Extempore* (as it permits the user to supply terms *extemporaneously*). Here the user supplies 'Republican' out-of-turn (see Fig. 1.21, top). This causes some of the hyperlinks presented on the root page (e.g., Hawaii), those which do not lead to the webpages of Republican congresspeople, to be pruned out (see Fig. 1.21, bottom). When used in conjunction with traditional browsing, the unsolicited reporting [5] in-

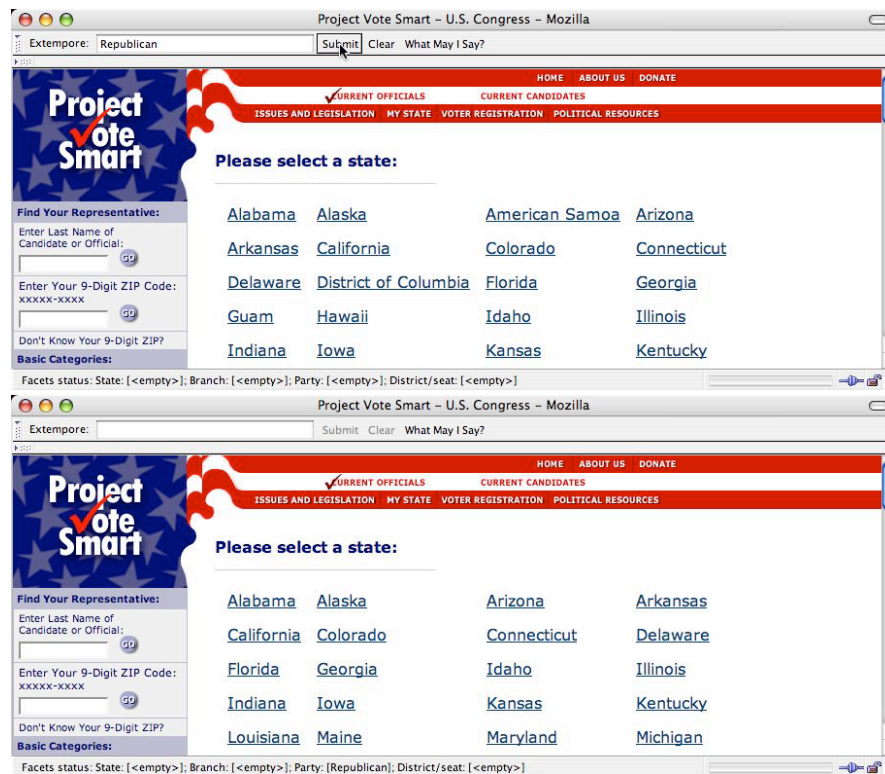


Fig. 1.21 Illustration of an out-of-turn interaction with PVS through the *Extempore* toolbar.

volved in supplying an out-of-turn input supports a simple form of mixed-initiative interaction [63] and can be viewed as an approach to integrating querying and browsing in information hierarchies [14].

In sites where each level of the hierarchy corresponds to a facet of information assessment, such as PVS, out-of-turn interaction permits the user to explore the facets in any order without the designer enumerating all possible paths of navigation. In hierarchies where each level does not correspond to a facet, such as Yahoo! and the Open Directory Project (ODP) at dmoz.org, out-of-turn interaction behaves more as a pruning operator and reveals to the user the portions of the taxonomy pertaining to their query. For example, ODP contains 16 top-level categories and a user starting from the homepage would be hard-pressed to know that only four (Home, Shopping, Business, and Regional) contain links to information about 'ice cream makers.' An out-of-turn interaction reveals these categories. In fact, the search feature provided in ODP is similar to out-of-turn interaction with the exception that ODP flattens the hierarchical structure in response to a query (see Fig. 1.22) whereas out-of-turn interaction preserves the hierarchical nature in order to retain context.

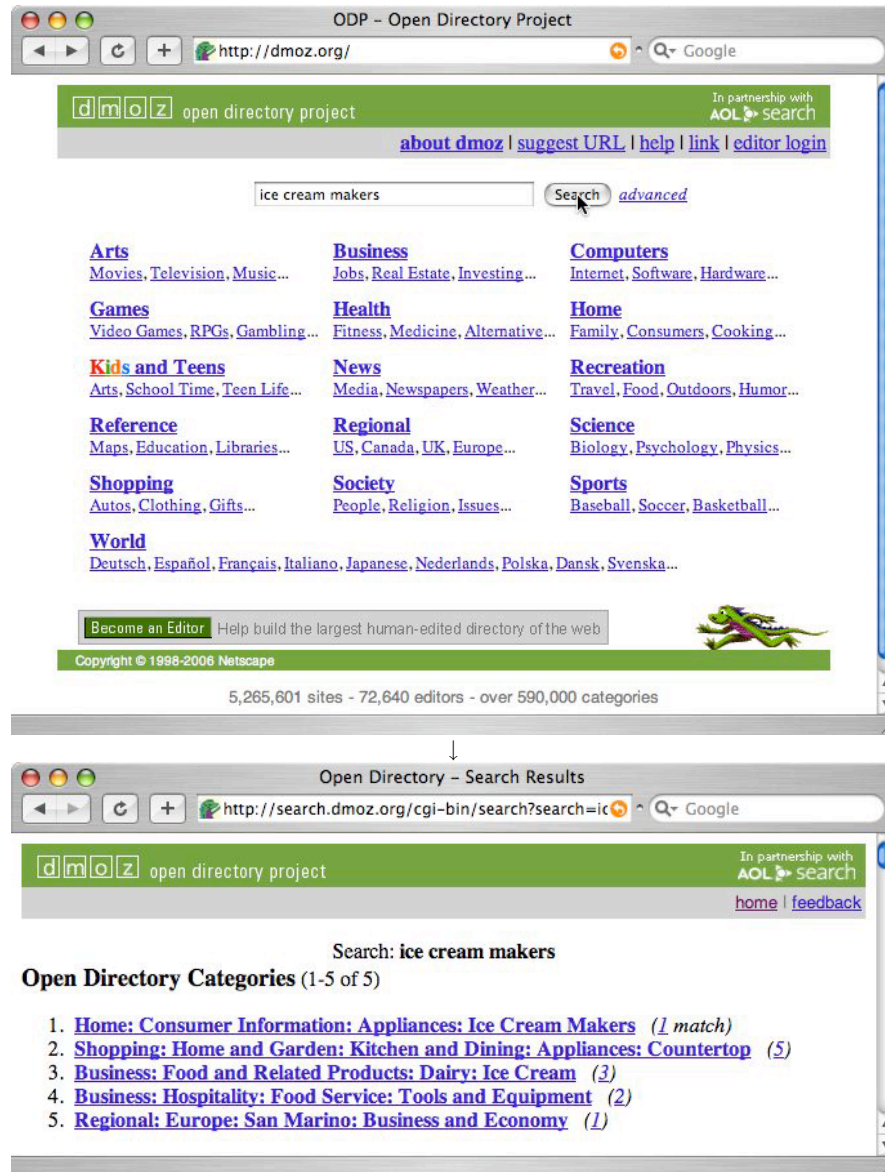


Fig. 1.22 (top) A query for 'ice cream makers' in the Open Directory Project and (bottom) its result as a flat list.

Notice that with the interpretation of out-of-turn interaction presented here, an unexpanded query will yield the same result as its expanded version and therefore query expansion here is simply a feedback mechanism to expose dependencies, unlike its use in traditional IR.

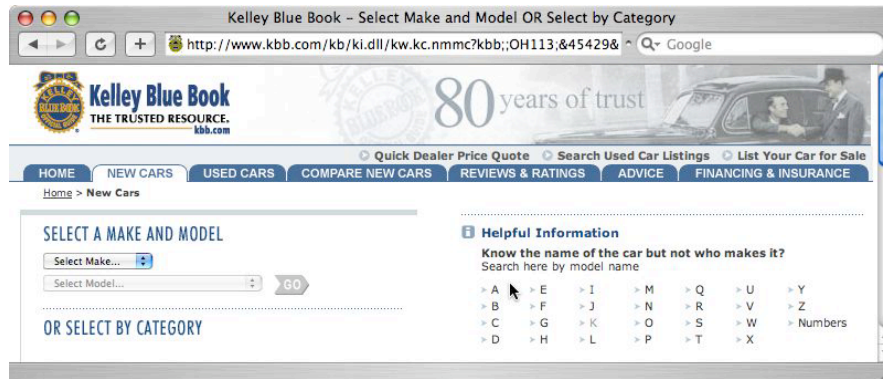


Fig. 1.23 Facility for automobile-make lookup by model in the online Kelley Blue Book.

There are other means of exposing dependencies underlying information hierarchies during information-seeking. For example, the Kelley Blue Book (KBB) online at kbb.com provides a facility for automobile-make lookup by model (see under heading titled ‘Helpful Information’ in Fig. 1.23) since FDs of the form ‘*model* \rightarrow *make*’ are implicit in the domain of automobiles. When browsing new cars in KBB, users are first asked to make a selection for automobile make (see under heading titled ‘SELECT A MAKE AND MODEL’ in Fig. 1.23). The lookup facility allows the user to search for the make of an automobile based on the model so that they can proceed with the information-gathering dialog on the left side of the window in Fig. 1.23. A more sophisticated example of support for dependency exploration is *Sony’s Advisor* facility available through sonystyle.com when browsing products such as digital cameras and camcorders.

1.6 Personalizing Faceted Search (Koren, Zhang)

1.6.1 Introduction

One of the primary ways users manipulate a faceted search interface is by refining their current query by clicking a facet-value pair from a list of possible system suggestions. How effective users are at finding their documents of interest is related to the quality of the query refinements suggestions. Traditionally, ad hoc approaches have been used to determine which values for a facet should be presented to the user

during interaction. One common approach is to simply display all available values for a facet. While this may be effective when the number of available values is small, this approach may overwhelm users when the number too large [77]. Another approach is to display only the first few alphabetically ordered values [36]. While this approach avoids overwhelming the user with many values, it arbitrarily biases the interface towards values earlier in the alphabet. A better method is to display the most frequent values for a facet. However this method is not user centric since the most frequent values are endemic to the corpus instead of the users.

This chapter focuses on a user centric approach to determine which values are most useful to users: Personalizing Faceted Search. Personalization allows the system to present the facet-value pairs that can help the user quickly find the document(s) that the current user is most interested. In order to determine which facet-values are most useful to a particular user, we analyze the distribution of values in corpus, and user's feedback on documents while using the system. With this knowledge, we can tailor the faceted search interfaces to individual users.

1.6.2 Related Work: Personalized Search and Filtering

The idea of personalizing search is not a new idea in the information retrieval community [43, 21, 12, 47, 29, 17, 74, 4, 85, 90, 81]. Dumais, Cutrell, Sarin and Horvitz automatically generate queries based on keywords within an email being read or composed by a user [28]. To improve retrieval results, Bharat treats the previous information requests from the user as the context of the current query [12], whereas Shen, Tan & Zhai use the preceding queries and clicked-document summaries as the context of the current query [74]. On the other hand, researchers have developed personal information integration environments that store a particular individual's heterogeneous information and the context of the information, providing content and context-based retrieval [2, 26, 31]. Rui et al. [57] explored biasing cosine similarity methods based on user feedback in order to retrieve more documents that were similar to a user's interests. Abrol et al. personalized semi-structured search interfaces by using a user's transactional feedback from his/her queries [1]. Shen et al. used implicit user feedback, such as query refinement and click logs, to customize a KL-divergence model for document retrieval [75]. Personal WebWatcher passively observed a user's browsing behavior in order to highlight links that matched the inferred task [60].

On the other hand, personalization is a heavily studied problem in the information filtering research community and the research can be traced back to 1970s. For example, **content-based adaptive filtering** studies the scenario in which a recommendation system monitors a document stream and pushes documents that match a user profile to the corresponding user. The user may read the delivered documents and provide explicit relevance feedback, which the filtering system then uses to update the user's profile using relevance feedback retrieval models or machine learning algorithms (e.g. Boolean models, vector space models, traditional proba-

bilistic models [67], inference networks [18], language models [22], Support Vector Machines, K-nearest neighbors clustering, neural networks, logistic regression, or Winnow [56, 92]). **Collaborative filtering** goes beyond merely using document content to recommend items to a user by leveraging information from other users with similar tastes and preferences. Memory-based heuristics and model-based approaches have been used [52, 25, 15, 48, 44, 40, 10, 3].

1.6.3 Personalization based on Collaborative Filtering

Faceted search interfaces share three characteristics. The interfaces present a number of facets along with a selection of their associated values, any previous search results, and the current query¹⁵. By choosing from the suggested values of these facets, a user can interactively refine the query. The interface also provides a mechanism to remove previously chosen facets, thus widening the current search space.

In personalized faceted search, the key problem is to rank facet-value pairs according to how helpful they are for a particular user to use for query refinement. Complicating this task is the fact that rarely, if ever, does the system have access to user ratings of on individual facet-value pairs. Instead, most of the existing faceted search systems, such as Amazon.com and Netflix.com, are designed so that users rate each individual document. This design trade-off is reasonable since the relevance/rating of a facet-value pair is not a well defined problem and thus hard for the user to provide. Besides, a user usually has seen a overwhelming number of individual facet-value pairs, and users rarely experience any particular facet-value pair in isolation. By rating the whole document, a user express a preference over many facet-value pairs simultaneously, especially for facet-value pairs that they may only be tangentially aware of.

Personalized faceted search is a comparatively new field that has not been well studied. Fortunately, we can develop personalized faceted search interface based on the earlier work in personalized search. For example, we can first use traditional collaborative filtering techniques to predict a user's ratings on unseen documents, and then propagate information from ratings on whole documents to individual facet-value pairs.

The basic assumption of collaborative filtering is that users that have similar preferences on some documents may also have similar preferences on other documents. Therefore the algorithm provides recommendations for a user based on the opinions of other like-minded users.

In collaborative filtering, users ratings over items are represented as a matrix A , where $A_{u,i}$ is user u 's rating on item i . Many collaborative filtering techniques have been proposed to predict the missing cells in the matrix [52, 25, 15, 48, 44, 40, 10, 3].

In this chapter, we first introduce two basic collaborative filtering techniques, then discuss how to go from document ratings to facet-value pair ranking.

¹⁵ For an interactive faceted search interface, the current query is the facet-value pairs the user has selected so far.

1.6.4 Item-item similarity or user-user similarity

There are two very commonly used collaborative filtering approaches: the first one compares each user to the other users, while the second one compares the items to each other. These are called *user-user similarity* and *item-item similarity* respectively. We describe Konstan et al.'s Pearson's correlation based user-user algorithm [52] below.

For each user, we calculate the average rating assigned by that user \bar{A}_u to all rated items. Each unknown rating is then estimated as the user's average rating, perturbed by sum of the difference between every other user's assigned rating and his/her average rating, weighted by the correlation among the commonly rated items of the current user to every other user.

Formally, this is stated:

$$A_{u,i} = \bar{A}_u + \sum_{v=1}^k \frac{w_{u,v}(A_{v,i} - \bar{A}_v)}{|w_{u,v}|} \quad (1.1)$$

where k is the number of users that have at least one rated item in common with user u , and $w_{u,v}$ is the Pearson's correlation between user u 's ratings and user v 's ratings. Recall that Pearson's correlation is:

$$\begin{aligned} w_{u,v} &= \frac{\sum_{j=1}^m ((A_{v,j} - \bar{A}_v)(A_{u,j} - \bar{A}_u))}{\sigma_v \sigma_u} \\ &= \frac{\sum_{j=1}^m ((A_{v,j} - \bar{A}_v)(A_{u,j} - \bar{A}_u))}{\sqrt{(\sum_{j=1}^m (A_{v,j} - \bar{A}_v)^2)(\sum_{j=1}^m (A_{u,j} - \bar{A}_u)^2)}} \end{aligned} \quad (1.2)$$

where σ_u and σ_v are the standard deviations in user u 's and user v 's ratings, and m is the number items that user u and user v have both rated.

Herlocker et al. [41] examined using other similarity methods such as Spearman's correlation, information entropy, mean-squared difference, and found they performed similar to Pearson's correlation.

In its simplest form, the item-item algorithm is similar to the user-user algorithm, only with the rows and columns exchanged. Item-item similarity can be extended to take into account the content of the items being rated. For example, Sawar et al. [71] estimated ratings by summing the ratings of the other rated items, weighted by the cosine similarity of the rated and unrated plain text documents.

1.6.5 Singular Value Decomposition

One problem with the collaborative filtering techniques discussed in section 1.6.4 is that the user-item matrix can become very sparse as the number of items and users increase. This sparseness can sometimes lead to poor predictions.

By creating a low dimensional approximation of the rating matrix, it is possible to improve the accuracy of the predicted ratings. One such technique that has been used successfully in the personalization and collaborative filtering domains is *singular value decomposition* (SVD) [73, 72, 30, 62, 11].

SVD, also known as *latent semantic indexing* (LSI) in the information retrieval community [24], works by combining rows and columns that are found to be strong correspondence. These correspondences are called *latent factors*.

Applying SVD to collaborative filtering task, we factor the $m \times n$ user-item matrix A into three smaller matrices U , Σ , and V . U is $m \times h$, Σ is a $h \times k$ diagonal matrix, and V is $n \times h$ matrix.

$$A \simeq U \Sigma V^T \quad (1.3)$$

The values along the main diagonal of Σ are the biggest h *singular values* of A in decreasing order. Each row of U and V contain orthogonal *singular vectors*. The vectors in U are known as the *left singular vectors*, while the vectors in V are the *right singular vectors*. Since we want a low dimensional approximation of the rating matrix, only the first $h < \text{rank}(A)$ singular values and singular vectors are used.

Typically, A is approximated as the product of two matrices:

$$A \approx (U(\sqrt{\Sigma})^T)((\sqrt{\Sigma})^T V) \quad (1.4)$$

We can view the first matrix as the hidden representation of the users, and the second matrix as the hidden representation of the items. With this approximation, the predicted rating for a user u on item i can be calculated as $A_{u,i} = \bar{A}_u + (U(\sqrt{\Sigma})^T)_u^T ((\sqrt{\Sigma})^T V)_i$.

1.6.5.1 Recommending facet-value pairs

There are many methods to propagate information from ratings on whole documents to individual facet-value pairs. One method is to assign facet-value pair x_i a score based on the expected rating a user u gives to a document d containing that facet-value pair:

$$\begin{aligned} f(u, x_i) &= E[R_u(d) | x_i \in d \in D] \\ &= \sum_{d \in D} R_u(d) P(x_i | d) \\ &\simeq \frac{\sum_{d \in D} R_u(d) I_{\{x_i \in d\}}}{\sum_{d \in D} I_{\{x_i \in d\}}} \end{aligned} \quad (1.5)$$

Where D is the set of documents selected by the current query, $R_u(d)$ is the rating of user u on document d , and I is the indicator function.

One simple way to suggest facet-value pairs for query refinement is presenting the top scoring facet-value pairs that are contained in the documents selected by the current query. While this is an attractive option, this approach can suffer when the top scoring values are redundant, where an extreme case is that if they co-occur in

the same documents. One elegant solution proposed by Chen and Karger [19] is to condition each suggestion on the assumption that none of the previous suggestions are relevant to the current query. For our purposes, this means that the $k + 1$ th suggestion is the top scoring value that is not contained in a document that contains any of the previous k suggested values.

To determine the order that the facets are presented to each user, a simple approach is using the average score of the suggested values for each facet for the null query and then fix order of the facets throughout the lifetime of the user interaction session(s).

1.6.6 Personalization using Content Based Filtering

In this section we provide a complementary approach to personalization based on the content of the documents. Motivated by relevance feedback retrieval models and content based adaptive filtering techniques, we focus on two statistical models: a model of the documents being searched, and a model of a user.

1.6.6.1 Document Model

While every faceted document has a set of facets associated with it, the number of values that each facet has in a particular document can vary a lot. We model this by expressing the number of values each facet has in a random document as a draw from an multivariate Normal distribution:

$$\langle n_1, \dots, n_K \rangle \sim MVN(\vec{\mu}, \Sigma) \quad (1.6)$$

where n_k contains the number of values for facet k in the document, $\vec{\mu}$ is a K dimensional vector containing the mean number values for each facet, and Σ is the corresponding $K \times K$ covariance matrix.

Each facet in a document has a certain semantic meaning that dictates the type and thus the probability distribution of the values that can be associated with it. Five common types of facets are: *nominal*, *ordinal*, *interval*, *ratio*, and *free-text*. Nominal facets take discrete and orderless values. The values to this type of facet can be modeled as draws from a multivariate Bernoulli distribution. Ordinal facets also take discrete values, but there is an implicit ordering to these values. An example of this would be field that identified the sensitivity of a document as being for “full release”, “limited release”, or “secret”. Interval facets can take any value on a defined range, as long as the range excludes an explicit zero point. Ratio typed facets on the other hand can take zero as a value. Values for ordinal, interval, and ratio facets can be modeled as draws from a normal distribution. Free-text facets allow arbitrary text to be associated with documents. Traditional statistical information retrieval

techniques represent each word in unstructured text as a draw from a multinomial distribution.

After identifying the type and the proper probability distribution over the values associated with facet, the probability of a document existing is simply the product of the probability of each desired value occurring for the appropriate facet.

Table 1.1 Facet Types and Distributions

Facet Type	Values	Example	Distribution	Prior
Nominal	Unique Tokens	Director	Multivariate Bernoulli	Multivariate Gamma
Ordinal	Repeatable Ordered Tokens	Critic's Rating (e.g. A, B, C, ...)	Normal	Normal
Interval	Repeatable Numbers Excluding Zero	Year of Release	Normal	Normal
Ratio	Repeatable Numbers Including Zero	Running Time	Normal	Normal
Free-Text	Repeatable Tokens	Synopsis	Multinomial	Dirichlet

1.6.6.2 User Model

Instead of estimating a user rating for a document, in this approach we estimate the probability of a document being relevant to a particular user. Similarly, instead of calculating a suitability score for a facet-value pair to a user, we estimate the probability of a particular facet-value pair appearing in a document relevant to a user. For simplicity and without loss of generality, we assume that a document is either *relevant* or *nonrelevant* to a user. From this we can estimate the probability that any document will be relevant to a particular user, and the probability that a particular facet-value pair x_k will be contained in a relevant or a nonrelevant document. This tuple is the user relevance model and is represented as:

$$\theta_u = \{P(\text{rel} \mid u), P(x_k \mid \text{rel}, u), P(x_k \mid \text{non}, u)\} \quad (1.7)$$

where $k = 1, \dots, K$.

These individual probabilities can be estimated from training data. For example, assume for a particular user u , there exists a set of training data $D_u = \langle D_{u,\text{rel}}, D_{u,\text{non}} \rangle$, where $D_{u,\text{rel}}$ is the set of documents marked by u as being relevant, and $D_{u,\text{non}}$ are the set of documents marked as nonrelevant. If the facet type is free text, the maximum likelihood estimation of θ_u is:

$$P(\text{rel} \mid u) = \frac{|D_{u,\text{rel}}|}{|D_u|} \quad (1.8)$$

$$P(x_k \mid \text{rel}, u) = \frac{1}{|D_{u,\text{rel}}|} \sum_{d \in D_{u,\text{rel}}} I_{x_k \in d} \quad (1.9)$$

$$P(x_k \mid \text{non}, u) = \frac{1}{|D_{u,\text{non}}|} \sum_{d \in D_{u,\text{non}}} I_{x_k \in d} \quad (1.10)$$

This setup is very similar to the commonly used relevance language model in information retrieval [95]. We leave it as an exercise for the readers to derive the estimation for other facet types.

As stated in section 1.6.1, it may take a while before enough user specific feedback information is gathered from a particular user, thus user could suffer from the so called “cold start” problem. To handle this problem and get a level of acceptable performance from the very beginning, a hierarchical Bayesian model is used and found success in user modeling experiments [94, 99, 97]. In this model, each individual user model is considered as a draw from a prior distribution that is common to all users. By using common prior, gaps in a particular user’s model can be filled in by using information from the community of users. To applying the hierarchical Bayesian modeling approach for personalized faceted search, each distribution for each facet needs a separate prior that is estimated from the training data of all users in the system. Table 1.1 suggests different priors for different facet types.

Based on document models and user models described above, one can generate faceted search interface in various ways. For example, we can rank the facet-value pair based on Equation .

1.6.7 An Ontological Approach

An alternative approach to personalize faceted search is using ontology created manually or automatically. Tvarožek and Bieliková [86] use the distance between values in a hierarchical ontology to measure similarity, and thus relevance to users. In cases where an explicit ontology does not exist, one can be automatically constructed [80, 13, 91, 84].

The technique works as follows: Let $L_u(x)$ be the relevance of a facet x to user $u \in \mathcal{U}$ as computed by the ontological similarity of the facet and the user model [7]. In this case, the ontological similarity is the reciprocal of the maximum number of links needed to transverse from each value being compared, to a value common to both.

Each user model is then compared to the other users in the system, in order to calculate the *cross relevance* of a facet x to a user. This value, $C_u(x)$, is the average of the relevance of the facet to each user, weighted by the similarity of the each user v to the current user u .

$$C_u(x) = \frac{\sum_{v \in \mathcal{U}} \text{similarity}(u, v) L_v(x)}{|\mathcal{U}|} \quad (1.11)$$

The *global relevance* of a facet, $G(x)$, is calculated as the average relevance of the facet to every user. The *static relevance* of facet to a user is a weighted combination of the cross relevance and the global relevance.

In order to make recommendation to a user regarding a specific query, the temporary in-session relevance of each facet x to the current user u is introduced. This value, $T_u(x)$, is simply the fraction of query refinements (i.e. clicks) that utilize a facet x of all refinements in the current search session. This number is combined with the static relevance of the facet to determine the current *dynamic relevance* of a facet to the user. Values from facets with the highest dynamic relevance are then suggested to the user for query refinement. Evaluated with the Factic system, the ontology based approach reduce the number actions required for users to find their documents of interest when compared to an un-personalized baseline [87] .

1.6.8 Evaluation Regime

Considering various personalized faceted search techniques, which one works better on a particular task? To compare personalization methods, an evaluation metric is needed. Traditionally user studies have been used to determine satisfaction with different user interfaces. While undeniably useful, user studies have some drawbacks. First, they are expensive to hold. A number of users must be gathered and then tested on the proposed system. This takes a nontrivial amount of time for even user studies with a moderate number of subjects. User studies also problematic when being used to evaluate personalized systems, as the test subjects may not interact with the system long enough for a sufficient user profile to be learned. This can lead to inconclusive, or possibly even contradictory results.

Koren et al. [53] proposed a complementary inexpensive evaluation metric based on calculating the expected utility to a user of a faceted search interface, through the use of simulated user interactions. This method allows designers to quickly compare various algorithms and determine which algorithms are the most promising. By using this method, or a similar one, designers can conduct fewer user studies by only submitting the top performing algorithms for an in depth user study. A similar approach has seen success when evaluating spoken dialog systems [55, 32].

The evaluation system works as follows: Assume that the goal of the search interface is to enable users to find their documents of interest with the least amount of effort. In order to measure this effort, the actions that a user can perform when interacting with the system are identified and the system is rewarded or penalized depending on what action is performed. A series of user interactions are simulated using a combination of real-user feedback and heuristics. The interface is then scored according to the expected total reward for an interaction session.

Given S user interaction sessions, the empirical utility of the interface can be estimated easily:

$$U = \sum_s^S \sum_t^{T_s} R(q_{s,t+1}, a_{s,t}, q_{s,t}) \quad (1.12)$$

Let us now define the reward function. As stated earlier, we assume that the goal of the interface is to allow users with the least amount of effort on their part. In order to measure this effort, the designer must identify which actions a user can perform at each step of the interaction. Koren et al. identified eight actions common to many personalized faceted search interfaces, along with the rewards the system receives for each. These actions are listed in table 1.2.

Table 1.2 User Actions and Rewards

Action	Reward
Select Facet-Value Pair	negative
De-select User Selected Facet-Value Pair	zero
De-select System Selected Facet-Value Pair	negative
View More Facet-Value Pairs	negative
Mark Document as Relevant	positive
Mark Document as Non-relevant	negative
View More Documents	negative
End Session	zero

Calculating $P(a | q_t, u, D)$ is much more problematic, since its functional form is unknown. Instead of employing real users through a user study, actions are simulated based on certain assumptions about how real users interact with faceted search systems. Without loss of generality, it is assumed that each simulated user is searching for exactly one *target document* and that the simulated user can recognize the document and the facet-value pairs that are indicative of that document. At each step of the search session, the simulated user scans the top ranked documents that match current query looking for the target document. If it is found, then it is selected and the session ends. If the target document is not found, the simulated user removes any facet-value pairs that are contained in the current query that do not match the target document. Once the query is cleaned of incorrect terms, the simulated user scans the list of presented facet-value pairs. A pair is selected by some method for inclusion in the query from the set of pairs that match the target document. If none of the suggested facet-value pairs match the target document, then a facet is chosen at random and all of its values are examined until a matching facet-value pair is found for inclusion. If no matching value can be found for any facet, then the simulated user scans through the complete list of returned documents until the target document is found.

When there are multiple matching facet-value pairs, deciding which one to include in the query can greatly impact how much additional searching is required to find the target document. Koren et al. suggested four possible selection methods. *Stochastic users* simply select one of the matching facet-value pairs randomly from a uniform distribution. *First-match users* scan the list from top to bottom, and select the first matching facet-value pair found. This heuristic is modeled after how users select matching documents from a ranked list of results. *Myopic users* select the matching facet-value pair that is contained in the least number of documents. With this method assumes that users are trying to reduce the search space as quickly as possible. *Optimal users* that perform actions that directly optimize the utility of

interface were also identified, but not examined in detail due to the complexity in searching for the optimal policy for the user to execute.

Although the simulated users differ from real users, the evaluation methodology does provide insight into understanding how various faceted interface design algorithms perform [53]. This evaluation method is neither better or worse than user studies. Instead, the approach serves to complement user studies by being cheap, repeatable, and controllable.

1.6.9 Conclusions

This section presented the problem of determining which facet-value pairs the system interface should provide to a user for query refinement was. In particular, we focus on personalized faceted search techniques that try to find facet-value pairs most useful to individual users. We introduced three major approaches, collaborative filtering based faceted search personalization, content based personalization, and ontology based personalization. We present a utility based evaluation framework for various faceted search interfaces, and the general idea is that the best interface should minimize the number/cost of interactions needed to find a document of interest.

1.7 Summary

This chapter demonstrated the breadth of design challenges and solutions in the realm of faceted search and dynamic taxonomies. Although there are some established standards in the user interface design for applications in this area, especially the emerging advanced navigation modes lack established, documented and usability evaluated approaches. We hope that future research will shed some light on the efficiency and efficacy of these novel approaches, in order to facilitate both design and usage of applications for faceted search and dynamic taxonomies.

References

1. Mani Abrol, Neil Latache, Uma Mahadevan, Jianchang Mao, Rajat Mukherjee, Prabhakar Raghavan, Michel Tourn, John Wang, and Grace Zang. Navigating large-scale semi-structured data in business portals. In *Proceedings of the 27th VLDB Conference*, pages 663–666, 2001.
2. E. Adar, D.R. Karger, and L. Stein. Haystack: Per-user information environments. In *Conference on Information and Knowledge Management*, 1999.
3. G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17:734–749, 2005.
4. Eugene Agichtein, Eric Brill, and Susan Dumais. Improving web search ranking by incorporating user behavior. In *SIGIR '06: Proceedings of the Annual international ACM SIGIR conference on Research and development in information retrieval*, 2006.
5. J. F. Allen, C. I. Guinn, and E. Horvitz. Mixed-Initiative Interaction. *IEEE Intelligent Systems*, Vol. 14(5):pp. 14–23, 1999.
6. G. Amato and C. Meghini. Faceted content-based image retrieval. In G.M. Sacco, editor, *DEXA Work. Dynamic Taxonomies and Faceted Search (FIND)*, pages 402–406. IEEE Computer Society, 2008.
7. Anton Andrejko, Michal Barla, and Michal Tvarožek. Comparing ontological concepts to evaluate similarity. In *Tools for Acquisition, Organisation, and Presenting of Information and Knowledge: Research Project Workshop*, pages 71–78, September 2006.
8. Holger Bast and Ingmar Weber. When you're lost for words: Faceted search with autocompletion. In Andrei Broder and Yoelle Maarek, editors, *SIGIR'06 Workshop on Faceted Search*, pages 31–35, Seattle, USA, August 2006. ACM.
9. Marcia J. Bates. The design of browsing and berrypicking techniques for the online search interface. *Online Review*, 13(5):407–424, 1989.
10. Robert M. Bell and Yehuda Koren. Lessons from the netflix prize challenge. *SIGKDD Explor. Newsl.*, 9(2):75–79, 2007.
11. Robert M. Bell and Yehuda Koren. Lessons from the netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 9(2):75–79, December 2007.
12. K. Bharat. Searchpad: Explicit capture of search context to support web search. In *Proceeding of 9th International WWW Conference*, 2000.
13. Wray Bintine and Henry Tirri. Multi-faceted learning for web taxonomies. *KD-NET 2002: 2nd Workshop on Semantic Web Mining*, 2002.
14. R. Bodner and M. Chignell. Dynamic Hypertext: Querying and Linking. *ACM CSUR*, Vol. 31(4es), December 1999. Article No. 15.
15. John S. Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. Technical report, Microsoft Research, One Microsoft Way, Redmond, WA 98052, 1998.

16. M. Bruls, K. Huizing, and J. van Wijk. Squarified treemaps. In *Proc. of Joint Eurographics and IEEE TCVG Symp. on Visualization (TCVG 2000)*, pages 33–42. IEEE Press, 2000.
17. P. Brusilovsky and C. Tasso. Preface to special issue on user modeling for web information retrieval. In *User Modeling and User Adapted Interaction 14 (2-3)*, 2004.
18. Jamie Callan. Document filtering with inference networks. In *Proceedings of the Nineteenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 262–269, 1996.
19. Harr Chen and David R. Karger. Less is more: Probabilistic models for retrieving fewer relevant documents. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '06)*, pages 429–436, New York, NY, USA, August 2006. ACM Press.
20. E. F. Codd, S. B. Codd, and C. T. Salley. Providing olap (on-line analytical processing) to user-analysis: An it mandate, 1993. White paper.
21. W. B. Croft. Relevance feedback and personalization: A language modeling perspective. In *Proceedings of Second DELOS Workshop: Personalization and Recommender Systems in Digital Libraries*, 2001.
22. W. Bruce Croft and John Lafferty. *Language Modeling for Information Retrieval*. Kluwer Academic Publishers, Norwell, MA, USA, 2003.
23. Raimund Dachsel, Mathias Frisch, and Markus Weiland. Facetzoom: a continuous multi-scale widget for navigating hierarchical metadata. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1353–1356, New York, NY, USA, 2008. ACM.
24. Scott C. Deerwester, Susan T. Dumais, Thomas K. Landauer, George W. Furnas, and Richard A. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science and Technology*, 41(6):391–407, 1990.
25. J. Delgado and N. Ishii. Memory-based weightedmajority prediction for recommender systems. In *ACM SIGIR'99 Workshop on Recommender Systems*, 1999.
26. Xin Dong and Alon Y. Halevy. A platform for personal information management and integration. In *Conference on Innovative Data Systems Research*, 2005.
27. Marian Dörk, Sheelagh Carpendale, Christopher Collins, and Carey Williamson. Visgets: Co-ordinated visualizations for web-based information exploration and discovery. *Visualization and Computer Graphics, IEEE Transactions on*, 14(6):1205–1212, 2008.
28. Susan Dumais, Edward Cutrell, Raman Sarin, and Eric Horvitz. Implicit queries (iq) for contextualized search. In *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 594–594, New York, NY, USA, 2004. ACM Press.
29. L. Finkelstein, E. Gabrilovich, Y. Matias, E. Rivlin, Z. Solan, G. Wolfman, and E. Ruppín. Placing search in context: The concept revisited. In *Proceedings of 10th International World Wide Web*, 2001.
30. Simon Funk. Try this at home. <http://sifter.org/~simon/journal/20061211.html>, December 2006.
31. Jim Gemmell, Gordon Bell, and Roger Lueder. Mylifebits: a personal database for everything. In *Communications of the ACM*, vol. 49, Issue 1, 2006.
32. Kallirroi Georgila, James Henderson, and Oliver Lemon. User simulation for spoken dialogue systems: Learning and evaluation. In *Proceedings of the 9th International Conference on Spoken Language Processing (INTERSPEECH-ICSLP 2006)*, 2006.
33. R. Godin, R. Missaoui, and A. April. Experimental comparison of navigation in a Galois lattice with conventional information retrieval methods. *International Journal of Man-Machine Studies*, 38(5):747–767, 1993.
34. M. A. Hearst. Design recommendations for hierarchical faceted search interfaces. In *Proc. SIGIR 2006, Workshop on Faceted Search*, pages 26–30, August 2006.
35. Marti Hearst. Uis for faceted navigation: Recent advances and remaining open problems. In *Workshop on Computer Interaction and Information Retrieval, HCIR*, Redmond, WA, October 2008.

36. Marti Hearst, Ame Elliott, Jennifer English, Rashmi Sinha, Kirsten Swearingen, and Ka-Ping Yee. Finding the flow in web site search. *Communications of the ACM*, 45(9):42–49, 2002.
37. Marti A. Hearst. Clustering versus faceted categories for information exploration. *Commun. ACM*, 49(4):59–61, 2006.
38. J. Heer, M. Agrawala, and W. Willett. Generalized selection via interactive query relaxation. In *Computer-Human Interaction (CHI)*, pages 959–968. ACM, 2008.
39. Jeffrey Heer and George Robertson. Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1240–1247, 2007.
40. Jonathan L. Herlocker, Joseph A. Konstan, Al Borchers, and John Riedl. An algorithmic framework for performing collaborative filtering. In *SIGIR '99: Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 230–237, New York, NY, USA, 1999. ACM Press.
41. Jonathan L. Herlocker, Joseph A. Konstan, Al Borchers, and John Riedl. An algorithmic framework for performing collaborative filtering. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 99)*, pages 230–237, New York, NY, USA, 1999. ACM.
42. Michiel Hildebrand, Jacco van Ossenbruggen, and Lynda Hardman. /facet: A browser for heterogeneous semantic web repositories. In *International Semantic Web Conference*, pages 272–285, 2006.
43. Haym Hirsh, Chumki Basu, and Brian D. Davison. Learning to personalize. *Commun. ACM*, 43(8):102–106, 2000.
44. Thomas Hofmann and Jan Puzicha. Latent class models for collaborative filtering. In *IJCAI '99: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 688–693, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
45. David Huynh and David Karger. Parallax and companion: Set-based browsing for the data web. (submitted to WWW 2009), 2009.
46. David F. Huynh, David R. Karger, and Robert C. Miller. Exhibit: lightweight structured data publishing. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 737–746, New York, NY, USA, 2007. ACM.
47. Glen Jeh and Jennifer Widom. Scaling personalized web search. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, 2003.
48. Rong Jin, Joyce Y. Chai, and Luo Si. An automatic weighting scheme for collaborative filtering. In *SIGIR '04: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 337–344, New York, NY, USA, 2004. ACM Press.
49. David Karger and Mc Schraefel. The pathetic fallacy of rdf, <http://swui.semanticweb.org/swui06/papers/karger/pathetic.fallacy.html>, 2006.
50. Amy K. Karlson, George G. Robertson, Daniel C. Robbins, Mary P. Czerwinski, and Greg R. Smith. Fathumb: a facet-based interface for mobile search. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 711–720, New York, NY, USA, 2006. ACM.
51. Georgi Kobilarov and Ian Dickinson. Humboldt: Exploring linked data. In *Linked Data on the Web Workshop (LDOW2008) at WWW2008*, Beijing, China, 2008.
52. Joseph A. Konstan, Bradley N. Miller, David Maltz, Jonathan L. Herlocker, Lee R. Gordon, and John Riedl. GroupLens: Applying collaborative filtering to Usenet news. *Communications of the ACM*, 40(3):77–87, 1997.
53. Jonathan Koren, Yi Zhang, and Xue Liu. Personalized interactive faceted search. In *Proceedings of the 17th International Conference on the World Wide Web (WWW '08)*, 2008.
54. Bill Kules. *Supporting exploratory web search with meaningful and stable categorized overviews*. PhD thesis, University of Maryland at College Park, 2006.
55. Esther Levin, Roberto Pieraccini, and Wieland Eckert. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transaction on Speech and Audio Processing*, 8(1):11 – 23, January 2000.
56. David Lewis. Applying support vector machines to the TREC-2001 batch filtering and routing tasks. In *Proceedings of the Eleventh Text REtrieval Conference (TREC-11)*, 2002.

57. Rui Li, Shenghua Bao, Yong Yu, Ben Fei, and Zhong Su. Towards effective browsing of large scale social annotations. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 943–952, New York, NY, USA, 2007. ACM.
58. Gary Marchionini. *Information seeking in electronic environments*. Cambridge University Press, New York, NY, USA, 1997.
59. D. R. Millen, J. Feinberg, and B. Kerr. Dogear: Social bookmarking in the enterprise. In R. E. Grinter et al, editor, *Conf. Human Factors in Computing Systems (CHI)*, pages 111–120. ACM, 2006.
60. Dunja Mladenić. Personal web watcher: Design and implementation. Technical report, J. Stefan Institute, Department for Intelligent Systems, Ljubljana, Slovenia, 1998.
61. M. Narayan, C. Williams, S. Perugini, and N. Ramakrishnan. Staging Transformations for Multimodal Web Interaction Management. In *Proc. of WWW'04*, pages 212–223.
62. Arkadiusz Paterek. Improving regularized singular value decomposition for collaborative filtering. In *Proceedings of the KDD Cup and Workshop at the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2007.
63. S. Perugini and N. Ramakrishnan. Personalizing Web Sites with Mixed-Initiative Interaction. *IEEE IT Professional*, Vol. 5(2):pp. 9–15, 2003.
64. P. Piroli and S. Card. Information foraging models of browsers for very large document spaces. *Proceedings of the Advanced Visual Interfaces Workshop, AVI '98*, pages pp. 83–93, 1998.
65. J. O'Regan R. Rensink and J. Clark. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, (8):368–373, 1997.
66. Ronald E. Rice, Maureen McCreadie, and Shan-Ju L. Chang. *Accessing and Browsing Information and Communication*. MIT Press, 2001.
67. S. Robertson and K. Sparck-Jones. Relevance weighting of search terms. In *Journal of the American Society for Information Science*, volume 27, pages 129–146, 1976.
68. Ann Rockley. *Managing Enterprise Content: A Unified Content Strategy*. New Riders Press, October 2002.
69. G. M. Sacco. Dynamic taxonomies: A model for large information bases. *IEEE Transactions Knowledge and Data Engineering*, 12(3):468–479, 2000.
70. Giovanni Maria Sacco. Dynamic taxonomies: A model for large information bases. *IEEE Trans. Knowl. Data Eng.*, 12(2):468–479, 2000.
71. Badrul M. Sarwar, George Karypis, Joseph A. Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on the World Wide Web (WWW '01)*, pages 285–295, 2001.
72. Badrul M. Sarwar, George Karypis, Joseph A. Konstan, and John Riedl. Incremental singular value decomposition algorithms for highly scalable recommender systems. In *Proceedings of the 5th International Conference on Computer and Information Technology (ICCIT 2002)*, 2002.
73. Badrul M. Sarwar, George Karypis, Joseph A. Konstan, and John T. Riedl. Application of dimensionality reduction in recommender system – a case study. In *Proceedings of the Workshop on Web Mining and E-Commerce Systems at KDD 2000 (WEBKDD 2000)*, August 2000.
74. Xuehua Shen, Bin Tan, and Chengxiang Zhai. Context sensitive information retrieval using implicit feedback. In *SIGIR '05: Proceedings of the Annual international ACM SIGIR conference on Research and development in information retrieval*, 2005.
75. Xuehua Shen, Bin Tan, and Chengxiang Zhai. Context-sensitive information retrieval using implicit feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '05)*, pages 43–50, New York, NY, USA, 2005. ACM Press.
76. Ben Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Transactions on Graphics (TOG)*, 11(1):92–99, 1992.
77. Vineet Sinha and David R. Karger. Magnet: Supporting navigation in semistructured data environments. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data (SIGMOD '05)*, pages 97–106, 2005.

78. Moritz Stefaner and Boris Müller. Elastic lists for facet browsers. In *DEXA '07: Proceedings of the 18th International Conference on Database and Expert Systems Applications*, pages 217–221, Washington, DC, USA, 2007. IEEE Computer Society.
79. Moritz Stefaner, Thomas Urban, and Marc Seefelder. Elastic lists for facet browsing and resource analysis in the enterprise. In *DEXA Workshops*, pages 397–401. IEEE Computer Society, 2008.
80. Emilia Stoica, Marti A. Hearst, and Megan Richardson. Automating creation of hierarchical faceted metadata structures. In *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT 2007)*. NAACL-HLT, 2007.
81. J. Teevan, S.T. Dumais, and D.J. Liebling. To personalize or not to personalize: Modeling queries with variation in user intent. In *Proceedings of the 31st annual international conference on Research and development in information retrieval*, 2008.
82. Jaime Teevan, Christine Alvarado, Mark S. Ackerman, and David R. Karger. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422, New York, NY, USA, 2004. ACM.
83. Jenifer Tidwell. *Designing Interfaces : Patterns for Effective Interaction Design*. O'Reilly Media, Inc., November 2005.
84. Daniel Tunkelang. Dynamic category sets: An approach for faceted search. In *ACM SIGIR '06 Workshop on Faceted Search*, August 2006.
85. Andrew Turpin and Falk Scholer. User performance versus precision measures for simple search tasks. In *SIGIR '06: Proceedings of the Annual international ACM SIGIR conference on Research and development in information retrieval*, 2006.
86. Michal Tvarožek and Mária Bielíková. Personalized faceted navigation for multimedia collections. In *Proceedings of the 2nd International Workshop on Semantic Media and Personalization (SMAP 2007)*, pages 104–109. IEEE Computer Society, 2007.
87. Michal Tvarožek and Mária Bielíková. Personalized faceted navigation in the semantic web. In *Proceedings of the International Conference on Web Engineering (ICWE 2007)*, pages 511–515, 2007.
88. Barbara Tversky and Julie B. Morrison. Animation: can it facilitate? *International Journal of Human-Computer Studies*, 57:247–262, 2002.
89. Wesley Willett, Jeffrey Heer, and Maneesh Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1129–1136, 2007.
90. S. Xu, S. Bao, B. Fei, Z. Su, and Y. Yu. Exploring folksonomy for personalized search. In *Proceedings of the 31st annual international conference on Research and development in information retrieval*, 2008.
91. Kiduk Yang, Elin K. Jacob, Aaron Loehrlein, Seungmin Lee, and Ning Yu. The best of both worlds: A hybrid approach to the construction of faceted vocabularies. In *Proceedings of the Asian Information Retrieval Symposium*, pages 301–304, 2004.
92. Y. Yang, S. Yoo, J. Zhang, and B. Kisiel. Robustness of adaptive filtering methods in a cross-benchmark evaluation. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2005.
93. Ka-Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. Faceted metadata for image search and browsing. In *CHI '03: Proceedings of the conference on Human factors in computing systems*, pages 401–408. ACM Press, 2003.
94. Kai Yu, Volker Tresp, and Shipeng Yi. A nonparametric hierarchical bayesian framework for information filtering. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '04)*, pages 353 – 360, New York, NY, USA, 2004. ACM Press.
95. Chengxiang Zhai and John Lafferty. Model-based feedback in the language modeling approach to information retrieval. In *Proceedings of Tenth International Conference on Information and Knowledge Management*, pages 403–410, 2001.

96. J. Zhang and G. Marchionini. Coupling browse and search in highly interactive user interfaces: a study of the relation browser++. In *JCDL '04: Proceedings of the 4th ACM/IEEE-CS joint conference on Digital libraries*, pages 384–384, New York, NY, USA, 2004. ACM.
97. Yi Zhang and Jonathan Koren. Efficient bayesian hierarchical user modeling for recommendation systems. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '07)*, New York, NY, USA, July 2007. ACM Press.
98. D. X. Zhou, N. Oostendorp, M. Hess, and P. Resnick. Conversation pivots and double pivots. In *Computer-Human Interaction (CHI)*, pages 959–968. ACM, 2008.
99. Philip Zigoris and Yi Zhang. Bayesian adaptive user profiling with explicit and implicit feedback. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management (CIKM 2006)*, pages 397 – 404, New York, NY, USA, 2006. ACM, ACM Press.