Fluid Views: A Zoomable Search Environment

Marian Dörk, Sheelagh Carpendale, Carey Williamson
Department of Computer Science, University of Calgary
Calgary, Alberta, Canada
{mdoerk, sheelagh, carey}@ucalgary.ca

ABSTRACT

We present Fluid Views, a web-based search environment designed to bridge overview and detail by integrating dynamic queries, semantic zooming, and dual layers. The most common form of search results is long ranked and paginated lists, which are seldom examined beyond the top ten items. To support more exploratory forms of information seeking, we bring together the notion of relevance with the power of visual encoding. In Fluid Views, results portray relevance via size and detail in a dynamic top layer and semantic similarity via position on a base map. We designed Fluid Views with temporal, spatial, and content-defined base maps for both textual and visual resources, and tested our prototype system on books, blogs, and photos. Interviews with library professionals indicate the potential of Fluid Views for exploring collections and exciting directions for future research.

Keywords

Information visualization, information seeking, Web search, zoomable user interface, level of detail, exploratory search.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: Misc.

General Terms

Design, Human Factors.

1. INTRODUCTION

With the growing abundance of digital information spaces in our lives, we rely on effective exploration tools. Search interfaces typically offer bottom-up details showing only a few of the many results, while visualizations commonly provide top-down overviews of large datasets. Despite many improvements in accuracy and relevance, search engines present results in much the same manner as fifteen years ago, as ranked lists of snippets and thumbnails. While the accuracy and immediacy of access are valuable, there is no sense of overview conveyed by a long ranked list. Lists of news items, books, and photos lack the semantic richness of a thoughtful layout, as manifested in a newspaper, book shelf, or photo album. A significant tension in information exploration lies between *detail* of items (as provided by search engines) and *overview* of a collection (as provided by visualizations)—two fundamentally different perspectives on information.

Extensive prior research has tried to narrow the gap between overview and detail. A prominent sequence of visualization and interaction research was concerned with improving access to detailed views via zooming [e.g., 3, 21]. However, zooming introduced a range of navigational issues [17], and was never applied to search. There has been considerable work on integrating visualization into search interfaces [e.g., 9, 14, 26], although mostly as a supplement. We address this gap, exploring how to support search more fully by drawing upon information visualization techniques.

Our goal is to support exploratory forms of information seeking by offering both overviews of an information space and direct access to individual results. We introduce *Fluid Views*, a zoomable search environment consisting of two coupled visualization layers (see Figure 1). The top layer is a dynamic view of information resources that are sized according to search relevance. The base map provides a subtle visual context and positions results using their temporal, spatial, or content-based similarity. Fluid Views is a functional web-based prototype, which we have applied to three information spaces: books, blogs, and photos. Library professionals have shared encouraging feedback about the potential of Fluid Views as well as indicating future directions.

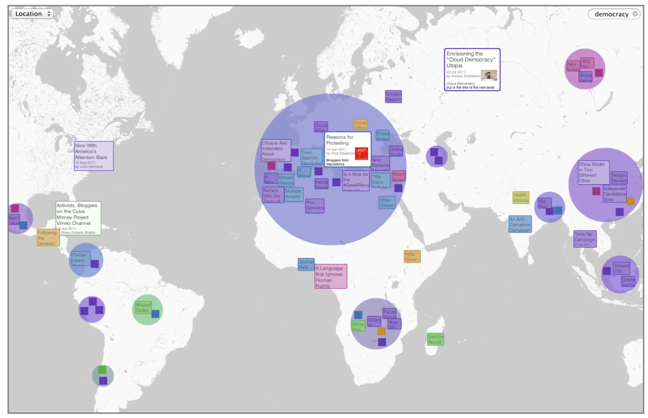

Figure 1: In the Fluid Views interface, items are positioned using similarity and adjusted in size and level of detail according to relevance.

AVI '12, May 21-25, 2012, Capri Island, Italy.

2. BACKGROUND

In the following, we give a brief summary of the research on information visualization for search, visual overviews, and zooming techniques as it pertains to Fluid Views.

Result Visualization. Some research has investigated using information visualization for the display of search results. These can be distinguished as preview or overview displays [13]. On one hand, results can be displayed as visual previews, for example, by blending thumbnails with text snippets and visualization [26]. On the other hand, overview visualizations have been used to provide context for results [6]. In order to help searchers assess the relevance of the results, small per-result visualizations can represent document-query similarities [14]. Another approach to adding interactive graphical overviews is to use the layout for displaying the results. Result thumbnails can be positioned not only as lists, but also as simple visualizations such as zoomable bar charts [20]. In our work, we seek to explore richer base visualizations with temporal, spatial, and content-based arrangements and investigate more closely the abstraction and display of items within these visualizations.

Visual Queries and Facets. Visualizations can also be used to formulate queries against datasets and documents. For example, dynamic queries can offer tightly-coupled filters for visualizations [1], and visual representations can indicate query term combinations [24]. More recently, there has been work at the intersection between faceted navigation and information visualization [9, 18, 25]. Visualizations for query formulation and faceted navigation allow for sophisticated filtering, however they typically separate the visualization from the results. That is, they enhance querying and filtering but they shift focus away from individual results.

Visual Overviews. The term 'overview' is used widely and variedly in information visualization [15]. Shneiderman's visual information-seeking mantra underlines its primacy: "Overview first, zoom and filter, then details-on-demand" [22]. Spatialization techniques have been developed to provide overviews of large high-dimensional data sets such as document collections. Landscape and galaxy metaphors circumscribe representations that utilize proximity of visual elements to indicate similarity between the underling resources [5, 27]. Spatialization is used, for example, to represent knowledge domains and scientific areas [2, 4]. Visual overviews such as galaxies, landscapes, or maps reveal overall patterns rather than individual resources. In contrast, we aim to integrate both abstract overviews of collections with literal details of individual resources.

Zoomable Interfaces. If overviews provide birds-eye perspectives, then zoomable interfaces are parachutes. Pad introduced the interface concept of an infinite zooming space for interacting with digital objects at varying levels of detail [21]. When zooming in Pad, resources reveal more and more of their content: from a title to a summary to an entire text. Pad++ expanded the original idea to exploring presentations, drawings, poetry, and Web browsing [3]. The idea of infinite zooming interfaces is highly relevant to Fluid Views. However, Pad and Pad++ had difficulties of zooming into blank, visually undefined space, which has been described as 'desert fog' [17]. Suggestions made to address this, such as the use of landmarks, critical zones, and multi-scale residue, inform our work via the constructs of base layers, visual ranking, and multi-scale clustering, respectively.

Interacting with Overviews. Orientation and perceptual smoothness have improved with advances in zooming interaction and animation, particularly by coupling panning speed with the zoom level [16]. Our approach, which considers factors other than scale for the appearance of visual elements, is more comparable to cue-based techniques discussed in a recent survey paper [7]. For example, in spatialization visualizations, search matches can be highlighted by colour [5] or halos [2]. In contrast, we use item size and level of detail to indicate search relevance. We seek to integrate representational and interactive aspects that have been considered separately in research on overviews and zooming.

3. TOWARDS ZOOMABLE SEARCH

While common keyword search has begun to offer increasingly satisfying results, the display of these results is still commonly the top-ranked items in a long list. This kind of information display works sufficiently well when searching for very specific items, however, it may not encourage more exploratory forms of information seeking [19]. The motivation behind this work is the duality of horizontal exploration (free-ranging lateral choices) and vertical immersion (targeted, focused direction) as complementary activities of many information seekers [10]. To support both of these search activities (broad and high-level vs narrow and low-level), our main design goals were:

- ease transition from overviews to detailed perspectives,
- display resources in context of their collection,
- scale results according to relevance or popularity, and
- provide clickable access to many results.

In summary, the aim is to expand the common top-ten ranked results to several hundred items that are visually and interactively accessible in a meaningful interface. To achieve this, we provide a search interface that integrates relevance-based ranking with scale-variant visual representations.

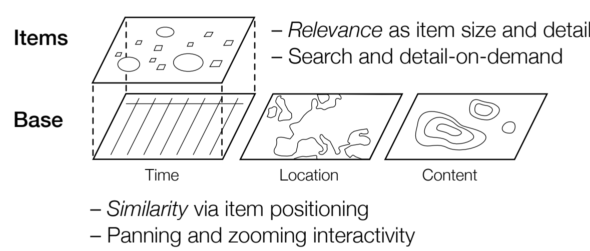

Figure 2: Fluid Views comprise two interconnected layers: the top layer features items at varying scales superimposed over a zoomable base layer, for example, representing time, location, and content.

As illustrated in Figure 2, the basis of Fluid Views is a *dual-layer approach* that allows for complementary scaling mechanisms as well as visual encodings. This layer-based approach is inspired by thematic maps, where content layers are superimposed onto geographic base maps. This idea is also related to hybrid visualizations that combine multiple techniques into an integrated view [e.g., 8, 23] as well as mapping mashups prevalent on the Web. Pad [21] and Pad++ [3] introduced the idea of a single zoomable layer, which triggered an entirely new way of thinking about interfaces, but also introduced challenges of navigating into blank spaces. We use a dual-layer approach to a zoomable search environment to alleviate these issues.

4. DESIGNING FLUID VIEWS

To explain our design decisions when creating Fluid Views as a dual-layer zoomable search environment, we describe the *base layers* and the *item layer*, and our use of clustering to integrate them into one interface.

4.1 The Base Layers

A base layer provides spatial semantics. It contains a map that supports continuous uniform scaling, similar to early zooming interfaces [3, 21] and as modeled by space-scale diagrams [12]. However, unlike Pad++, our base layer holds a visual background that covers the entire layer. This backdrop provides visual landmarks in the zooming environment, positional semantics for the items, and ambient orientation within the zoom space [17]. The visual background is kept subliminal: recognizable if needed, without detracting from the items being displayed. Homogeneous scaling of the entire base map reveals regions that are appropriate for a given scale and conceals them as they become too large or too small to display. We created three example base layers where search items can be positioned based on similarities in time, location, and content.

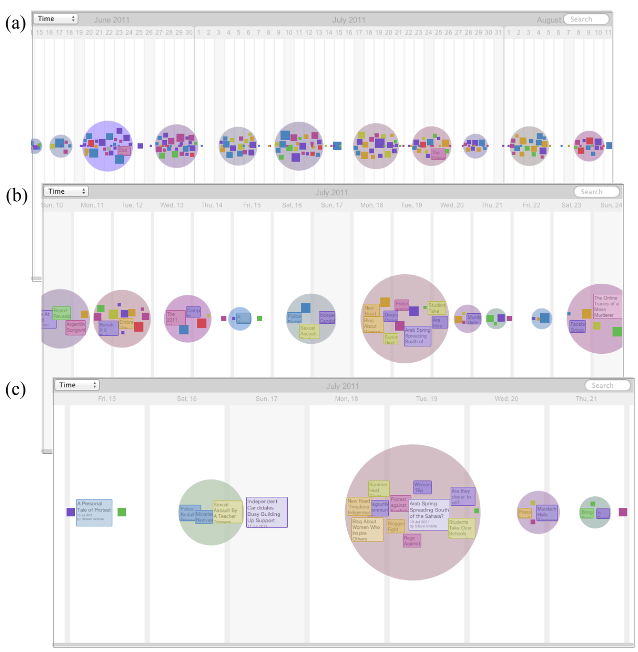

Figure 3: The time scale is changed from (a) two months, to (b) two weeks, and (c) seven days.

The *time map* has a single linear but zoomable dimension of time. It displays visual areas for common temporal units such as years, months, days, or hours progressing left-to-right horizontally. Colours chosen to delineate the time regions are light grey, dark enough to be clearly readable, and light enough to avoid interfering with the actual information items. To keep the time map clearly readable, it only shows two temporal levels at a time. These two levels are arranged to indicate nesting of the smaller unit within the larger one. As shown in Figure 3, areas representing months contain those representing days, which in turn contain hours when zoomed in further. Labels on the upper edge name the periods currently being shown and adjust depending on the width of the respective temporal area. Every zooming and panning operation gradually extends, contracts, or shifts the time window, and transitions the visual units into place.

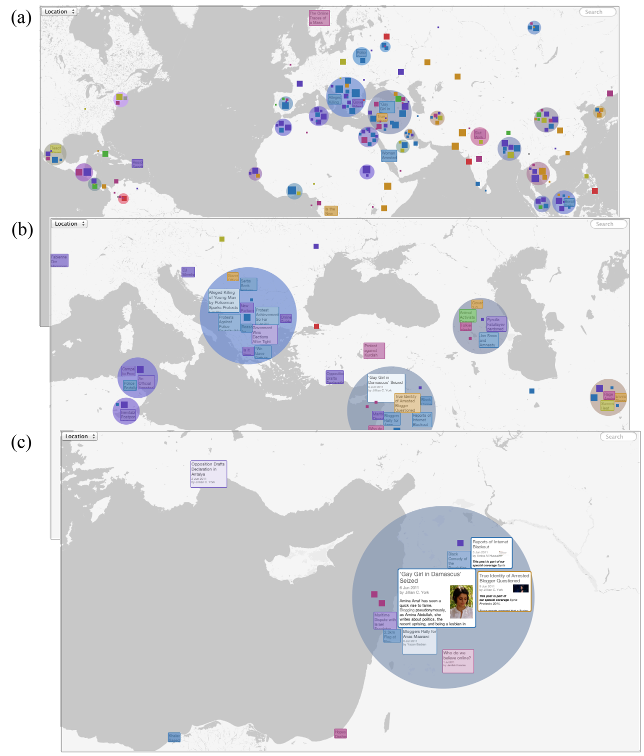

Figure 4: Zooming from a continental view (a) to the Mediterranean (b), and finally a regional close-up (c). With increasing geographical detail in the map, the displayed blog articles are continuously refined.

The *location map* provides a two-dimensional interactive geographic map. For simplicity, the map primarily indicates the natural formations and boundaries, such as the contours of continents, rivers, and oceans (see Figure 4). White areas represent land, while light grey is used to indicate water bodies such as rivers, lakes and oceans. With this design choice, land-based items are more likely to appear on a bright background, since the location map uses geospatial coordinates for item positioning. A spatial filter determines the items displayed in the view based on the map.

The *content map* is a general N-dimensional base layer. It uses multi-dimensional scaling as a spatialization technique to represent similarities within an information collection. Locations with a high concentration of items appear as bright amorphous shapes, while darker shades of grey are used for regions with no resources. This approach parallels the sparsity of resources on water bodies in the location map. The content map serves as a generic map that can make use of any content-defined dimensions, such as words from textual resources or colours from images, as the basis for similarity determination. In Figure 5, digital photographs

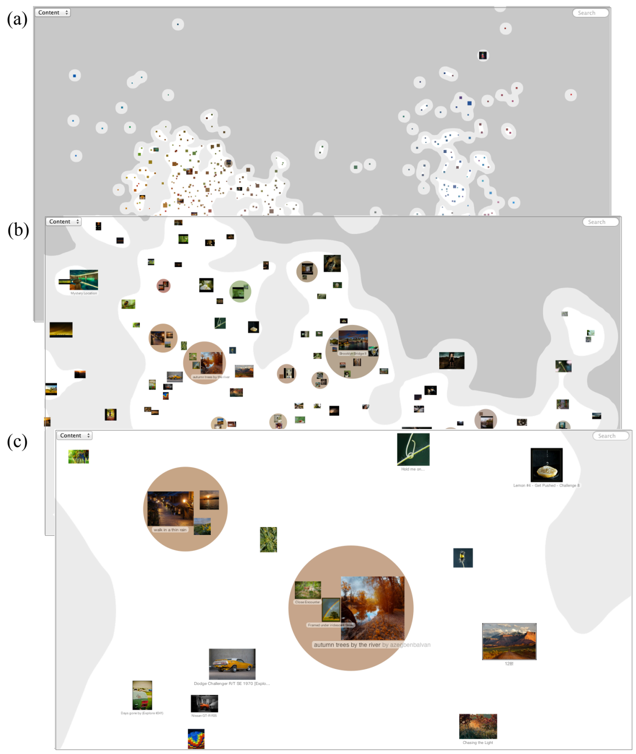

Figure 5: When zooming into the content map from (a) to (c), the base layer gets gradually magnified and the items, here photographs, are appropriately filtered and displayed with more detail.

are positioned based on the similarity of dominant colours. The spatialization method generates output values ranging between -1 and 1 in each of the X and Y dimensions, which are then used to determine the positions of items.

4.2 The Item Layer

While the base maps are designed to provide context, the item layer is designed to offer the information items for the viewer's attention. First, each item is positioned according to its metadata and the semantics of the current base map. Next, the items are sized according to their relevance. This is calculated either according to the item's general popularity or according to its relevance to the current search query. An item's relevance is encoded by its screen real estate (i.e., pixel area), and thus how much detail it can reveal. In our current design, we support two types of resources as items: visual and textual. The visual resource type is primarily designed for images, such as photographs or paintings, while textual resources are more suitable for articles and books. Hybrid media are also possible: an image may have a title or description, and a text may have an accompanying image.

The item layer has two main interaction methods: search queries and direct manipulation. As a textual search is entered, the interface dynamically ranks resources according to relevance, which can result in varying scale changes for the displayed items. In the absence of any search terms, items are sized according to their popularity, as indicated by number of comments, views, or editions. Clicking on an item immediately reveals its contents in full detail, using

Figure 6: Items scale from small dots to large previews. While discrete steps exist for showing certain aspects of an item, size and detail changes to image (top) or text (bottom) displays are continuous.

a suitably large portion of the screen. Furthermore, items respond to the interactions of the base map, namely panning and zooming. These interaction techniques provide additional ways to adjust the level of detail in the visual representations. In general, panning and zooming produce uniform display changes, while search and detail operations trigger more varied but highly selective changes.

Regardless of the interaction methods used, all display changes are carried out as animated transitions to aid comprehension. For this, the challenge is to provide semantic zooming functionality as a family of multi-level representations that gradually shift from small to large with sensible increases of detail. The representations of visual and textual resources are handled differently at intermediate scales to exploit the corresponding modality.

The smallest representation of a *visual resource* uses the dominant colour extracted from the image data. While the use of colour can be regarded as an abstract representation, it is based on the visual content of the resource. When an item's size exceeds a certain threshold, an actual thumbnail replaces the square. With continuous scale changes, the image data is dynamically replaced with images of appropriate display resolution (see Figure 6, top row). Labels of visual resources include a title with gradually-increasing font-size and, beyond certain size intervals, additional metadata.

The colour used for a *textual resource* is based on themes extracted from the overall collection of items. The colour is used for the smallest dot, as well as the border and background in intermediate representations. As the size of the item gradually increases, the title is augmented with the author's name, the date of publication, and eventually the actual text (see Figure 6, bottom row). Once a textual resource reaches a certain size, one can scroll through the entire document within the interface. While all resources can be viewed in high detail within the interface, the titles of resources are also hyperlinks to the respective Web pages.

The search function is associated with the item layer, since it exclusively changes the scale of information resources. The search box allows basic Boolean expressions and simple term searches. By placing '+' and '−' in front of terms, items that do or do not have a certain term are highlighted. If no operators are used, the search terms are treated as a logical disjunction (OR). Item relevance increases with more matching terms. This results in a more nuanced display that emphasizes items that satisfy all query terms, but also

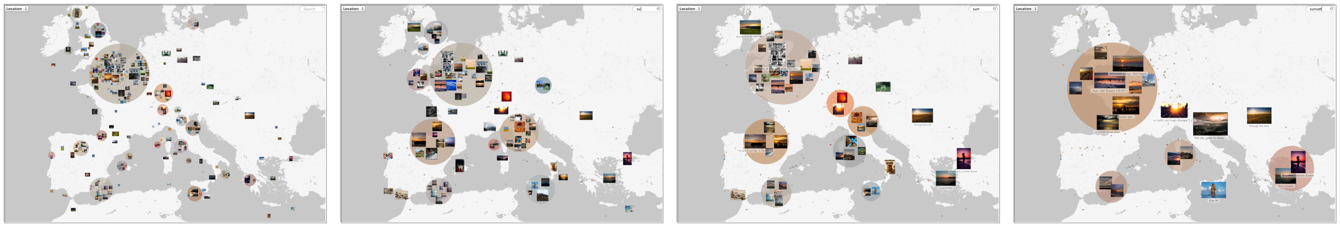

(a) No query (b) 'su' (c) 'sun' (d) 'sunset'

Figure 7: When changing search terms, item sizes and details gradually change according to relevance scores.

reveals items that have some terms. While terms are being typed into the search box, the top layer gradually changes the sizes for all items accordingly (see Figure 7). Items that do not match any search term are reduced to their minimal size releasing additional screen space for relevant items.

4.3 Clustering

To avoid overlaps when many items are positioned at close locations, items are grouped into circular clusters when close together (see Figure 8, left). The position of a cluster is the weighted average of the positions of its members. The colour of the cluster is a (slightly brighter) average of the colours from the contained items. The size of a cluster circle directly corresponds with the aggregated pixel area of its constituent members, which reflects their relevance values. Since cluster sizes are based on aggregate relevance, it is difficult to compare clusters with regard to cardinality. However, because the clusters in Fluid Views do not conceal members, the number of elements per cluster remains visually discernible.

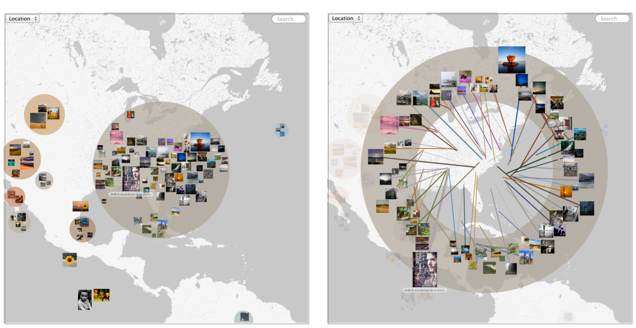

Figure 8: To see the exact positions of cluster members, a cluster can be opened following a transition arranging the items around the cluster and displaying edges to their actual place on the base map.

Special interaction techniques are provided for revealing more detail within clusters. When a cluster is clicked, its items are moved towards the periphery of the cluster and returned to the size they would assume if not part of a cluster. As shown in Figure 8, right, an edge from each item on the periphery indicates its original position on the base map not unlike excentric labelling [11]. In this way, clusters can be thought of as lenses that allow items to be viewed at two levels of detail. Specifically, they reduce overlap so that all items can be recognized, and they facilitate the examination of detailed positions. These detail operations on clusters are also similar to portals in Pad [21], but they are embedded and integrated with the displayed objects.

4.4 Integrating the Two Layers

While most search interfaces keep item size constant in a simple list and associate relevance with the position in the list, the Fluid Views interface uses item size to convey relevance, freeing position to be used as a visual variable that can represent similarity relationships. The visual arrangement of items via size and position constitutes a form of abstract representation that is blended with varying levels of literal representation of items. Analogously, clusters themselves are hybrid formations that merge literal aspects of individual items into abstract representations of clusters.

To achieve this combination of an abstract representation (overview) with the literal search results items (detail), the two layers in Fluid Views provide two distinct functions. The base layer provides visual context for the items by providing a subtle texture of a given facet, as well as mapping and zooming operations. The top layer serves as a dynamic result space that displays the items with varying sizes and details, which can be manipulated through search and detail operations. In order to make good use of screen space, items are sized according to their current relevance while also taking into account the size of the screen.

The visualization sequence in Fluid Views can be summarized by a 6-step process:

1. *Positioning:* set dimensions and extent of base map.
2. *Selection:* determine which items are in view using position information from the base map.
3. *Ranking:* calculate relevance values based on item popularity or search matches.
4. *Spacing:* allocate screen space and set item sizes.
5. *Clustering:* group items based on their proximity.
6. *Draw:* display, change, or hide items and clusters.

Panning and zooming operations provided by the base map trigger all steps, while search query changes trigger the last four steps. Detail-on-demand operations only change the display of a single item or cluster (step 6).

5. EXAMPLE SCENARIOS

We have applied Fluid Views to books, blogs, and photographs, as exemplars for a wide range of resource types. For each collection, we outline a possible exploration scenario.

Books on Time Map. From Open Library,[1] we selected the 500 fiction books with the most editions. Year of first edition, subject place, and subject headings were used to position the books in the time, location, and content layers

[1] http://openlibrary.org/

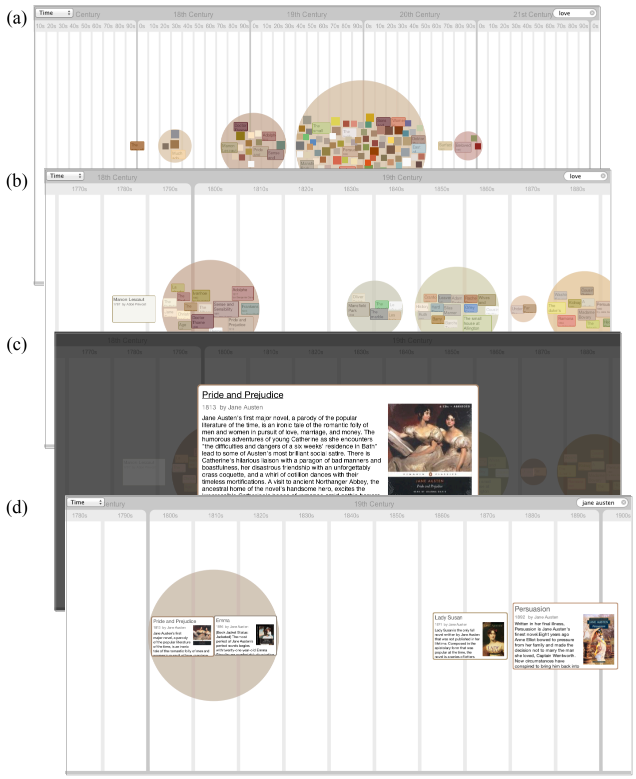

Figure 9: The scale in the time map is adjusted from five centuries (a) to one century (b). Clicking on a book brings up a detail view (c) and a search reveals other books by the same author (d).

respectively. A person interested in romance novels may enter 'love' as an initial search term (see Figure 9). The resulting view shows fewer books with varying sizes, aggregated in clusters (a). The person then zooms into the medium-sized cluster at the beginning of the 19th century (b), where the book 'Pride and Prejudice' catches the person's attention. This triggers the information seeker to click on the book to reveal the author, description, and cover image (c). Remembering the author's name from a discussion with a friend, the person changes the search query to 'Jane Austen' revealing three other books by the same author (d).

Blog Articles on Location Map. We extracted 500 articles from the community blog Global Voices.[2] Time of publication and tags are directly in the RSS feeds, and geographic coordinates were extracted using a web service.[3] A person interested in the Middle East may start out with a high-level view of many recent blog articles positioned on a world map according to primary geographical topics (see Figure 10). After zooming in slightly towards the region (a), the searcher may enter a general term such as 'peace' to explore recent peace developments in the wider region (b). Seeing a relatively large cluster in the middle of the map, the viewer clicks on it and is able to see the more detailed positions of the contained blog articles (c). The searcher may continue their exploration by selecting individual articles, zooming further in, or changing the search query.

[2]http://globalvoicesonline.org/

[3]http://geonames.org/

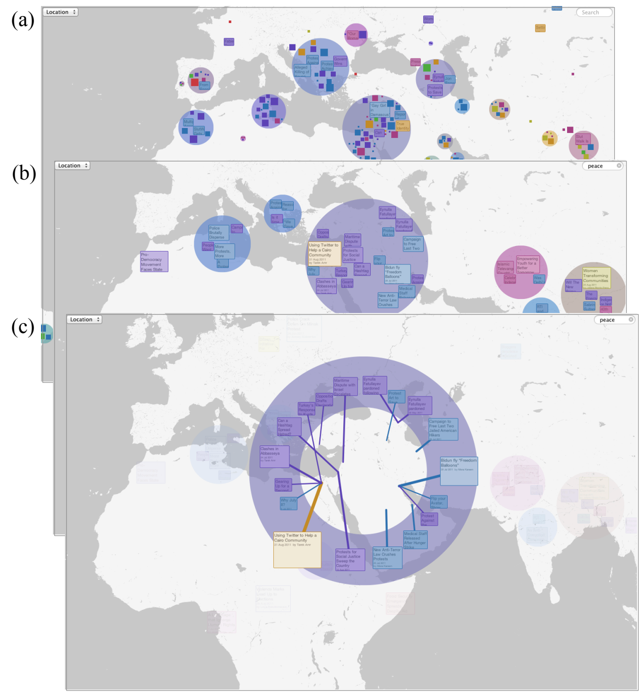

Figure 10: Looking for 'peace' in blog articles about the Middle East, North Africa, and South Europe.

Photos on Content Map. We retrieved 500 photos from Flickr's most interesting photographs for a day.[4] The photo's upload time, location where it was taken, and dominant colours are used for the base maps. A person interested in finding inspiring photos might enter 'nature' as a search term (see Figure 11). This filters the view showing more items on the left side of the display (a) towards which the person zooms (b). Intrigued by the photo of a squirrel, the person zooms in further (c) and notices two photos of a bumblebee and a curlew in close proximity. The person moves the semantic content map towards these other photos and then resets the search terms to see other related photos (d).

6. EXPERT INTERVIEWS

To better understand how Fluid Views could support everyday search practices, we conducted a qualitative study in a university library. The main goal was to learn about the potential of Fluid Views in the context of current developments in a library. The participants (4 female, 2 male) are all trained library professionals with different roles, including instruction, reference work, and administration, as well as content and technology development. Our study was conducted in two parts: part one was a short explanatory demo of Fluid Views answering questions until the participants were satisfied that they understood Fluid Views; part two was individual semi-structured interviews, which each lasted between 15-45 minutes. The questions were formulated to engage in an open critique-based conversation. Example questions were: "Could you explain how library patrons find resources today? What role could a system like Fluid Views play in their search practices? Do you see prob-

[4]http://www.flickr.com/

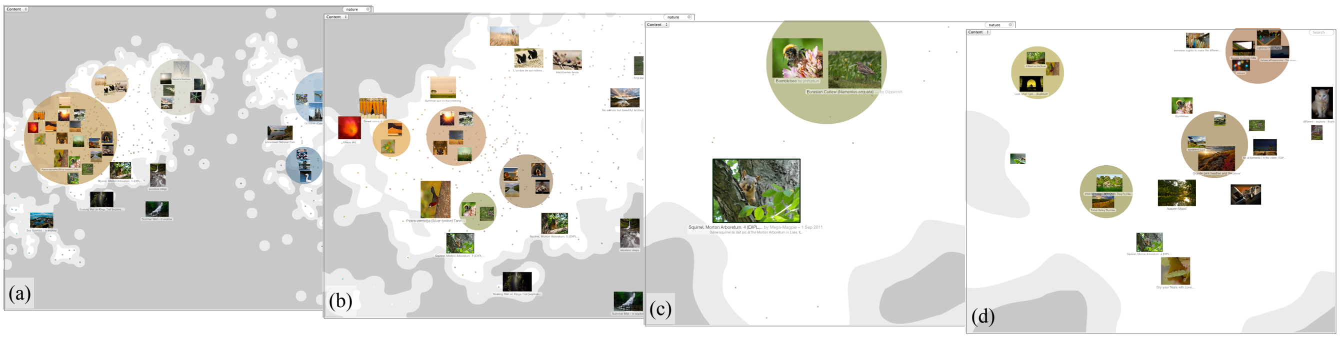

Figure 11: Photos are arranged by colour similarity. Having entered 'nature' as a query (a), the viewer gravitates towards large cluster on left (b), zooms into a particular photo (c), and views nearby photos (d).

lems, limitations, or challenges with Fluid Views, or similar visualization interfaces?" From our interview data, we extracted the main themes, which are discussed below.

6.1 Current Library Search Practices

When describing current information-seeking practices at the library, the two main themes were the book shelf and the catalog's search interface for exploring information.

The Diminishing Role of Shelves. Almost all participants referred to a recent move of the library collection into a high-density storage facility, which opened a few months before the interviews took place. According to several participants, this transition raised challenges for many library patrons, who perceived this change as taking away "the ability to experience serendipity [when] browsing the collection" [P1]. In comparison to shelves, which inherently associate similar items, participants described the library's catalog search interface as lacking a "sense of scope" or "physical sense of the collection" [P2]. Since many catalog items are stored off-site, scholars miss the ability to wander around in the collection [P4]. One participant viewed the information density of search interfaces as inferior to shelves [P3].

Search and Exploration. Most participants discussed different search strategies, and the role of result lists. One participant noted that many people start with a wide spectrum of resources, and then filter them. However, when paging through result lists, fatigue can terminate the exploration at an arbitrary point [P3]. Participants discussed the effect of lengthy linear result lists on the search experience, for example, when doing a subject search: "For the most part you need to go through page by page to really see the scope of what's showing up there. [...] It might be number 2, 4, 22 and then 50 – so you don't get a sense that it actually is a focused area of interest" [P6]. Most participants felt that known-item search was well supported by the library's tools. In contrast, there are "poor mechanisms for browsing" [P3]. That is, support is lacking for "people that aren't so goal driven, but really just want to be exposed to information, just want to see it" [P4].

6.2 Potential of Fluid Views

When asked about the possible role of Fluid Views for information seeking, participants noted several advantages for exposing relationships and refining searches.

Relationships. One of the recurring themes was how the Fluid Views interface exposes relationships that are difficult to see in conventional search interfaces. Participants noted that Fluid Views could help "getting a view of how much material is around" [P2] and to "see the 400 or 500 hits [and] how this is broken up" [P3]. Another participant contrasted visualization with conventional book shelving: "Library browsing has always been very linear, when it has been possible at all. Books stand beside each other, and they can only be at one place. So this breaks this up and would really encourage this cross-disciplinary view [...]" [P2]. Two others highlighted the fact that information seekers sometimes benefit from encountering information that they didn't know existed, possibly because it is outside of their domain, or using a different terminology [P1, P5]. In this respect, the Fluid Views interface gives a searcher "the ability in one view to get a sense of relationships" by showing the relevance of resources in related domains [P6].

Refinement. Besides the visual representation of resources, participants also noted the power of visualization for refining a search. For example, narrowing a given search with visualizations can be done "more quickly than trying to think of other search terms in your head" [P1]. In particular, concepts that are difficult to express in words, such as physical location or time spans, become possible since "visually it is very easy to capture" [P4]. Several participants noted that Fluid Views visualizations would also add value for "focused" searches [P6] by getting a sense of the results that are retrieved by a specific search. One participant summarized the dual value of Fluid Views as "a nice fusion between the flexibility of list based systems and getting closer to the kind of information density that [...] a shelf would provide" [P3].

6.3 Limitations and Challenges

The recurring concern for most participants was how Fluid Views would handle a large collection, like a library catalog. To address scalability, two participants suggested to start from a small set, such as a "filtered point" [P3] or a "curated collection" [P4]. In addition, participants, especially those concerned with content curation and technology development, mentioned incomplete metadata as a great challenge.

A range of suggestions made by participants involved integrating additional types of data. For example, two participants mentioned circulation activity as a potential source for item relevance. Two participants suggested augmenting the display with an indication of the availability of resources. Searchers who intend to check out a book would appreciate "a layer for what is available right now and right here" [P6].

7. CONCLUSION

Search interfaces offer increasing levels of access and accuracy, but often obscure relationships provided by traditional ways of organizing information. Considering that information seekers operate at varying levels of exploration and immersion [10], it can be beneficial to blend collection overviews and item details, and reduce the effort in shifting between the two. However, existing search interfaces expose limited numbers of results in relevance-ranked lists with little recognition of data similarities or relations. On the other hand, the key idea of information visualization is revealing patterns and relationships in large datasets. To enable a richer, more exploratory information-seeking experience, we combined search ranking and visual encoding techniques.

We have presented a search interface concept that integrates two layers representing items at varying levels of detail on top of a data facet. This dual-layer approach unifies information relationships and search by positioning items according to similarity and sizing to reflect their relevance. We have designed and built Fluid Views as a zoomable search environment with three base maps for time, location, and content. In summary, our contributions are:

– presenting search results spatially, with clickable access to details, providing an overview of larger groupings,

– providing a subliminal underlying layer that offers semantics for the positioning of the search results,

– integrating these layers in a search environment that combines dynamic queries and semantic zooming, and

– offering a functional interface for textual and visual resources along three facets (time, location, content).

To explore the potential of Fluid Views, we applied the technique to books, blogs, and photos and interviewed library professionals. Their feedback pointed out how Fluid Views was interesting to them because it addressed several current issues in library search. In their interest in whether Fluid Views could be deployed they raised questions about the scalability and extensibility of the Fluid Views approach.

8. ACKNOWLEDGEMENTS

We thank our colleagues from the InnoVis group and iLab at the University of Calgary for valuable advice. We are also grateful to the study participants and anonymous reviewers for their constructive feedback on this research. Funding was provided by SMART Technologies, NSERC, and AITF.

9. REFERENCES

[1] C. Ahlberg and B. Shneiderman. Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays. In *Proc. of ACM CHI*, pages 313–317, 1994.

[2] K. Andrews et al. The Infosky Visual Explorer: Exploiting Hierarchical Structure and Document Similarities. *Information Visualization*, 3(4):166–181, 2002.

[3] B. Bederson and J. Hollan. Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proc. of ACM UIST*, pages 17–26, 1994.

[4] K. Börner, C. Chen, and K. Boyack. Visualizing Knowledge Domains. *Annual Review of Information Science and Technology*, 37(1):179–255, January 2005.

[5] M. Chalmers. Using a Landscape Metaphor to Represent a Corpus of Documents. *Spatial Information Theory: A Theoretical Basis for GIS*, pages 377–390, 1993.

[6] E. Clarkson, K. Desai, and J. Foley. ResultMaps: Visualization for Search Interfaces. *IEEE TVCG*, 15(6):1057–1064, Nov/Dec 2009.

[7] A. Cockburn, A. Karlson, and B. Bederson. A Review of Overview+Detail, Zooming, and Focus+Context Interfaces. *ACM Computing Surveys*, 41(1):1–31, 2008.

[8] C. Collins, G. Penn, and S. Carpendale. Bubble Sets: Revealing Set Relations with Isocontours over Existing Visualizations. *IEEE TVCG*, 15(6):1009–1016, 2009.

[9] M. Dörk, S. Carpendale, C. Collins, and C. Williamson. VisGets: Coordinated Visualizations for Web-based Information Exploration and Discovery. *IEEE TVCG*, 14(6):1205–1212, Nov/Dec 2008.

[10] M. Dörk, S. Carpendale, and C. Williamson. The Information Flaneur: A Fresh Look at Information Seeking. In *Proc. of ACM CHI*, pages 1215–1224, 2011.

[11] J.-D. Fekete and C. Plaisant. Excentric labeling: dynamic neighborhood labeling for data visualization. In *Proc. of ACM CHI*, pages 512–519, 1999.

[12] G. Furnas and B. Bederson. Space-scale Diagrams: Understanding Multiscale Interfaces. In *Proc. of ACM CHI*, pages 234–241, 1995.

[13] S. Greene, G. Marchionini, C. Plaisant, and B. Shneiderman. Previews and Overviews in Digital Libraries: Designing Surrogates to Support Visual Information Seeking. *Journal of the American Society for Information Science*, 51(4):380–393, 2000.

[14] M. Hearst. TileBars: Visualization of Term Distribution Information in Full Text Information Access. In *Proc. of ACM CHI*, pages 59–66, 1995.

[15] K. Hornbæk and M. Hertzum. The Notion of Overview in Information Visualization. *International Journal of Human-Computer Studies*, 69(7-8):509–525, 2011.

[16] T. Igarashi and K. Hinckley. Speed-dependent Automatic Zooming for Browsing Large Documents. In *Proc. of ACM UIST*, pages 139–148, 2000.

[17] S. Jul and G. Furnas. Critical Zones in Desert Fog: Aids to Multiscale Navigation. In *Proc. of ACM UIST*, pages 97–106, 1998.

[18] B. Lee et al. FacetLens: Exposing Trends and Relationships to Support Sensemaking within Faceted Datasets. In *Proc. of ACM CHI*, pages 1293–1302, 2009.

[19] G. Marchionini. Exploratory search: From finding to understanding. *Comm. of the ACM*, 49(4):41–46, 2006.

[20] Microsoft. PivotViewer. `http://www.microsoft.com/silverlight/pivotviewer/` (Retrieved 2012-03-18).

[21] K. Perlin and D. Fox. Pad: An Alternative Approach to the Computer Interface. In *Proc. of ACM SIGGRAPH*, pages 57–64, 1993.

[22] B. Shneiderman. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proc. of IEEE Symposium on Visual Languages*, pages 336–343, 1996.

[23] B. Shneiderman and A. Aris. Network visualization by semantic substrates. *TVCG: Transactions on Visualization and Computer Graphics*, 12(5):733–740, Sep/Oct 2006.

[24] A. Spoerri. InfoCrystal: A Visual Tool for Information Retrieval. In *Proc. of IEEE VIS*, pages 150–157, 1993.

[25] M. Stefaner and B. Müller. Elastic Lists for Facet Browsers. In *Proc. of IEEE DEXA*, pages 217–221, 2007.

[26] J. Teevan et al. Visual Snippets: Summarizing Web Pages for Search and Revisitation. In *Proc. of ACM CHI*, pages 2023–2032, 2009.

[27] J. Wise et al. Visualizing the Non-Visual: Spatial Analysis and Interaction with Information from Text Documents. In *Proc. of IEEE InfoVis*, pages 51–58, 1995.