Interactive Visualization of Multiple Query Results

Susan Havre[1], Elizabeth Hetzler[1], Ken Perrine[1], Elizabeth Jurrus[1], Nancy Miller[2]
Battelle Pacific Northwest Division[1]
P.O. Box 999
Richland, WA 99352
(509) 375-6948
Battelle Memorial Institute[2]
505 King Ave.
Columbus, OH 43201-2693

[susan.havre,beth.hetzler,kenneth.perrine,elizabeth.jurrus]@pnl.gov,millern@battelle.org

Abstract

This paper introduces a graphical method for visually presenting and exploring the results of multiple queries simultaneously. This method allows a user to visually compare multiple query result sets, explore various combinations among the query result sets, and identify the "best" matches for combinations of multiple independent queries. This approach might also help users explore methods for progressively improving queries by visually comparing the improvement in result sets.

1. Introduction

The results of a query on a document collection are typically presented as an ordered list of document titles and/or a brief text summary. To rank the query results so that the most relevant documents are listed first, a search engine must calculate a measure of each document's relevancy to the query. This relevancy calculation may be the most computationally expensive step of the query process, and yet it is typically used only to order the returned documents. Some search engines show relevancy values to tell the user the degree of similarity of each document to the query. We believe the user may also find valuable information in the pattern of the relevance values—such as multiple documents with the same relevance value, a set of documents very relevant to the query, or a majority of documents very dissimilar to the query. However, the textual display of these values makes it difficult to understand such patterns.

This problem is further compounded when a user makes multiple related queries. The user may want to see not only the result set of the original query but also result sets based on components of the query. This would give the user some idea of the contribution of the query parts to the full query. Often a user will reformulate a query to try to improve the result set. Or, finding a relevant document, a user may want to submit a query by example. However, there are no good tools to compare multiple query result sets to evaluate the impact of query components or successive reformulations.

These needs can be solved by a query visualization tool. Such a tool may also be used to directly compare other types of query sets. For example, a researcher interested in evaluating relevancy measurement methods could compare the results of alternate methods on the same query and document set. While users may be aware of the variability of Internet search engine results, they cannot easily compare the results from different engines or even from the same engine on different dates. A visual approach could let them more directly compare and explore such results. As another example, a user might want to search a collection of images, where multiple queries would represent various ways of calculating similarity to a target image, such as similarity based on color, texture, orientation, etc. In some situations a combination of loosely or unrelated queries is needed. For example, lawyers preparing for a trial often seek case precedents. They look for these precedents based on multiple facets. If each facet is represented by a query, the visualization and interaction might help a user identify cases matching all or the majority of facets. Likewise, an inventor looking for patents similar to his invention might make a query for each key feature of his invention and search for existing work that matches the majority of his features.

In this paper, we present *Sparkler*, a portion of a prototype system that visualizes the relevancy of multiple queries to a collection of documents.

2. Approach

Sparkler is designed as an enhancement to a suite of text analysis tools with alternate views. For example, the tool suite can provide visualization of trends, overall similarity and clustering, and other metadata patterns. Sparkler allows a user to make multiple queries and compare the result sets within the integrated environment of other tools and views.

The visual portrayal we propose is independent of document/query representation method and of similarity measure. For this particular experiment, we use a vector space model to represent both the documents and the query statements, as described in [7]. For each part of the multi-query, we compare the vector representation of the query to the vector representation of each document in the collection. Numerous similarity functions have been proposed for such comparisons (for example, see [1]). Our prototype currently uses a simple Euclidean distance measure. The calculated distance/similarity between an individual query and all the documents in the collection comprise a ranked query result set. Our goal is to present multiple query result sets to the user in a simple, interactive way.

For portraying rank score in a query result set, various graphical display methods could be considered, such as icon size, color scale, and position. One of the most effective methods for perceptual comparison is position on a line [1], [2], [6]. In our proof-of-concept prototype, we use this method to present similarity between a query and the documents. We represent the query as a triangular icon at one end of an invisible line. In the simplest version, the documents in the collection are shown as small dots arranged along the line (see Figure 1a) where their distance from the query icon reflects their relevance.

This presentation provides a hint at the distribution of the results set, but overplotting obscures some document icons and hides information about the result set as a whole. By using a second dimension to spread the icons, the distribution pattern becomes more obvious as shown in Figure 1b. Icons of documents with similar relevance values are positioned the same distance from the query icon but spread along an arc centered on the line in a variation of a simple histogram. The spreading allows access to individual icons that are occluded in the one-dimensional variation. The prototype allows the user to control both the spacing of the radial spread and the granularity of the histogram bins.

To see multiple queries simultaneously, we place the query icons in the center of the graphic with the linear arrangement of document icons for each query radiating outward forming a sparkler. Each query and its document icons form an arm of the sparkler. Using this structure, the icons for documents closest to each query are closest to the center of the display. We call these *inner icons*. The icons of documents least similar to a query are located near the farther end of the query's arm. We call these the *outer icons*.

Figure 2 shows a zoomed-in view of two expanded sparkler arms, corresponding to Queries 1 and 2. The cluster of query icons appears near the upper left of the image. The sparkler arms radiate outward from the query icons. The query indices appear clockwise of the query and result icons. Each expanded sparkler arm reveals a distribution pattern of the document icons; we call this pattern the *result profile*. For brevity, we refer to the "result profile for Query X" as "Profile X".

The two result profiles in Figure 2 differ slightly in position and shape. Profile 1's innermost icon is closer to its query icon than Profile 2's. However, the bulk of Profile 2 is closer to its query icon than Profile 1. The

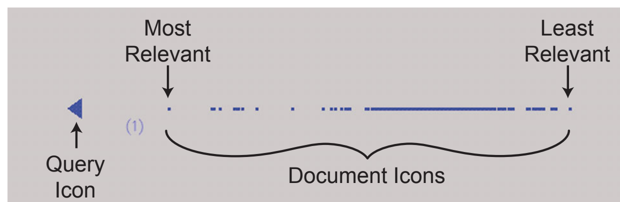

a.

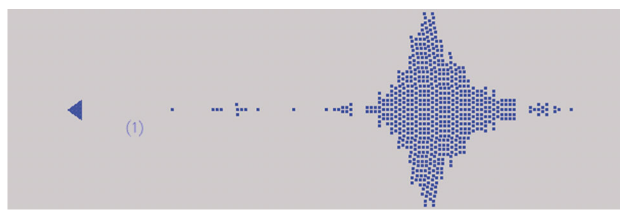

b.

Figure 1. Document icons are plotted in relation to the query icon (a) overplotted along the invisible query line and (b) spread radially

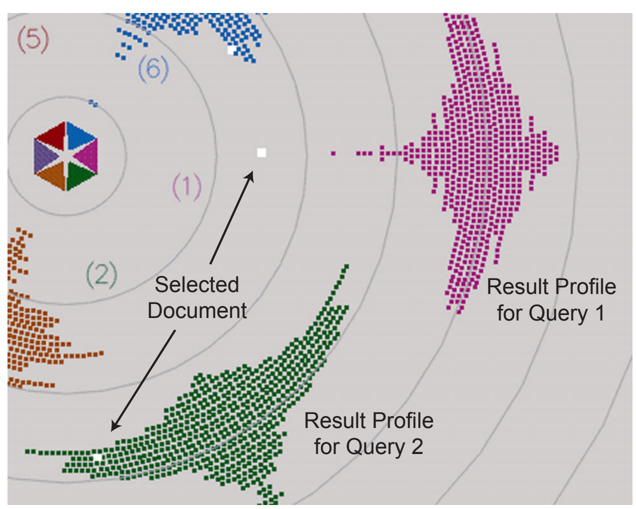

Figure 2. Zoomed-in view of two expanded result profiles and one selected document

shorter distance between the inner and outer icons of Profile 2 indicate less variation in relevance. When a document icon is selected within a profile, all icons across the profiles for that document are highlighted in white. In Figure 2, we have selected the innermost icon of Profile 1. The icon for that document is also highlighted in Profile 2; the document is not as closely related to Query 2. Circular grid lines make it easy to compare distances.

Figure 3 is a Sparkler visualization of 568 visualization patents. The queries comprise six concept phrases from a potential patent. A user might want to use such a set of queries to search for existing work, called "prior art," before submitting a patent application. The six result profiles are arranged in clockwise order starting from the positive x-axis. By default, the scale of the profiles is calculated for the combination query set: we find the closest match among all profiles and place the corresponding document icon at a fixed position, and then use that distance to scale all the other positions. For comparison among multiple query profiles, we would want to make the scale more explicit, fix it across displays, or provide a controllable "shrinkage" mechanism.

We find it helpful to apply a log function to spread the icon positions outward, reducing the occlusion of icons closest to the center. This spreading emphasizes gaps defining icon groups of similarly relevant documents. In Figure 3, the inner icons of Profile 6 appear separated into groups of distinct levels of relevance. This pattern suggests likely break points for selecting groups of similarly relevant documents rather than an arbitrary fixed number of most relevant documents. In usability tests, users demonstrated an intuitive understanding of this.

The prototype shows the title of the last selected document at the bottom of the display (not shown in the figures). The document text can be retrieved on demand.

For combining multiple query results, we define a special "bull's-eye" selection mode. By clicking anywhere in the display, the user defines the radius of a bull's-eye, shown as a red circle in Figure 3. The system automatically selects all documents that satisfy the bull's-eye selection rule. The default selection rule requires that a document must fall inside the bull's-eye ring on each result profile. The user can loosen the rule, for example, so that a document will be selected if it falls inside the ring for at least three profiles. The user can disable profiles so the bull's-eye selection rule no longer applies to them. When a profile is disabled, the query index number has a slash across it and the unselected icons along the profile darken in color. Selected documents in disabled profiles are still highlighted. In Figure 3, Profile 5 has been disabled and the bull's-eye rule set so that any document with icons inside the ring for at least three profiles is selected. Thus, any patent with high-similarity value to at least three of our query concepts except Query 5 (which is disabled) is highlighted.

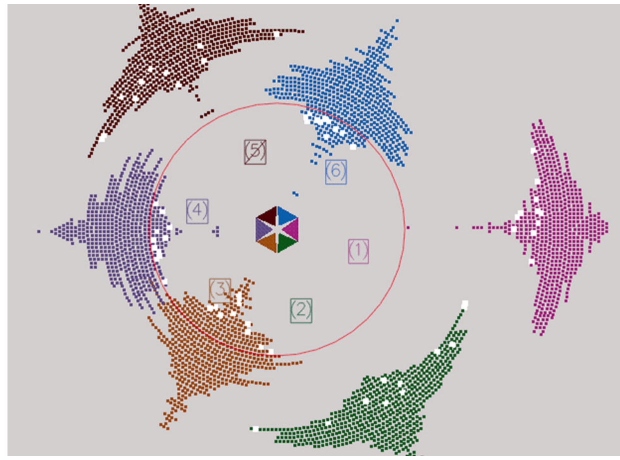

Figure 3. Result profiles of six queries with some documents selected by the "bull's-eye"

3. Related Work

Many visualizations leverage the benefits of circular displays as an effective method for portraying multi-dimensional data. For example, spider graphs [14], kiviat diagrams, multiple superimposed star plots (sometimes called circular parallel coordinates), Neighborhood Explorer [10], and Starstruck [5] all bear visual similarity to Sparkler. These displays seek either to portray similarity among multi-dimensional data objects, so that the spatial pattern of the circular positions is important, or to portray performance along a number of measures, where each arm represents a particular measure and the goal is to optimize the combination of measures. The unique contribution of this paper is twofold: first, by applying this kind of presentation to query results, we not only provide additional information on a single query result set but also enable explicit comparison among the multiple query result sets. Second, applying this display method to this kind of problem lets us leverage key perceptual characteristics. For example, by placing the query icons in the center of the display and explicitly showing relationships among the collection and each query, we are using both proximity and line position to show vector-based similarity. Also, the metaphor lets us place the most important items in the center of the display; if a user focuses on the inner icons of one arm, the inner icons of all other arms are near the center of the user's visual field.

We certainly are not the first to propose visualizing multiple query result sets. Veerasamy & Belkin [13] propose an approach that shows the strength of each individual query term for each retrieved document in a tabular layout with documents ordered by overall rank. The key words or phrases of the query are arranged from top to bottom on the y-axis. Along the vertical column for each document, bars are sized to show the strength of the query words or phrases in that document.

TileBars [3] uses a similar philosophy at a finer level of detail to show relevance to multiple term sets within the narrative flow of documents. Each document is algorithmically partitioned into a number of sections. Horizontal bars represent individual documents, with bar length indicating the document's length and partitions explicitly shown. Each term set in the query is represented as a row; for example, for three queries, there are three rows in each bar. Tiles along each row are colored from white (low relevance) to black (high relevance); by scanning these tiles, the user can identify the document passages most relevant to term combinations in the query.

In VIBE [9], query terms or term sets become Points of Interest, represented by icons in a 2D visualization. The retrieved documents are placed within the convex hull created by these Points of Interest, using a vector sum algorithm. Users can move the Points of Interest and see which documents respond to resolve ambiguities. WebVIBE, based on the VIBE approach but with a magnet pull metaphor, has been applied to the issue of comparing Internet search engines [8]. VIBE has some similarities to our approach: calculating relevance to each query part and placing the documents most relevant to all queries in the center of the image. The Sparkler visualization differs in that the relationship between each document and each query is explicit in the presentation, while VIBE represents the combination of similarity calculations for each document. Also, in VIBE, a document related to only one Point of Interest is placed near the edge of the display, while in Sparkler, the document is close to the center for the query to which it relates, even if it is an outer icon for the other result profiles.

Lyberworld [4] places queries on the outside of a Relevance Sphere. Query results are positioned in the sphere to represent the combination of attractions between the document and the various queries. The user can move the query locations along the sphere surface to resolve ambiguities.

InfoCrystal [11] is a visual multi-query approach that can be applied to both Boolean and vector-based queries. The queries are represented as points on a regular polygon (triangle for three queries, square for four, etc.) For Boolean queries, a set of icons is generated inside the polygon that represents all possible Boolean combinations of the queries. The icon shape represents the number of queries included and its location indicates which queries are included. The user can then select the appropriate icons to generate queries and group such polygons to create complex combinations. A variation applies InfoCrystal to vector-based queries. Again, a given document's location within the polygon is determined by a combination of its relevance to the query documents. There is no explicit representation of the individual relevance scores.

4. Comparing Query Results to a Truth Set

This section explores multiple query results using a standardized document collection from the third Text Retrieval Conference (TREC-3). In addition to multiple document collections, TREC-3 provides a set of query topics and assigns relevance judgments that declare each document relevant, or not relevant, to each topic. For our experiment, we select one collection and one topic using the relevance judgments to define a truth set. We then compare the multiple query result sets to each other and to the truth set.

The corpus selected is the Associated Press (AP) 1990 corpus. We isolate news stories from a 7-day period from June 24 to June 30 and select TREC topic 108 (Japan Protectionist Measures). The relevance judgments for topic 108 during this period define a set of 17 relevant documents out of the 2116 total documents.

4.1 Multi-query formulation

For our study, we formulate several alternate query statements based on information provided with TREC Topic 108. Topic 108's title, summary, narrative, and the first seven concepts are listed below:

Topic 108 elements:

1. Japanese Protectionist Measures (Topic 108 Title)
2. Japanese policies or practices which help protect Japan's domestic market from foreign competition (Topic 108 Summary)
3. Japanese law or regulation a governmental policy or administrative procedure a corporate custom or a business practice which discourages or even prevents entry into the Japanese market by foreign goods and services A document which reports generally on market penetration difficulties but which does not identify a specific Japanese barrier to trade is NOT relevant (Topic 108 narrative)
4. Japan (Topic 108, first concept)
5. Ministry of International Trade and Industry MITI Ministry of Foreign Affairs
6. protectionism protect
7. tariff subsidy quota dumping obstruction retaliation
8. structural impediment product standard
9. trade dispute barrier tension imbalance practice
10. market access free trade liberalize reciprocity

From these we select a set of words and create queries from the individual words and various combinations that we guess might be useful:

1. Japan
2. trade
3. barrier
4. protection
5. measures
6. Japan trade
7. Japan barrier
8. Japan protection
9. Japan trade barrier
10. Japan trade protection
11. Japan trade measure
12. Japan protection measure

Figure 4 shows the sparkler for our 12 queries. The result profiles were calculated according to the approach described in Section 2. Profile 1, "Japan," is positioned along the positive x-axis. The other query arms are arranged clockwise around the sparkler center.

4.2 Initial observations

The result profiles vary. The profiles for the three-word queries (9–12) are all close to their query icons, indicating that several documents are similar to the queries. Profile 3 is flat and wide, while Profile 10 is elongated and narrow. The positions of the result profiles also vary. Profile 3 is positioned much farther from the query icons than other profiles. Profiles 10 and 12 are close to their query icons.

One might expect that, in general, the result profiles of the single-word queries would be positioned further from the center than those of multiple-word queries. Certainly, the position of the "barrier" profile (3) meets that expectation. It appears that "barrier" is weakly related to the document collection. However, the "trade" profile (2) is surprisingly close to its query icon indicating that trade is highly relevant in some of the news stories.

The two-word "Japan barrier" profile (7) is not very close to its query icon. In fact, it appears to be a reflection of the "Japan" profile (1). This seems plausible because we know from Profile 3 that "barrier" does not have much similarity to this particular document collection.

4.3 Exploration

To start our exploration, we switch to bull's-eye mode and focus on the five single-word queries. Noting that Query 2 ("trade") has profile icons closest to their query

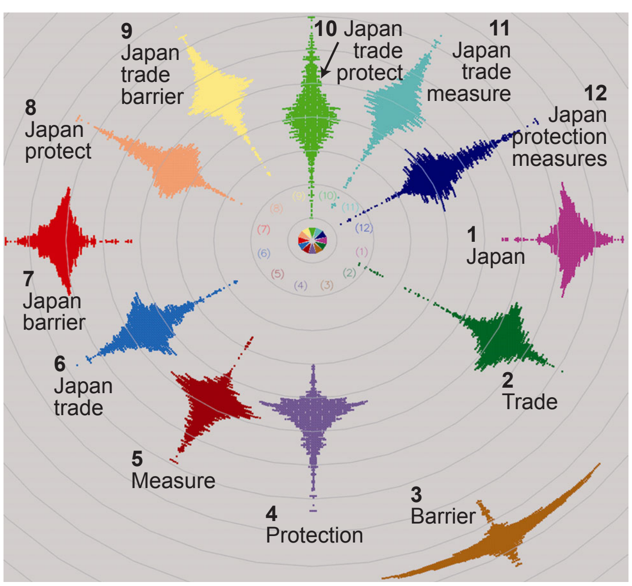

Figure 4. Queries and the respective result profiles for the AP 1990 corpus

icon, we disable all but the "trade" arm. Figure 5 shows the bull's-eye dropped on a wide gap. The 13 selected icons represent documents most relevant to "trade."

Looking at the location of the selected icons in the other profiles, we see they are dispersed. In the "Japan" profile (1), these documents appear among the outer icons. This is true for all result sets for queries not containing the word "trade." Looking at the document titles and text, we see the highly relevant stories cover trading in the futures market for soybeans, corn, sugar, and oil.

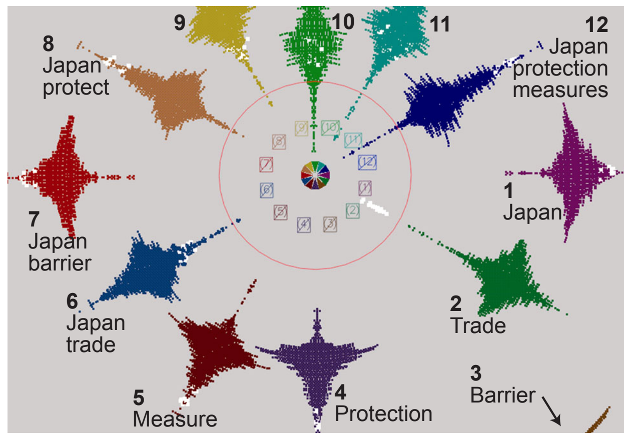

Figure 5. We select document icons inside the ring along the Query 2 result profile. These same documents are also highlighted in the other sets.

Because we are looking for documents specifically related to "Japan," we decide to look at that result profile instead. We disable Query 3 and enable Query 1, dropping the bull's-eye at the first natural gap and capturing seven documents as shown in Figure 6 on the next page. The selected news article titles are listed here, ranked by the order they occur in the Query 1 result profile:

1. Japan, U.S. Seek Final Agreement on Structural Reforms
2. Structural Trade Talks Extended
3. Trade Talks Expected to Continue into Wednesday
4. Japan, U.S. Trade Negotiators Begin Final Round of "Structural", 0265
5. Japan's Public Spending to Overshadow Trade Talks
6. Japan, U.S. Trade Negotiators Begin Final Round of "Structural", 0282
7. Negotiators on Final Lap of Japan-US Trade Talks

These 7 documents were among the 17 judged relevant by TREC. We cannot evaluate the ordering since documents were marked only relevant or non-relevant. These documents also occur as the most relevant to Queries 5 through 12. They group naturally—distinct from the next closest document in several profiles. If the purpose of our search is to retrieve just the most relevant documents, this might be a place to stop the exploration.

Because we know there are 17 relevant documents, we continue by dropping the bull's-eye at the next gap in

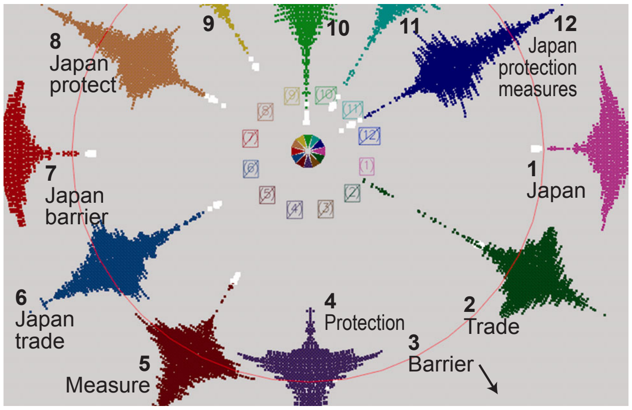

Figure 6. We select document icons inside the ring along the Query 1 profile, placing the ring at the first natural gap in the "Japan" profile

Profile 1, "Japan" (see Figure 7); 24 documents are selected. We have tight grouping of these documents along the inner icons of several profiles. Profiles 6, 7, 11, and 12 have the best homogeneity—that is, the least mixing of selected and unselected documents. The 24 selected documents include all but one of the TREC relevant documents.

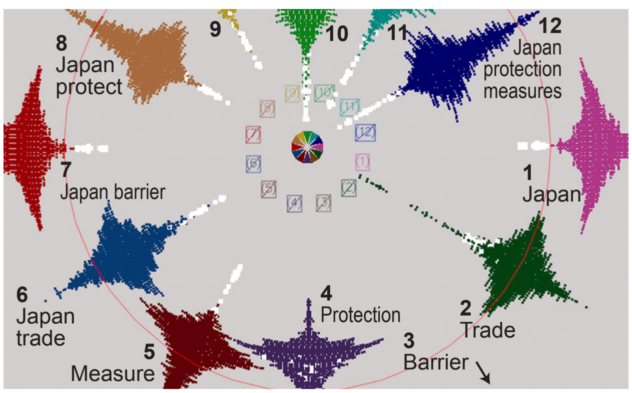

Figure 7. We extend the ring to the second gap in the "Japan" profile, selecting 24 documents

For most search purposes, we would likely stop at this point. We recognize that the narrative for topic 108 would make it hard to exactly match the truth set. The narrative states that a document is not relevant unless it identifies a specific trade barrier, a subtlety not recognized by our method.

A user might want to examine outliers in the selection set for additional insights. The connecting web in Figure 8 shows consensus across several profiles on the weaker similarity of some outliers. The user could decide to remove these from the selection set or to explore them for possible interesting material or additional query concepts.

This example demonstrates several interesting points about the information conveyed by the visualization.

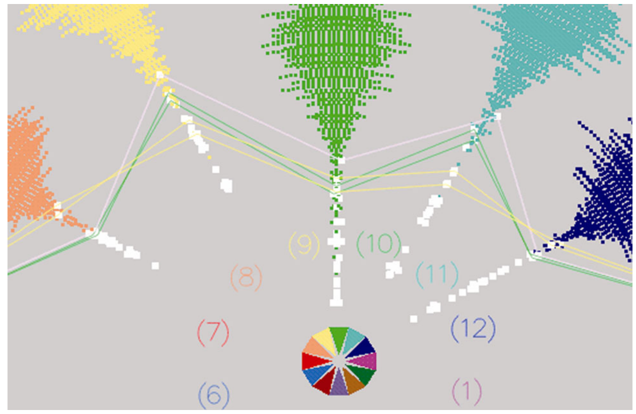

Figure 8. The webs connecting the five furthest outliers for Profile 12 show that the same documents are also the outliers in other profiles

First, there is good consensus on the relevant documents across all the three-word queries (9–12), the two-word queries (6–8), and two of the one-word queries (1, 5). In addition, even though the results for the "Japan" query are overall further from the center (indicating a lower relevance score), a ranked list of them would provide relatively consistent results with the three-word queries for this particular example. As another observation, consider the case where a user might start with a three-word or longer query and have the system automatically not only run the combination (e.g., Query 10) but also present results for the individual components (e.g., Queries 1, 2, and 4). This might tell the user that there are other documents in the collection strongly related to Japan and trade and to protection, including protection of the spotted owl, and workers in the mining and timber industries. If these are interesting to the user, he or she might proceed to investigate these topics as well. This kind of serendipitous exploration has been noted as desirable; however, as systems seek to make query results best reflect a user's need, enabling serendipitous discoveries in an unobtrusive manner is a challenge [12].

5. Comparison of Internet Search Results

An alternate use for this visual presentation is to compare results from different search methods, different vectorizations, different distance calculation methods, etc. To illustrate this idea, we present an example using a simple search with Internet search engines. We submitted the words *automatic summary* to five search engines: AltaVista, Excite, Go, Google, and Lycos. We captured the top 30 results for each engine, pooled these documents, and then eliminated duplicate results (that is, identical files, regardless of URL). In the end, there were five result lists of 29, 22, 18, 30, 27 URLs, respectively (often a single search engine returned duplicates). We call these the "top" returns.

Because not all search engines provide relevance scores, we used the return order to establish rank and applied a constant offset to separate the positions of document icons along each query arm. Where duplicates occurred in a list, the top-ranked occurrence was left in place and the lower ranked one eliminated. For each arm, the remaining documents (those not in the top returns) were assigned to an arbitrary location at the outer end of the arm.

In comparing multiple search engines, a user might want to find all documents consistently appearing in the top returns. If we use the bull's-eye capability to identify documents selected by all five search engines, only one document qualifies (Figure 9). It discusses a recent workshop on automatic summarization.

An alternate comparison might seek the set of URLs in the top set of two or more engines. We used the bull's-eye to select all documents on at least two arms; Figure 10 shows the result. By looking at the distribution, we can see that the top-ranked five or six documents for Excite, Go, and Google are included in the top 30 of at least one other engine. AltaVista's top nine do not appear in any other top 30. Only two out of Lycos' top five returns are confirmed by another engine.

6. Usability Testing

We conducted an informal usability evaluation with two male and two female users, varying in age from under 30 to over 45. One user was colorblind. None of the users were familiar with multi-query visualizations, including the Sparkler visualization; one was familiar with vector-based queries in general. The version tested displayed document icons with overplotting (Figure 1a).

Each user received a brief demonstration of the prototype features and interactions using data and queries entirely different from the evaluation sets. Each user was encouraged to interact with the Sparkler prototype and the demonstration data until he or she felt comfortable with it. The users were also given ordered lists of titles corresponding to the results displayed in the sparkler.

We asked each user to perform a series of simple tasks, designed to check how well they understood the prototype. All correctly used the proximity along the sparkler arms to identify document relatedness to a query and compare multiple query results. They also correctly used the bull's-eye capability to select documents that are equally related to all active queries. However, they were not able to use a document's position on the various sparkler arms to predict the document subject matter. When asked to select highly relevant documents, they used the spacing breaks along the query arms as cutoffs, although they did not articulate this as their reasoning.

In answering a questionnaire, three of four participants felt that Sparkler helped them understand the relatedness of documents to a single query and to multiple queries. They felt the visualization added to their understanding of a ranked list of titles and indicated that they would likely use it. The fourth user found it difficult to judge relative distances to the center, even with the aid of a grid overlay and the bull's-eye.

All were interested in the visible differences between the individual concept queries and the combined query. They were also interested in comparing results, for example, between the combined query and the bull's-eye with different combinations of the individual concepts. Most were interested in exploring apparent outliers to understand why they were or were not part of a result set. The users suggested that the sparkler would let them come up with better query results through improved queries.

We ran a second formative usability test with two professional information analysts, whose jobs focus strongly on finding and analyzing information from large

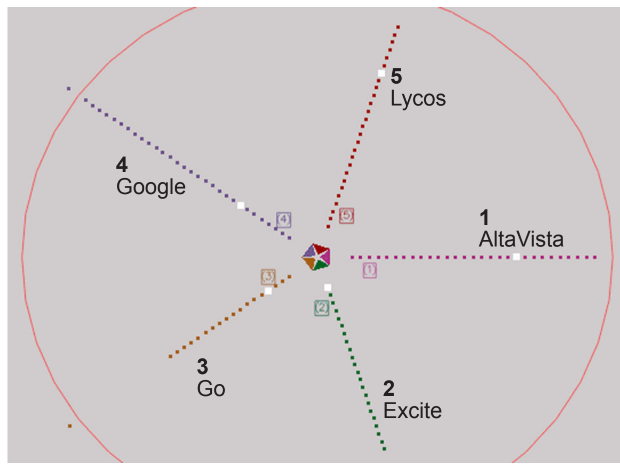

Figure 9. Only one document appeared on the top-30 lists of all search engines

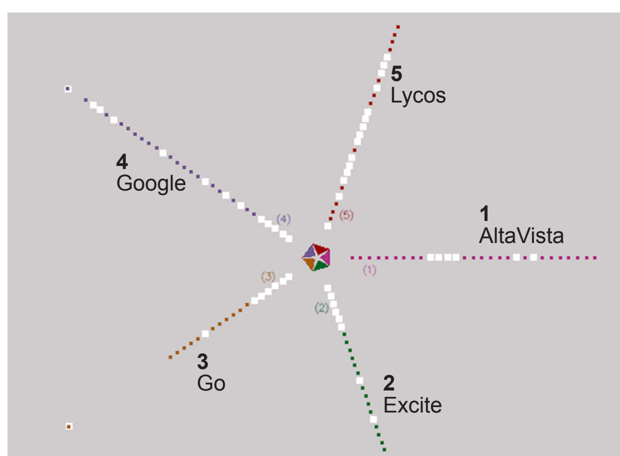

Figure 10: Documents that fall within the bull's-eye for at least two queries are distributed along each arm in distinct patterns

databases of text. They each had extensive experience with Boolean queries and less experience with ranked vector-based queries. Both users took some time to fully understand the sparkler and its possible implications. They cited several possible uses, such as suggesting new queries (together with a title list), combining queries, and interpreting relevance of results. The analysts found it useful to see the sparkler in conjunction with a similarity-based dot plot display, and a list of titles. They provided several suggestions, including a progressive disclosure capability. Both strongly preferred the result profiles to the linear display with overplotting, because they showed where the bulk of the relevancy scores fell, and suggested cutoff points for relevant documents. Overall, they found the sparkler added significantly to their understanding of a ranked query list and indicated that they would use it if it were integrated into an analysis tool. Both analysts were interested in the relationships between the single-word queries and the multi-word queries, and the differences in document rank produced by adding more query terms. They observed that they were gaining insight into how the query engine performs.

7. Conclusions and Future Work

We have introduced a visualization of query results where the distance between icons representing the query and one or more documents is based on a similarity/ relevance measure. This visualization can depict multiple query results in a single display.

The Sparkler visualization allows the user to easily see the ranked order of relevance of documents to a query. In addition, the distribution of the icons forms a relevance pattern that may reveal useful information. For instance, the natural gaps and distribution profile along an arm suggest ways to partition the result set.

The visualization can be applied to several classes of user problems, such as finding documents related to several distinct but related queries, exploring various combinations of queries, progressively improving queries, and comparing alternate distance/similarity measures or alternate document/query representations.

Our plans for future work focus first on a number of user interaction improvements, such as adding the capability to add/delete/edit queries interactively, to reorder them, and to allow more user control over the sparkler spacing options. We need to improve usability by making the query text more accessible. For this paper, we have supplied query titles near the ends of the arms. We need to implement this capability in the prototype.

We would like to further explore the visualization to better understand its utility for various query tasks. We want to understand scalability implications, such as knowing the maximum number of queries a user can handle perceptually. Finally, we would like to apply Sparkler to queries in other media, such as images.

8. Acknowledgments

We gratefully acknowledge the contributions of our colleagues at Battelle to the development and testing of the Sparkler visualization. Special thanks go to Jim Thomas, Lou Martucci, Lucy Nowell, Harlan Foote, Renie McVeety, Rik Littlefield, Sharon Eaton, and Wanda Mar. Battelle’s Information Technology Platform funded this research.

9. References

[1] Carswell, C. Melody. (1992) “Choosing Specifiers: An Evaluation of the Basic Tasks Model of Graphical Perception”, *Human Factors*, 34(5), 535-554.

[2] Cleveland, W. S. and McGill, R. (1985) “Graphical Perception and Graphical Methods for Analyzing Scientific Data”, *Science*, 229(August), 828-833.

[3] Hearst, M. (1995) “TileBars: Visualization of Term Distribution Information in Full Text Information Access”, in *Proc. ACM SIGCHI ’95*, 59-66.

[4] Hemmje, M., Kunkel, C., and Willett, A. (1994) “LyberWorld – a Visualization User Interface Supporting Fulltext Retrieval”, in *Proc. Seventeenth Annual International ACM SIGIR Conference* (TREC-1), 1-20.

[5] Hetzler, B. Whitney, P., Martucci,, L., Thomas, J. “Multifaceted Insight Through Interoperable Visual Information Analysis Paradigms”, in *Proc. IEEE Symposium on Information Visualization*, InfoVis '98, October 19-20, 1998, Research Triangle Park, North Carolina, 137-144.

[6] Mackinlay, J. (1986) “Automating the Design of Graphical Presentations of Relational Information”, *Transactions on Graphics*, 5(2), 110-141.

[7] Miller, N., Wong, P, Brewster, M, & Foote, H. (1998) “TOPIC ISLANDS™ - A Wavelet-Based Text Visualization System”, in *Proc. IEEE Visualization ‘98*, ACM Press, 189-196.

[8] Morse, E. *Evaluation of Visual Information Browsing Displays*, PhD Dissertation, Univ. of Pittsburgh, 1999.

[9] Olsen, K. A., Korfhage, R. R., Sochats, K. M., Spring, M. B., & Williams, J. G. (1993). “Visualization of a Document Collection: the VIBE System”, *Information Processing & Management*, 29(1), 69-81.

[10] Spence, Robert. (2001) *Information Visualization*. ACM Press, New York, 85-88.

[11] Spoerri, A. (1993). “InfoCrystal: A Visual Tool for Information Retrieval” in *Proc. Visualization '93*, IEEE Computer Society Press, Los Alamitos, 150-157.

[12] Toms, E. (1998). "Information Exploration of the Third Kind: The Concept of Chance Encounters." Presented at Information Exploration Workshop for SIGCHI '98.

[13] Veerasamy, A. and Belkin, N. “Evaluation of a Tool for Visualization of Information Retrieval Results”, in *Proc. SIGIR ’96*, 85-92.

[14] Wilkinson, L. *The Grammar of Graphics*, Springer-Verlag, New York, 1999.

[15] Zobel, J. (1998) “How Reliable Are the Results of Large-Scale Information Retrieval Experiments?,” *Proc. SIGIR ’98*, 307-314.