Scalable Visual Reasoning: Supporting Collaboration through Distributed Analysis

William A. Pike, Richard May, Bob Baddeley, Roderick Riensche, Joe Bruce, and Katarina Younkin
Pacific Northwest National Laboratory
{william.pike, richard.may, bob.baddeley, rmr, joseph.bruce, katarina.younkin}@pnl.gov

ABSTRACT

We present a visualization environment called the Scalable Reasoning System (SRS) that provides a suite of tools for the collection, analysis, and dissemination of reasoning products. This environment is designed to function across multiple platforms, bringing the display of visual information and the capture of reasoning associated with that information to both mobile and desktop clients. The service-oriented architecture of SRS facilitates collaboration and interaction between users regardless of their location or platform. Visualization services allow data processing to be centralized and analysis results to be collected from distributed clients in real time. We use the concept of "reasoning artifacts" to capture the analytic value attached to individual pieces of information and collections thereof, helping to fuse the foraging and sense-making loops in information analysis. Reasoning structures composed of these artifacts can be shared across platforms while maintaining references to the analytic activity (such as interactive visualization) that produced them.

KEYWORDS: analytical reasoning; knowledge management; visual analytics; mobile applications; service-oriented architecture.

1. INTRODUCTION

The benefit of visualization as a knowledge discovery aid is well established. As a means for rapidly generating insight from large data collections, visual tools can be indispensable. Yet a central problem in visualization remains understanding – or even extending – the place of visual depictions in the broader knowledge construction and communication process. The recent identification of a field of "visual analytics" indicates a growing recognition that visualization involves not just the display of information, but the entire process of information discovery, analysis, and dissemination.

A number of models have been created to describe the process of information analysis [1, 2]. These models commonly include a foraging loop and a sense-making loop. The foraging loop involves the gathering and processing of data. The sense-making loop includes the development of understanding about the data and the creation of a finished analytic product. Each loop can contain many cycles and need not be mutually exclusive – advancing to the sense-making loop does not preclude a return to foraging.

Another model [3] uses three states (explore, enrich, exploit) to describe the technique used by analysts to find high-value information artifacts The exploration process is a widening of the collected material in an effort to ensure high-value artifacts have been brought into the analytical space. The analytical space is then enriched by weeding out the items that are deemed of lower value. After the analytical space has been sufficiently reduced, the high-value material is exploited (read and analyzed). As with the foraging and sense-making loops, these three states are not mutually exclusive. It is routine for an analyst to iteratively cycle between the states while performing an analysis.

Tools designed to help the analytical process must ensure that the high-value artifacts are captured and retained through to the "exploit" phase. However, in contemporary practice there is typically a divide between the tools that analysts use to gather information (such as search engines), tools to organize information (such as web browser bookmarks or text documents), tools to explore information (such as text visualization), and tools to record findings (such as concept maps or slide presentations). When using visualization to explore data, there can be a particular disconnect between the cognitive process of discovery and the ability of the tool to record these discoveries so they can be shared, evaluated, and revised.

0-9785699-1-1/07/$25.00 ©2007 IEEE.

The work described here endeavors to create a visual analytic environment that situates visualization within a collaborative knowledge construction workflow, easing the process of bringing new information into visual tools and recording the knowledge that results. The collaborative visual environment is designed to bridge the assembly, analysis, and dissemination of information and to support these activities across the range of client devices information that analysts use. This environment, the Scalable Reasoning System (SRS), is a toolkit that provides teams with a graphical workflow mechanism for recording evidence, assumptions, hypothesis, and other reasoning structures. Integrated into this workflow are dynamic visualization components for exploring user-created text data sets *in situ*, allowing the insight generated from visual exploration of data to be directly integrated into an emerging problem-solving process. These visualizations can be created on both desktop and mobile clients. Reasoning structures created in SRS can be shared with other users on other devices – either in real-time or asynchronously – and converted into reports that maintain references to the process of visual discovery that created them.

2. RELATED WORK

There is growing recognition in the fields of visualization and Computer-Supported Cooperative Work that collaborative tools must actively help their users construct and share knowledge. MacEachren [4] identifies three roles for visualization in support of collaborative activity: visualization can be the object of collaboration (as might be the case with shared maps), it can support dialogue, and it can support coordinated activity. Frequently, the role of visualization in collaborative sense-making is limited to its ability to diagrammatically represent reasoning – these visualizations can variously play roles as dialogue support tools and, in the sense that their representation of this dialogue helps further it, as objects of collaboration. Visual argumentation languages such as that in ClaiMaker [5], for instance, can provide structured vocabularies for articulating claims and mechanisms to annotate data so as to encode relations between text documents and the concepts expressed therein. Compendium [6] and Codex [7] follow more "unstructured" models, and are able to capture reasoning discourses – both personal and collective – in the contexts in which they occur through dialogue mapping that accommodates a wider array of concept and relationship types. The Oculus Sandbox [8] is an analytical sense-making system that combines computational linguistic and analytical functions. In [9], analysts can record individual perspectives using a graph-based interface and can fuse multiple perspectives to find areas of agreement or uncertainty. Suthers [10] demonstrates that these graph-based visual representations of reasoning processes can effectively facilitate mutual knowledge construction.

Our intent with SRS, however, is to capture knowledge *in situ*; that is, during the course of its construction. The goal is not just to support dialogue, but to create a space for collaborative discovery. To this end, SRS merges visual argumentation languages like the above with data visualization interfaces, bringing convergence to the typically separate acts of analysis and argumentation. Personalized knowledge maps [11] have been proposed as one mechanism for linking structure observed in data to structure inferred by human analysts; these maps provide an example of how data-driven visualizations and user-driven knowledge exposition can interact, but are not designed to support capture of the discovery and reasoning process over time and across multiple visualizations. On the other hand, methods for capturing entire work sessions (e.g., [12]) allow analysts' work practices to be effectively stored, shared, and replayed, but at too low a level (typically at the scale of individual application events) to be meaningful in a cooperative knowledge construction task. Moreover, few extant methods scale across platforms so cannot support distributed teams that include field personnel collecting and analyzing data in place.

The fields of knowledge representation and cognitive science offer a supporting scaffold for our efforts to integrate analysis and exposition. Knowledge management tools, of which visual reasoning aids are a class, should reflect the situated work practices of their users [13] and accommodate the dialogical, manipulative nature of exploration [14, 15]. Communities of practice need tools that help them test, refine, and implement emergent, experience-based, solutions [16]. Indeed, others have recognized a need for what might be called "creativity support systems" – tools to help assemble information, analyze it, collaborate with others, and distribute resulting products [17]. (Further detail on the theoretical motivation behind our work can be found in [18]). Empirical studies of such collaborative inquiry in investigative settings, notably [19], have indicated user desire for (a) the ability to link information artifacts to the reasoning processes in which they figure, as an aid to keeping a community informed about the state of its knowledge, (b) the ability to generate "big picture" reports, and (c) minimizing ontological complexity. Some extant tools accomplish a few of these objectives: [20] provides a utility for describing observations, beliefs, and uncertainty associated with online documents using semi-formal semantic markup, and [21, 22] capture analytic events in context – although this context is not yet used as a reasoning aid.

3. BUILDING A SCALABLE REASONING SYSTEM

Our Scalable Reasoning System is designed to support collaborative inquiry by combining collection of source evidence, visual argumentation, and text visualization. The primary components of SRS are outlined below, with emphasis on how each contributes to the goal of fusing the foraging and sense-making loops into a single environment that supports the recursive construction of knowledge across devices and users.

3.1. Service-oriented visualization and analysis

Central to supporting collaborative visual analysis in distributed teams is the need to accommodate the multiple interfaces and interaction techniques each team member might use. For instance, a team of SRS users might include operations center personnel using large-screen displays, field personnel using mobile devices such as PDAs or smartphones, and decision makers or other stakeholders who receive reports from these personnel and need the ability to understand and evaluate the team's findings.

Rather than build special-purpose visual tools for multiple platforms, we take the approach of creating visualization and analysis *services* that separate computationally intensive data processing from the display of resulting visual depictions. These services also manage the collection of new information from the field, providing a single point of entry for getting data onto a visual display, and provide a common user experience across devices.

In our current implementation, a centralized data store represents each data artifact (described in more detail in the following section) as an XML document and associated schema; in this fashion data from multiple sources with different schemas may be stored in the same repository. Applications interact with this data via abstractions of two key elements of the SRS visual analytics sub-system: "Harvests" and "Bubbles." Harvests represent collections of artifacts that have been passed through a text-processing algorithm to identify key topics and relations among them. A harvest process includes:

- Parsing relevant text from the supplied artifacts.
- Running a clustering algorithm, modified from [23], to produce a two-dimensional spatial representation of the documents clustered and arranged by a set of similarity measures.
- Storing the results of the parsing and clustering in the database.

Bubbles, described further in Section 3.4, represent clusters of related information artifacts as determined by the harvesting process. Bubbles simplify a depiction of text relationships by reducing a cluster of related artifacts to a single content bubble that supports easier drill-down style interaction.

An important feature of the analysis service model is the ability to incorporate application-specific code that is used during the text parsing stage of a harvest. For instance, it is possible apply custom "Text Extractor" and "Text Scrubber" objects to text documents prior to harvesting. A Text Extractor object preferentially selects fields from the XML body of a data object based on that object's schema (allowing documents with different schema to be returned as the result of a single query). Text Scrubbers operate on the resultant extracted texts, and apply application-defined transforms (for example, filtering out specific words, providing translation, expansion of known acronyms, etc).

There are four primary benefits to the service-oriented analytic model. First, multiple synchronous applications can interact with these services at once, allowing many analysts to use the same data, and allowing users to collaborate remotely over a visualization without having to also distribute the data that produced it. Second, visualization can be brought to devices that otherwise lack the processing power to compute the necessary data transforms. Third, this single visualization broker allows users to easily share visualizations – essentially, letting someone else see what they saw. The services simply provide clients with the instructions to draw a particular visualization, so in sharing a visualization a user simply shares the service call that produced a depiction of interest. Finally, running all user interaction against a single service set enables new notification tools; a user can establish a subscription that allows them to be alerted when another user has created information of interest.

3.2. Implementing reasoning artifacts

Every information object – whether text, image, multimedia, or collection thereof – is managed by the SRS services as a "reasoning artifact". These artifacts are the basic elements from which analyses are composed. Our taxonomy of reasoning artifacts, adapted from [24], includes the following elements:

- *Source information*: An individual unit of data that does not yet have an analytic role, but is germane to an analysis.

- *Evidence:* Source information may turn into evidence – that supports or refutes an explanation, for instance – when it has utility in a reasoning strategy.
- *Assumption:* Important to identifying user biases and mental models that color an analysis, assumptions include such information as informed conjectures about unknowns.
- *Argument:* Individual pieces of evidence and assumption may be linked into argument structures that reflect analytic judgment.
- *Hypothesis*: Evidence, assumptions, and arguments combine to form hypothesized explanations for observations.

Within these categories are specializations, such as causal arguments, and groupings, such as temporal and spatial patterns. SRS users tag artifacts with one of these categories, and the artifact's role can be associated with a specialized schema in the SRS backend.

One interface to the Scalable Reasoning analysis environment is a visual, web-based client for managing information collections and marshalling them into hypotheses that can be disseminated. This interface is shown in Figure 1. Every interaction with this client involves background communication with one or more SRS services. Here, reasoning artifacts are represented as rectangular "widgets." Both the web based client and the mobile client (introduced in Section 3.5) have capabilities for collecting and exploring these artifacts that are suited appropriately for the environment. The web based client collects information primarily from the shared database, from web pages, from manually entered observations, and from local files. The mobile client takes advantage of the camera, video and voice recording, and GPS coordinates to create a rich array of reasoning artifacts that can then be shared. Because these artifacts are all stored centrally, each environment has access to all the artifacts and can render them according to the capabilities of each platform.

3.3. Building thought chains

As information is gathered, the analyst can package lower-level artifacts together into sets representing higher-order collections, and can indicate that particular artifacts act as associations between others (for instance, a causal model might attribute an item of evidence to a spatial pattern). Figure 2 depicts a section of a sample analysis in the domain of pandemic flu as it appears in the workspace. Each artifact in the analysis is represented as a rectangular widget that may contain other widgets nested within. Here, the "Drug resistance" artifact describes a hypothesis (indicated by the light bulb icon at lower right) about the possibility of the H5N1 influenza strain not responding to vaccines; this artifact is founded on three sub-elements (boxed icons within it) that could be opened to explore the hypothesis further. To create these hierarchies, users simply drag-and-drop reasoning artifacts onto each other. This packaging behavior helps users collapse potentially deep analyses into a single visual element. SRS simultaneously displays these hierarchical relations alongside associative relations, shown as green arrows connecting artifacts. The drug resistance hypothesis, for instance, is linked to a scenario about an H5N1 pandemic through an assumption about exposure to mutated virus strains (encoded in the arrow

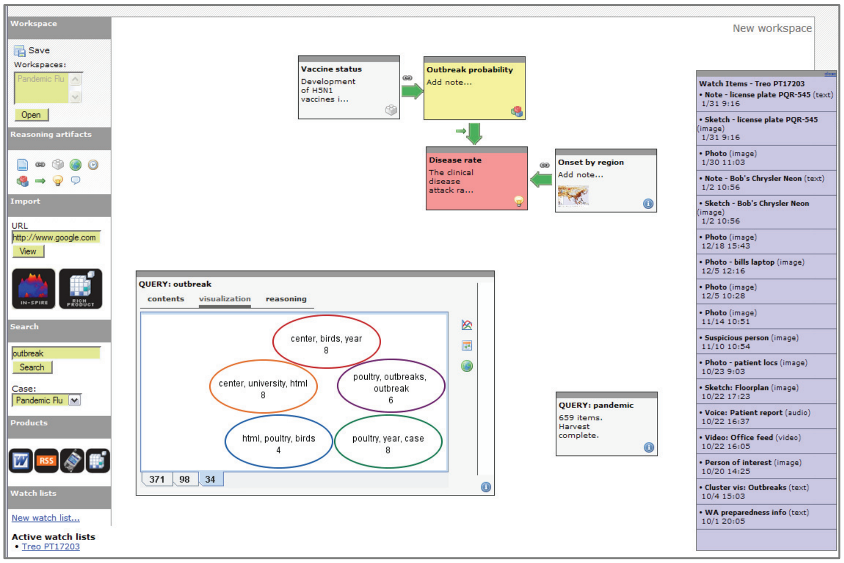

Development of H5N1 vaccines i...

The clinical disease attack ra...

659 items. Harvest complete.

Watch Items - Treo PT17203
- **Note - license plate PQR-545** (text) 1/31 9:16
- **Sketch - license plate PQR-545** (image) 1/31 9:16
- **Photo** (image) 1/30 11:03
- **Note - Bob's Chrysler Neon** (text) 1/2 10:56
- **Sketch - Bob's Chrysler Neon** (image) 1/2 10:56
- **Photo** (image) 12/18 15:43
- **Photo - bills laptop** (image) 12/5 12:16
- **Photo** (image) 12/5 10:28
- **Photo** (image) 11/14 10:51
- **Suspicious person** (image) 11/10 10:54
- **Photo - patient locs** (image) 10/23 9:03
- **Sketch: Floorplan** (image) 10/22 17:23
- **Voice: Patient report** (audio) 10/22 16:37
- **Video: Office feed** (video) 10/22 16:05
- **Person of interest** (image) 10/20 14:25
- **Cluster vis: Outbreaks** (text) 10/4 15:03
- **WA preparedness info** (text) 10/1 20:05

Figure 1. SRS web workspace, with reasoning diagram composed of reasoning artifacts and a "query" artifact opened to reveal a visualization of query results.

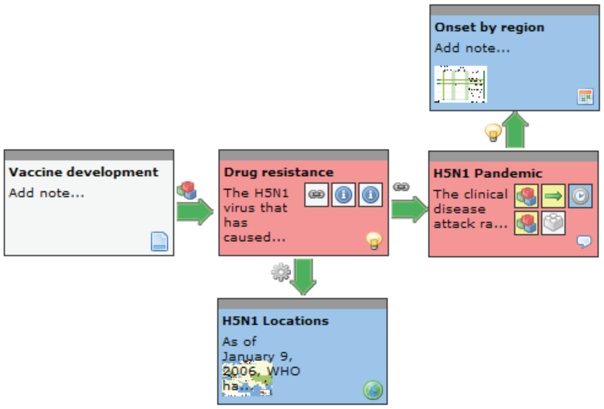

Figure 2. Sample analytic pathway for an avian flu analysis.

linking the two nodes; although not depicted in this view, the assumption could be "unpacked" to examine its structure in detail). The pandemic scenario, in turn, defines the analyst's estimate of a temporal pattern of the onset of infection by geographic area.

Any reasoning artifact can be turned into a connector to assert a relationship between two other artifacts. In this manner the reasoning diagrams produced in SRS represent not just "flat" concept maps, but concept maps whose nodes and edges represent potentially deep evidence packages (a connector, for instance, might be created out of an entire sub-analysis, which the analyst compresses to view as a single edge in the current depiction).

Supplementing each artifact depicted in a reasoning graph is a numerical value that can be used to perform mathematical reasoning and analysis over a hypothesis network. This value is a confidence rating, which ranges from strong confidence the artifact is true to strong confidence the artifact is false. The user can manually set the confidence of each artifact, or specific rules can set confidence values by default (i.e. photographs are automatically given a high confidence because they are a primary source).

On links between artifacts another value is specified (Figure 3). This value is the probative force, or strength of the relationship. Its range is from strongly supportive to strongly refutive. In other words, the value represents how much one artifact supports or refutes another artifact.

These ratings fulfill the basic requirements for a modeling engine. Both Bayesian and Dempster-Shafer network modeling techniques can be applied to this model and used to determine the likelihood of a line of reasoning and analysis of competing hypotheses. The implementation of such an inferential engine, while outside the scope of the present work, would allow reasoning structures to be updated automatically when user assessments change.

Analytic processes are captured in SRS in two ways. First, a versioning mechanism tracks the activity related to each artifact, so as items are connected, disconnected, re-connected elsewhere, and their confidence values changed, the history of this testing is maintained and auditable. Using asynchronous connections to the SRS servers, a user's interaction with the client workspace is continually serialized in the background; the state and content of each reasoning artifact is automatically updated in the SRS database each time a user manipulates it. Second, as the following section describes, the reasoning environment is designed to reduce the computational and cognitive distance between information analysis and the reasoning process by allowing dynamic analytic visualizations to run inside an artifact.

3.4. Organizing collections across queries and visualizations

Until now we have discussed just the visual argumentation aspects of SRS. Via SRS services, these argumentation structures are synchronized in real time with the SRS data store, allowing collaborating analysts to share reasoning structures easily. But to help close the gap between foraging and sense-making, SRS enables users to explicitly record how their manipulation of a data visualization (foraging) results in new knowledge (sense-making). A single artifact might represent a collection of thousands of documents returned in response to several web queries, along with a collection manually created by the user. Instead of simply serving as a container for these documents in the workspace and relegating their analysis to a separate environment, document clustering

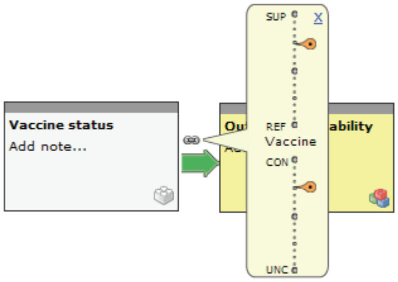

Figure 3. Setting confidence and probative force values for a relationship artifact.

techniques using the services introduced in Section 3.1 are integrated into the workspace.

Document clustering provides a mechanism for communicating the similarity in content between information artifacts – the more they have in common, the closer they appear in the information space. For large document sets, such a space can quickly become overcrowded, making it less valuable to the user. Clustering documents in a hierarchical fashion and providing only cluster-level information (versus document specific information) sufficiently reduces the complexity of the visualization while maintaining the ability to navigate the set according to document content.

Figure 4 shows a cluster visualization composed of bubbles representing topic clusters in an artifact of interest. Each bubble is labeled with the major topics that define it and the number of documents it contains. Selecting any bubble "opens" it, drilling down to create a new visualization of the document relationships within just that subset. A row of tabs along the bottom of the visualization serves as a breadcrumb trail that allows the user to quickly traverse the navigation history that brought them to a particular point in the visualization. The number of clusters shown at any one time is adjustable by the user to show more or less topical granularity. Any artifact in a reasoning diagram that contains "sub-artifacts" within it can be visualized with this tool. (In the case of a query, the query itself is the "artifact", and the documents that the query returns are sub-artifacts; in the case of a user-created information package, an artifact might be a container called "relevant documents" and the sub-artifacts are all the other artifacts that the user has dragged onto this node to package them together.) Like any other reasoning artifact, queries are represented as visual nodes in the workspace. Opening a query node to reveal its contents also reveals controls to select a mechanism for visualizing the contents.

More than just in-place visualization of the items contained in a reasoning diagram, fusion of foraging and sense-making is accomplished by allowing the user to physically extract insight-generating clusters of documents from a visualization and tie them into an emerging reasoning structure. In a cluster visualization, a user might drill-down to identify a collection of key documents, and then use those documents as evidence to underlie a developing hypothesis. To accomplish this, the user simply drags the relevant cluster out of the visualization and onto the surrounding workspace, where it takes on the appearance of a new artifact while maintaining the history of visual exploration that produced it. These extracted collections can then be wired into a reasoning diagram using connectors, or they can be dragged onto each other to create new collections that represent the union of information discovered through exploration of separate visualizations. The extraction of high-value information supports the "exploitation" reasoning phase identified in [4]; document clusters representing lower-value collections can then be pruned away by deleting a bubble, enriching the information space. Explicit computational support for these analytic acts, the ability to have several unique visualizations open in the web client at once, and the ability to create new fused collections that represent bits and pieces of each, are significant contributions of the SRS approach.

3.5. Mobility: Collection, analysis, and dissemination

One goal of the Scalable Reasoning System is to create analytic tools that scale down to a mobile agent without significant loss of the capabilities that exist on the desktop as well as create tools that make collection of reasoning artifacts in the field simple and intuitive, and collaboration with other agents seamless and synergistic. Leveraging the same visualization and analysis services that power the web client, our mobile SRS environment essentially targets three facets of a mobile user's needs: data collection, data viewing/analysis, and search.

We have implemented a variety of collection capabilities, including text, GPS location, voice, photo, video, and sketch. Upon collection of the data, it is automatically uploaded to the central server to allow immediate collaboration, harvesting, and analysis. Mobile search

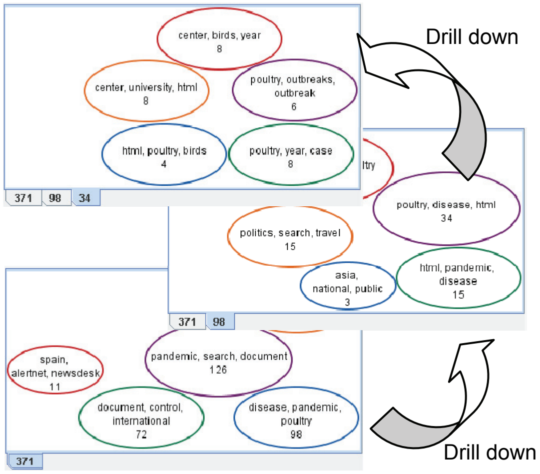

Figure 4. Drilling down into a cluster visualization reveals the content structure of documents therein.

tools are suited to a variety of situations. For example, a location search, which will return documents tagged within a given radius of a point, is ideal for the mobile user interested in things that are happening at or near their current location. Other searches will run text analysis tools on the database to extract topics and allow the user to drill down to the documents of interest by choosing successively specific topics. These searches return visualizations almost identical to the desktop renderings, though suited for better visibility on mobile devices (Figure 5; clusters are converted to 1-dimensional blocks where blue bar indicates "size" of bubble).

All data collected by the mobile device is automatically tagged with the time, the location (if available), and an identifier for the device, allowing a collaborator on the desktop to get updates when the mobile user adds content. This is of particular use to a dispatcher watching as multiple agents at a scene acquire evidence.

We are doing further work to facilitate the automatic transfer of information back to the mobile device, through multimedia messaging or other means. A primary concern with mobile users is avoiding data overload; while pushing information from the mobile device to the SRS knowledge base is comparatively straightforward, there are significant usability issues in ensuring that only the most relevant information is distributed to mobile personnel. SRS users can create "watch lists" that serve as location- or topic-based subscriptions. These subscriptions act as standing queries that notify users of new information or changes in other users' confidence assessments.

3.6. Sharing analytical reasoning spaces

Because the central SRS services mediate production of and interaction with data visualizations (in their capacity as harvesting and clustering routines) and reasoning strategies (through capture and serialization of the artifacts a user creates), it becomes possible to share complete and editable depictions of problem-solving processes between analysts accessing SRS through different clients. Descriptions of the workspace are saved in such a manner that they can be reconstructed on the device as well as on other devices. Devices will only render as they are able, so representation will be different across platforms, but the content will be the same. This ability allows agents in the field to contribute to the analysis and decision making process, and also gives a user the freedom to work on a workspace while traveling and switch to another platform at the desktop without changing context. Users have access to the reasoning spaces of other users, as well, so that collaboration and information sharing on multiple levels is enabled.

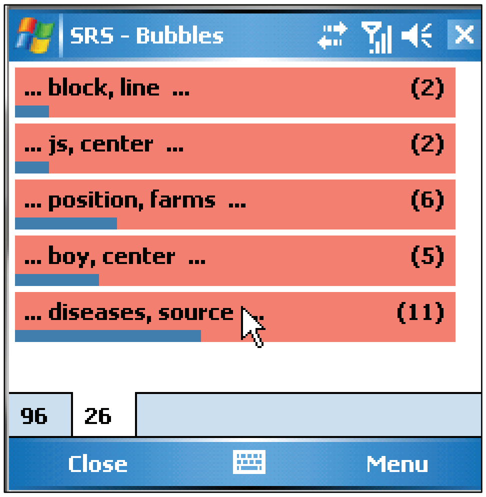

Figure 5. Conversion of cluster visualization as shown in Fig. 4 to a "one-dimensional" view appropriate for a small-screen device.

In circumstances where the veracity of a particular reasoning component comes into question, or beliefs evolve, it is possible to recreate the analysis in light of the updated information. Any analyst can retrieve a prior workflow created in SRS, on any device, and can update it with new analysis.

4. CHALLENGES FOR SCALABLE VISUAL REASONING

One of the central issues in creating visual analytic environments that support collaborative exploration and sense-making is the need for infrastructural support. Our approach of simple visualizations that offload much of the processing to centralized servers requires a coordinated event model that guides communication between client and server. While such a model is feasible for systems on the scale of SRS, it becomes a barrier when trying to integrate several visual tools into a single collaborative framework. We anticipate that as these "lightweight" visual interfaces become more widely adopted, communities will work toward standards for interoperable visualization services. Many of our visualizations, for instance, rely on the Scalable Vector Graphics (SVG) standard, so any service that produces SVG output could potentially be integrated, although this standard just accommodates the visual display and not the interaction model. A further concern is that in highly mobile environments, the centralized infrastructure we currently employ may be infeasible. In these circumstances, we may work toward enabling peer-to-peer collaborative analysis in place of our client-server approach, as well as encrypted communication over insecure networks.

When manipulating shared visual objects, issues of concurrency and conflicting updates become important. Other work [e.g., 25] has addressed the problem of concurrency of view in collaborative visualization. In our case, version control plays the dual role of keeping each user's modifications to shared reasoning artifacts separate, while also tracking the different roles that versions of the same artifact play in different contexts. We are currently exploring visual techniques that help users understand the evolution of reasoning structures over time.

Our system also supports semantic scalability in visual exploration. For users to be able to collaborate in the process of visual discovery, they need to be able to share more than just the data depictions; users need to be able to communicate the meaning each attaches to aspects of these visualizations. Thus, visualizations must scale from depicting low-level data relationships to the higher order knowledge that emerges from these relationships. Our reasoning diagram approach is one way to capture the association between data and knowledge. However, explicit recording of semantic relationships can require users to be more introspective than they have either the comfort or the time to be. Ideally, reasoning strategies would be captured during collaborative visualization with as little user intervention as possible. As the research community develops a better understanding of the relationship between manipulating a visualization and the construction of resulting knowledge, it may become possible to infer analytic activity through passive observation of user interaction alone.

5. CONCLUSIONS

We have introduced a collaborative visualization approach based on a technique we call Scalable Reasoning. Our goal is to create visual analysis services that can be deployed on a range of devices, including mobile systems with traditionally limited graphical and interaction capabilities. Beyond scaling across devices, our system also supports the capture of analytic insight that visualizations generate through a graphical argumentation structure founded on a taxonomy of reasoning artifacts. Insight can be recorded in situ, in the context of the visualizations that produce it, and can be shared between users both synchronously, by pushing artifacts to another user's device, and asynchronously, by retrieving and reusing a prior reasoning strategy from the SRS data store.

The service-oriented analytic model has four primary benefits: 1) many analysts can interact with the same data,; 2) visualization can be brought to devices that otherwise lack the processing power to compute the necessary data transforms; 3) a single visualization broker service allows users to easily share visualizations; and 4) running all user interaction against a single service set enables new notification tools.

We are currently engaged in a pilot deployment of SRS with a regional law enforcement organization. This deployment is evaluating the impact of bringing visual analyses to its mobile users and the benefit of enabling those uses to share visual representations with their office-bound colleagues. We are also expanding the mobile visualization work to bring a lightweight reasoning capture interface to mobile users, allowing them to record their investigation process in place.

SRS represents a growing recognition that computational aids to analytical work must incorporate richer representations of the reasoning process. As more examples of systems that support introspection and shared reasoning, our hope is that communities of practice will increasingly see these systems as necessary to effective and efficient collaboration.

REFERENCES

[1] J. Bodnar, "Warning analysis for the information age: Rethinking the intelligence process," Joint Military Intelligence College, Center for Strategic Intelligence Research, Washington, DC 2003.

[2] P. Pirolli and S. Card, "Information foraging," Psychological Review, vol. 106, pp. 643-675, 1999.

[3] E. Patterson, E. Roth, and D. Woods, "Predicting vulnerabilities in computer-supported inferential analysis under data overload," Cognition, Technology, and Work, vol. 3, pp. 224-237, 2001.

[4] A. MacEachren, "Moving geovisualization toward support for group work," in Exploring Geovisualization, J. Dykes, A. MacEachren, and M.-J. Kraak, Eds. Amsterdam: Elsevier, 2005, pp. 591-609.

[5] G. Li, V. Uren, E. Motta, S. Buckingham Shum, and J. Domingue, "ClaiMaker: Weaving a semantic web of research papers," in The Semantic Web - ISWC 2002: First International Semantic Web Conference. vol. 2342, I. Horrocks and J. Hendler, Eds. Berlin: Springer-Verlag, 2002, pp. 436-441.

[6] S. Buckingham Shum, "Sensemaking on the Pragmatic Web: A hypermedia discourse perspective," in 1st International Conference on the Pragmatic Web, 2006.

[7] M. Gahegan and W. Pike, "A situated knowledge representation of geographical information," Transactions in GIS, vol. 10, pp. 727-749, 2006.

[8] W. Wright, D. Schroh, P. Proulx, A. Skaburskis, and B. Cort, "The sandbox for analysis: Concepts and evaluation," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Montreal, 2006, pp. 801-810.

[9] S. Brennan, K. Mueller, G. Zelinsky, I. Ramakrishnan, D. Warren, and A. Kaufman, "Toward a multi-analyst, collaborative framework for visual analytics," in IEEE Symposium on Visual Analytics Science and Technology 2006, Baltimore, MD, 2006, pp. 129-136.

[10] D. Suthers, "Collaborative knowledge construction through shared representations," in Proceedings of the 38th Hawaii International Conference on System Sciences (HICSS-38), Waikoloa, Hawaii, 2005.

[11] J. Novak, "Helping knowledge cross boundaries: Using knowledge visualization to support cross-community sensemaking," in Proceedings of the 40th Hawaii International Conference on System Sciences, Waikoloa, Hawaii, 2007.

[12] S. Li and A. Hopper, "What you see is what I saw: Applications of stateless client systems in asynchronous CSCW," in Fourth International Conference on Computer Science and Informatics (CS&I'98), October 23-28 1998, Research Triangle Park, NC, 1998.

[13] U. Schultze and R. J. Boland, "Knowledge management technology and the reproduction of knowledge work practices," Journal of Strategic Information Systems, vol. 9, pp. 193-212, 2000.

[14] F. Nake and S. Grabowski, "Human-computer interaction viewed as pseudo-communication," Knowledge-Based Systems, vol. 14, pp. 441-447, 2001.

[15] S. Dustdar, "Caramba - A process-aware collaboration system supporting ad hoc and collaborative processes in virtual teams," Distributed and Parallel Databases, vol. 15, pp. 45-66, Jan 2004.

[16] A. de Moor, "Patterns for the Pragmatic Web," in Conceptual Structures: Common Semantics For Sharing Knowledge, Proceedings. vol. 3596, 2005, pp. 1-18.

[17] B. Shneiderman, "Codex, memex, genex: The pursuit of transformational technologies," International Journal Of Human-Computer Interaction, vol. 10, pp. 87-106, 1998.

[18] W. Pike, R. May, and A. Turner, "Supporting knowledge transfer through decomposable reasoning artifacts," in Proceedings of the 40th Hawaii International Conference on System Sciences, Waikoloa, Hawaii, 2007.

[19] R. Carvalho, J. Williams, I. Sturken, R. Keller, and T. Panontin, "InvestigationOrganizer: The development and testing of a Web-based tool to support mishap investigations," in Proceedings of 2005 IEEE Aerospace Conference, 2005, pp. 1-10.

[20] Y. Gil and V. Ratnakar, "An interactive tool for capturing information analysis and decision-making," in Knowledge Engineering and Knowledge Management, Proceedings of the thirteenth international conference, Siguenza, Spain, 2002, pp. 37-42.

[21] K. Wideroos and S. Pekkola, "Presenting the past: A framework for facilitating the externalization and articulation of user activities in a desktop environment," in Proceedings of the 39th Annual Hawaii International Conference on System Sciences (HICSS'06) Track 7, 2006, p. 148a.

[22] P. Cowley, L. Nowell, and J. Scholtz, "Glass Box: An instrumented infrastructure for supporting human interaction with information," in Proceedings of the 38th Annual Hawaii International Conference on System Sciences, 2005, p. 296c.

[23] B. Hetzler and A. Turner, "Analysis experiences using information visualization," IEEE Computer Graphics and Applications, vol. 24, pp. 22-26, 2004.

[24] J. Thomas and K. Cook, "Illuminating the Path: The Research and Development Agenda for Visual Analytics," Los Alamitos, CA: IEEE Press, 2005, p. 200.

[25] S. Ryu, H. Kim, J. Park, Y. Kwon, and C. Jeong, "Collaborative object-oriented visualization environment," Multimedia Tools and Applications, vol. 32, pp. 209-234, 2006.