

# Entity Workspace: an evidence file that aids memory, inference, and reading

Eric A. Bier<sup>1</sup>, Edward W. Ishak<sup>2</sup>, Ed Chi<sup>1</sup>

<sup>1</sup> Palo Alto Research Center, Inc., 3333 Coyote Hill Road, Palo Alto, CA 94304, USA  
{bier, echi}@parc.com

<sup>2</sup> Columbia University, Department of Computer Science, 1214 Amsterdam Ave., 450 CS  
Bldg, New York, NY 10027  
ishak@cs.columbia.edu

**Abstract.** An intelligence analyst often needs to keep track of more facts than can be held in human memory. As a result, analysts use a notebook or evidence file to record facts learned so far. In practice, the evidence file is often an electronic document into which text snippets and hand-typed notes are placed. While this kind of evidence file is easy to read and edit, it provides little help for making sense of the captured information. We describe Entity Workspace, a tool designed to be used in place of a traditional evidence file. Entity Workspace combines user interface and entity extraction technologies to build up an explicit model of important entities (people, places, organizations, phone numbers, etc.) and their relationships. Using this model, it helps the analyst find and re-find facts rapidly, notice connections between entities, and identify good documents and entities to explore next.

## 1 The Need for a Better Evidence File

An intelligence analyst working on a case may encounter hundreds of documents. Each document, in turn, may contain any number of facts relevant to the case. The number of facts can quickly become too large to be held in human memory alone. For this reason, analysts often use a notebook or *evidence file* to record the facts uncovered so far. This evidence file may be an electronic document into which the analyst places snippets of text and hand-typed notes about the case.

Electronic documents have a number of capabilities that make them suitable for use as evidence files. For example, they are easy to edit and read. They can hold both textual and pictorial material. They support full text search. Finally, as the analyst's product is generally an electronic document, evidence stored in an electronic document is easy to re-purpose for use in the final report.

However, as used today, electronic documents lack the capability to help analysts make sense of the facts captured in them. Although the computer system knows that the analyst has selected certain pieces of text, it doesn't know which entities mentioned in this text are important, nor how they are related to each other.

In this paper, we describe a software evidence file called *Entity Workspace* that helps the analyst explicitly represent information about relevant entities and their

relationships. Entities of interest to intelligence analysts include people, places, organizations, telephone numbers, addresses, weapons, bank accounts, and so on. With this information in hand, Entity Workspace is able to offer the analyst several kinds of assistance during the analytical process. In particular, Entity Workspace helps the analyst: (1) re-find important facts quickly, usually without the need for explicit searching or querying; (2) “connect the dots” by visualizing relationships between entities, supporting inference tasks; (3) discover new information by recommending good documents and entities to study next, based on entities that the analyst has previously marked as being of high interest. (4) Finally, it accelerates the process of reading through new documents by highlighting key terms automatically.

Entity Workspace can provide this kind of assistance to analysts because it builds up a model of documents, entities, and entity relationships that we call the *Document-and-Entity Graph*. The analyst does not directly view or edit this graph. It is constructed automatically as a side effect of organizing information in the workspace.

Section 2 of this paper describes the drag-and-drop user interface that Entity Workspace provides for capturing entities and their relationships. With that user interface in mind, section 3 describes the Document-and-Entity Graph data structure that is built behind the scenes to represent the information that the analyst has captured. Section 4 then describes a user interface technique called entity quick click that allows the analyst to record entity information rapidly. Section 5 describes the Entity Workspace tool as a whole, including its four different panels. Section 6 then describes several ways that Entity Workspace supports and accelerates the analytical process. The final sections present related work, our plans to extend and test Entity Workspace, and our conclusions.

## 2 Rapid Entity Capture from Large Document Collections

We first describe the user interface features of the Entity Workspace. These user interface features enable analysts to rapidly capture entity relationships and other related information from a large document collection.

One of the main functions of Entity Workspace is to record the names of important entities and relationships between them. This is done using a part of the tool called the Workspace Pane (or simply “the workspace”), which is used as follows:

An analyst is looking for facts relevant to a case. When the analyst finds a relevant fact in a document, she may wish to capture that fact in the evidence file. With Entity Workspace, the process of capturing facts begins by copying entity names from the document into the Workspace Pane. Entity names can be copied using traditional copy-and-paste or drag-and-drop. Better methods of capturing entity names are also available, as will be described later.

### 2.1 Workspace Objects

Suppose that an analyst is working on a terrorism case and has encountered the (fictitious) report in Fig. 1. This report and other examples in this paper are drawn from “The Sign of the Crescent,” an intelligence analysis tutorial developed by Frank

Hughes, a professor at the Joint Military Intelligence College. Some of the important entities in this report are highlighted. If the information in this report looks relevant, the analyst may wish to record important entities, such as dates, names of people, names of banks, account numbers, names of businesses, and phone numbers. Any number of entities can be added to the workspace.

1) **Report Date: 1 April, 2003.** FBI: Abdul Ramazi is the owner of the **Select Gourmet Foods** shop in **Springfield Mall, Springfield, VA.** [Phone number 703-659-2317]. **First Union National Bank** lists **Select Gourmet Foods** as holding account number **1070173749003.** Six checks totaling **\$35,000** have been deposited in this account in the **past four months** and are recorded as having been drawn on accounts at the **Pyramid Bank of Cairo, Egypt** and the **Central Bank of Dubai, United Arab Emirates.** Both of these banks have just been listed as possible conduits in **money laundering schemes.**

Fig. 1. A fictitious FBI report that mentions a person, a company, and several other entities

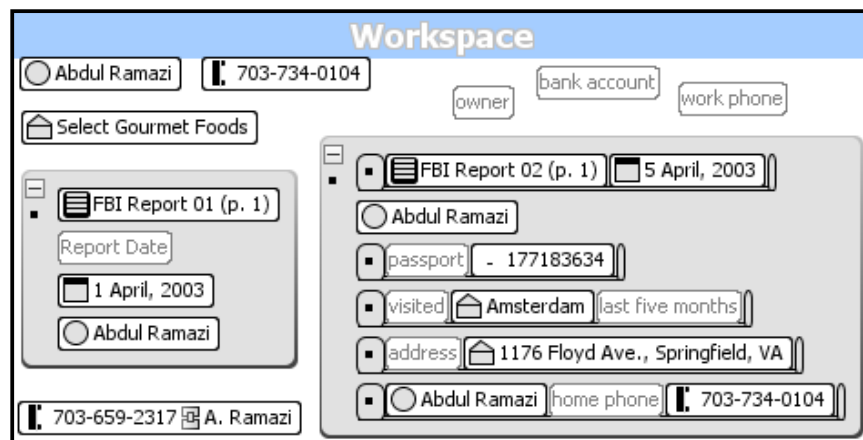


Fig. 2. An entity workspace containing (top to bottom, left to right): entities, comments, a group, a group containing, statements, and an entity coupled to another entity.

Each entity recorded in the workspace is displayed as a separate visual object, called an *entity object*. Each entity object can be independently dragged and positioned anywhere in the workspace. The upper left corner of Fig. 2 shows three entities, representing a person, a telephone number, and a shop. Each entity is displayed as a rounded rectangle containing text describing the entity. A small icon at the left end of each entity indicates its type, if it is known. For example, a circle indicates a person, a calendar icon indicates a date, a house icon indicates a place, and a telephone receiver icon indicates a telephone number.

In addition to objects representing entities, the analyst can add arbitrary text strings to the workspace representing comments, labels, descriptions of relationships, categories, and so on. Each one is called a *comment object*, and the analyst sets its text by copying text from a document or typing. The upper right corner of Fig. 2 shows three comment objects, labeled “owner”, “bank account”, and “work phone”. Comments look like entities, but have no icon inside and have gray text and border.

Finally, an analyst can add a representation of a particular page of a document in the analyst's collection, called a *document page object*. Examples of these in Fig. 2 have labels beginning with "FBI Report". The analyst can drag a document page object into the workspace from the Entity Inspector and System Suggestions panes, which are described below. Document page objects are also added to the workspace automatically under certain circumstances. These objects record where a particular fact was discovered. Clicking on a document page object causes the corresponding page to be displayed in a document viewer so the analyst can review it.

## 2.2 Snap-Together Relationships

Entity objects, comment objects, and document page objects can all be dragged to any position in the workspace. However, a special function is triggered when an object is placed in close vertical proximity to another object. In this case, our system creates a new composite object, called an *entity group*, and places both objects in the new group. Like an entity, a group is displayed as a rounded rectangle and is made large enough to contain all of its child entities and is displayed with a non-white background, as shown in the middle left and lower right of Fig 2. In our system, a group is also displayed with a collapse button in the upper left hand corner (to request a condensed view).

The analyst places objects together in a group to indicate that they are related to each other in some way. For example, the left hand group in Fig. 2 indicates that Abdul Ramazi is connected in some way to the date April 1, 2003. Dragging additional objects near to an existing group places these objects into the group. A group can contain any number of objects. The group at the left edge of Fig 2 has four objects: a document page object, a comment, and two entity objects.

The analyst indicates a stronger relationship between two objects by placing them together in a statement. This is done by dragging one object in close horizontal proximity to a second object. Our system responds by creating a new composite object, called an *entity statement* and places both objects into it. The statement is shown visually by placing the objects side-by-side and displaying "book end" shapes to the left and right of them and an underline below them.

Statements can be placed in groups, just as non-compound objects can. The right hand group of Fig. 2 shows a group that includes five statements and one non-compound object. The analyst can add any number of objects into a statement before, after, or in between existing objects to make a longer statement. Like groups, statements can contain a mixture of different object types including entity, comment, and document page objects.

An object can appear in the workspace any number of times. For example, the entity "Abdul Ramazi" appears twice in the right hand group of Fig. 2 (on the second and last lines). An object can appear in multiple groups, in multiple statements, both in and out of groups, and in and out of statements. Each appearance of the object is called an instance, where all instances of an object represent the same real world object. Having multiple instances of an object allows the analyst to express its relationships to any number of other objects in multiple contexts.

The system makes it easy to copy any object instance. The new copy, like the original, can then be dragged, placed in a group, placed in a statement and so on.

### 2.3 Coupling Objects

Some objects are so closely related that it is helpful to ensure that they are always displayed together. For example, the analyst may decide that whenever a certain phone number appears on the screen, the name “Abdul Ramazi” should appear next to it. To do this, the analyst points to both objects in turn and gives a “coupling” command. Thereafter, every current and future instance of this phone number will appear followed by a link icon and the name, like the entity in the lower left corner of Fig. 2, helping the analyst remember the connection.

### 2.4 Indicating Degree of Interest

Each entity in the workspace is associated with a number that indicates the degree to which that entity is of interest to the analyst. The degree of interest is a number from 1 to 4, where 1 means low interest and 4 means high interest. Entities are initially given a rating of 2 or “possible interest”. Our system varies the display properties of an entity based on its interest rating. For example, in our system, items rated 1 are in blue italics, items rated 2 are in black plain face, items rated 3 are in red plain face, and items rated 4 are in red bold face. Fig. 3 shows examples. The analyst clicks on the icon at the left of each entity object to cycle through the degree-of-interest values.

### 2.5 Entity Quick Click: Entity Extraction Enables Rapid Copy-and-Paste

Intelligence analysts often need to produce results quickly, so they need efficient tools. In order to make Entity Workspace practical, we needed a faster way to record entities than traditional copy-and-paste or drag-and-drop. Entity Workspace includes two accelerated methods for recording entities and their relationships that we call *entity quick click* and *linguistic grab*.

Entity quick click is described elsewhere [2]. We summarize it here, and describe its modification for intelligence analysis. With this technique, the system runs an information extraction algorithm over each document to be displayed. Information extraction finds entities such as person names, company names, organization names, telephone numbers, addresses, and currency amounts. When the document is displayed to the analyst, these entity phrases are highlighted, as they are in Fig. 1.

Next, the analyst presses a key (such as Shift) to initiate entity quick click mode. Then, the analyst positions the cursor over a desired phrase and clicks on it. The phrase is copied into the entity workspace. In addition, the newly added entity is placed into an entity group that includes a document page object corresponding to the page in which the analyst just clicked; this entity group becomes the current entity group. By clicking on additional phrases, the analyst adds more entities to this group. The result looks like the left hand group in Fig. 2. By holding down different keys (e.g. Shift and Control), the analyst can add the selected phrase to the current entity group as a comment object instead of an entity. In this way, the analyst can add entities and comments to the entity workspace at a cost of one cursor motion and one mouse click per addition. This is much faster than traditional copy-and-paste which requires roughly four times as many steps.

## 2.6 Linguistic Grab: Automatic Relationship Extraction between Entities

Another technique for rapid copying is *linguistic grab*. With this approach, the system runs both an entity extraction process as described above and additional information extraction operations, based on computational linguistics, to find relationships between entities. The analyst selects a passage to import. As with entity quick click, the system creates a new entity group in the workspace and places into it document page objects representing the page(s) that the analyst has selected. It then places into the group an entity object for each extracted entity. In addition, the system adds a statement object to represent each extracted relationship found between two or more entities. For example, in the document of Fig. 1, the information extractor concludes that “Abdul Ramazi” has a “person-company” relationship with the “Select Gourmet Foods shop” and creates a statement that includes the person name, the company name and a comment to indicate the nature of the relationship.

## 3 The Entity Workspace Tool

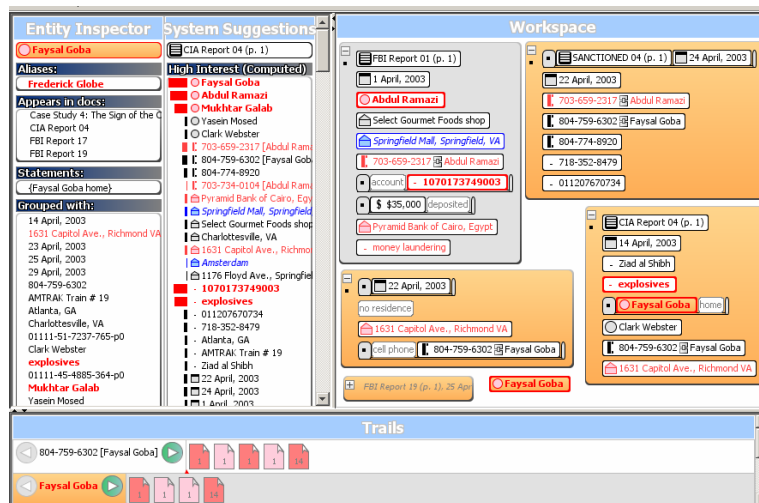


Fig. 3. The four panels of the Entity Workspace Tool

Having described some of the user interface features for recording entities from document collections, we now describe the entire Entity Workspace tool. As shown in Fig. 3, the tool includes four panels. Counter-clockwise from upper right they are: the Workspace panel, the System Suggestions panel, the Entity Inspector panel, and the Trails panel. We next describe the purpose of each panel.

The *Workspace Panel* serves as the entity workspace described above. It includes representations of entities, comments, document pages, groups, and statements. Entities can be added to this panel by dragging them from the Entity Inspector or

System Suggestions panel, or by copying them from the document reader (shown in Fig. 1) using copy-and-paste, drag-and-drop, entity quick click, and linguistic grab.

The *System Suggestions Panel* shows a list of all of the entities that are represented in the workspace, ordered according to how closely related each is to high-interest objects. This display is sorted first by the type of entity (e.g. all entities of type “person” precede all entities of type “phone-number”), and then by closeness score. The top of this panel displays a document page object representing the page that is currently open for reading in the document reader.

At any time the analyst may select an entity in the workspace by clicking on it. The *Entity Inspector Panel* responds by displaying information that is known about that entity, including the text string that represents it, any text strings that serve as aliases for it, any documents that mention it, any statements in which it participates, other entities that appear in the same groups with it, and other entities that appear in the same documents with it.

When an entity or set of entities is selected, the system finds all documents in the repository that mention that set of entities and displays information about these documents in a horizontal band in the *Trails Panel*. It displays one icon for each such document, using icon color and other graphical properties to indicate facts about the document, such as whether or not the analyst has read it, how many pages it has, and how highly the analyst rates it.

## 4 A Graph Model of Gathered Intelligence

As the workspace grows, with the addition of entities, statements, and groups, the system creates a data structure, called the document-and-entity graph or DE Graph, to represent the relationships indicated by the arrangement of objects. The DE Graph is an undirected graph whose nodes are the entity objects and document page objects described above. An edge is placed in the DE Graph between any two objects that are known to be related according to the workspace information.

In particular, an edge is placed between all pairs of objects that share a group in common or that share a statement in common. In addition, edges may be added in the DE Graph because of information available outside the workspace. For example, an edge may be placed between two objects if they are mentioned near each other in one or more documents. Fig. 4 shows two entity groups next to their associated DE Graph.

If an object is not inside of any groups or statements, then no edges attach to it (as for the “explosives” object in Fig. 4). Comment objects are not included as nodes in the graphs, but may be used to label edges.

Each edge may be associated with a number representing its “weight”. Edges are given higher weight if the workspace indicates that the objects joined by the edge are more closely related. For example, edges joining objects that share a statement are weighted more heavily than edges that share only a group. Edges joining objects that share multiple groups or statements may be weighted more heavily than edges joining objects that share only a single group or statement. In the case that the graph includes information about co-occurrence in a collection of documents, edges between objects that are explicitly linked in the workspace (because they share a group or statement)

may be weighted more heavily than edges joining objects that are linked only because they are mentioned in the same document.

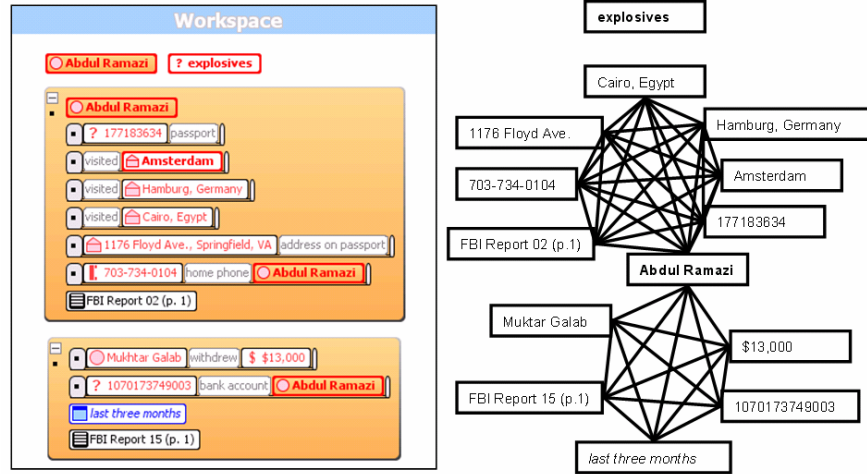


Fig. 4. A workspace and its corresponding Document-and-Entity Graph

The system uses the DE Graph to provide the analyst with memory aids, recommendations, inference aids, and support for reading, as described below. Roughly, given this weighted graph, we compute closeness scores that describe how related entities are to each other. The system makes use of semantic and association networks to do this conceptual search. This method is based on computational models of *Spreading Activation* [Anderson84] used in the study of relevancy in human memory. Starting from entities that the analyst has rated as being of high interest, activation spreads from one node to another through the weighted edges of the graph. The result is a score for each entity that represents its computed degree of interest, given how connected it is to entities already known to be of interest.

## 5 Improving Memory, Inference, and Reading

As soon as an analyst has recorded some information into the entity workspace, the system can begin to offer the analyst help in remembering facts, making inferences, and finding new information more rapidly. These capabilities are described next.

### 5.1 Augmenting Working Memory

Entity Workspace uses the DE Graph and the panels described above to help the analyst remember what has already been learned from the document collection. In particular, each visible entity serves as a reminder of a person, place, phone number or other concept that has been encountered. Each statement and group represents one



or more relationships between the entities. When the analyst selects an entity, all instances of that entity in the Workspace panel highlight in orange, as shown for the selected entity “Faysal Goba” in Fig. 3. In addition, all groups that contain an instance of that entity also highlight in orange. This allows the analyst to find such groups at a glance. By examining the other entities that are mentioned in these same groups, the analyst is quickly reminded of the set of entities that are directly linked to the selected entity.

Likewise, the Entity Inspector panel shows all of the directly linked information about an entity in a single column of information that is easy to scan. The Trails panel takes the analyst through all documents that mention the entity. The System Suggestions panel lists all of the entities recorded so far in the workspace.

When the analyst clicks on a document page object in any panel, the corresponding document page is immediately displayed, to help the reader review the page.

## 5.2 Supporting Rapid Inference

Our system also provides tools to help the analyst discover connections between entities, even if those connections are indirect. For example, when two entities are coupled, as described above in the section on Coupling Objects, it becomes possible to see connections between entities that require synthesizing information from multiple sources. For example, say that the analyst is recording information from a document describing a phone call from “703-659-2317” to “804-759-6302”. The analyst may use entity quick click to put these two phone numbers in a group. If the analyst has previously coupled the phone numbers to a Person A and a Person B respectively, the resulting entity group will immediately display the names of both people, allowing the analyst to infer that Person A has called Person B.

As a second example, suppose the analyst wishes to see if there is any known connection between some Persons C and D. The analyst can answer this question by *selecting both entities*. In the case of multiple selections, our system displays the two selected entities in distinctive ways (such as with distinctive colors). For example, the first selection may be in orange and the second selection may be in magenta. Each group that mentions only the first entity is highlighted in its corresponding color. However, groups mentioning *both* entities are given a distinctive color, such as green for example. In the resulting display, the analyst can rapidly find those groups that contain information that relates the two selected entities to each other.

As a third example, the analyst may wish to view all entities connected to a given entity by paths of length two. In this case, the analyst may activate the *Highlight Linked Entities* operation. In this case, the system highlights a selected entity and all entities that are directly linked to it, while dimming all other entities. By looking at dimmed entities that are in the same group with a highlighted entity, the analyst can see immediately the entities that are at two removes from the selected entity.

## 5.3 Providing Reading Aids

Entity Workspace supports an analyst’s reading processes in three different ways. In particular, it provides reading recommendations, rapid reading through all of the

documents on a topic, and keyword highlighting within each document to aid skimming. We describe each of these kinds of support further in this section.

The System Suggestions panel rank orders all entities of each type as a way of recommending to the analyst which entities are worth reading more about. The rating of each entity is determined using a spreading activation algorithm over the DE Graph. The initial activation of each node is computed from the degree of interest score set on the corresponding object by the analyst, as described above. After spreading activation, the nodes with the highest scores will be those that were explicitly rated highly by the analyst, those that are linked from highly rated nodes by the shortest paths, those with the most highly weighted links to highly rated nodes, those with multiple paths to highly rated nodes, or a combination of these factors. If the analyst decides to follow a recommendation, she can click on the name of the recommended entity. The system responds by highlighting related entities in the Workspace and Entity Inspector panels and related documents in the Trails Panel, making it is for the analyst to learn more about the recommended entity.

Whenever the analyst selects an entity or set of entities in the Workspace Panel, the system constructs a search engine query based on the selected entities and runs it against all of the documents in the analyst's document collection. The resulting documents are rank ordered based on the degree to which they match the query and are then displayed, best first, in the Trails Panel next to a description of the query, as shown in Fig. 3. For example, one trail in this figure corresponds to a search for a person named "Faysal Goba" and another for Faysal Goba's telephone number. The analyst can press forward and backward buttons (shaped like left- and right-pointing arrows in the figure) to visit all of the documents in the trail, in turn, much as one would flip through a stack of paper documents. This user interface helps the analyst skim for important facts.

As each document is displayed, words and phrases in the document are highlighted automatically. Several kinds of highlighting are used in order to call attention to the most important facts. In particular, all entities found by the entity extractor are highlighted in a color of relatively low salience, such as yellow. By contrast, phrases that have been given a high degree of interest rating by the analyst are highlighted in a more salient color, such as red. Finally, the system gives a distinctive appearance to phrases that match the current search query; these are highlighted with an orange background, which is larger in area than the other two forms of highlighting.

## 6 Related Work

Our work on Entity Workspace builds on other recent work related to recording evidence, spatial hypertext, automatic highlighting, automating inference, reading recommendations, and reading through multiple documents. We survey some of this literature here.

Entity Workspace is similar in many respects to the Sandbox component of Oculus nSpace [Wright2005]. Oculus Sandbox allows the analyst to copy-and-paste text snippets from a document into a work area. The analyst can then drag snippets around to group them and organize them. Our system differs in that the analyst identifies and

records specific entities and entity relationships from documents, in addition to human-readable snippets, and allows the analyst to associate a degree of interest with each one. The resulting machine-readable data structure allows our system to provide the analyst with additional help.

The automatic spatial grouping used in Entity Workspace's snap-together relationships was inspired by the Niagara system described by Good [Good2003]. In Niagara, the analyst can group text snippets by moving them close together. Entity Workspace extends this approach by supporting two different kinds of grouping that result, respectively, from moving objects close together in vertical or horizontal directions, and by building the Document-and-Entity graph to represent all of the relationships in the workspace.

Several recent systems employ automatic highlighting of text to aid reading and skimming. For example, the Scent Highlights component of the 3Book system automatically highlights words related to a query and sentences containing them to direct the reader's attention during skimming. Likewise, the Reader's Helper [Graham1999] highlights phrases judged to be similar in concept to a reader's topic of interest. Entity Workspace applies these ideas to intelligence analysis by highlighting both automatically-extracted entities and also phrases that have been given a high degree of interest rating by the analyst.

A number of authors have reported on systems that automate the process of making inferences for intelligence analysis. Coffman et. al describe the use of subgraph isomorphism to find suspicious patterns in a graph of entities and relationships [Coffman2004]. SRI's Link Analysis Workbench [Wolverton2003] searches for entities in a graph that match a pattern of suspicious behavior either exactly or approximately. By contrast to these automated approaches, Entity Workspace provides user interface tools for analysts to directly aid inference, based on whatever information that analyst is viewing at any given moment.

Several previous systems support a reader in deciding what documents to read next. Woodruff et al. described a Reading Recommender that analyses the relationships based on textual similarity and co-citation between a set of documents and a list of documents read so far, and recommends new documents to examine. Our own CorpusView system [Bier2004] displays a visualization showing at a glance the most highly rated unread documents, which act as an implicit recommendation. Entity Workspace builds on these approaches in two ways. First, because the analyst assigns degree of interest values to individual entities, recommendations are based on a relatively rich model of the analyst's interests. Second, the system recommends both documents to read and also specific entities to learn more about.

Finally, the idea of reading through a "trail" of documents at high speed is inspired by one of the aspects of the Memex system [Bush45] as envisioned by Vannevar Bush sixty years ago. We have tuned the idea to modern intelligence analysis by providing a visualization of a set of trails, each of which corresponds to a query about an entity or set of entities. The visualization that we provide is similar in some respects to Oculus TRIST [Jonker2005]; both systems show an icon per document and use graphical presentation to distinguish read and un-read documents, for example. Entity Workspace trails differ in that they are created automatically as a result of interactions in the entity workspace.

## 7 Conclusions and Future Work

The design of Entity Workspace is intentionally a compromise. We wish to elicit more information from analysts during the analytic process, so that our tools can provide them with more assistance. On the other hand, we want to ensure that the effort required from the analyst to give us this information is very small, so that the overall effect is to improve the speed and quality of analysis. Thus, while we are sometimes tempted to ask the analyst, for example, to specify in detail the exact relationship of each pair of related entities, we have resisted this temptation in favor of keeping the user interactions relatively quick and simple. This tradeoff appears to be successful. With only a little extra effort, the analyst constructs an evidence file that captures entity names, multiple appearances of entities, and coarse information about how entities are related. This information is enough to allow us to provide user interface tools that support memory, inference, and reading.

The design of Entity Workspace has been influenced by interactions with four intelligence analysts providing sample problems, sample document sets, and feedback over the course of a year. These interactions have included videotaped sessions where an analyst would work a sample problem using the tool. While the results have been informal, these sessions have given us confidence in an approach based on recognized entities and in specific features of the system such as snap-together relationships, quick click coping of entities, and entity coupling to aid inference.

Acting on feedback from analysts, we are extending Entity Workspace to hold more information, so that it can handle cases involving larger sets of entities, relationships, and documents. We are also adding to Entity Workspace facilities more commonly associated with document editors, such as the ability to copy and paste information out of the workspace, print the workspace, search through the workspace, or organize the material into sections.

We are currently hardening this technology so we can deploy it to analysts and see if it will improve their performance on real analytical tasks.

This research was funded in part by the Advanced Research and Development Activity NIMD and ARIVA programs (MDA904-03-C-0404).

## References

1. Anderson, J. R., Pirolli, P. L. Spread of Activation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 10 (1984): pp. 791—798.
2. Bier, E, Good, L., Popat, K., and Newberger, A. A document corpus browser for in-depth reading. *Proceedings of the Joint Conference on Digital Libraries (JCDL)*, 2004, 87-96
2. Bier, E. A., Ishak, E. W., and Chi, E. Entity quick click: rapid text copying based on automatic entity extraction. Submitted to the Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2006), 2006.
3. Bush, V., "As We May Think." *The Atlantic Monthly*, July 1945. Reprinted in *Interactions*, 3(2), 1996, pages 35-67.
4. Chi, E. H., Hong, L., Gumbrecht, M., Card, S. K. ScentHighlights: highlighting conceptually-related sentences during reading. *Proceedings of the 10th International Conference on Intelligent User Interfaces*, 2005, pages 272-274.

5. Coffman, T., Greenblatt, S., and Marcus, S. Graph-based technologies for intelligence analysis. *Communications of the ACM*, Volume 47, Number 3 (Mar. 2004), 45-47.
6. Good, L. E., *Zoomable User Interfaces for the Authoring and Delivery of Slide Presentations*. PhD dissertation, Department of Computer Science, University of Maryland, October 27, 2003. <http://hdl.handle.net/1903/262>
7. Graham, J. The Reader's Helper: a personalized document reading environment. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, 1999, pages 481-488.
8. Jonker, D., Wright, W., Schroh, D., Proulx, P., and Cort, B. Information triage with TRIST. 2005 International Conference on Intelligence Analysis.
9. Kozima, H. and Furugori, T. Similarity between words computed by spreading activation on an English dictionary. Proceedings of the 5<sup>th</sup> Conference of the European Chapter of the Association for Computational Linguistics, EACL-93, pages 232-239
10. Wolverton, M., Berry, P., Harrison, I., Lowrance, J., Morley, D., Rodriguez, A., Ruspini, E., and Thomere, J. LAW: A workbench for approximate pattern matching in relational data. In *The Fifteenth Innovative Applications of Artificial Intelligence Conference (IAAI-03)*, 2003.
11. Wright, W., Schroh, D., Proulx, P., Skaburskis, A., and Cort, B. Advances in nSpace – the sandbox for analysis. Poster at the 2005 International Conference on Intelligence Analysis.