

The Sandbox for Analysis – Concepts and Methods

William Wright, David Schroh, Pascale Proulx, Alex Skaburskis and Brian Cort
Oculus Info Inc.

{bill.wright, david.schroh, pascale.proulx, alex.skaburskis, brian.cort}@oculusinfo.com

ABSTRACT

The Sandbox is a flexible and expressive thinking environment that supports both ad-hoc and more formal analytical tasks. It is the evidence marshalling and sense-making component for the analytical software environment called nSpace. This paper presents innovative Sandbox human information interaction capabilities and the rationale underlying them including direct observations of analysis work as well as structured interviews. Key capabilities for the Sandbox include “put-this-there” cognition, automatic process model templates, gestures for the fluid expression of thought, assertions with evidence and scalability mechanisms to support larger analysis tasks. The Sandbox integrates advanced computational linguistic functions using a Web Services interface and protocol. An independent third party evaluation experiment with the Sandbox has been completed. The experiment showed that analyst subjects using the Sandbox did higher quality analysis in less time than with standard tools. Usability test results indicated the analysts became proficient in using the Sandbox with three hours of training.

Keywords

visual analytics, information visualization, sense making, gestures, user-centered design

ACM Keywords

H.5.2 [Information Interfaces]: User Interfaces. I.3.6 [Methodology and Techniques]: Interaction Techniques.

INTRODUCTION – TRIST AND THE SANDBOX

As part of the Novel Intelligence from Massive Data (NIMD) research program [2], new interactive, information visualization techniques are being investigated which tightly couple massive data, software agents and the analyst’s exploration task. A break-through in finding novel intelligence is believed possible if all the components are combined in a system of systems.

During the NIMD program, several new interactive interface techniques have been developed and tested that utilize information visualization’s ability to amplify human cognition by increasing mental resources, reduce search time, enhance recognition of patterns, allow perceptual inference, allow monitoring and provide a manipulable medium [5, 15, 16]. These techniques have been combined into an integrated cognitive workspace where intelligence analysts see, and interact with, more information, more quickly, with more comprehension. This space is called “nSpace” and is the combination of the multi-dimensional linked views for triaging massive data found in TRIST, “The Rapid Information Scanning Tool” [13], with the visible and flexible cognitive mechanisms of the Sandbox.

nSpace is a unifying human information interaction interface and an integration platform. The nSpace ambition is to be a system of systems, performing in combination with computational linguistics, reasoning and information retrieval technologies, ensuring use of the most effective technologies for the task at hand. This paper focuses on the Sandbox component and discusses its rationale, concepts and capabilities.

COGNITIVE TASK ANALYSIS FOR THE SANDBOX

In preparation for the Sandbox investigations, three types of cognitive task analysis (CTA) studies have been completed including structured interviews [25], activity analysis of data logged by the NIMD Glass Box [7] and content analysis of selected Glass Box (GB) observational data.

From the structured interviews, conducted with fourteen analysts, it was clear that working with a local knowledge base or shoebox of pertinent information including references and annotations is a common daily practice. The following are illustrative interview excerpts:

- “I put the information in piles on the desk sorted by topics and arranged in time. The topics are what I will write about. I make notes on anything and everything”
- “Research is gathering bits. Analysis is what it might add up to.” It can take two weeks to two months to bring together disparate raw data into a set of connections and patterns of activity. The patterns can continue to be refined and developed.
- “Organization is half the battle. I use ppt to organize my thoughts. e.g. Five hypotheses and 100 pieces of

evidence, incl. indexing, writing notes and reasoning. I can spend three or four days compiling results."

- "The basic concept of analysis includes rolling up, drilling down and pivoting. You rollup from lots of detail, clutter, maps. Show results in the same space. You drill down to the needle, or by clicking through. Filtering is important: by geography, time, hierarchy. Pivoting is picking a thing that becomes the focus. Find out more on this element. Spawn a new search."

A recently completed ethnographic study of the intelligence community [14] made similar observations during structured interviews:

- "First, I print the daily traffic I'm concerned with; then I lay out all of the relevant stuff in front of me on my desk or the floor; then I start looking for threads."
- "I'm looking for links and patterns. Once I figure out the pattern, I can figure out where to look next."
- "I come up with a few scenarios and see what the evidence supports."

Analysis of analyst activity data logged by the NIMD Glass Box [6, 7] was also done. The Glass Box (GB) provides an instrumented infrastructure that captures data about what an analyst is doing during information retrieval, reading, analysis and reporting activities. Events logged include keyboard/mouse, window (active window, active application, location on screen, etc.), file save, copy/paste, Web browser events (URLs, queries, and query results) as well as events for those tools with logging capabilities. Screen captures operate at one image per second.

With a focus on the "analysis methods" activities, GB activity data was processed to gain insight about the "costs" incurred by analysts, the "micro-products" produced as well as friction points. Costs include, for example, time spent by tool and by activity. Micro-products include items like number of references, pastes and documents saved.

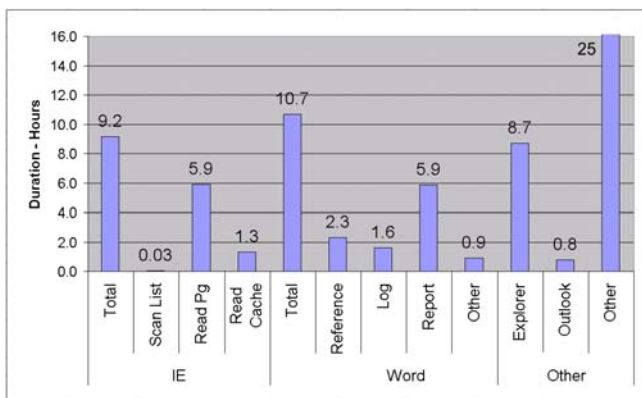


Figure 1. Cost Profile. Times in hours for Average Task.

Figure 1 shows, for two analysts and nine tasks, the average time spent using current tools and processes. Almost equal time was spent in Internet Explorer (IE) doing information retrieval, in Word doing both reference saving and analysis,

and in the File Explorer organizing / navigating directories and files as well using the desktop. Word documents were categorized as reference documents, a log of observations/notes and reports. Note that the File Explorer category is a "catch-all" and includes "gaps", i.e. no observed activity. The "Other" category includes applications such as Reference Manager, Adobe pdf reader, Netscape browser and the Glass Box itself.

Of additional interest, are higher level observations about "content" objects. Observations, for example, about how many hypotheses were made, or how many elements of evidence were marshaled for assertions, or how many trains of thought were noted. These kinds of observations are possible but require human coding of content. For one analyst across two tasks, observations about process and the content were extracted by watching the GB screen capture videos. Using a framework of guidelines, judgments were applied to categorize document content into assumptions, evidence, reference and other content objects. Counts of content objects for one task is shown in Figure 2. Note that

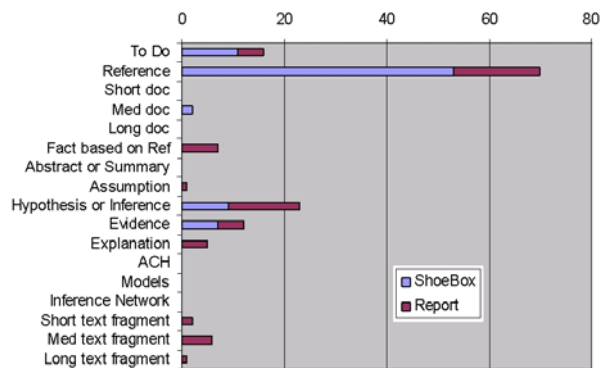


Figure 2. Counts of Content Produced by an Analyst for One Task.

the "Shoebbox" was a log file called notes.doc and the "Report" was the final delivered file Bulletin.doc. Of interest is that there seems to be no observable use of analytical methods such as ACH (analysis of competing hypotheses), models, or inference networks. The main analytical method being used seems to be evidence marshalling. Tying the pieces together appears to have been done in the analyst's mind and/or as part of the report creation, or was done offline and unobserved. In fact, the analyst printed at least part of the shoebbox to more easily refer to it while writing the report.

Observations suggest evidence marshalling and synthesis is particularly difficult. To get the "big picture" by looking at many pages of text, the analyst relies heavily on memory to connect the dots. Many analysts decide to print their shoebbox, spread the pages to see most of them at once, and also make generous use of highlighter ink. There were few analytical methods observed. Considering the relatively short term tasks studied and the difficulty of building say inference networks or other analytical charts in Word, it is not surprising. In the end, this seems to indicate that

analysis (hypothesizing, fleshing out hypotheses with evidence that proves or refutes, corroborating, evaluating the diagnosticity of the evidence, and assessing which hypotheses are most likely, etc.) is something that happens mostly in the analyst's head. Researchers [8] have warned about the consequences of doing analysis primarily in one's head. The analyst is more prone to cognitive biases. It may not be as easy to collaborate and to communicate about the analytical process and how the conclusions were reached than if analytical methods were used and so made the thinking explicit.

RELATED WORK

Post It Notes are a common medium for analysis (Figure 3a). They are inexpensive, easy to use, flexible in terms of what can be captured on them and how they can be organized. Different background colors, size, even shapes are available, which can be used to encode metadata about the content of the note. Of course, an unusable critical mass of notes is reached fairly quickly, and content is difficult to scan, edit, re-arrange, update or transfer.


Concept Map software quickly captures and organizes thoughts with hierarchical linked text [11]. In *MindManager* [17], Figure 3b, typing can be started anywhere in the thinking space. Text formatting (size, color, font, style) encodes some metadata. Thoughts can be placed anywhere but the dominant organization mechanism is a hierarchical tree with the root node centered in the middle of the space. Visual vocabulary and interactions are limited. An automatic layout enforces order but arrangements are not flexible. Even for tree-like data structures, usability decreases as the volume of data increases.




3a. Post It Notes



3b. MindManager



3c. i2 Analyst's Notebook



3d. Visual Links

Figure 3. Related Systems.

Link analysis software, such as the *Analyst's Notebook* [12] and *Visual Links* [24], Figures 3c and 3d respectively, has

diagrammatic visual representations and is mainly used for visualizing connections (e.g. transactions, phone calls, 'is-related-to' relations, etc) between various types of entities. The icons and brief labels facilitate scanning. Over-simplification is possible as the full content and context of information is hidden. It seems better suited as a report tool than a thinking tool since it does not encourage various alternative thinking. As it is time consuming to layout the information, analysts may get attached to a particular layout and be reluctant to try different organization schemes that could provide different insights. Furthermore, assumptions and evidence are not easily distinguishable.

Also of interest, the Data Mountain [22] investigated 2D and 3D spatial layouts to take advantage of human spatial cognition and spatial memory when working with large numbers of documents. Flatland [19] augmented white boards allowing users to get white space more easily without having to remove content from view, to improve the readability of common white boarding artifacts such as to-do lists and drawn maps, and to retrieve past content using time tags and context queries. Finally, information foraging theory [20, 21] assumes that people, when possible, will modify their strategies or the structure of the environment to maximize their rate of gaining valuable information. The Sandbox work presented here is focused on the "sense-making loop" and the "exploiting" process of the exploration-enrichment-exploitation stages of foraging.

SANDBOX PERFORMANCE OBJECTIVES

The first objective for the Sandbox is to be a flexible and expressive place for visual thinking and working with evidence. Analysts need to be able to quickly and easily place, arrange, re-arrange, group, emphasize, highlight and compare. Brainstorming, generating hypotheses, marshalling evidence, collaborating and reporting, all need to be supported. It must be flexible enough to support many types of analysts and analytical styles. Manipulation and organization of information should be direct and tactile. Intuitive iconic representations of entities, concepts and relations constitute a common, shared visual vocabulary. It should be a visual thinking space for considering evidence in multiple frames, designed for user driven tactics in evidence marshalling and sense making like "people-places-organizations-events" and "put-this-there" cognition.

Another objective is for the Sandbox to support best analytical practices without imposing additional training requirements or cognitive and procedural strain. There are a variety of analytical practices including organize-conceptualize-hypothesize, the application of scientific method, competing ideas as well as the clear and explicit representation of reasoning. Additional analytical methods now in use include [3, 8, 9]:

- Generate hypotheses before searching for evidence;
- Multi-dimensional analysis of process, organization,

- Inference networks;
- Analysis of competing hypotheses;
- Use multiple strategies simultaneously for hypothesis generation and evidence marshalling;
- The seven methods of evidence marshalling;
- Situational logic. Constructing scenarios with evidence; and
- Toulmin argumentation.

Supporting a larger volume and complexity of analysis is also an objective. The Sandbox needs to scale to massive data. Analysts already work with large amounts of complex data, and need to be able to work with even larger amounts. More evidence, issues and options need to be considered.

A key objective for the Sandbox is to provide a measurable increase in performance. More and higher quality analysis must be able to be performed in less time. As a first step, the aim is to support the same quality of analysis and reporting in less time. The longer term objective is to increase productivity over the whole analytical workflow with fewer steps, higher quality results, in less time.

Sandbox performance objectives also include supporting multiple simultaneous tasks as well as providing an initial capability for collaboration. Analysis is not sequential, and because analysts are engaged in multiple simultaneous tasks, it must be easy to switch contexts. Collaboration is a creative working mode and operating the Sandbox with multiple analysts at the same time would allow shared perspectives and shared construction of lines of thought.

SANDBOX PROTOTYPE CAPABILITIES

Introduction

The Sandbox is a work space that supports visual thinking, providing alternatives to paper or text editors for analysis activities such as hypothesizing, fleshing out hypotheses with evidence, corroborating, grouping, annotating and prioritizing. Explicit thinking helps ensure more rigorous thinking and thus reduces the impact of some cognitive biases. Explicit visual thinking increases an analyst's cognitive span [5], and also makes it easier to collaborate by making visible the nature and structure of the analysis.

An analyst would typically start an assignment by brainstorming and making notes in the Sandbox. As shown as the top of Figure 4, questions would first be posed. Information retrieved using TRIST (e.g. images, snippets, documents and entities as shown in Figure 5) is placed in the Sandbox. This information (observations, evidence and references together with the analyst's own notes) is arranged, linked and grouped according to topics and issues as shown in Figure 4. Evidence is organized into meaningful patterns. The analyst develops an explicit understanding by iteratively collecting, arranging and writing. Formal analytical methods assist in weighing the strength of assessments and understandings. Analysis

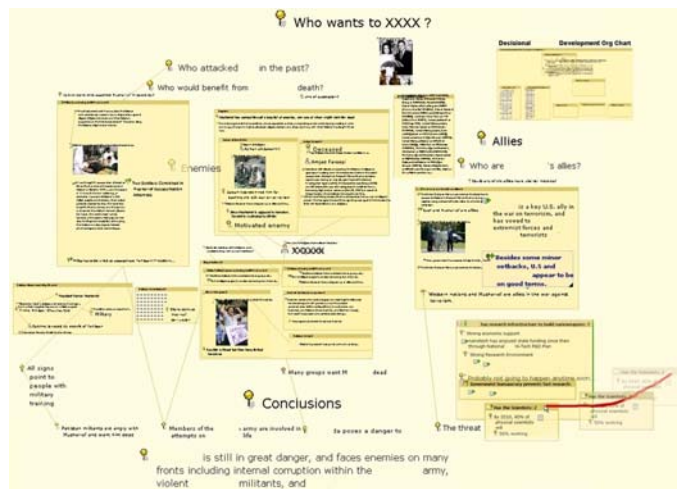


Figure 4: The Sandbox

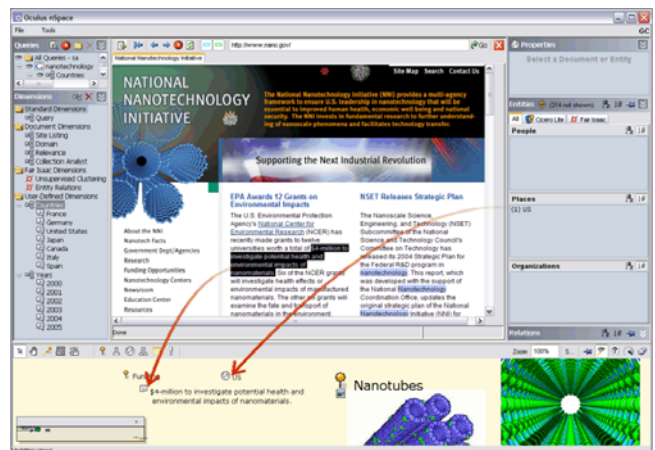
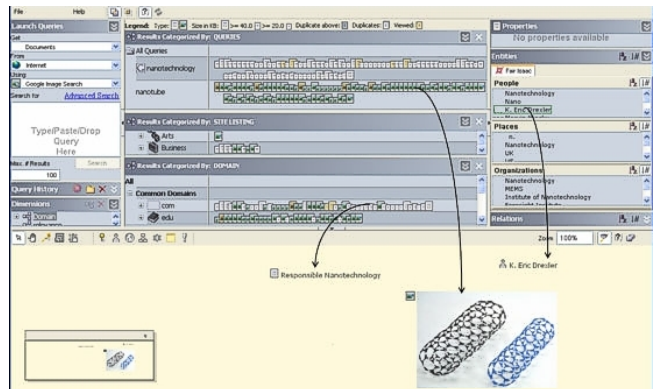


Figure 5: Placing Objects from TRIST into the Sandbox

model templates can be applied to a set of evidence to provide a new point of view. Assertions with pro/con “gates” make explicit the supporting / refuting evidence. The analyst can walk through a Sandbox, highlighting items of importance, to review and collaborate with a colleague. At the end, the analyst can make a report by dragging and dropping from the Sandbox to MS Word. Information is automatically formatted and references are maintained.

Source Attribution and Context

The analyst can save any relevant information, including documents, snippets, images, tables, etc. by dragging them into the Sandbox from TRIST as well as MS Word, MS Explorer, IE and other systems. TRIST uses linked views and multi-dimensional frameworks to speed scanning and identification of relevant reference materials. Dimensions can be system or user defined and can be properties, for example, such as “time”, “country”, “source” etc. TRIST also includes entity extraction dimensions as well as document clustering for content-generated dimensions.

When relevant information is saved by dragging it into the Sandbox from TRIST, references are saved too. No time consuming manual tracking of sources is required. After selecting the information in the Sandbox, its source is highlighted in TRIST, making it easy to check context and scan all the metadata associated with the source. So for example, if country, time and source are used as a multi-dimensional framework, highlighting a piece of evidence in the Sandbox will highlight that same evidence in that TRIST multi-dimensional framework. This approach also eliminates the context switches that take place when collecting information and later going back its source.

One piece of evidence, or *all* the evidence used in an analysis argument can be selected and then seen in the TRIST framework. The analyst can quickly see, for instance, if most of the evidence in an assessment is old or new, or from one or several countries, or from one or multiple sources. The source context of the evidence is easily accessible.

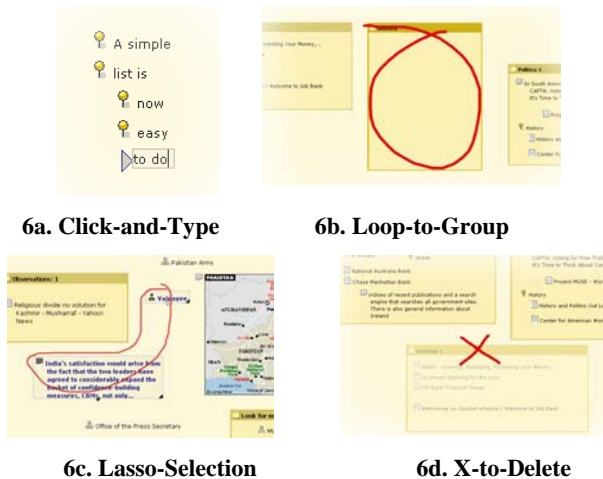


Figure 6: Simple Gestures Help Interactions with Ideas

Put-This-There Visible Flexible Cognition

Direct manipulation and annotation are used to build and express meaning. Analysts can write anywhere, group anything, place anything anywhere, add and edit entities (e.g. people, organizations, places) and add and edit text snippets. Whole documents as well as images can also be placed there and used as references. Links allow connections to be made among items.

Thinking happens quickly and must not be frustrated by a burdensome medium difficult to operate. A minimum of “buttonology” and no dialog boxes or forms are used to create, edit, place, arrange, re-arrange, group, emphasize, highlight and compare. A few simple, naturally expressive gestures link thinking with action. Examples are in Figure 6

Powerful Finger

Level-of-Detail (LOD) can be changed dynamically with the ‘powerful-finger’ mechanism. A simple gesture, the red upward stroke in Figure 7, over an area of interest will progressively add more details (e.g. seeing more text) or scale the item to add emphasis up to the desired effect. The revealed details can either be persisted (if for example that information is particularly relevant to the current train of thoughts) or be temporary and be allowed to fade back (if for example the details were quickly highlighted for emphasis while collaborating with a colleague).

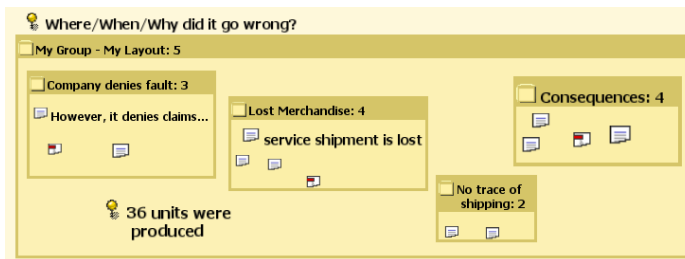


Figure 7. Powerful Finger Gesture Adds Emphasis.

Analytical Model Templates

Models of, for example, a manufacturing process, provide a structured framework to think about subjects of interest and events. A process model helps explain what is happening, why it is happening and what can be done to change it. Many analysts use and invent models or templates to analyze situations. Models help share and harness analytical knowledge and experience. In addition, different models help the analyst exercise a different point-of-view to help think out of a standard mindset or pre-disposition. Finally, when an analyst has a process model in mind, or in the Sandbox, what they don’t know jumps out at them. Gaps are evident.

The Sandbox provides automatic evidence layout using “templates” of analytical frameworks. Normally, evidence (e.g. snippets, observations, documents) is organized in a specific way to support the analyst’s put-this-there cognition, as shown in Step1 of Figure 8 with an analysis of a fictional shipping problem. Step 1 is as the analyst thinks of the situation. But to apply another point of view, the analyst opens a new template of, for example, a process model, Step 2, and the existing evidence is automatically placed in the new layout using continuous animation to maintain some sense of context and change. Placement of evidence reflects similarity or closeness of fit as shown in Step 3. Evidence that does not fit is placed outside and might prompt edits to the template/model, or rethinking of



Step 1. Exploration Workspace.



Step 2. Blank Manufacturing Process Template.



Step 3. Manufacturing Process Automatically Filled with Items from Exploration Space.

Figure 8. Using Templates.

the analysis. Placement inside the template could also reflect certainty of fit (e.g. center of node vs. on the edge). The template reveals a gap in the “Package” stage and the analyst wonders if there’s a problem with labeling. The Sandbox allows the creation and use of models to support various concurrent ways of thinking about a problem without adding any additional cognitive strain. The purpose is to lower the cost of looking at evidence from multiple perspectives, trying different models, which would counter rigid mindsets biases while highlighting the strengths and deficiencies of all models considered. For example, an organizational chart from one context may be applied to a different context to see if an organization under study may be more or less than expected.

Other applications of this technique include having many such layouts placed side-by-side, so the analyst can draw conclusions such as which model best fits the data. Another analysis method would be to use the linked selection mechanism across multiple models, to select evidence, to see where it is used and so reveal the diagnosticity of particular evidence items.

Document similarity using context vector clustering from Fair Isaac [4] is used to implement templates. Snippets are scored and then assigned to the closest cluster. The system can use lightly supervised methods. Training is with keywords, snippets and sample documents. Subsequent “learning” and improvement is possible with “experience”.

Dragging Hand Gesture Through Evidence Gates

The explicit representation of assertions and how evidence supports or refutes each assertion is another mechanism that has been implemented to support best practices in analysis.

Assertions make explicit the points the analyst is trying to prove/ disprove. In the Sandbox, an assertion is represented as a group. For the title, a statement is entered that needs to be proved true or false. The assertion group has “Support and Refute Gates” along the sides. The act of binding evidence to an assertion is done fluidly by dragging the evidence to the assertion group through the appropriate gate. The supporting gate is along the right edge and the refuting gate is on the left of the group box. The gates light up red or green when evidence is brought into the assertion. The green supporting edge is lit up in Figure 9. In the assertion, supporting evidence items receive a green plus sign. Refuting items receive a red negative sign. Neutral objects are dragged through the top or bottom of the assertion.

‘Refute Gates’ encourage the application of the scientific method, where the focus is on disproving hypotheses. Nesting assertions creates inference networks. In addition, at the top-left corner of an assertion, a graphical cumulative ‘score’ is displayed for an ‘at a glance’ reminder of the quantity and weight of evidence found so far. Weight or significance of contribution is set by the analyst by repeatedly dragging the evidence through a gate. (Right clicking a plus/minus sign also brings up a five point selection scale for significance.) The cumulative score takes into account all evidence including nested assertions.

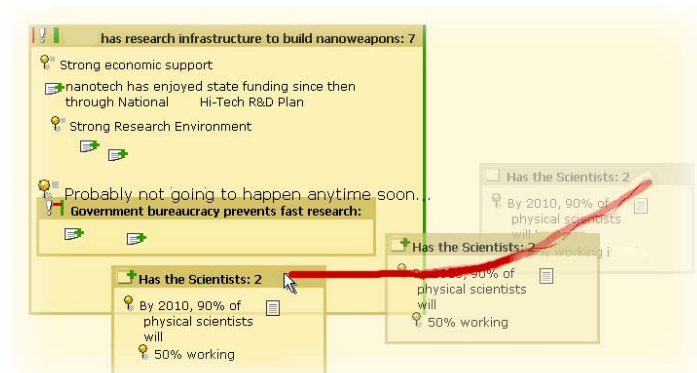


Figure 9. Assertions with Evidence Gates.

Assertions can provide a quick visual understanding of the components of the analysis and how they interact. Assertions also allow useful meta data about evidence to be easily encoded during the gesture.

Using Evidence From Other Systems

Observations contained in snippets of text from documents with source attribution is one type of evidence. Another type of evidence may be found with computational assistance. For example, agent-based simulation models may reveal an unexpected correlation in system behaviors, or an hypothesis generation system may suggest a new possibility. The Sandbox allows analysts to drag and drop data elements to and from specialty analytical systems. Once the analyst is finished using the specialty system, any conclusions, notes and supporting evidence can be returned to the Sandbox where it can be used in a larger structured

argument or analysis. A thumbnail screen shot of the specialty system is also added. An example is shown in Figure 10 for a specialty system called GeoTime that analyses events in time and geography [15]. Clicking on the thumbnail invokes the specialty system, in it's own windowing system, with the full data and application context associated with that conclusion.



Figure 10: Integrating Evidence from a Specialty System

Managing Space

Using the Sandbox, one objective is to allow the analyst to deal with many more plausible scenarios, at once, than could have been done before. To support a higher volume of information and thoughts, in a 'put-this-there' environment, without hindering productivity, means addressing the mechanical cost of arranging items when the aim is simply to create or delete space for example. This is accomplished with a simple gesture encoding the command and its parameters. An amount of space is created or deleted in proportion to the size of the gesture at the location and in the direction performed. The sandbox items are animated to their new location and collision detection with friction is used to insure that everything will stay visible and that the order of layout will be maintained.

In Figure 11, an insert gesture is used to add space. Items are animated away from the area while maintaining relative positions as well as continuity of context for the user. A gradual dampening is applied so that items further away move less until no movement is done at all. With collision detection, similar to Data Mountain [22] and QuickSpace [10], another way to create space is to grab one item and then use it to knock other items away.

Scalability, LOD and Navigation

Each item in the Sandbox can be seen at multiple levels of detail (LOD) individually. At the lowest LOD, intuitive iconic representations are compact and easier to recognize and scan than text. Nothing needs to be hidden. As an example, entity icons encode key attributes of gender and profession as shown in Figure 12.



Figure 11: Creating Space with Editor Gestures.



Figure 12. Entity Icons Encode

LOD can be changed dynamically with the powerful finger gesture. Thumbnail level of detail for groups minimizes a group while still preserving the semantic layout of information. Examples of various levels of detail are shown in Figure 13.

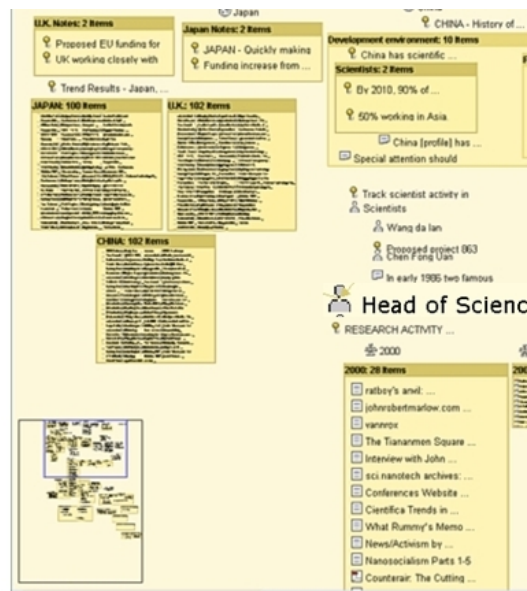


Figure 13: Thumbnail (Lower Left) and Levels of Detail

When an object that is outside the visible sandbox area is linked to a visible and selected object, an icon of the non-visible object is displayed in a white frame around the Sandbox window frame.

Navigation using the Thumbnail Map for the whole Sandbox allows quick directed movement within larger Sandboxes. Clicking in the Thumbnail Map or dragging the

indicated window frame in the map moves the viewer accordingly. There is also a Scroll Tool to drag the screen in any direction, and the middle mouse button zooms in or out of the Sandbox workspace. The Scroll Tool is best used for small movements, while the thumbnail makes it easy to move across large portions of the sandbox in one click.

Concept Maps

Some search engines [1] use concept map structures and models of user interests to specify or modify queries that support the analysis. Concept maps are automatically generated in the Sandbox on any text using a text-to-concept-map algorithm accessed via an integration protocol. The resulting map is displayed in a callout transparent bubble next to the source text. The map is graphically editable by the analyst. Emphasis can be adjusted by interactively scaling a node or link. Alternative terms can be selected with an ontology also integrated via the protocol (see later Technical Architecture section).

EXAMPLE WORKFLOWS IN NSPACE

With a new tasking, the analyst might start brainstorming in the Sandbox, using prior and tacit knowledge to generate hypotheses or alternative scenarios, to jot down key questions and an analytical strategy or plan. Divergent thinking, crucial at this stage, is unrestricted in the Sandbox. The analyst can click on the Sandbox and type annotations anywhere. Thoughts can flow freely, and be quickly organized to create more structure on the fly. Notes can be placed close to an object or can be attached to it by dragging the note on top of the object.

To help investigate hypotheses and find answers to questions, information retrieval is done with TRIST. (Working with two or even three monitors, one for TRIST, one for TRIST's integrated document viewer and one for the Sandbox, saves time and maintains task context in view.) As relevant information is found in TRIST, it can be immediately integrated into the Sandbox "big picture", by either binding it to various concept, thought or argument structures (i.e. an information object) with links or groups, or creating new structures to account for that evidence. Information objects (notes, snippets, documents and images) can be recursively associated with container information objects like entities (person, place, organization, other), groups, assertions and links.

The analyst arranges information within the Sandbox to encode meaning. For example, supporting details are placed with higher level concepts. Related issues are placed next to each other. Overview issues might be placed at the top. Miscellaneous details at the bottom. Perhaps the analysis strategy is placed on the right. Follow-on issues might be on the left. The Sandbox does not force the analyst into a particular model or mode. Using a layout of text (words, phrases, snippets, lists), entities, links and groups within groups, the analyst can describe, think and reason about a situation by building higher level cognitive

mechanisms such as processes, organizations, relationships among strengths/weaknesses, input/output factors, etc., whatever construct helps organize and provide insight. Items more important to the analyst can be emphasized in size. Meaning is built and expressed through direct manipulation of thoughts and information. A diagram of the situation and assessment emerges that consists of words, text, groups, entities and links. In the Sandbox, evidence is marshaled and sense is made of it.

This flexible, visual cognition is supported through ease of spatial arrangement. Drawing a circle around objects selects them for moving. Drawing a second circle around selected objects groups them. Items or groups can be moved and dropped on other groups. Groups can also be moved with collision detection active and so knock other groups out of the way when the group is placed. Editor gestures insert and delete Sandbox space by animating objects out of the way or bringing them closer together.

More formal analytical methods are supported. Assertions with evidence gates combine supporting and refuting evidence and include an at-a-glance indicator of strength. Assertions can be combined into an analysis of competing hypotheses (ACH) framework. Visual cues indicate diagnosticity of the ACH evidence. An evidence item in many alternatives is blurred while an item uniquely supporting one alternative hypothesis is sharp. Analytical models, encapsulating expert knowledge and best practices, can be quickly applied and un-applied to a situation to provide another point-of-view. To reduce the cognitive biases associated with a particular mindset, exploring different perspectives is a key workflow. Trying to fit the collected evidence to different models might reveal gaps or different possible interpretations of the evidence.

Tight integration with TRIST supports fluid iterative discovery, blurring the boundaries between information retrieval and analysis. For example, entities in the Sandbox can be dragged to TRIST to form an entity query, and of course TRIST query results can prompt new Sandbox hypotheses. Any objects found in TRIST and saved to the Sandbox will always be linked to TRIST. Selecting a Sandbox object quickly shows which query it came from and the corresponding contexts (i.e. position in various dimensions). This linked viewing with TRIST helps provide credibility assessment of evidence in the Sandbox.

At any stage, the big picture in the Sandbox provides at a glance what needs to be proved or disproved, where evidence is lacking, and helps keep the collection and analysis process focused and efficient. Interim results can be interactively reviewed with colleagues. Final reports can be generated by moving selected elements to MS Office, or by making an interactive final report in the Sandbox itself.

TECHNICAL ARCHITECTURE

nSpace hosts two client applications: TRIST and the Sandbox. The applications are written in Java, and connect

to other graphical interfaces, analytical technologies or middle-tier system through standard protocols. An underlying scene graph management and rendering engine is used that provides access to OpenGL graphics acceleration and the latest video hardware. This allows Sandbox interactions and animations to scale smoothly from a just few objects to a large complex analysis.

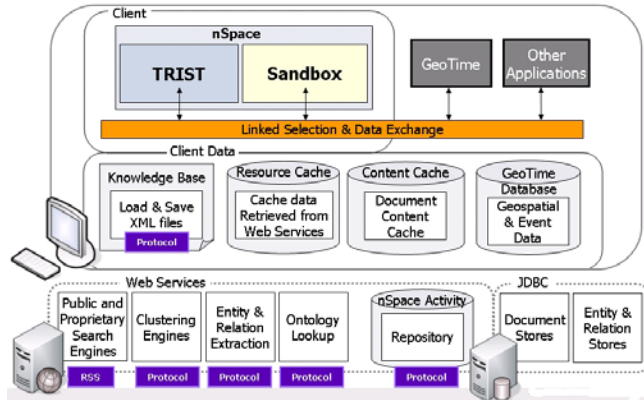


Figure 14: Technical Architecture

As shown in Figure 14, nSpace has a multi-tier architecture for scalability and ease of deployment. Existing open XML standards and new Web Services interfaces encapsulate services in each tier to provide scalability, modularity and data processing functionality. The client application uses a local XML repository to store its knowledge base. In addition, all analyst activity is published to a repository on the services tier (and also to any other services that wish to subscribe to user analytical activity).

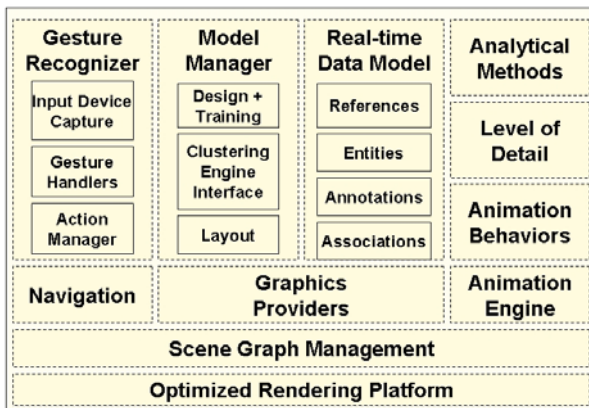


Figure 15: nSpace Sandbox Organization

Internal Organization

As shown in Figure 15, the Sandbox is arranged with an underlying scene graph management system together with specific data models, view generators and interaction management. The graphics providers implement business logic, visual vocabularies and controls for interacting with data objects. Associations provide the mechanism for linked views. Navigation is achieved via the gesture system or through interactions with the thumbnail renderer. The model manager implements the analytical templates.

Gestures

A single and multi-stroke gesture recognition system monitors the different kinds of interactions the user can make with the Sandbox using various input devices (mouse, stylus, keyboard, tablet) and defines a simple small set of mechanisms that serve as intuitive controls. Points are captured, a smoothing algorithm is applied, the curve is scored and then matched. Feedback to the user on the match is provided as the gesture is done. Normal interactions (e.g. picking, scaling, dragging) use the same gesture pipeline but have standard windows handlers.

Protocol

Integration with local and remote computational linguistic, knowledge and analytical services is through an nSpace Web Services protocol. The Protocol defines a knowledge representation format and a standard, open XML interface and extensible messaging scheme for exchanging data. Different components (e.g. search engines, ontologies, text-to-concept-map services, entity extraction) can be “plugged-in” as required.

EVALUATION EXPERIMENT

An independent third party, the National Institute of Standards and Technology (NIST) Visualization and Usability Group (VUG) performed an evaluation experiment of nSpace-Sandbox in March 2005 using an intelligence analysis task-oriented methodology [18].

Four analysts participated in the evaluation experiment. Three hours of standardized training were provided by NIST. A usability test was administered afterwards. Each analyst worked for eight hours on a well defined, representative analytical task. At the end, questionnaires were administered and debriefings held.

This evaluation with four real users is preliminary but gives an early indication of usability and utility. All four subjects reported being able to work more quickly and to do a higher quality of analysis with the Sandbox than with MS Word.

Additional experiments are planned to further assess the performance impact of the Sandbox in particular and the integrated NIMD system as a whole. These experiments will include comparison with a baseline analytical task, allow more precise observations of task timings, and will also include peer assessment of analytical work products to allow comparison of product quality with the baseline.

CONCLUSION

The Sandbox is an analytical sense-making system. Its’ objective is to be a fluid, flexible medium of analysis and expression that combines human insight with computational linguistic and analytical functions. Analysts sort through, organize and analyze large numbers of bits of data extracted from a variety of sources. The Sandbox offers flexible organization support for thoughts and information relevant to the analysis, encourages making thinking explicit, facilitates the application of various analytical methods,

scales to larger amounts of information and thinking, makes it easier to keep more of that information and thinking available to human working memory at once, and increases analytical performance by eliminating friction points and overhead that steal time from analysis. The system includes innovative interaction and visualization techniques including rapid gestures for placing, moving, grouping and marking information, use of miniatures and level of detail, templates for building visual models of information, machine learning methods for grouping information, visual assessment of evidence and methods for integrating computational analytical systems into a common interface.

In a preliminary experiment, four analysts said they were able to increase the quality and the speed at which they do analysis. They were able to operate the Sandbox functions after three hours of training.

Future work includes further evolving of Sandbox human information interaction capabilities, additional experiments and the continued collaboration with other NIMD researchers to develop a robust integrated system for analysis that combines massive data, computational analytics, human information interaction components and the analyst, all in a system of systems.

Acknowledgement

The Sandbox was inspired by Dr. Lucy Nowell's vision of Sandcastles. This study was supported and monitored by the Advanced Research and Development Activity (ARDA) and the National Geospatial-Intelligence Agency (NGA) under Contract Number NMA 401-02-C-0032. The views, opinions, and findings in this report are those of the authors and should not be construed as an official DOD position, policy, or decision, unless so designated by other official documentation. The authors wish to thank all ARDA staff for their support and encouragement.

References

1. Alonso, R. and H. Li, Combating Cognitive Biases in Information Retrieval, In Proc. *International Conference on Intelligence Analysis*, 2005.
2. ARDA Novel Intelligence From Massive Data, NIMD, http://www.ic-arda.org/Novel_Intelligence/, 2002.
3. Bodnar, J.W., Warning Analysis for the Information Age: Rethinking the Intelligence Process, *Joint Military Intelligence College (JMIC)*, December, 2003.
4. Caid, W. and Pu Oing, System and Method of Context Vector Generation and Retrieval, *United States Patent* 5,619,709, 1997.
5. Card, Stuart, J. Mackinlay, B. Shneiderman Readings in Information Visualization, *Morgan Kaufman Publishers*, San Francisco, CA., 1999.
6. Cowley, P., L. Nowell, J. Scholtz, Glass Box: An Instrumented Infrastructure for Supporting Human Interaction with Information, In Proc. *Hawaii International Conf. on System Sciences*, 2005, 296.3

7. Hampson, E., and P. Cowley, Instrumenting the Intelligence Analysis Process, In Proc. *International Conference on Intelligence Analysis*, 2005.
8. Heuer, Richard, Psychology of Intelligence Analysis, *Center for the Study of Intelligence*, 1999.
9. Hughes, F. and D. Schum, Discovery-Proof-Choice, The Art and Science of the Process of Intelligence Analysis – Preparing for the Future of Intelligence Analysis, *JMIC Joint Military Intelligence College*, 2003.
10. Hutchings, D. and J. Stasko, QuickSpace, Short Paper, *ACM CHI Conference*, 2002, 802-803.
11. Institute for Human Machine Cognition, University of West Florida., Cmap Tools, <http://cmap.ihmc.us/>
12. i2 Limited., <http://www.i2inc.com>
13. Jonker, D., W. Wright, D. Schroh, P. Proulx and B. Cort, Information Triage with TRIST, In Proc. *International Conference on Intelligence Analysis*, 2005.
14. Johnston, R., Analytic Culture in the U.S. Intelligence Community, Center for the Study of Intelligence, *Government Printing Office, Pittsburgh, PA*.
15. Kapler, T. and W. Wright, GeoTime Information Visualization, *Information Visualization Journal*, Palgrave Macmillan, 4,(2), Summer 2005, 136-146.
16. Larkin, J. and H. Simon, Why a Diagram is (Sometimes) Worth Ten Thousand Words, *Cognitive Science*, 11(1), 1987, 65-99.
17. Mindjet Inc., <http://www.mindjet.com>
18. Morse, E., M. Potts Steves and J. Scholtz, Metrics and Methodologies for Evaluating Technologies for Intelligence Analysts, In Proc. *Conference on Intelligence Analysis*, 2005.
19. Mynatt, E., T. Igarashi, K. Edwards, A. LaMarca, Designing an Augmented Writing Surface, *IEEE Computer Graphics and Applications*, July 2000, 55-61.
20. Pirolli, P. and S. Card, Information Foraging in Information Access Environments, In Proc. *SIGCHI Conference on Human Factors*, 1995, 51-58.
21. Pirolli, P. and S. Card, The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis, In Proc. *International Conference on Intelligence Analysis*, 2005.
22. Robertson, G., M. Czerwinski, K. Larson, D. Robbins, D. Thiel and M. van Dantzich, Data Mountain, In Proc. *ACM Symposium on UIST*, 1998, 153-162.
23. Rose, Russ, ChairP1000 Committee, P1000 Report, *Office of Research and Development, CIA*, 1996.
24. Visual Analytics Inc., <http://www.visualanalytics.com>
25. Wright, William and Kapler, Thomas, Speaking with Analysts – Observations of Current Practices with Massive Data, submitted to *Journal of Intelligence Community Research and Development*.