

Design of Zoomable User Interfaces And Lens Tools For Intelligence Analysis

Raymond D. Rimey, William N. Klarquist, Christian H. Debrunner *

Lockheed Martin Space Systems Company, Astronautics Operations
P.O. Box 179, Mailstop 4370, Denver, CO 80201
Raymond.D.Rimey@lmco.com

Abstract

This paper introduces the ideas of zoomable user interfaces and lenses and explains how these ideas can benefit the design of future intelligence analysis workstations.

1.0 Introduction

A zoomable user interface (ZUI) is based on an infinite drawing surface in which graphical objects can have any real-valued scale. The elevated prominence of scale leads to several new user interface design concepts. One of these concepts, called a lens, is an object (widget) that sits on the drawing surface and provides a selective view onto some part of the drawing surface. Usually a lens' view looks on the area directly underneath the lens, but more generally the view can look on any area of the surface at any zoom, in which case the lens is referred to as a portal. Lenses are a natural choice for viewing and interacting with things on a zoomable surface. Another important aspect of zoomable user interfaces is that people's highly developed spatial reasoning, recall and manipulation skills can be utilized extensively in the interface design.

This paper explores how zoomable user interfaces and lenses can benefit the design of future intelligence analysis workstations. The most obvious mapping of these new concepts is to present the

user with a single geo-referenced world surface with embedded imagery and intelligence data (at various scales). This paper also touches on other mappings of zoomable spaces, such as time, task structure, and tool palettes. The presentation includes illustrative examples from within a research prototype for an image analysis workstation that we are building to explore these user interface concepts and associated interactive target recognition technologies and semi-automated image understanding technologies. The paper also discusses how these user interface design concepts address specific needs of the image or intel analyst.

Zoomable user interfaces have many potential applications beyond military image analysis, related intelligence analysis domains, command post systems, etc. Map-oriented application domains include commercial geographic information systems (GIS), planetary science database systems, and personal interactive maps such as a downtown attraction and street map, a restaurant guide, a museum map, etc. Other uses are database information browsers, zoomable web pages, or individual zoomable images within an internet shopping catalog.

Our own research, however, focuses on the military image analyst domain. Our project and the associated prototype system are called INIMEX (INteractive IMage EXploitation). The top-level goals for this research are as follows:

- Give the image analyst ready access to geo-registered context information: maps, current and historical imagery, reports, and collateral data such as signals intelligence, etc.

* This work is sponsored by DARPA, contract number F33615-97-C-1097, monitored by the U.S. Air Force Research Laboratory.

- Give the image analyst an intuitive and efficient user interface so the analyst can incorporate large amounts of information into an analysis task.
- Give the image analyst intuitive methods for navigating within a 2D world space, and for navigating through a spatio-temporal world space. The analyst must maintain a good sense of spatial context.
- Give the image analyst intuitive methods for selectively viewing and navigating through this information. The analyst must maintain a good sense of task-specific information context.
- Give the image analyst insight into and control over interactive or semi-automated analysis aids. This project focuses on (a) interactive aids for cueing and identification of vehicles utilizing model-based vision, and (b) semi-automated aids for wide area search and facility monitoring utilizing image understanding technology.

Our user interface research approach strives to avoid the drawbacks of a classic command-based WIMP (windows, icons, menus, pointer) interface design. Current image and intel analyst workstations provide exactly this type of user interface. Such workstations include classic electronic light tables (ELTs) from many commercial vendors, research prototypes for next-generation image exploitation workstations utilizing image understanding technologies (such as the RADIUS testbed [3] and Target Jr. [6]), and other more comprehensive intelligence analysis workstations (such as the Army's All Source Analysis System).

The general appearances of these interfaces are all somewhat similar and familiar: a packed menu bar, numerous levels of sub-menus, several modes, numerous choices of icon bars or palettes, multiple-combination keyboard and mouse shortcuts, and a central data viewing window that operates like a classic drawing program. These interfaces require a long time to learn and can be complex to use even for common tasks. The user must memorize a huge number of visible command choices and navigate through the interface to find the remaining commands that are not made directly visible. The interface is command-based, which requires multiple steps (e.g., select object, select operation, apply, repeat on same object until done),

and requires large spatial/context switches (move from command to object concept, move from one part of screen to another), and requires that all commands for all objects be organized within the top-level user interface. Modes add further interface complexity and can confuse the user. The interested reader is directed to [7] for an excellent discussion of the deficiencies of command- and WIMP-based user interfaces.

In contrast, our user interface research emphasizes direct manipulation within the single window that presents the geo-referenced workspace. This is not just a data viewing window; The entire INIMEX user interface co-exists with the data in this workspace. INIMEX provides one infrequently used global menu (providing collaboration facilities), and currently provides one pop-up menu that is common to all lens tools. All other interface elements are direct manipulation tools, usually lenses, that reside on the geo-referenced surface. New instances of tools are created on the geo-referenced surface via a zoomable tool palette, which provides several novel methods for a user to select or create needed tools.

This paper is organized as follows. Section 2 introduces the system concept presented to the user, in which all display and all interaction is performed within a single geo-referenced surface. Section 3 discusses zoomable user interface design. Section 4 covers lens tool design in detail. Section 5 summarizes our work on zoomable tool palettes. Section 6 summarizes our novel results to date.

2.0 Pure Geo-referenced Interface

The conceptual model that a user has of the INIMEX workstation is a single geo-referenced surface, viewable at any scale, in which all data is embedded with an appropriate position, shape and scale. Every piece of data has a location in the geographic coordinate system (latitude and longitude) and has some associated notion of size and shape. In particular, all images are warped so their pixels are properly geo-registered. This user model is sometimes called a sandbox. Many types of data objects may be contained in a sandbox, as illustrated in Figure 1, including for example: maps, area delimitations, lines of communication, site models, signals intelligence reports, electronic intelligence reports, various types of imagery,

moving target track data, force level analysis, analyst reports, and sensor platform tasking.

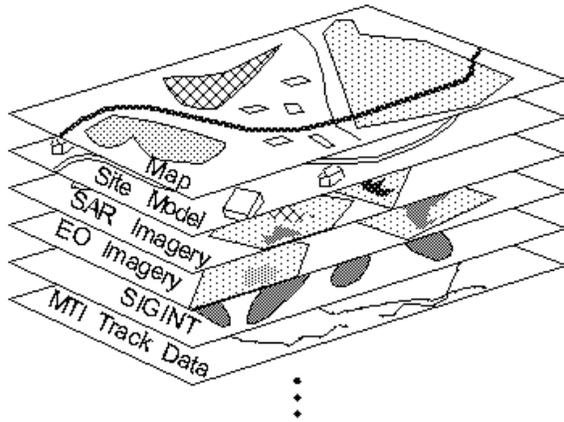


FIGURE 1. The sandbox geo-references all data and presumes that the user will select, view and work with only a portion of the data during a specific analysis task.

It is important to emphasize that INIMEX has no user interfaces (display or interaction) outside of the sandbox. Many emerging intel analysis systems include a map display on which some data (though rarely images) can be overlaid, but not all data can be overlaid on the map, other data is accessed and displayed using separate windows with their own interfaces, and each window display is surrounded by areas containing menus and palettes.

Methods for selectively accessing, viewing and utilizing the sandbox data are at the heart of the user interface thrust of the INIMEX system design. There are simply so many data objects that they can not all be viewed at once, and only a small portion of them will be relevant to any given analysis task. The user also needs the ability to see only certain details of a single object, and to visualize such information in different ways.

3.0 Zoomable User Interfaces

INIMEX utilizes a single infinite drawing canvas (called the INIMEX surface), which contains objects such as text, graphics and widgets. In addition, this canvas has the special property that it can be viewed using any real-valued zooming factor. The standard mouse bindings currently used for INIMEX allow the user to pan using the left mouse button and to smoothly zoom in and out with the middle and right mouse buttons. Dragging during a

zoom results in a smooth simultaneous pan and zoom motion.

Zoomable user interface design can utilize the following fundamental ideas, made possible by the availability of a zoomable display surface. Illustrative examples are presented later in this section:

- **Navigation:** Smooth zooming and/or panning is an intuitive method for navigating through a large work surface. The system can offer animated moves, in which the system automatically zooms out to an overview and pan/zooms to a new location.
- **Structured Information Organization:** A collection of information objects can be structured for presentation in both scale (zoom) and space.
- **Selective Presentation of Detail:** A range of visible detail about one object can be structured within the scale dimension around that object. The presentation of more/less detail is often coupled with spatial zooming in/out on an object, and is called semantic zooming.

Humans possess highly developed spatial abilities, which can be utilized extensively in the design of a zoomable user interface. These spatial abilities include navigating to a specific point/area of a populated surface, recalling the position or context of an object on the surface (e.g., for search), reasoning, and direct manipulation of objects.

Zoomable user interfaces are relatively new. Information visualization problems often involve zoomable data views, and some work on 3D user interfaces exists. Zoomable web-page content is emerging. The best known (2D) zoomable user interface research involves the Pad++ system [1]. INIMEX utilizes Pad++, and we are collaborating with the Pad++ researchers.

3.1 Smooth Zoom and Animated Pan/Zoom Movements

The smooth pan/zoom capability of the INIMEX system and the overall ease-of-use for spatial browsing are simple, basic strengths often overshadowed by other more novel aspects of the INIMEX system. The ability to perform smooth zooming, smooth panning, and combined smooth pan/zoom movements is of fundamental importance in an image analysis workstation and in other

map-based military intelligence analysis applications. Current systems provide smooth panning (or an approximation), and most provide jump zooming by factors of two of isolated image displays or of relatively plain map displays.

The image analyst utilizes zooming for basic navigation through a geo-referenced space -- for example to move from analyzing something at one side of a large industrial facility to something at the other side, perhaps one mile away. The continuity provided by an intuitive ability to smoothly pan and zoom greatly enhances an analyst's ability to maintain spatial context during analysis tasks.

The animated two-step pan/zoom capability, presently unique to Pad++ and INIMEX, is a powerful technique that further helps the image analyst maintain spatial context during analysis tasks. This animated move transforms the display from one position/scale to another by first zooming the display out so that both the start and end positions are visible, then panning over toward the end position, and finally zooming in to the final zoom. During this process the user gets a clear idea of the relative positions of the starting and ending locations.

3.2 Organizing Information In Scale

Scale adds a third dimension to the space in which user interfaces can be designed. Every piece of text, graphical object, or widget in a zoomable graphics display system generally has a minimum and maximum screen size at which it is visible. A widget remains active no matter what zoom the widget is viewed at, as long as the widget is visible. Figure 2 shows a simple example of how objects can be organized in scale space and made selectively visible. Here one object, a slide of text, is organized "behind" another slide (specifically in-between two letters on the first slide). This figure also shows some zooming interface widgets, a button and slider that remain active over a range of zoom values. (Many zoomable user interface concepts are most compelling through dynamic demonstrations. However, this paper must rely on static snapshots of display screens.)

The INIMEX application domain adds some constraints on how information can be organized. Since the user is presented with a single geo-referenced surface, objects tied to the real world must

reside at or near their geo-referenced location on the surface. Therefore pages of information should not be haphazardly placed within the INIMEX geo-referenced surface solely for the sake of presentation organization.

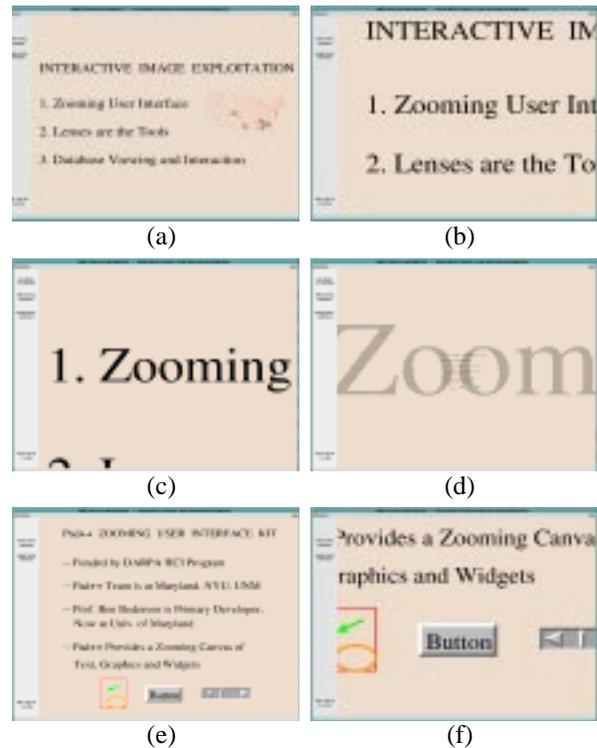


FIGURE 2. Sequence of screens showing how information can be organized in the zoom dimension, and showing widgets that are operable within a range of zoom values.

The user himself can utilize the scale dimension to organize tools on the work surface. For example, if the user has a specific set of tools he uses for analysis task X at location A in an industrial facility, he can leave instances of those tools on the surface, exactly where the tools will be used, and visible only over a range of zoom views. A different set of tools used at a different part of the facility, or used for an overview of the facility, etc., can similarly be left in other places on the work surface. Tools left on the INIMEX surface are remembered from one session to the next

While the image analyst will spend most of his or her time in the geo-referenced world performing analysis tasks, the user interface could utilize additional, clearly distinct, zoomable spaces for other purposes. For example, Section 5 summarizes our idea of a zoomable tool palette, which includes tool palettes organized in a task sequence structure.

Other zoomable surfaces/structures that we have been exploring are for analyst reports, briefings, and temporal analysis problems.

3.3 Semantic Zooming

The analyst's ability to maintain informational context during an analysis task is greatly influenced by what information is displayed and how it is displayed as the analyst works on a task. Work is harder when extraneous information is displayed (screen clutter) and necessary information is not displayed in the most useful manner.

In semantic zooming, an object's visual representation and the information it conveys changes in a coordinated fashion along with the amount of spatial zooming of the object. Typically the amount of detail conveyed will increase as the view of an object is zoomed in more.

A simple example of semantic zooming causes different resolution raster maps to automatically fade in and out as the analyst smoothly zooms in from an overview of a large area to a close-up of a small area. Of course any one map being displayed will smoothly zoom over its visible range. The effect is that the analyst always sees the best map for the immediate zoom level. Existing intel workstations typically require the user to select a specific map for display from a tabular list.

Vector map data consists of point, polyline and polygon data and associated attributes. Semantic zooming can be utilized to selectively display map elements (e.g., display increasingly smaller road types when zooming in) or to selectively display attribute values (e.g., zooming into a mountain peak displays that peak's attributes).

Semantic zooming can be extremely useful for pigeon-holing military deployment information into INIMEX's geo-referenced world. For example, in Figure 3, as the user zooms in on a specific target vehicle in a synthetic aperture radar (SAR) image, increasing amounts of vehicle information are displayed. The information displayed here reflects automated analysis performed by the MSTAR target identification system [5]. Initially a number of target identification objects are visible, displayed as boxes, which indicate only that vehicles have been detected there. Further zooming

shows: the highest scoring target identification hypothesis, the ranked list of identification hypotheses, the addition of a bar graph depicting the matching scores for each hypothesis, and finally numerical labels for the bar graph. Note that as an alternative to semantic zooming, each one of the visual representations for a single target identification object could alternatively be provided by a separate lens (see Figure 11).

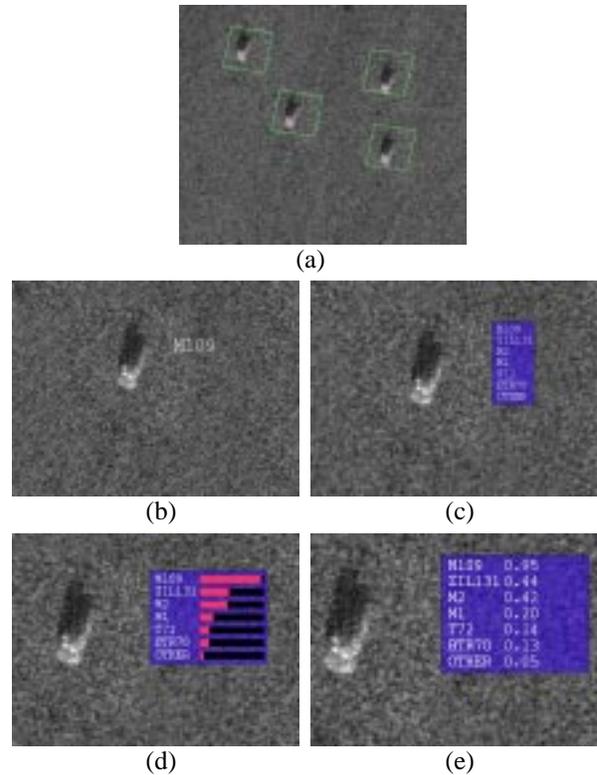


FIGURE 3. This sequence of display fragments shows one form of semantic zooming of the results from a automated target identification system that analyzed the SAR image shown in the background. These display fragments are from inside the view of a lens that is too large to be visible here.

Many different types of data objects are typically available within a military intel analysis workstation (e.g., signal intelligence hits, activity reports) and typically contain a range of detailed attributes that are natural candidates for semantic zooming. For example, semantic zooming of standard military symbology denoting unit structure and control measures (i.e., area boundaries) for a deployed force would depict the symbology corresponding to different echelon deployments (e.g., company, battalion, regiment, brigade, etc.). Once a formation of enemy vehicles is detected and verified with the aid of tools like the target id lens in Figure 3,

the analyst might want to see the situational context of that formation, which a zoomed-out military symbology lens can provide, as in Figure 4(c).

4.0 Lenses

Lenses are user interface tools that provide a window with different views of the INIMEX data surface. Lenses reside on the INIMEX surface. A lens shows a particular visual representation of a particular subset of objects that overlap the lens, and the user interacts with objects through a lens, or in some cases interacts with the lens through various kinds of controls embedded in the lens itself. What a lens displays is called a view, since it is a view onto the INIMEX surface.

Lenses provide the following benefits in the design of user interfaces:

- Lenses reduce graphical display clutter and informational display clutter. Instead of using global overlays, a lens can be scaled to cover a portion of the display or a smaller lens can be scanned over an area of interest. Multiple kinds of data can be viewed simultaneously, without clutter problems, by using multiple lenses.
- Lenses help the user maintain spatial and information context. A lens and its view are always visible within the context of whatever is displayed in the global screen display. Selective use, placement and scanning of lenses allows a user to view exactly the necessary data throughout the process of doing an analysis task.
- New lenses can be constructed by stacking and optionally gluing multiple lenses together. The user can (interactively) create a complex lens from primitive lens elements to display any combination of information.
- Lenses provide real-time response to the user, even for operations that involve intense computation. For example, a lens can perform image enhancement in real time. This is possible because a lens typically shows data for a small area, which can be processed in real time. Large areas can be quickly scanned by a small lens.
- Tools can be left on the surface for future use at that area. Since military analysts often view their assigned analysis sites at regular time intervals and perform the same basic analysis

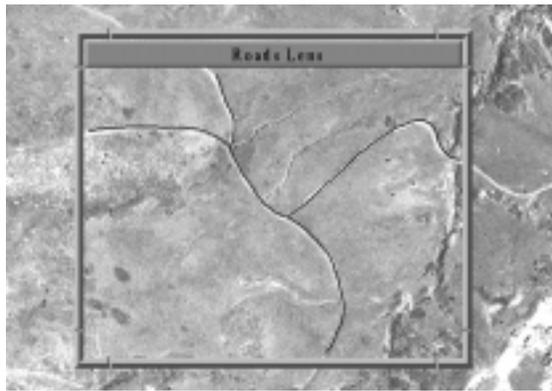
task, they can leave custom lens sets pre-positioned at each analysis area.

- Tools residing on the surface, with a geospatial position and size, are highly intuitive to use during analysis tasks that involve geo-referenced data in the spatial vicinity of the tools. Since such tools sit at the work point, rather than off to the edge of the screen, they minimize the shifting of spatial attention and context. With large displays (e.g., 3x4 foot), large motions become impractical.
- Lenses are a natural fit for spatial database queries and visualization, and the iterative construction of complex database queries.

The current concept of a lens was pioneered by Xerox PARC [2] and Pad/Pad++ [8] researchers. The Pad/Pad++ researchers developed a more generalized concept of a lens called a portal, which is a lens that shows a view at any zoom of any portion of the Pad++ surface. This group went on to focus on other zoomable user interface ideas, e.g. [1]. (Each lens in IMINEX is implemented using a Pad++ portal widget. This paper commonly uses the term lens, but when the portal widget's view is not looking at a position below the widget then we term the tool a surface portal.) The Xerox based group constructed a number of lenses, tested their use in drawing applications, database query/visualization, and some map viewing uses, and they also explored transparent palettes of lenses called toolglass sheets [4,9]. More familiar but less general versions of the lens and toolglass sheet concepts existed long ago in the form of pixel magnifier lenses, image view lenses, and transparent tear-off tool icon palettes. Other related concepts were fish-eye and non-linear magnification lens views.

4.1 Layer View Lenses

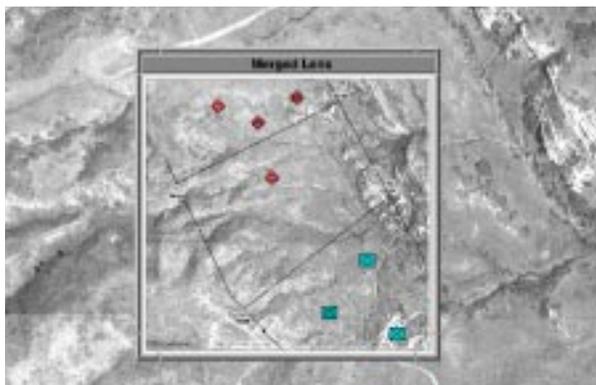
Many types of geospatial data are most naturally organized as multiple layers of different data types. For example, many kinds of map data consist of line graphics, and many military application domains involve graphical overlays describing such things as the deployment of friendly forces. This type of data can be quite dense visually, and applications typically provide a user interface to turn each specific layer on or off within the application's global display. Screen clutter is a fundamental problem.



(a)



(b)



(c)

FIGURE 4. Lenses can display a layer of data such as line graphics depicting (a) roads or (b) hydrology. Lens (c) is a “merged” combination of several lenses, including military symbology, with magnification.

Several lenses that display map data are depicted in Figure 4. Here the map data is organized into layers and each lens simply displays all the data within a specific layer. Clicking the mouse through a lens of this type toggles whether that layer of data is displayed globally on the main screen.

It is crucial for an image analyst to maintain a feeling for the spatial context in which he is working. Analysts often prefer to see a reference map or

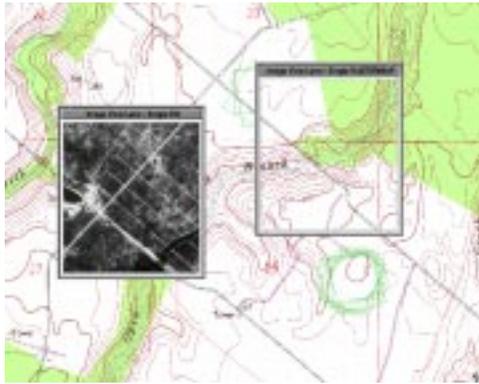
overview image in the background of the screen in order to help maintain an idea of where they are located in the world. Unfortunately there can only be one background at a time, and such a background can confuse the appearance of the primary data the analyst is studying. Two lenses, a rastermap viewing lens and an image viewing lens, provide simultaneous display of both types of data, as illustrated in Figure 5. An important aspect of this example is that the user can scan one of the lenses around on the screen, for example by following a road with the rastermap lens, and thus simultaneously combine a (registered) view of the overview image and the rastermap.

4.2 Simple Interaction Through Lenses

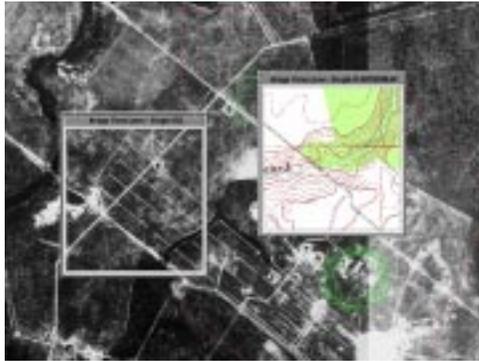
User interaction through lenses can take any number of forms. The common theme for all such interaction is direct manipulation, either of graphic objects that directly represent objects in the world, or of special graphical objects associated with a world object and representing subparts or properties of that real-world object. A simple example of the former is a lens that displays image footprints (filled polygons), which the user can interact with. An example of the latter is a lens that overlays a hypothesized geometric model over a vehicle, and provides graphical handles to change the aspect of the vehicle, the turret angle, the presence of external fuel tanks, etc.

Figure 6 shows INIMEX’s footprint lens. This lens displays a filled polygon -- the so-called footprint -- for each image. Footprints are displayed in a stacking order that reflects the stacking order of the corresponding images if they were all displayed at once. As the mouse moves over different footprints, each footprint is highlighted to provide visual feedback. The “zoom in” and “zoom out” mouse buttons send a footprint (and the associated image) to the front and back of the stacking order respectively. Clicking the left mouse button over a footprint toggles the visual representation (outline or pixel) of an image object. When an image object is in the outline mode (the default) only the polygonal outline of the image is visible, and when it is toggled to pixel mode the image pixels are visible.

Mensuration, the ability to measure real-world distances (meters typically), is a basic requirement for the image analyst. Types of distances can range



(a)



(b)



(c)

FIGURE 5. (a) A rastermap is displayed globally (by clicking through the rastermap viewing lens) while an image view lens lets the user see specific details in an EO image. (b) The EO image is displayed globally while a rastermap lens provides a selective view of the map. (c) The user utilizes both lenses, freeing up screen space for other more important data or tools.

from the dimensions of a vehicle, building, or facility to the movement of a force element on the battlefield. Since everything in INIMEX is georeferenced, distances on the screen are always proportional to distances in the real world.

INIMEX provides a simple mensuration lens, as shown in Figure 7. The user drags the endpoints of

a line segment -- through the lens -- and the lens always shows the length of the line segment. Other mensuration lenses could present shapes (e.g., polygon) and display relevant measurements.

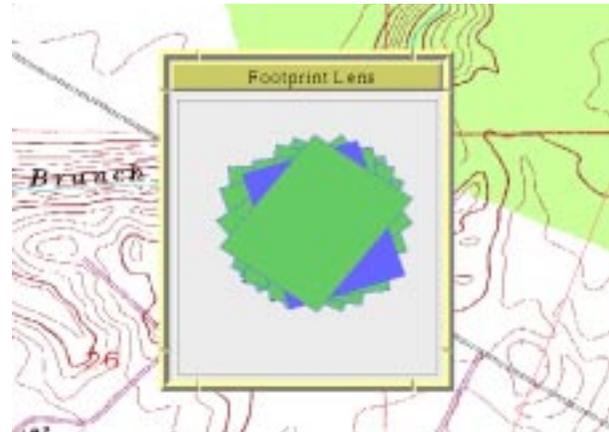


FIGURE 6. The footprint lens is used to re-arrange the stacking order of displayed images, and to turn the display of individual images on and off.

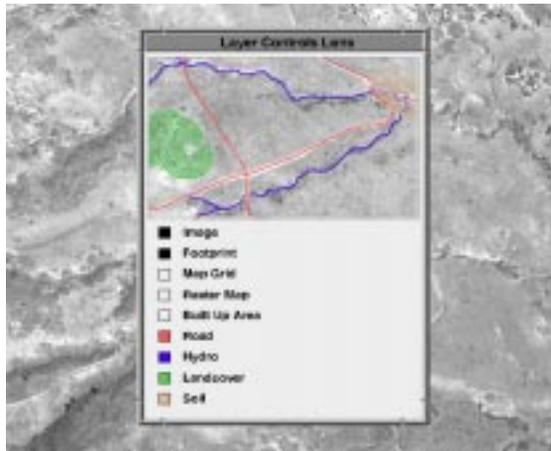


FIGURE 7. This mensuration lens shows the length of a line segment manipulated by dragging the line's endpoints through the lens.

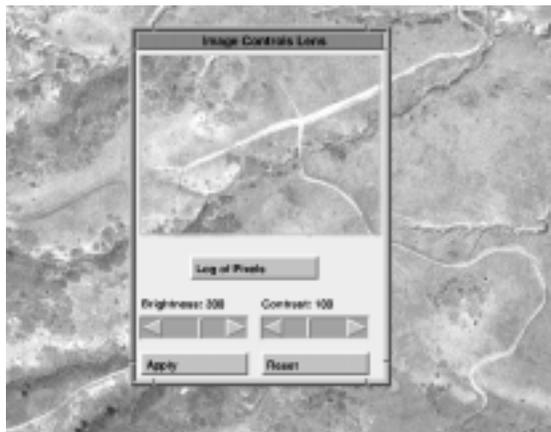
Lens interfaces can utilize conventional widgets placed inside the frame of the lens but distinct from the view (portal widget) inside that frame. Figure 8 shows two examples. In both cases the buttons, checkboxes and sliders affect only the view of the lens, and clicking through the lens applies those changes to the global view or entire image.

Studying changes over time is a fundamental element of image analysis, and the analyst needs a variety of tools to support this requirement. A rather basic need is an ability to peruse imagery for a specific area according to when the images were

taken. The interface provided for perusing imagery in most current image exploitation systems typically consists of a tabular list of image file names and their acquisition dates.



(a)



(b)

FIGURE 8. (a) The layer controls lens provides a classic interface, a set of checkboxes, for viewing multiple data layers. (b) The image controls lens provides sliders to modify factors such as brightness and contrast for one image.

Figure 9 shows one of INIMEX’s tools for perusing imagery, a lens that shows geo-registered footprints of the relevant subset of images and incorporates a separate zoomable timeline populated with checkboxes corresponding to the image objects. The top part of the lens acts just like the footprint lens described previously, with the addition that as the mouse passes over a footprint the corresponding checkbox on the timeline is also highlighted. The bottom part of the lens shows a zoomable view into a timeline space, which supports semantic zooming over different timeline scales (i.e., hour, day, week, month, year). As the mouse passes over a checkbox on the timeline the

footprint of the corresponding image object is highlighted, and clicking on the checkbox has the same effect as clicking on the footprint, meaning the image object is toggled from outline mode to pixel display mode.

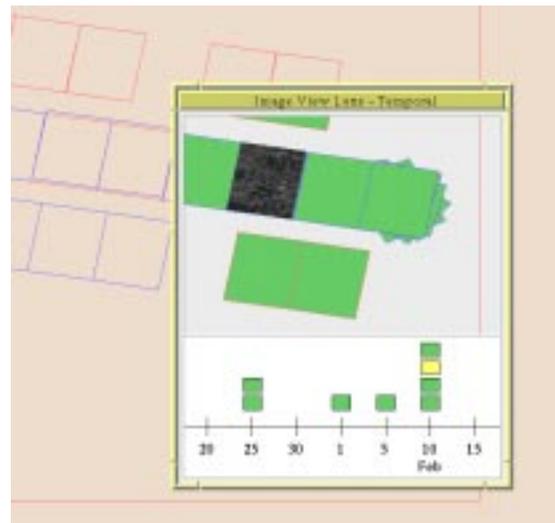


FIGURE 9. This lens allows the analyst to view and peruse imagery along the time dimension.

A family of lenses, similar to this timeline-based image selection and viewing lens, provides a somewhat complex interface to a needed capability within INIMEX, rapidly perusing imagery in a search area and perusing target identification results for individual vehicles in that imagery. Lenses in this family overlay a display of specific data layers associated with the selected image and provide different methods for interacting with the overlaid data.

The current collection of timeline-based lenses were built to satisfy an immediate need for perusing target identification results. We have been working to replace this monolithic-style lens with a combination of stacked primitive lenses. The basic idea, in this example, is to have a separate and generic timeline lens that is stacked underneath a target identification lens.

4.3 Lenses That Show Different Visual Representations of an Object

One of the more striking and useful aspects of a lens is that different lenses can show different graphical representations of an object.

The ability to add graphical annotations to an area in the world or to individual images is a basic capability for any image exploitation workstation. Annotations by an image analyst usually take the form of text and line graphics of the sort provided by a very simple drawing program. Figure 10 shows a simple example. Here a user has begun to create some freehand annotations, which is done by drawing through an ink lens. The lenses in the figure interpret freehand annotation as specific shapes. Clicking on an object through such a lens converts the primary representation of the graphical object to the respective shape.

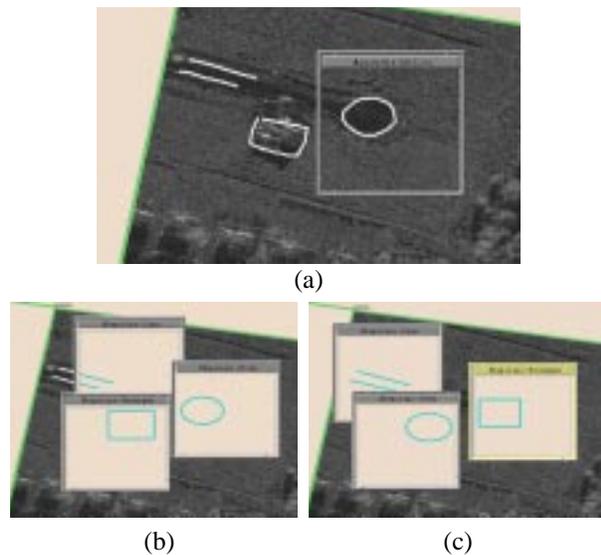


FIGURE 10. Illustration of how complex symbolic annotations can be created by interpreting hand drawn annotations through different lenses. Clicking on a shape through a lens resets the visual representation to that lens' shape.

The INIMEX application domain contains many types of objects for which the user may want to see different visual representations at different times. For example, a set of lenses can provide different visual representations of a vehicle identification hypothesis, as shown in Figure 11.

Zooming our perspective out from an individual vehicle to a collection of vehicles inside one image, the collection of vehicle hypotheses inside an image constitutes a different kind of hypothesis object. Here the image analyst would most naturally view the information in this object as a summary report for that one image. Figure 12 shows some lenses that generate different types of reports.

INIMEX is developing a research prototype of a future end-user image exploitation workstation that includes target cueing and identification technology. Part of the interface will be an interactive ATR capability that allows a user to view and manipulate 3D geometric models matched to a specific vehicle underneath a lens, and to view a SAR image predicted for that model. Figure 13 shows a 3D model rendering lens, built by integrating OpenGL with Pad++ and INIMEX.

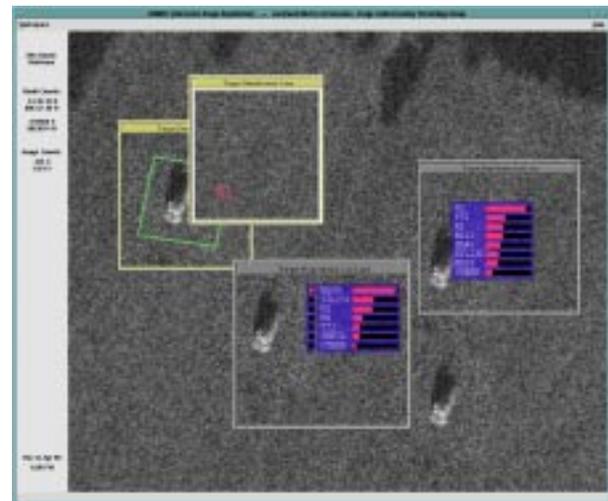


FIGURE 11. This set of lenses provides different visual representations of the vehicle identification hypothesis data for vehicles.

4.4 Lenses Whose Displays Involve Intense Computation

The user's conceptual model for a lens is that the lens provides a view of (or operates on) only the data directly underneath the lens. Most image processing operations require intense numerical computations that seldom can be accomplished in anything close to real time, especially when images have as many as ten thousand rows and columns.

Lenses enable an analyst to view image processing results in real time, such as an image enhancement. This is possible because the lens actually only performs the operation on the portion of the image visible underneath it. By scanning a small-sized lens over areas of interest the analyst can perform analysis operations over a large area in real time.

The target identification hypotheses generated by the MSTAR system are currently run off-line at the time that imagery is first ingested into the INIMEX system. Depending on the specific system configu-

ration and desired output, the MSTAR system can require up to many minutes or even hours to process a single SAR image. Some useful but more limited MSTAR functions run at speeds close to those needed for interactive uses and we are working to speed up some functions.

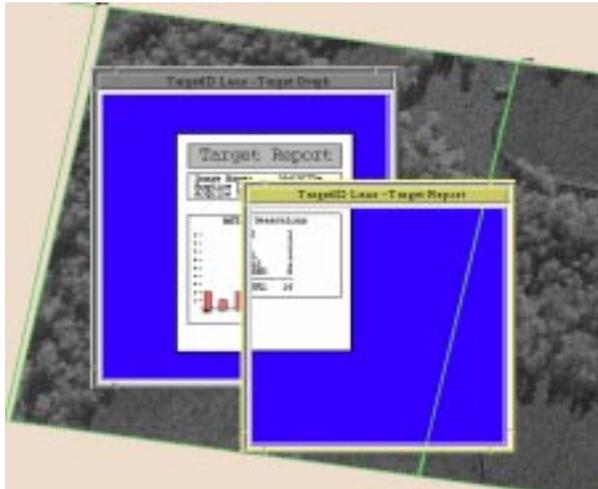


FIGURE 12. The vehicle identification results for one image can be summarized in different ways such as a tabular list, a bar graph, or a detailed report.

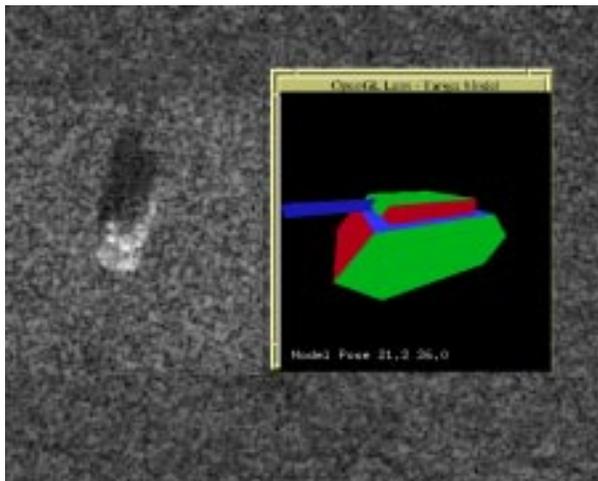


FIGURE 13. Initial version of lens that displays an OpenGL rendered view of a vehicle model associated with a specific target identification.

4.5 Lenses That Interact With Each Other

The set of lenses available to the user must provide a huge variety of functions. How to provide this variety of functionality in the user's mind and hands is a fundamental design question. At one end of the spectrum, the user could be provided with a set of complex, monolithic lenses. Each lens would

provide a number of related functions, selectable through various kinds of option selections built into the lens. At the other end of the spectrum, the user could be provided with a set of primitive lenses, several of which must be stacked to compose a complex function. The monolithic approach to lenses is straightforward, but the stacked primitive approach to lenses is less familiar and raises a number of implementation, interface design and usability questions.

Conceptually, a stack of lenses implements a sequence of operations on a set of input objects: The bottom-most lens selects a set of input objects (overlapping the lens' view), applies an operation to the set, and outputs a set of objects. The objects output from a lens are rendered and may be passed as input to a lens stacked above.

Certain types of lenses perform operations that are trivially described, and thus several lens instances of that type can be merged to create a new lens that performs a combination of the original operations. One example is a stack of image processing lenses, where each lens operation is essentially specified by a parameterized operators. A stack of these lenses can be merged into a single new lens with an operation that is simply the concatenation of the original operators. Another example is the merging of layer view types of lenses, which display everything in a particular set of data layers. In this case the merged lens concatenates the layer sets from the individual lenses (retaining the proper display order of the layers). Figure 14 shows two examples of merged layer view lenses.

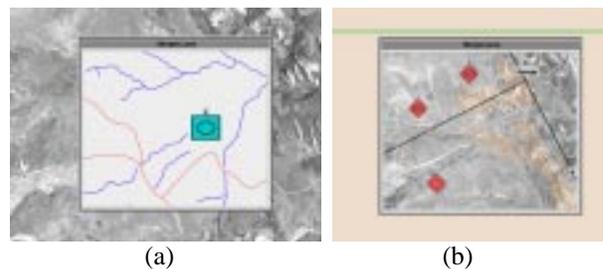


FIGURE 14. Each of these lenses is a merger of several layer view lenses: (a) friendly units, roads and hydrology (all on a blank background), and (b) enemy units, control measures, soils and imagery (providing the background).

Examples of more sophisticated stacked lenses include: (1) a temporal range selection lens under an image exploitation lens under a plotting lens, (2)

a site model lens stacked over two adjacent image selector lenses resulting in a 3D site model being properly projected onto the selected imagery and the nominal map background, and (3) multiple database query/visualization lenses stacked to compose a complex database query.

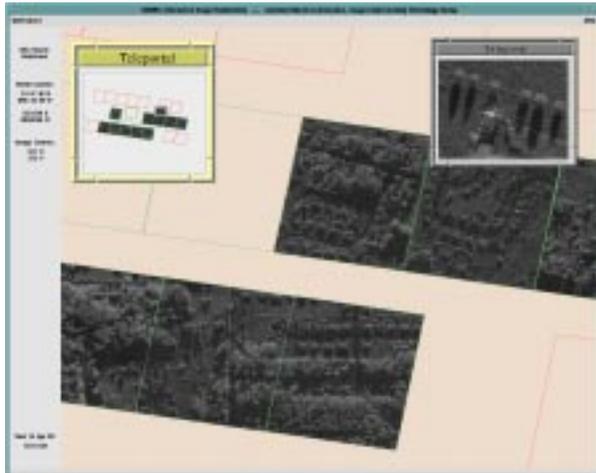


FIGURE 15. The teleportal tool is essentially a Pad++ portal widget with window dressing and one mouse binding. The portal displays a view of a bookmarked area, and clicking the mouse through the portal causes the main display to be transformed to view the bookmarked area.

4.6 Portals

Recall that a *surface portal* is a generalization of a lens that provides an independent view of any area of the Pad++ surface at any zoom. The view may show any set of the display layers on which all Pad++ objects reside. Figure 15 shows an example of a surface portal used within INIMEX, a kind of spatial bookmark called a teleportal. The teleportal shows a view of the bookmarked area, and clicking through the teleportal causes the main display (itself a view) to be transformed via an animated pan/zoom to the bookmarked area. Once there, the teleportal view is reset to the position and zoom before the teleport motion was initiated. Normally, all lens and surface portal tools in INIMEX reside on the geo-referenced surface. The teleportal has a different property enabled, called stickiness [1], that effectively causes the tool to be glued to the glass of the display screen as the main display's view changes. If the teleportal was not sticky it would remain at the previous geo-location, an option available to the user through the popup property menu shared by all lenses.

Other than teleporting to a bookmarked location, a surface portal can simply serve as a convenient visual reference. For example, one can provide a thumbnail overview of the current site under analysis, or one can provide a close-up of specific vehicles located in some imagery from a week ago and suspected of being the same vehicles found today at a different location on the battlefield.

5.0 Zoomable Tool Palettes

The functionality provided by conventional user interface elements (such as a menubar, menus, and tool palettes) can be provided in other ways within a zoomable user interface. For example, INIMEX utilizes a new zoomable space of tool palettes, called a toolspace, instead of a menubar. The toolspace contains multiple redundant tool palettes. One is a collection of tool palettes structured so as to sequence the user (especially newer users) through common analysis tasks. As shown in Figure 16, a graph-like display shows each step as a node in a graph -- called the task graph. Each node in the task graph contains a specialized tool palette, access to a description of what to do and how to do it for that step, and a "next" button that performs an animated pan/zoom motion to the next node/step in the task graph.

6.0 Summary

Our goal is to develop innovative, simple and intuitive user interface design elements for future image exploitation systems and other intelligence analysis systems. This paper has summarized the results of our first round of exploration and prototyping, which focussed on basic issues of zoomable user interface (ZUI) design and lens design. The prototype system has been very well received by a number of groups who have seen it. So far, INIMEX has demonstrated the following novel claims:

- First image exploitation application using a zoomable user interface.
- First illustration of sophisticated lenses for image exploitation.
- First coherent set of lens tools for the image analyst.

- [5] J.C. Mossing, T.D. Ross, "An Evaluation of SAR ATR Algorithm Performance Sensitivity to MSTAR Extended Operating Conditions," Proceedings of SPIE, Vol. 3370, SPIE'98 Algorithms for SAR Imagery V, April, 1998.
- [6] J.L. Mundy, "Image Understanding Research at GE," Proceedings of the Image Understanding Workshop, 1997, Morgan Kaufman, pages 425-430.
- [7] J. Nielsen, "Non-Command User Interfaces," CACM, 36(4):83-99, April 1993.
- [8] K. Perlin, D. Fox, "Pad: An Alternative Approach to the Computer Interface," Proceedings of SIGGRAPH'93, pages 57-64.
- [9] M.C. Stone, K. Fishkin, E.A. Bier, "The Movable Filter as a Interface Tool," Proceedings of CHI'94, pages 306-312.