


Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments

By Nicholas F. Polys*, Seonho Kim and Doug A. Bowman



This paper describes our recent experimental evaluation of Information-Rich Virtual Environment (IRVE) interfaces. To explore the depth cue/visibility tradeoff between annotation schemes, we design and evaluate two information layout techniques to support search and comparison tasks. The techniques provide different depth and association cues between objects and their labels: labels were displayed either in the virtual world relative to their referent (Object Space) or on an image plane workspace (Viewport Space). The Software Field of View (SFOV) was controlled to 60° or 100° of vertical angle and two groups were tested: those running on a single monitor and those on a tiled nine-panel display. Users were timed, tracked for correctness, and gave ratings for both difficulty and satisfaction on each task. Significant advantages were found for the Viewport interface, and for high SFOV. The interactions between these variables suggest special design considerations to effectively support search and comparison performance across monitor configurations and projection distortions. Copyright © 2006 John Wiley & Sons, Ltd.

Received: 6 October 2006; Accepted: 6 October 2006

KEY WORDS: 3D interaction; visual design; usability testing and evaluation; Information-Rich Virtual Environments (IRVE); multimedia information systems; human factors

Introduction

Increasingly designers, engineers, scientists, and students require 'Integrated Information Spaces' where spatial, abstract, and temporal data are simultaneously available and linked. To address this problem in our work we are developing Information-Rich Virtual Environments (IRVEs). An IRVE combines the capabilities of *virtual environments* and *information visualization* to support the integrated exploration of spatial, abstract, and temporal data. IRVEs are therefore concerned with information design and interaction techniques that enable both the independent and combined navigation and comprehension of these different data types.^{1,2} In other words, they require optimized visualization configurations to achieve the most accurate and complete mental model of the data.

Virtual environments (VEs) can provide users a greater comprehension of spatial objects—their perceptual properties and their spatial relations. Perceptual information includes 3D spaces that represent physical or virtual objects and phenomena including geometry, lighting, colors, and textures. As users navigate within a rich virtual environment, they may need access to the information related to the world and objects in the space (such as name, function, attributes, etc.). How to effectively present this related information is the domain of Information Visualization, which is concerned with improving how users perceive, understand, and interact with visual representations of abstract information.³

This enhancing abstract (or symbolic) information could include text, links, numbers, graphical plots, and audio/video annotations. Both perceptual and abstract information may change over time reflecting their temporal aspects. This work examines what integrated visualization capabilities are necessary for users to gain a full understanding of complex relationships in their heterogeneous data, and to create advantageous research,

*Correspondence to: N. F. Polys, 660 McBryde Hall, Virginia Tech, Blacksburg, VA 24061-0106, USA. E-mail: npolys@vt.edu

design, and decision-support applications. By taking account of how humans build their cognitive models and what perceptual predispositions and biases are in play, designers can take steps to minimize or leverage their effect. This line of inquiry has been termed the 'inverse problem of design' by Joseph Goguen⁴ and 'Information Psychophysics' by Colin Ware.⁵ More recently Thomas and Cook,⁶ recast this problem as 'Visual Analytics,' referring to this science of facilitating analytical reasoning through interactive visual interfaces. Regardless of the formulation, it is clearly an open challenge to understand what constitutes a 'good' or effective design in interactive, information-rich applications.

The challenge is to provide a set of organized, multi-dimensional representations that aid users to quickly form accurate concepts about and mental models of the system they are studying. In IRVEs, the information design problem can be summarized as:

Where and how should enhancing abstract information be displayed relative to its spatial referent so that the respective information can be understood together and separately?

Support for search and comparison tasks is fundamental to the success of any IRVE interface and there has been little research into how different display techniques impact user performance, satisfaction, and perceived difficulty. In addition, there is little understanding of these techniques' properties on different screen sizes and spatial resolutions. Information design for IRVEs presents many rendering and layout challenges including managing the display space for visibility, legibility, association, and occlusion of the various data.⁷

This paper provides a thorough treatment of our recent experimental evaluation of IRVE layout spaces across single and tiled nine-screen displays for search and comparison tasks. The goal of this evaluation was to understand the effectiveness of an annotation's Layout Space across different display sizes and different Software Fields of View (SFOVs). In this work, we focus on *depth* and *association* cues between labels and their referents and devise a controlled experiment aimed at understanding the effective parameters of IRVE information design for two viewing platforms and two viewing projections.

Specifically, we are interested in the perceptual cues provided by two different layout spaces and their tradeoffs for performing fundamental types of tasks across different monitor configurations (1 and 9) and different projection distortions (60° or 100° of vertical angle). The results of this empirical evaluation provide insight into how IRVE information design tradeoffs

impact task performance and satisfaction and what choices are advantageous under various rendering distortions. In addition, this evaluation addresses the problem of how designers should consider the transfer of IRVE interfaces between screen-sizes (e.g., a single-monitor and multiple LCD monitor display; Figure 1).

Background

Information Design

The integration of text and image stimuli has been studied as it relates to the comprehension of multimedia presentations. Faraday and Sutcliffe's work^{8,9} supports the claims of Chandler and Sweller¹⁰ that co-references between text and images can improve the comprehension of complex instructional material. For example, the integration of text and image information in multimedia presentations has resulted in better user recall than images only. While proximity between object and label may be a useful heuristic for static or 2D imagery, it is not clear how it should be applied to dynamic renderings of VEs.

Feiner *et al.*¹¹ enumerated locations for the display of information labels in Augmented Reality (AR). The display locations described by Feiner *et al.* were organized for AR design paradigms. For IRVEs, we must adapt the terminology to incorporate synthetic environments on desktop and immersive devices. We characterize display locations according to a user's perspective and what coordinate space the information resides in: abstract information may be located in *object space*, *world space*, *user space*, *viewport space*, or *display space*. Conceptually, these layout spaces align with those described in Barrilleaux¹² for Java3D.

Bell *et al.*¹³ developed an impressive strategy for dynamically labeling environment features on the image plane for mobile AR. They used a Binary Space Partition tree (BSP) to determine visibility order of arbitrary projections of the scene. From visible surface data for each object the frame, a view-plane representation is then used to identify each visible object's rectangular extent. The algorithm identifies 'empty' spaces in the view-plane and draws the annotation (such as a label or image) in the nearest empty space by depth order priority.

Display Parameters

Both resolution and physical size of a display play an important role in determining how much information




Figure 1. Example display configurations used in this experiment. Both images show the same world with an Object Space layout: Single screen with 100 degree SFOV (top) and nine-screen with 60 degree SFOV (bottom).

can or should be displayed on a screen.¹⁴ Tan *et al.*¹⁵ found evidence that physically large displays aid user's performance due to increased visual immersion. Swaminathan and Sato¹⁶ examined the advantages and disadvantages of large displays with various interface settings and found that for applications where information needs to be carefully studied or modified, 'desktop' settings are useful, but for collaborative, shared view and non-sustained and non-detailed work, a 'distance' setting is more useful. This work orients our design claims and evaluation to the paradigm of the single-user desktop workstation. Mackinlay and Heer¹⁷ proposed seam-aware techniques to perceptually compensate for the bezels between tiled monitors. In contrast, our system rendered views naively, splitting images across monitors as though there were no bezels.

While a number of studies have examined the hardware and the display's (physical) Field of View (e.g., Dsharp display¹⁸), less is known about the performance benefits related with the Software Field of View (SFOV) and VEs. However, Draper *et al.*¹⁹ studied the effects of the horizontal field of view ratios and simulator sickness in head-coupled VEs and found that 1:1 ratios were less disruptive than those that were far off. There is also a good body of work on SFOV in the information visualization literature, typically with the goal of overcoming the limitations of small 2D display spaces. Furnas, for example, introduced generalized Fish-Eye views²⁰ as technique that allows users to navigate data sets with 'Focus-plus-Context.' Gutwin's recent study²¹ showed that fisheye views are better for large steering tasks even though they provide distortion at the periphery.

Information-Rich Virtual Environments (IRVEs)

Across a wide variety of domains, analysts and designers are faced with complex systems that include spatial objects, their attributes and their dynamic behaviors. In order to study and understand these systems, these users require a unified environment to explore the complex relationships between their heterogeneous data types—*integrated information spaces*. The problem of integrated information spaces arises from the fact that users must comprehend the nature and relationships within each data type individually as well as the relationships between the data types.

Information-Rich Virtual Environments (IRVEs) have been proposed as a solution to the problem of integrated information spaces.² IRVEs start with realistic perceptual and spatial information and enhance it with abstract and temporal information. In this way, the virtual space serves as a context for the methods of VEs and InfoVis to be combined and enable a unified interface for exploring space and information.

IRVE Task Space

For IRVEs, we can readily identify four categories of tasks: exploration, search, comparison, and pattern finding. Exploration and search requires navigating the world space and the information space and recognizing that a perceptual object has abstract information associated with it. Comparing requires examining abstract and perceptual data that are interrelated but may have unique structures. Finding patterns and trends refers to the recognition and extraction of higher-level forms in perceptual and abstract information.

In IRVEs, these activities will be performed both *within and between* the spatial, abstract, temporal information. As we have noted, there are fields whose particular interest is examining performance within these respective types (i.e. VEs and Information Visualization). For example, Schniederman²² enumerated a set of generic information visualization tasks that provide a structure to design interface features. For VEs, Bowman *et al.*²³ propose a set of generic tasks whose requirements we design for. For the problem of integrated information spaces, it is crucial that we examine how these tasks can be supported *between* information types. Therefore, we introduce the notion of a *task-information mapping*:

- Spatial to Abstract Information (S → A): In this case, the user wishes to use the VE as an index into the abstract information. An example task is details-on-demand, in which the user desires to retrieve abstract information related to a given element in the VE space.
- Abstract Information to Spatial (A → S): In this case, the user wishes to proceed from abstract to VE. This can enable the user to control the VE through abstract information. For example, by selecting data in a separate abstract display, users could highlight desired objects in the VE, filter uninteresting objects, or automatically travel to a related location.

To begin this project, we critically examined existing approaches to information layout in IRVEs for Details-on-Demand interaction. On most desktop VEs, selection of virtual objects is accomplished via raycasting from the mouse pointer into the scene. In response to selection interaction, the system response is to display or highlight the abstract information related to that object. For example, a user may select a virtual object and toggle its information as ‘on’ or ‘off.’ These annotations may contain any type of information from text to images to interactive graphs and windows, but *where* that information is displayed and *how* it is visually associated to its referent object is the subject of our investigation.

IRVE Information Design Space

In an IRVE, there may be a wealth of data and media types embedded-in or linked-to the virtual space and objects. Users require interfaces that enable navigation through these various types and while a number of applications have designed interfaces for this purpose, they are typically ad hoc and specific to the application. Clearly, a more rational design approach is required.

We have presented an information design-space for IRVEs that provides a means to separate information design concerns into three principal dimensions: Visual Attributes, Layout Attributes, and Aggregation Attributes.⁷ Visual Attributes refer to the visual properties of a label such as typography, color, and transparency. Layout Attributes refer to how the label is drawn in relation to its referent. The Aggregation dimension defines the level of abstraction that is represented in the annotation (label). The focus of this evaluation however is to understand the role of Layout Attributes on user performance across screen-sizes and projection distortion.

In IRVEs, annotations may reside in a number of coordinate systems—what we term the ‘Layout Space.’ As mentioned above, these are: World, Object, User,

Viewport, and Display spaces. Depending on the IRVE display technique used, each of these spaces may present a range of depth cues that are consistent between the annotation and referent. In addition to the various depth cues that may be provided by a Layout Space, we can also consider the relation of object and annotation on the image plane in terms of Gestalt Association cues. Thus we have a two-dimensional design space for Layout Attributes, shown in Table 1. IRVE Display techniques of High Association and High Occlusion reside in the top left corner. IRVE Display techniques of Low Association and Low Occlusion reside in the lower right corner.

The evaluation we describe in this paper seeks to understand the visual design tradeoffs that exist concerning the layout of abstract information in relation to its referent object in the virtual environment. Specifically, we are interested in the occlusion-association tradeoff. This tradeoff occurs because the stronger the visual cues of association between object and label, the more of the spatial environment is occluded; the less visual association, the less occlusion. This tradeoff can be summarized by the following design claim:

More consistent depth cues and Gestalt cues between annotation and referent

- (+) May convey more information about the relation between annotation and referent (i.e. less ambiguity)
- (–) May cause more occlusion among scene objects and annotations and therefore less visibility of each information type

The two layout spaces we examine in this research are termed: ‘*Object Space*,’ in which annotations are displayed in the virtual world relative to their referent object, and ‘*Viewport Space*,’ in which annotations are displayed in a planar workspace at or just beyond the image plane. In Object Space, annotations representing abstract information are spatially situated in the scene, which provides depth cues such as occlusion, motion

	Common region	Proximity	Connectedness	Common fate	Similarity
Object	x	x	x	x	x
World	x	x	x	x	x
User	x	x	x	x	x
Viewport	x	x	x	x	x
Display	x	x	x	x	x

Table 1. IRVE layout attributes: Layout space × association cues

	Proximity	Connectedness	Common fate	None
Occlusion	○	○	○	
Motion parallax	○	○	○	
Relative/size/perspective				
None		V	V	

Table 2. Depth and gestalt cues presented by Object (O) and Viewport (V) space layouts used in this experiment

parallax, and linear perspective consistent with the referent object; in addition, the annotation and referent are visible in the same region of the screen (Gestalt proximity).

The Viewport space, in contrast, is a 2D-layout space at or just beyond the near-clipping plane. As such, annotations and geometry in the Viewport space is rendered last and appears as over-layed on top of the virtual world’s projection. Annotations in Viewport Space typically do not provide depth cues consistent with their referents, but do provide guaranteed visibility and legibility of the annotation. The relationship of perceptual cues in the conditions tested is shown in Table 2. IRVE Display techniques of High Association and High Occlusion reside in the top left corner. IRVE Display techniques of Low Association and Low Occlusion reside in the lower right corner.

Object Space

One existing layout technique, termed ‘Object Space,’ is to locate the abstract information in the virtual world and in the same coordinate system as its referent object. By co-locating the enhancing information with its referent object in the virtual space, this technique provides depth cues that are consistent with the referent object such as occlusion and motion parallax. If the object is moved or animated, the label maintains its relation to the object, giving a tight visual coupling between annotation and referent. Labels may be located in the object’s coordinate system through a number of layout algorithms including: relative position, relative rotation, bounding box, screen-bounds, and force-directed methods.

We have previously described software objects that encapsulate a number of IRVE layout behaviors.^{7,24,25} These ‘Semantic Objects’ allow the specification of multiple levels of detail for both objects and their labels,

which enables proximity-based filtering on either type. In addition, Semantic Object labels can be: billboarded to always face the user and maintain upright orientation, connected to the object with a line, and scaled by user distance through a number of schemes (none, periodic, and continuous).

In Gestalt terms, Object Space can provide strong association cues including connectedness, proximity, and common fate.²⁶ However, there are some limitations to Object Space, especially for search and comparison tasks. For example when using Object Space layouts, not all labels may be visible at once and spatial maneuvering may be required to make them visible and legible. In addition, when comparing abstract information that is rendered as a graph for example, the effects of the perspective depth cue can make comparison difficult. Figure 2 shows an example of the Object Space layout technique used in a 3D-cell model.

Viewport Space

To address the limitations of Object Space layouts, we designed and implemented a new IRVE interface we call the ‘Viewport Workspace’ where a selected object’s label is toggled into a Heads-Up-Display at the image plane where is it always visible regardless of the user’s position and viewing orientation. In the software definition of our interface, we maintain a pixel-agnostic stance and scale and locate labels according to parameters of the environment’s projection (rendering). Labels are sized and located in world units relative to the specified Software Field of View (SFOV) and the distance to the near-clipping plane.

In the Viewport Workspace, labels can also be connected to their referent objects with lines extending into the scene. The layout of labels in the 2D-Viewport space is managed by a parameterized BorderLayout-Manager script. Like its Java Swing inspiration, the

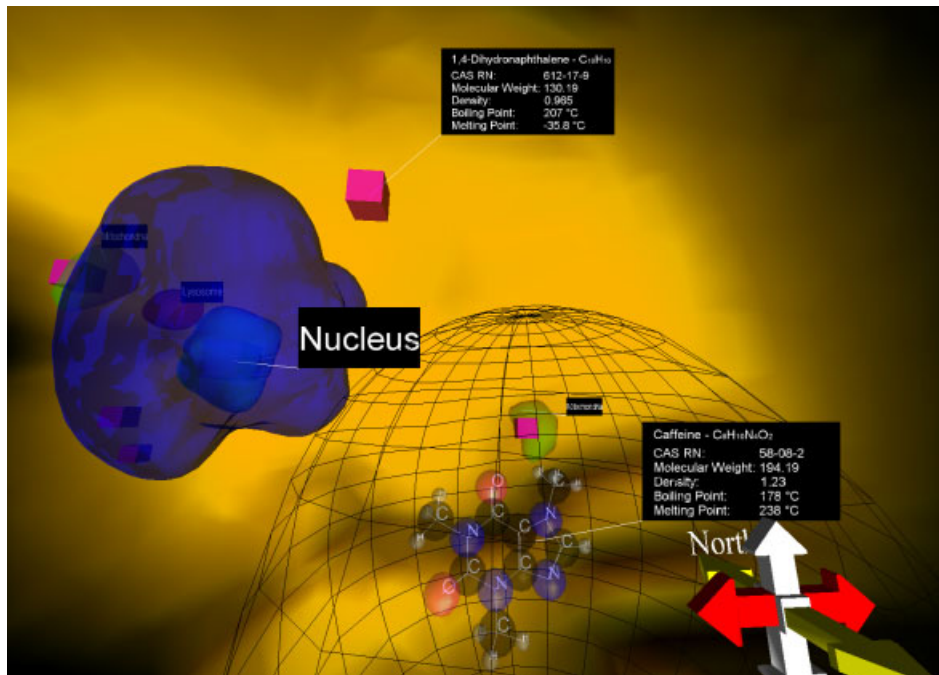


Figure 2. The object Space IRVE layout technique.

BorderLayoutManager divides the layout space into four regions or containers: North and South, which tile horizontally across the top and bottom, and East and West, which tile vertically on the right and left sides.

The Viewport Space BorderLayout we defined can be specified with container capacity and the fill order for the four directions using the BorderLayoutManager. The location of any given label is determined by the order in which it was selected. If the user does not like a label's location, they can click and drag it to a new location in the Viewport Workspace. Figure 3 shows an example of the Viewport Space layout technique used in a 3D-cell environment.

By providing a pixel-agnostic layout space and manager at the image plane (layout positions are not described in pixels), we can easily scale labels and containers to the size of the display and projection. For example, we may only be able fit a half dozen labels legibly in one container on a single-screen display. However when we render that same interface on a nine-screen display, the labels scale proportionately and also become larger. Using our Viewport Space approach, we can adapt the label scale and container capacity to display the labels at an equivalent pixel size as on the single-screen. On a nine-screen display and holding pixel size constant to the value on a single-screen, we can

get approximately three times as many labels in one container.

System Design

A recent study examined exploration and search tasks in immersive IRVEs using Head-Mounted Displays (HMDs).³¹ The study compared combinations of two navigation techniques and two layout techniques for textual information. The two techniques for annotation labels were: 'in the World' (Object Space) or in a 'Heads-Up-Display (HUD)' (Viewport space). The two navigation techniques were HOMER³² and Go-Go.³³ For naïve search, the HUD technique was significantly better for both navigation types and a significant advantage to a combination of HUD and Go-Go navigation was demonstrated. However, it is not clear how such results will transfer to desktop or large-screen VE's where system input and output channels are so different.

The questions we set out to answer with this experiment include:

- Is a layout space with guaranteed visibility better than one with tight spatial coupling for certain tasks?

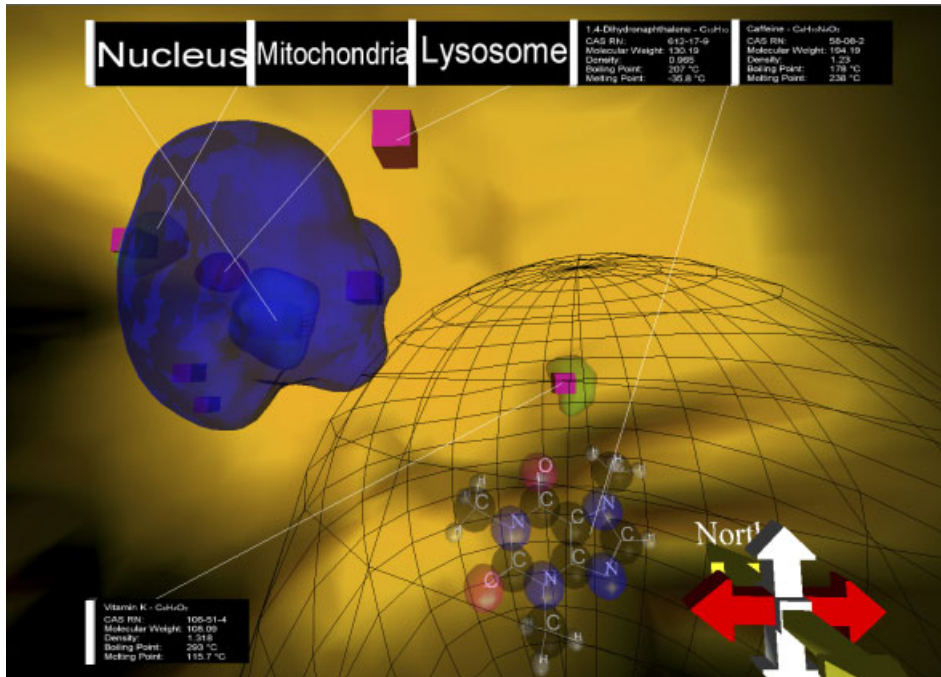


Figure 3. The viewport space layout technique.

- Do the advantages of one layout space hold if the screen size is increased?
- Do the advantages of one layout space hold if the SFOV is increased?

Field Of View

In understanding how humans perceive a virtual environment on a particular display, the concept of Field of View (FOV) is essential. For desktop displays we can describe at least two important kinds of FOVs: the Display Field of View (DFOV) and the SFOV. DFOV refers to the amount of visual angle that the physical display surface occupies in the user's visual field—a nine-screen display offers approximately three times more DFOV angle than a single-screen when viewed from the same distance. For example, a 17-inch monitor viewed from 65 cm provides a 22.5° vertical DFOV; three stacked 17-inch monitors viewed from the same distance provides a 61.7° vertical DFOV. It follows that a larger DFOV will require larger saccades and head movements for users to traverse it visually.

The SFOV on the other hand, refers to the viewing angle of the camera on the virtual scene, which is rendered (projected) onto the display surface. Larger

SFOV values create fish-eye effect while smaller values create tunneled, telescoping effects. We decided to fix SFOV to two levels for our experiment: 60° vertical SFOV (which approximately matched the nine-screen DFOV) and 100° vertical SFOV to assess any impact on the performance of search and comparison tasks.

Formative Evaluation

An informal pilot study was performed to understand how users perceive and interact with our IRVE interfaces in different SFOVs across the different monitor conditions. The goal of the formative evaluation was to find initial values of SFOV and drag mappings for the full study. Users were given a large-scale virtual model of the Giza plateau and given 7–10 minutes to learn the navigation controls of the VRML browser. On standards-compliant engines for VRML/X3D the SFOV is default of 45° (0.785 radians) measured along the shortest screen projection (typically the *vertical*).

When they were comfortable with the navigation interface, the initial designs of Object Space and Viewport Space annotation layouts were presented to two users from the study pool, each on both screen configurations. The layout techniques were presented in

a cell environment like those used in the full study. Users used the up and down arrow keys to dynamically increase or decrease the SFOV as desired.

Pilot Results

In the cell environment, novice users were able to tolerate much higher SFOVs than we had anticipated. The average across all interface layouts and display sizes was 90.5° (1.58 radians) vertical. On the single-screen the average SFOV was 4.1 times the DFOV, while on the nine-screen, it the average was 1.4 times the DFOV. While this is an interesting interaction, there is not enough statistical power to draw any real conclusions here. The user tendency to high SFOVs is also interesting because in a cell environment there are few, if any, sharp edges or 90° angles.

More interesting perhaps were user strategies with a dynamic SFOV control. In the Object Space layout, Users increased the SFOV to gain overview spatial information and also increased the SFOV to recover detail abstract information (when it was just out of view, for example). In addition, Users decreased the FOV to focus in or telescope to targets on the screen projection; however users sometimes confused reducing the SFOV to actually navigating to the target.

In the Viewport Space layout, users increased the SFOV control to gain overview spatial information and then had to decrease it to make detail abstract information legible. User's association of annotation to its referent appeared to have a strong temporal component. For example, when looking up information, users commonly oriented to labels' appearance or disappearance on screen as result of selection/de-selection rather than tracing the connection lines between objects and labels. This suggests that the interactive Gestalt cue of 'Common Fate' is a strong association cue in Viewport Space. Also, users did not identify the dragging affordances of the annotations on the Viewport workspace (even though the cursor changed) without instruction.

Design Impact

The initial data and observations were used to improve the IRVE layout prototypes for the final study. This included choosing two levels of SFOV condition that were higher than the VRML default, but within the preferred range demonstrated in the pilot. Once the target SFOVs values were chosen, all mouse drag

mappings were calibrated between the interface layouts on all display sizes and SFOVs. In addition, we added a white 'handle' bar to the Viewport Space labels to emphasize their druggability.

User Study

To test the relative effectiveness of our IRVE layout spaces across displays and task types, we designed an experiment to test the following hypotheses:

- *H1*: With its guarantee of visibility and legibility, the Viewport workspace should provide an advantage for search tasks as well as tasks involving the comparison of abstract information. The Viewport workspace does not provide the consistent depth cues of occlusion and motion parallax; thus tasks involving spatial comparisons may be difficult.
- *H2*: We hypothesized that the increased display size and corresponding spatial resolution of the nine-panel display will be advantageous for tasks where exhaustive search and comparison is required because more information panels (annotations) can be displayed at once.
- *H3*: Higher software FOV will aid search tasks by including more of the scene in the rendering. However, higher software FOV will hinder spatial comparison tasks due to fish-eye distortion.

Participants

Participants were drawn from the international graduate population of the College of Engineering. There were 11 males and 5 females. Ten of the 16 subjects wore glasses or contacts and all of the subjects used computers daily for work. 81.25% of the subjects also used computers daily for fun and the remainder used them several times a week for this purpose. All subjects had at least a high-school level familiarity with cell biology. Two subjects were known to be actively working as Research Assistants on bioinformatics projects and they were assigned to different display groups.

Subjects self-reported their familiarity with computers: 87.5% reported 'very familiar' with the remainder reporting 'fairly familiar.' 31.25% of the subjects reported not having used a 3D VE system before. Of those that had, 63.6% had used immersive systems such as HMDs or a CAVE; the remainder had used desktop platforms only, typically for 3D games.

Equipment

We used a large display system consisting of nine tiled LCD monitors supported by five dual-head peripheral component interconnect (PCI) high-end graphics cards on a 2.5GHz Pentium 4 PC. With the support of Microsoft Windows XP operating system's advanced display feature, we could easily create an illusion of a single large screen without using any special software and hardware. The dimension of the nine-screen display is 103.6 cm × 77.7 cm in physical size with 3840 × 3072 = 11 796 480 pixels. The dimension of the small normal display is 34.5 cm × 25.9 cm in physical size with 1280 × 1024 = 1 310 720 pixels. Subjects were seated at a distance of 60–70 cm from the screen with their heads lined up to the center monitor.

Content and Domain

Environments built for the study were based on a 3D model of a cell and its constituent structures (e.g., nucleus, mitochondria, lysosomes). These objects provided landmarks within the cell and a basis for showing spatial relationships such as 'next-to,' 'inside-of,' etc. All cellular structures were defined with different levels of detail so that from far away they appeared semi-transparent, but within a certain distance were drawn as wireframes. In this way, when a user got close enough to structure, they could pick (select) objects inside of it.

For each trial, a set of 3D molecules were shuffled and arbitrarily located in the various structures of each cell including the cytosol; these were the targets for the search and comparison tasks. In each cell environment there was a nucleus, a nucleolus, three mitochondria, two lysosomes, and 13 small molecules (all organic compounds with a molecular weight of less than 195). Since molecular scales are a few orders of magnitude smaller than cellular scales, molecules were represented by pink cubes when the user was far away; the molecular structure was drawn when the user got within a certain distance. Each cell structure was labeled with its name and each molecule's label included its name, formula, molecular weight, and melting and boiling point.

The choice of a cell model as the content setting was made for a number of reasons. First, there is a wealth of organic chemistry data suitable for IRVE visualization^{27,28} and its natural representation is in a biological context. Second, in these contexts there is no horizon and requirement for physical constraints such as gravity,

landmarks and targets are distributed in all three dimensions making effective annotation layout and navigation a challenge. Third, education researchers have shown improved student performance by augmenting science lectures with desktop VEs including the 'Virtual Cell' environment for biology and the processes of cellular respiration.^{29,30} It is our hope that our interface design lessons may be directly applied to biomedical research and education software.

For each task, landmarks and targets in the cell model were shuffled to insure naïve search for every trial. Regardless of independent variable conditions, each environment had identical mappings of mouse movement to cursor movement and picking correspondence was maintained for navigation, selection, and manipulation interactions in the virtual environment. In addition, all environments included an identical HUD compass or 'gyroscope widget,' which helped the user maintain their directional orientation within the cell. All interface components are realized entirely in VRML.

Information Design Conditions

In both the Object Space and Viewport Space layouts, labels were always drawn 'up' regardless of the user's orientation. All labels were connected to their referent objects with a drawn white line (Gestalt connectedness). In both the Object Space and Viewport Space layouts, label size was determined by the minimum label size for text legibility on a 1280 × 1024 display, in this case 206 × 86 pixels for six lines of the Arial San Serif font.

In the Object Space conditions, labels were located relative to their referent object and offset orthogonally by a set distance. When toggled on, Object Space labels were periodically scaled according to user distance. This scaling was established to guarantee a minimum size for legibility from any distance. The actually rendered label size (when viewed head-on) could vary between 1 and 1.2 times the pixel area of a label in the Viewport condition depending on the distance to the object. In making this choice, we remove the depth cue of relative size in favor of access to the abstract information contained in the label. The depth cues of occlusion and motion parallax remain.

Because Object Space labels are co-located with objects in the virtual world, they are subject to the same magnification and distortion problems as other objects in the periphery of the projection. As a result, a label may appear to stretch (keystone) and scale as it moves away from the line of sight.

In the Viewport Space condition, we used the BorderLayoutManager described above; fill order was set to ['N', 'S', 'E', 'W']. The minimal legibility sizing meant that five labels could fit in any given container on the single-screen. As mentioned previously, when a Viewport Space is rendered on a nine-screen display, its projection is simply scaled up. In order to understand how the properties of larger screens affect usability, we decided to keep the label's pixel size constant. This means that we could now fit 15 legible labels in a given container on the nine-screen. While we realize this may be a confound to some degree, it allows us to ask the question of if we can improve Viewport performance by leveraging the larger screen size (with constant spatial resolution).

Tasks

In order to test how our IRVE layout techniques impact usability for search and comparison, we define four kinds of tasks (below). The task types are denoted by the following convention: [IRVE_TaskType: information-Criteria → informationTarget].

IRVE Search Tasks [S:] require subjects to either:*

- Find a piece of abstract information (A) based on some perceptual/spatial criteria (S). Task example [S:S → A]: What molecule is just outside of the nucleolus?
- Find a piece of perceptual/spatial information (S) based on some abstract criteria (A). Task example [S:A → S]: Where in the cell is the molecule with a molecular weight of 85.15?

IRVE Comparison Tasks [C:] require subjects to either:*

- Compare by some spatial criteria (S) and determine an abstract attribute (A). Task example [C:S → A]: 'Find the lysosome that is closest to a mitochondria. What is the melting point of the molecule in the lysosome?'
- Compare by some abstract criteria (A) and determine a spatial attribute (S). Task example [C:A → S]: 'Where in the cell is the molecule with the lowest melting point?'

Experiment and Method

We used a mixed design for this experiment (Figure 4). Subjects were randomly divided into two groups for the

N=16; time, accuracy, difficulty, satisfaction		<i>Within Subjects</i>		
			SFOV = 60	SFOV = 100
<i>Between Subjects</i>	1 screen	Object space	S:S→A S:A→S C:S→A C:A→S	S:S→A S:A→S C:S→A C:A→S
		Viewport workspace 1	S:S→A S:A→S C:S→A C:A→S	S:S→A S:A→S C:S→A C:A→S
	16 conditions = 16 environments per subject;			
	9 screen	Object space	S:S→A S:A→S C:S→A C:A→S	S:S→A S:A→S C:S→A C:A→S
Viewport workspace 2		S:S→A S:A→S C:S→A C:A→S	S:S→A S:A→S C:S→A C:A→S	

Figure 4. Experimental design.

between-subjects independent variable, which was the display size: one group performed all tasks on the single-screen display configuration and one group performed all tasks on the nine-screen display configuration. There were two within-subjects independent variables of two levels each: layout technique (Object or Viewport Space) and SFOV (60° or 100° vertical). For each condition, users were given one of each of the four task types mentioned above. Thus a total of 16 trials were presented to each subject in a counterbalanced order.

Users were introduced to each control mode of desktop VE navigation under the Cortona VRML browser. The metaphor was fixed to 'FLY' and users were educated and guided on how to use the plan, pan, turn, roll, go-to, restore, for control in the virtual world. Users were given the *Kelp Forest Exhibit* virtual environment, which is a 3D model of a large saltwater tank at Monterrey Bay Aquarium.³⁴ Users were directed to do things like 'fly into the tank; turn to your right 90° , is that a shark? Pan up to the surface; now down to the bottom; turn around; follow that diver ...' For the navigation portion of training, subjects took anywhere from 4 to 10 minutes to affirm that they felt comfortable with the controls.

Subjects were then given a sample 3D-cell environment with all the common landmark structures they would see in the experiment. In this environment, they were shown how to toggle object labels and how the cellular structures and molecules behaved depending on their proximity. Finally, they were instructed on the nature of the tasks. When users affirmed that they felt comfortable with the cell environment (typically 3–5 minutes), the experiment began.

In each trial, users were timed and recorded for correctness. In addition, they were asked to rate their satisfaction with the interface for that task and the level of difficulty of the task on a scale of 1–7. One part of each of three Cognitive Factors tests was given to each subject before the experiment began: Closure Flexibility (Hidden Patterns), Spatial Orientation (Cube Comparisons), and Visualization (Paper Folding).³⁵ This was intended to help understand the role of individual differences in utilization or preference of the various interfaces.

Results

For each trial, the dependent variables collected were: time, correctness, and user ratings of satisfaction and difficulty. A General Linear Model was constructed for these results to determine any significant effects and

interactions of the various experimental conditions to these metrics of usability. A *post hoc* analysis of the cognitive test scores using an independent samples *t*-test revealed that there was no significant difference between the two groups in terms of cognitive test scores.

Some general observations are notable. First, most users tended to search serially through the space in a lawn-mower pattern and used position controls more often than orientation controls. Across layout spaces, some users tended to select, read, and deselect objects along the way rather than keep them visible and travel on. In general, this strategy results in less visual clutter but required repeated re-selection if they did not immediately recall the information. After one or two experiences with a more exhaustive search, users typically adopted the strategy of leaving selected annotations visible until they occluded or distracted from their search.

There was a significant main effect on user accuracy across all tasks for the layout technique. The Viewport interface (mean = 84.7%) performed better than the Object space layout (mean = 75.6%) at $F_{1, 12} = 6.134$; $p = 0.029$. This result agrees with our first hypothesis and makes sense because with Viewport space, all active labels are visible and the HUD facilitates comparison. Because label size was controlled across levels, we know this is not a difference arising from legibility.

There was a significant interaction between display size and layout technique ($F_{1, 12} = 5.587$; $p = 0.036$): the single-screen group performed better under the Viewport interface, but the nine-screen performed better under the Object space layout (Figure 5). The large display and Viewport interface combination requires a large saccade and head movement in order to follow a connector line between an object and its label in another part of the display. In contrast, on the small display, there is little or no head and eye movement, connector lines are shorter and a given number of labels may be divided into more than one container. An additional advantage that Object space might have on the large display is that there is less occlusion between labels on the large displays.

There was also a significant interaction between layout, SFOV, and display. On the single-screen display, both techniques were roughly equivalent at small SFOV, but at large SFOV the Viewport interface provided a significant advantage. Figure 6 depicts the interaction of these three variables where $F_{1, 12} = 4.798$; $p = 0.049$. We attribute this difference to the high distortion of 100° SFOV on a 24° DFOV, a condition where Viewport space is more effective.

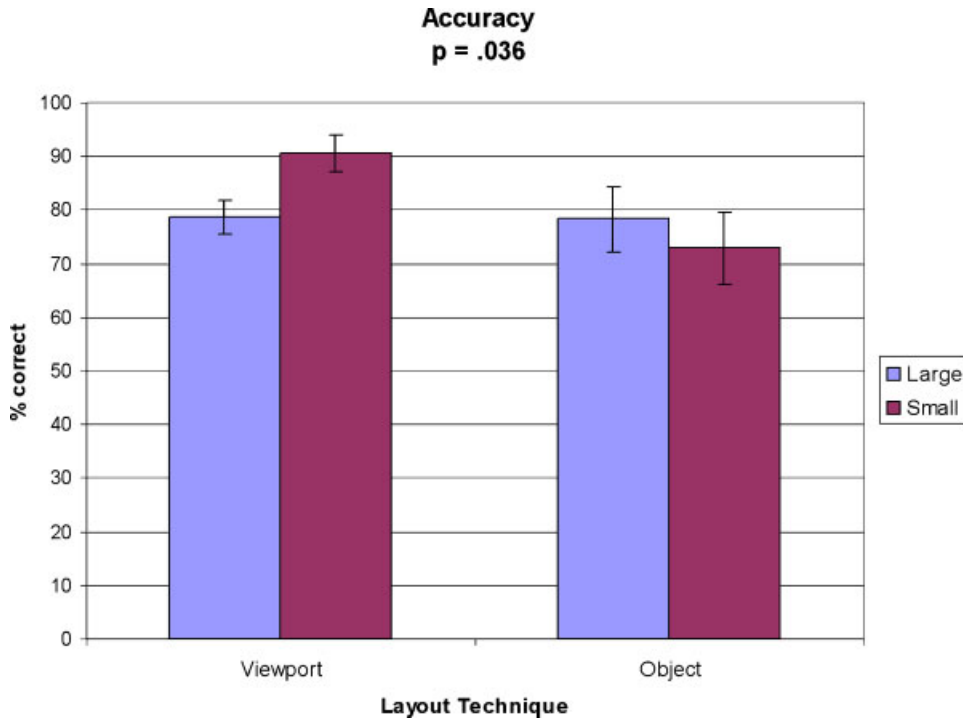


Figure 5. Interaction of Display size and Layout Technique on accuracy.




Figure 6. Interaction of layout technique, display size, and SFOV variables on accuracy.




Figure 7. Main effect of SFOV for search task accuracy.

Task-specific Results

For search tasks, there was a significant main effect for SFOV ($F_{1, 14} = 7.56; p = 0.016$) with the high SFOV being more accurate (95.3%) than the low SFOV (81.3%). This result, which is shown in Figure 7, can be explained because with a high SFOV, users can see more of the scene in the projection at any given time. For comparison tasks, small SFOV was significantly more accurate and this was a main effect ($F_{1, 14} = 4.61; p = 0.05$). This result also aligns with our hypotheses that comparison tasks (especially those on spatial criteria) may suffer under visual distortion.

The interaction of Layout and Display variables was mostly due to relative performance on comparison tasks (Figure 8). Here, Layout and Display were a significant combination $F_{1, 14} = 13.44; p = 0.003$. The Object space outperformed Viewport on the large display by percent correct (71.9 vs. 62.5), but the effect was reversed on the small display (62.5 vs. 87.5). Pairwise t -tests reveal that the significant difference between layouts on comparison tasks arises on the small display ($t(7) = 5.292; p = 0.001$).

Time

Subjects were timed from their first input event until the time they gave an answer they were confident in. The

sum time to complete all 16 tasks was longer for the nine-screen group than the one-screen group (32% longer), and this difference was almost significant ($t_{1, 14} = 0.184; p = 0.091$). There are a few interpretations for this result; the most obvious being the slower frame-rate with the nine-screen rendering (typically 1.2 fps vs. 6.7 fps during travel).

In addition, the physical size of the nine-screen display required users to make more mouse and head motion than when using a single-screen. In order to account for these differences, subsequent analysis was based on an 'adjusted time' for each group where the fastest possible completion time for a given trial was subtracted from each subject's recorded time for that trial. It should be noted that the effects described here were significant regardless of whether raw or adjusted time was used.

Time performance across tasks and displays carried significant main effects for both Layout technique and for SFOV. Figure 9 shows that the Object space interface (mean = 127.7 seconds) took longer than the Viewport interface (mean = 101.4 seconds); $F_{1, 12} = 5.244; p = 0.041$. The low SFOV of 60 (mean = 131.2 seconds) also took longer than the 100 SFOV (mean = 97.9 seconds) with $F_{1, 12} = 11.805; p = 0.005$ (Figure 10). This follows our general hypothesis that the Viewport interface would be advantageous over the Object interface and that larger SFOVs would be advantageous over smaller

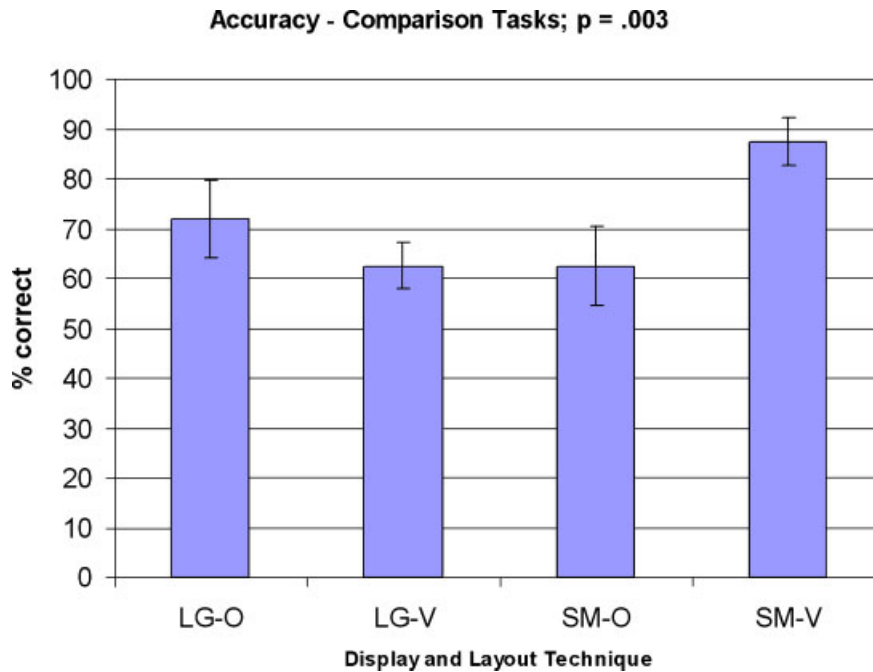


Figure 8. Interaction of Screen-size and layout technique on comparison task accuracy.

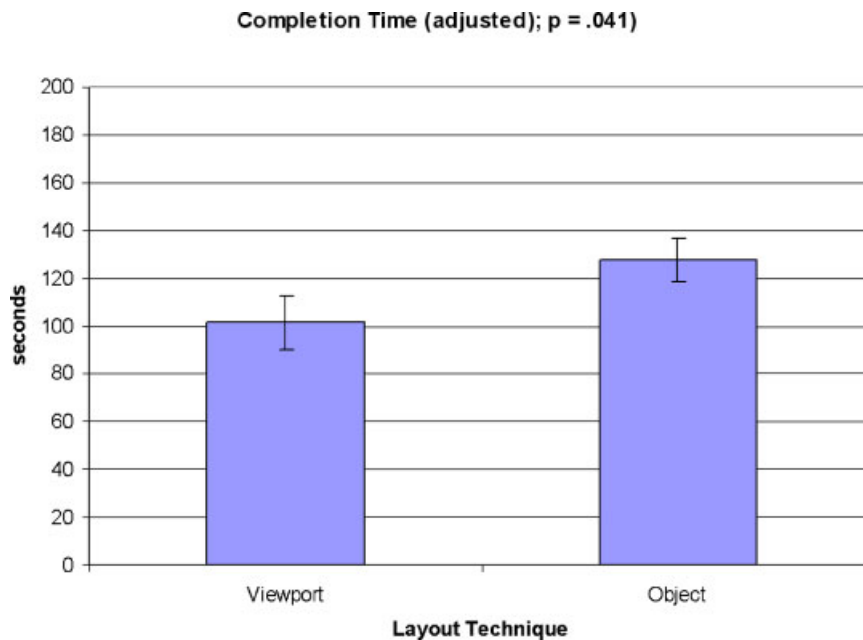


Figure 9. Main effect of Layout Technique on completion time.

SFOVs. This result is true of both search and comparison tasks.

There was also a significant interaction between the Layout and the SFOV ($F_{1, 12} = 19.094$; $p = 0.001$) variables. On low SFOVs of 60, the Object space technique

took longer than Viewport, whereas on 100 SFOV the Object space was slightly faster than Viewport (Figure 11). For the tasks we tested, it is clear that a 60 SFOV is a poor performer and in addition, this was a particularly poor combination with Object space layouts

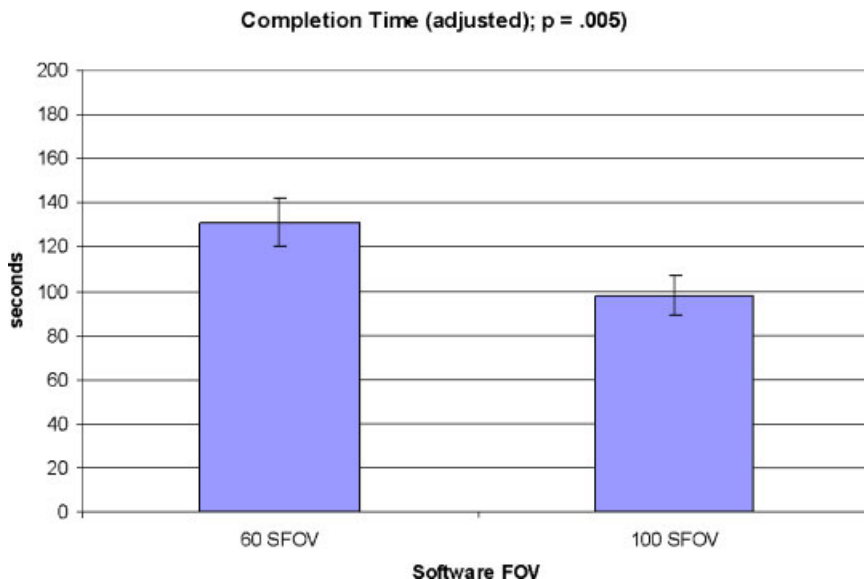


Figure 10. Main effect of SFOV on completion time.

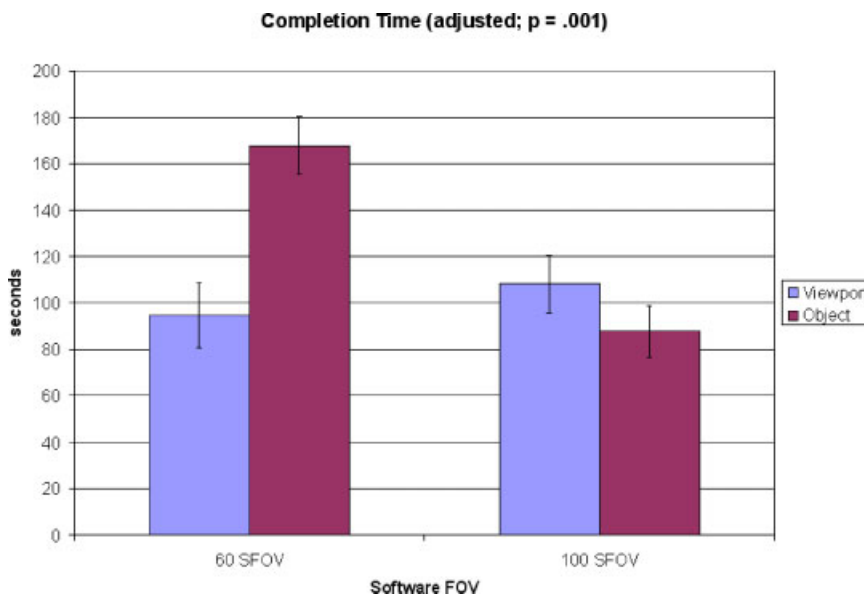


Figure 11. Interaction effect for SFOV and layout technique on completion time.

as the user is required to perform a lot of viewpoint manipulation to get the label into the viewing frustum.

Satisfaction and Difficulty

Results on these qualitative metrics are what we would expect from knowing about the relative performance of the layouts and SFOVs obtained by objective measures.

The subjective results actually followed the pattern for Time performance. For example, subjects rated the Viewport interface more satisfying ($F_{1, 12} = 5.788$; $p = 0.033$) and the Object space layout most difficult ($F_{1, 12} = 35.396$; $p < 0.001$). Subjects also rated the low SFOV as more difficult than the high SFOV and this difference was significant ($F_{1, 12} = 5.330$; $p = 0.040$).

There was also an interaction between layout technique and SFOV for both qualitative metrics. While

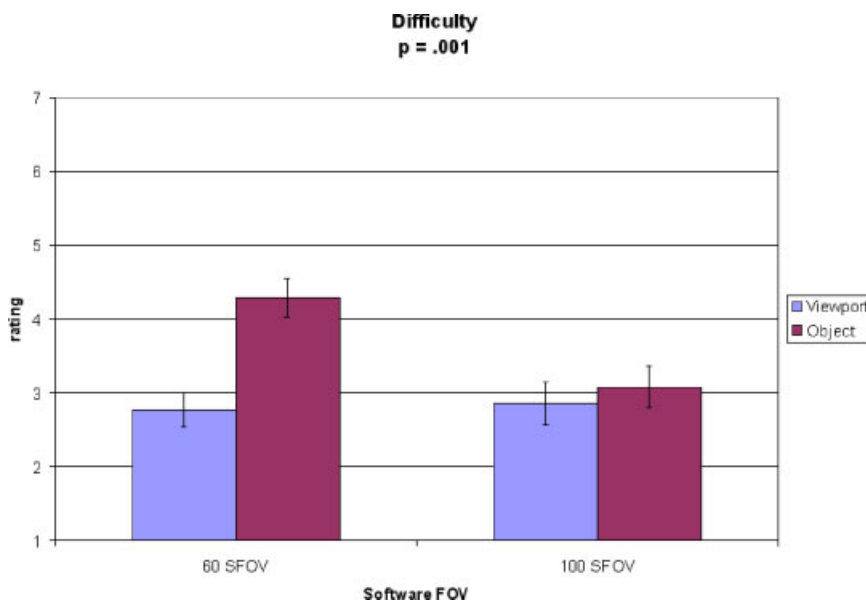


Figure 12. Interaction of layout technique and SFOV on user difficulty rating.

both interface types were rated similarly on the large SFOV conditions, in the small SFOV conditions subjects preferred the Viewport workspace ($F_{1, 12} = 8.007$; $p = 0.015$) and it was perceived as less difficult ($F_{1, 12} = 17.684$; $p = 0.001$). Figure 12 depicts this relationship.

Conclusions

Interface designs for IRVEs such as those used in cell biology research and education can benefit from a better understanding of the role of depth and association cues in supporting search and comparison tasks. In such environments, objects may be distributed in all three dimensions and there may not be a horizon or gravity constraint on navigation. The challenge facing designers and developers is, understanding the relationship of their information design choices (such as layout space) to the usability of their applications. For example, *Where and how should enhancing abstract information be displayed relative to its spatial referent so that the respective information can be understood together and separately?* The design problem is further compounded when considering the transfer of design layouts across rendering platforms.

In this study, we explored the relative performance of two IRVE layout spaces for search and comparison tasks in a desktop context. The first was an annotation layout scheme where the labels were co-located with their referent objects in the virtual scene in what we call

Object Space. While this technique provides a tight spatial coupling (via depth cues and Gestalt proximity) between the annotation and its referent object, annotations may not be fully visible because of other occluding objects in the scene. To guarantee visibility regardless of position or orientation in the VE, we developed an IRVE layout component that manages annotations on a HUD just beyond the near-clipping plane (Viewport Space). This study investigated the information design tradeoff between the spatial coupling guarantee or the visibility guarantee provided by annotation labels in either Object or Viewport layout spaces. In addition, we asked if the relative advantages of a layout space holds when the scene is rendered on a large screen or under large projection distortion.

Object versus Viewport Space

The first set of conclusions regards the usability of our IRVE layout techniques on a common single-screen setup. 'We asked: Is one layout space with guaranteed visibility better than one with guaranteed tight spatial coupling for certain tasks?' The results of this experiment showed that overall the Viewport interface outperformed Object space layouts on nearly all counts of accuracy, time, and ratings of satisfaction and difficulty across tasks. In other words, for the set of tasks performed, tight-spatial coupling of annotation to its referent (Object Space) was not as advantageous or

preferable as the consistent visibility provided by an image plane layout (Viewport Space).

This result suggests that the development and evaluation of a richer set of Viewport Space layout capabilities (such as the X3D Layering and Compositing Components) would be worthwhile. If the tight spatial coupling provided by Object Space layouts is deemed necessary (for aiding spatial comparison tasks for example), consider further refining Object Space designs including managed or emergent layout schemes to provide consistent visibility.

Single and Nine-screen Configurations

One of the main drawbacks to using our interfaces on the nine-screen display was the slower frame-rate. The VRML browser we used in the study did not work with the operating system to manage the hardware rendering with multiple video cards and displays. When the browser was enlarged to 1.5×1.5 screens or greater, the application switched to a software-rendering mode, which seemed significantly slower. However, the differences in time to completion across display configurations (due mainly to rendering speed) were not statistically related to task performance. We also found no statistically significant effect of display configuration on user accuracy.

The second research question we posed was: 'Do the advantages of visibility or tight spatial coupling hold if the screen size is increased?' Display size interacted with both Layout and SFOV variables for accuracy. The worst performing combination was the Object Space with a high SFOV on a small display. The best performing combination was the Viewport Space with high SFOV on a small display. However, on the large display, high SFOV the Object Space outperformed the Viewport Space.

With the tight spatial coupling, Object Space annotation schemes render the annotation with the rest of the scene. Annotations end up on the image plane nearby their referents—they provide the additional depth cues of occlusion and motion parallax and the additional Gestalt association cues of proximity with their referents. We can postulate that the advantage of the tight spatial coupling of Object Space only comes into effect when there is enough screen size (DFOV) to avoid the occlusion problem. Also, on the large screen size, tight spatial coupling means that users do not need to perform large saccades or head movements to see and read the annotation.

In examining the transfer of the Viewport BorderLayout interface design across display configurations, we

can that say the successful transfer of an interface to a larger display is not simply a matter of scaling. On the large display, our Viewport Space design had the capacity for three times as many annotations. However on the large display, ergonomics require special consideration. The BorderLayout Viewport Space annotations began in the N container, which was above the line of sight at the top edge of the nine-screen display. This made frequent reference fatiguing for users. There is substantial work to be done in exploring Viewport Space annotation designs, especially for large displays. This work suggests that design and management choices for image-plane-interface layouts may be different depending on the size of the display.

Software Field of View

The third research question this study addresses is: 'Do the advantages of visibility or tight spatial coupling hold if the SFOV is increased?' Preliminary results indicated that for the cell environment, users had a high tolerance for large SFOVs, but that the tolerance was much less on the large display. In the study overall, users significantly rated low SFOV conditions more difficult; the differences in satisfaction ratings between SFOVs was not significant. Because we cannot compare subjective metrics between subject groups, the relationship between DFOV and SFOV remains an open research question.

Our study results showed that overall our two SFOVs levels did not significantly affect accuracy performance. However, higher SFOVs were advantageous for time especially on search tasks, but negatively impacted accuracy especially on comparison tasks. This result supports our hypotheses about the benefits of a high SFOV for search tasks (by showing more of the scene in the periphery) and liability of a high SFOV for comparison tasks (by distorting a scene object's spatial location). It suggests that designers may consider modifying the SFOV dynamically depending on the user task.

Summary

Reflecting on the implications of these results, we can answer our original hypotheses and substantiate the following IRVE information design claims:

- Overall, the guaranteed visibility of Viewport Space offered significant performance and satisfaction advantages over the tight spatial coupling of Object Space annotation layouts. The effect was especially pronounced in the single-screen monitor configuration.

- The advantages of our Viewport Space layout did not transfer cleanly or scale iso-morphically up to the larger nine-screen configuration. On the large display condition for example, tight spatial coupling (Object Space) was more effective for accuracy across tasks but especially for comparison.
- Higher SFOVs decreased search time because they render more of the scene in the projection. Higher SFOV increased spatial comparison times because of fish-eye distortion.

The results of this evaluation contribute to our understanding of a fundamental layout space tradeoff in IRVEs. For example, they demonstrate that the visibility of the respective information types may be the more important design criteria—more important than visual or spatial association between information types. In addition, they provide initial guidance as to the challenges of designing integrated information spaces that are portable across display sizes and distortions. Still, the relationship between interface usability, SFOV and DFOV is an open research question; for example, what are the thresholds of screen size or projection distortion where various techniques break down and others become advantageous?

This evaluation has shown that there are display conditions where both Object and Viewport layouts can be advantageous. Designs and capabilities for both Object and Viewport layouts must be improved. For example, successful, portable IRVEs will require better text rendering facilities, layering and compositing functionality as well as support for pixel-agnostic layout mechanisms for the image plane. Future work includes continued design and evaluation of IRVE information display techniques to further examine the role of depth and association cues for common tasks.

ACKNOWLEDGEMENTS

The authors thank Chris North for his advisement and lab space, Robert Ball for his support and use of the nine-panel display, and the graduate students of CS 6724 Seminar in 3D Interaction for their helpful comments. Finally the authors thank Ron Kriz and Patrick Shinpaugh for their continued support through VT's University Visualization and Animation Group.

References

1. Bolter J, Hodges LF, Meyer T, Nichols A. Integrating perceptual and symbolic information in VR. *Computer Graphics and Applications IEEE* 1995; **15**(4): 8–11.


2. Bowman D, North C, Chen J, Polys N, Pyla P, Yilmaz U. Information-rich virtual environments: theory, tools, and research agenda. In *Proceedings of VRST*. ACM Press: Osaka, Japan, 2003; 81–90.
3. Card S, Mackinlay J, Shneiderman B. *Information Visualization: Using Vision to Think*. Morgan Kaufmann: San Francisco, 1999.
4. Goguen J. Information visualizations and semiotic morphisms. <http://citeseer.ist.psu.edu/goguen00information.html>, UCSD. 2000.
5. Ware C. Design as applied perception. In *HCI Models, Theories, and Frameworks: Towards a Multidisciplinary Science*, Carroll JM (ed.). Morgan-Kaufmann: San Francisco, 2003.
6. Thomas JJ, Cook KA. 'A visual analytics agenda'. *IEEE Computer Graphics and Applications* 2006; **26**(1): 10–13.
7. Polys NF, Bowman DA. Desktop information-rich virtual environments: challenges and techniques. *Virtual Reality* 2004; **8**(1): 41–54.
8. Faraday P, Sutcliffe A. Designing effective multimedia presentations. In *Proceedings of CHI*. ACM Press: Atlanta, GA, 1997; 272–278.
9. Faraday P, Sutcliffe A. Making contact points between text and images. *Multimedia*. ACM Press: Bristol, UK, 1998; 29–37.
10. Chandler P, Sweller J. Cognitive load theory and the format of instruction. *Cognition and Instruction* 1991; **8**: 293–332.
11. Feiner S, Macintyre B, Haupt M, Solomon E. Windows on the world: 2D windows for 3D augmented reality. In *Proceedings of UIST*. ACM Press: 1993; 145–155.
12. Barrilleaux J. *3D User Interfaces with Java3D*. Manning Publications Co: Greenwich, CT, 2001.
13. Bell B, Feiner S, Hollerer T. View management for virtual and augmented reality. In *Proceedings of UIST*. ACM Press: 2001; 101–110.
14. Wei B, Silva C, Koutsofios E, Krishnan S, North S. Visualization research with large displays. *Computer Graphics and Applications IEEE* 2000; **20**(4): 50–54.
15. Tan D, Gergle D, Scupelli P, Pausch R. With similar visual angles, larger displays improve performance. *CHI* 2003; **5**(1): 217–224.
16. Swaminathan K, Sato S. Interaction design for large displays. In *ACM Interactions*. January, ACM Press: 1997.
17. Mackinlay J, Heer J. Wideband displays: mitigating multiple monitor seams. In *Proceedings of CHI*. ACM Press: Vienna, Austria, 2004; 1521–1524.
18. Czerwinski M, Smith G, Regan T, Meyers B, Robertson G, Starkweather G. Toward characterizing the productivity benefits of very large displays. In *Proceedings of INTERACT*. 2003; 9–16.
19. Draper MH, Viirre ES, Furness TA, Gawron VJ. Effects of image scale and system time delay on simulator sickness within head-coupled virtual environments. *Human Factors* 2001; **43**(1): 129–146.
20. Furnas GW. Generalized fisheye views: visualizing complex information spaces. In *Proceedings of CHI*. ACM Press: 1986; 16–23.
21. Gutwin C, Skopik A. Fisheye views are good for large steering tasks. *Proc. CHI* 2003; **5**(1): 201–208.
22. Shneiderman B. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings of IEEE Visual Languages*. Boulder, CO: 1996.
23. Bowman D, Kruijff E, LaViola J, Poupyrev I. *3D User Interfaces: Theory and Practice*. Addison-Wesley: Boston, 2004.

24. Polys NF, Bowman DA, North C. Information-rich virtual environments: challenges and outlook. *NASA Virtual Iron Bird Workshop*, NASA Ames, <http://ic.arc.nasa.gov/vib/index.php>. 2004.
25. Polys N, Bowman D, North C, Laubenbacher R, Duca K. PathSim visualizer: an information-rich virtual environment for systems biology. In *Proceedings of Web3D*. ACM Press: Monterey, CA, 2004; 7–14.
26. Goldstein EB. *Sensation and Perception* (5th Edn). Brooks/Cole Publishing: 1999.
27. Polys NF. stylesheet transformations for interactive visualization: towards web3D chemistry curricula. In *Proceedings of Web3D*. ACM Press: St. Malo, France, 2003; 85–90.
28. Murray-Rust P, Rzepa Henry S, Wright M. Development of chemical markup language (CML) as a system for handling complex chemical content. *New Journal of Chemistry* 2001; 618–634.
29. McClean P, Saini-Eidukat B, Schwert D, Slator B, White A. Virtual worlds in large enrollment biology and geology classes significantly improve authentic learning. In *Proc 12th International Conference on College Teaching and Learning (ICCTL-01)*. J. A. Chambers. Jacksonville, FL, Center for the Advancement of Teaching and Learning 2001; 111–118.
30. North Dakota State University World Wide Web Instructional Committee (NDSU WWWIC), <http://www.ndsu-nodak.edu/instruct/mcclean/wwwic>. 2004.
31. Chen J, Pyla P, Bowman D. Testbed evaluation of navigation and text display techniques in an information-rich virtual environment. In *Proceedings of Virtual Reality*. IEEE: Chicago, IL, 2004; 181–188.
32. Bowman DA, Hodges LF. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *Proceedings of Interactive 3D Graphics*. ACM Press: 1997; 35–38.
33. Poupyrev I, Weghorst S, Billinghamurst M, Ichikawa T. The go-go interaction techniques: non-linear mapping for direct manipulation in VR. In *Proceedings of UIST*. ACM Press: 1996.
34. Brutzman D. Teaching 3D modeling and simulation: virtual kelp forest case study. In *Proceedings of Web3D*. ACM Press: Tempe, AZ, 2002; 93–101.
35. Ekstrom RB, French JW, Harman HH. *Manual for Kit of Factor Referenced Cognitive Tests*. Educational Testing Service: Princeton, NJ, 1976.

fication from 1999 to the present, Nicholas has also conducted numerous courses on the design and implementation of real-time networked 3D and techniques for 3D User Interfaces. His research interests include Human-Computer Interaction and Information Architectures for visualizing large heterogeneous data sets. He did his BA in Cognitive Science from Vassar College in 1996 and received his PhD in Computer Science and Human-Computer Interaction from Virginia Tech in 2006. Dr. Polys now works for the IT and Research Computing Division building visual analysis tools that advance science, medicine, and education.



Seonho Kim is a PhD student in Department of Computer Science at Virginia Tech. He did his BS in Department of Computer Science at Yonsei Univ, Korea, MS in Department Computer and Communication Engineering at POSTECH, Korea, and is working as a researcher in Korea Institute of Science and Technology Information (KISTI). His research interests include information system, digital libraries, and human computer interaction.



Doug A. Bowman is an Associate Professor in the Center for Human-Computer Interaction and the Department of Computer Science at Virginia Tech, where he directs the 3D Interaction Group. His research interests cluster around the areas of three-dimensional user interfaces and immersive virtual environments. Along with three colleagues, he is the author of the recent book *3D User Interfaces: Theory and Practice*, and he is also serving as the General Chair for IEEE Virtual Reality 2007. Dr. Bowman received his undergraduate degree from Emory University and his MS and PhD from the Georgia Institute of Technology.

Authors' biographies:



Nicholas F. Polys has worked in the creation and management of digital assets and virtual environments for more than 9 years. A key contributor to the X3D speci-