

Worlds-in-Wedges: Combining Worlds-in-Miniature and Portals to Support Comparative Immersive Visualization of Forestry Data

Jung Who Nam*

Computer Science & Engineering
University of Minnesota

Krista McCullough†

Computer Science & Engineering
University of Minnesota

Joshua Tveite‡

Computer Science & Engineering
University of Minnesota

Maria Molina Espinosa§

Computer Science & Engineering
University of Minnesota

Charles H. Perry¶

USDA Forest Service
Northern Research Station

Barry T. Wilson||

USDA Forest Service
Northern Research Station

Daniel F. Keefe**

Computer Science & Engineering
University of Minnesota




Figure 1: The worlds-in-wedges technique is applied here to visualizing three forests from the Forest Inventory and Analysis dataset.

ABSTRACT

Virtual reality (VR) environments are typically designed so users feel present in a single virtual world at a time, but this creates a problem for applications that require visual comparisons (e.g., forest scientists comparing multiple data-driven virtual forests). To address this, we present Worlds-in-Wedges, a 3D user interface and visualization technique that supports comparative immersive visualization by dividing the virtual space surrounding the user into volumetric wedges. There are three visual/interactive levels. The first, worlds-in-context, visualizes high-level relationships between the worlds (e.g., a map for worlds that are related in space). The second level, worlds-in-miniature, is a multi-instance implementation of the World-in-Miniature technique extended to support multivariate glyph visualization. The third level, worlds-in-wedges, displays multiple large-scale worlds in wedges that act as volumetric portals. The interface supports navigation, selection, and view manipulation. Since the techniques were inspired directly by problems facing forest scientists, the interface was evaluated by building a complete multivariate data visualization of the US Forest Service Forest Inventory

and Analysis public dataset. Scientist user feedback and lessons from iterative design are reported.

Keywords: worlds-in-miniature, 3D user interface, presence, comparative visualization

Index Terms: Human-centered computing—Virtual reality—;—Human-centered computing—Interaction techniques—; Human-centered computing—Scientific visualization—;—Human-centered computing—Geographic visualization—.

1 INTRODUCTION

Many data visualizations benefit from the sense of presence that is possible with immersive visualization in virtual reality (VR). In VR, users can analyze spatial relationships in data relative to their own bodies. They can also “experience” their data rather than just view it. For example, immersive visualizations make it possible for planetary geologists to stand near a martian crater and analyze slopes, areas in shadow, or other land features using the same spatial reasoning skills they would at a field site on Earth [9]. In VR, archaeologists can use spatial reasoning while analyzing artifacts to naturally refer back to memories of working on the actual dig site (e.g., remember we found a similar lamp to this one *in the trench over there* [while pointing]) [1]. Using a multi-modal, immersive speech-giving interface, scholars of ancient rhetoric can literally *feel* what it was like to give a speech at the hillside of the Pnyx in ancient Athens, recreated under different data-driven conditions for the assembly size and time period [14]. In all of these examples, the sense of presence enabled by VR adds value to the data visualizations, but for VR to be a truly valuable data visualization tool it must do more than immerse users within their data.

One critical task for many visualizations to support is visual

*e-mail: namxx054@umn.edu

†e-mail: mccullough.k.n@gmail.com

‡e-mail: jtveite@gmail.com

§e-mail: molin200@umn.edu

¶e-mail: charleshperry@fs.fed.us

||e-mail: barrywilson@fs.fed.us

**e-mail: dfk@umn.edu

comparison [11, 13]. Almost every data-driven scientific study or scholarly work involves some form of comparison, but, as Kim et al. report in their recent survey [13], comparison is an understudied problem in 3D visualization, especially for VR-based visualization.

In VR, whether used for data visualization or some other application, the tradition is to immerse oneself in a single world at a time, but this approach is known to be problematic for making visual comparisons. As Munzner puts it, “eyes beat memory” [18]. Juxtaposition (side-by-side), superimposition (semitransparent or other forms of overlays), and explicit encoding (when differences are possible to compute explicitly) are three useful, fundamental approaches to visual comparison that, in theory, are all preferable to trying to remember the details of a first dataset while looking at the second, or third, or fifth!

Our specific research is motivated by a collaboration with forest scientists who study data collected by hand by surveyors who inventory each tree within sample plots spread throughout the United States [19]. More than 20 variables are collected to describe the species, size, health, and other parameters of each tree. This data is used, in part, to develop carbon estimates of the nation’s forests, which impact climate policy. Using this data, we have developed algorithms to generate data-driven virtual reconstructions of plots throughout the U.S. Like the planetary geologists on Mars, this enables forest scientists to see and experience their data in a life-size scale that is intuitive, drawing upon their own lived experiences visiting the forests. However, the most interesting thing to do with this data is to compare the plots. How do we do this in VR? It does not make sense to superimpose the plots because they do not share the same spatial reference frame. If not superimposition, then perhaps juxtaposition (side-by-side), but how would this desktop-based technique translate to immersive visualization in VR?

To contribute to understanding these research questions, we introduce a novel interaction and visualization technique called Worlds-in-Wedges. Our approach divides the virtual space surrounding the viewer into pie-slice shaped volumetric wedges and fills each wedge with a different virtual world. By stacking the wedges around a common axis and placing the tracked user at the center, it is possible to create an immersive juxtaposition of spatial data, where the user may look around, using head-tracked stereoscopic graphics to compare multiple worlds. For example, Fig. 1, shows a comparison of forests in the West, Midwest, and East of the US.

This core idea can be enhanced with several extensions. It is important to have a strong spatial awareness of one’s location in each of the worlds. So, we utilize a multi-instance version of the classic World-in-Miniature (WIM) technique [21, 25] to provide a 3D map of each world. We also extend each WIM to function as a multivariate visualization. To convey how the worlds relate to one another (i.e., each plot’s location within the U.S.), we provide a worlds-in-context view. Finally, we consider how to support standard VR tasks, such as navigation and selection in this new style of multi-world immersion.

The specific contributions of this research are:

- A novel VR visualization and 3D user interaction technique for making simultaneous visual comparisons of multiple virtual worlds at real-life scale.
- A volumetric portal metaphor that combines portal and magic lens techniques to support simultaneous immersion in multiple virtual worlds.
- An extension of the traditional World-in-Miniature technique to demonstrate how multi-instance WIMs constructed with multivariate glyphs can be used for immersive visualization.
- Design, results, user feedback, and discussion for a case study based on our motivating application to visualizing the unique US Forest Inventory and Analysis dataset [19].

2 RELATED WORK

Our research builds upon related work in both data visualization and virtual reality.

2.1 Experiential Analysis Using VR

Architectural walk-throughs and other first person experiences have long been considered the killer apps for VR. Applications in this style construct data-driven worlds at life-size scale and make it possible for the user to experience the world as in real life. We can walk through historical sites that no longer exists [1, 2, 14] or even places like Mars, where it is impossible or difficult to be [9, 20]. This is a style of data analysis that is often overlooked since it is less quantitative than other forms. In experiential data analysis, the *feeling* of what it is like to speak to a crowd of 10,000 ancient Athenian citizens is just as important as the ability to accurately measure the dimensions of an architectural reconstruction [14]. This style of first-person, experiential data analysis is what we aim to support, but, unlike these prior works, Worlds-in-Wedges brings this style of analysis to the problem of simultaneous visual comparisons.

2.2 Comparative Visualization Using VR

Comparison is a crucial task for data analysis [10, 11, 13], but it is not well understood how to facilitate comparative analysis for datasets that are best displayed in VR [13]. A survey published in 2017 found that only 2 out of 41 visualizations for comparing spatial 3D and 4D data utilize VR and points to the need to investigate more in the area. The survey identifies four fundamental approaches for comparative visualization: Superimposition, Juxtaposition, Interchangeable (over time), and Explicit Encoding. One of the two VR comparative visualization examples, a system called Bema [14], supports comparison of multiple historical phases of archaeological reconstruction via an Interchangeable approach. The user is immersed in one world at a time and uses an interactive timeline to switch between worlds via a fade-in-fade-out transition. This worked, in part, because the three historical phases were spatially co-located. In contrast, Worlds-in-Wedges uses what would be called a Juxtaposition approach. The theoretical advantage with this approach is that the data is simultaneously visible, so users rely upon their eyes, not just their memory, to make comparisons. Also, juxtaposition does not require datasets to be able to be spatially registered so that they can be drawn on top of each other in a useful way. Like our work, the other VR example of comparative visualization, a perceptual study by Coffey et al. [5], uses a Juxtaposition approach, but the data is quite different. They compare animated abstract visual representations of neck motions (i.e., biomechanics data), which are small in scale. Our technique is tailored toward comparing life-size virtual worlds, where the impression is that the viewer is immersed in the data rather than seeing data as a floating object. This style of difference (i.e., *looking-in* vs. *looking-out* tasks) is potentially a critical perceptual distinction for visualization in VR environments [7]. Taking this, along with our desire to include first person VR immersive *experience* as part of the goal of the visualization, we are motivated to explore how comparative visualization in VR might work in the context of the need to “look out” at multiple life-size data-driven worlds.

2.3 Portals and 3D Magic Lens Metaphors

As in the portal metaphor in VR [4, 15, 24], in Worlds-in-Wedges, another virtual world is only visible through a three-dimensional “window”. However, rather than treating a portal as a window made of virtual glass, our interface does not clip the objects that are between the user and the window frame. This volumetric approach to portals is similar to the volumetric lenses introduced by Viega et al. [28] and the 3D-Box technique introduced by Kunert et al. [15]. However, these prior examples work with small-scale volumes, such as hand-held tools for applying filters to data within a small volume or as references for travel by teleporting. Our volumetric space is




Figure 2: The technique is composed of three linked layers, which can be interpreted differently depending upon the dataset. At the worlds-in-wedges level, each wedge always presents an immersive view of one world. At the worlds-in-miniatures level (the discs floating in the middle of the space), miniature versions of each world are displayed using data-driven glyphs or other representations that are helpful in a bird's eye view. Finally, at the worlds-in-context level, the relationships between the different worlds that can be activated in the display are visualized. On the left, in our forestry example, the context is spatial – each world corresponds to a plot at a specific latitude and longitude where the forest service has collected data, and worlds for three such plots are currently active. In the middle, we imagine an application to a different dataset where the relationship between multiple worlds is temporal rather than spatial. Here, multiple worlds are contextualized on a timeline, as in multiple historical reconstructions of an archaeological site. On the right, we imagine a more abstract relationship. Here, some high dimensional datasets (e.g., mathematical functions, parameterized simulations) are related through a 2D plot that displays a dimensionally reduced version of the data.

much larger. Also, our wedges are not finite volumes. Aside from the far clipping plane in the graphics renderer, there is no back face to the volume, allowing the user to look out to the world's horizon.

2.4 World-in-Miniature Techniques

Traditional WIM interfaces provide a miniature version of a virtual world that acts as a 3D map as well as an interaction space [21, 25]. Many extensions to the technique have been introduced to support features, such as scaling and scrolling [31], automatically selecting optimal views [27], and new ways of interacting with the WIM [6, 16, 25, 29]. Early work demonstrated how multi-instance WIMs can be constructed in order to depict multiple virtual worlds or multiple locations (e.g., different rooms) within the same world [21], a concept that is quite similar to our use of WIMs. The difference in our approach is linking each WIM instance with a corresponding immersive view (in a wedge). Our approach also incorporates another visual layer, placed on the bottom of WIMs, to convey how these WIMs are related to another. Finally, rather than rendering an exact miniature version of each world, we render an abstract data visualization of the world, where the terrain is intentionally simplified and realistic trees are replaced with multivariate data glyphs.

3 THE WORLDS-IN-WEDGES TECHNIQUE

This section describes the visualization and interaction techniques that make up the new approach to immersive comparative visualization we call the Worlds-in-Wedges technique. The technique requires a head-tracked stereoscopic VR display and a 6 degree-of-freedom tracked wand that is held in the dominant hand. As is common in today's commercially available VR hardware, the wand must have a primary button (e.g., a trigger button) as well as a touchpad or joystick to serve as a secondary button and provide 2D analogue input. Input from the non-dominant hand is limited to a button, which can be provided by another VR controller or a low-cost device, such as a wireless mouse. Our current implementation uses the HTC Vive VR display and its two controllers.

3.1 Visual Layout: A Linked, Three-Level Visualization

The Worlds-in-Wedges technique is composed of three-levels of visuals, as diagrammed in Figure 2.

Level 1: Worlds-in-Context

The worlds-in-context layer visualizes the relationship between the different virtual worlds. The specific relationship to display is determined by the comparison task at hand. In this paper, we fully develop the forestry example, which is based on comparing different physical locations. In this case, it makes sense to contextualize the different locations on a map. However, to assist others who might wish to apply the technique to other non-spatial datasets, we suggest two other possible mappings in Fig. 2. For comparing changes of a historical site at different points of time, the Worlds-in-Context layer could be realized as a timeline, where each world would be represented with a marker placed along the timeline. For comparing more abstract data, the Worlds-in-Context layer could be realized using a 2D plot, for example a scatter plot or a plot of a dimensionally reduced version of the data.

Level 2: Worlds-in-Miniatures

The worlds-in-miniatures layer is a multi-instance version of the classic World-in-Miniature VR interface [21, 25]. As in the traditional WIM, each WIM is a 3D map of the corresponding virtual world that can also be used as an interaction space. There is a visual indicator within the miniature world for the user's current location. However, rather showing an avatar of the user, our WIM shows a rendering of the wedge that is displayed in the next layer. We also reinterpret WIMs as a data visualization tool that complements the immersive display by displaying multivariate data-driven glyphs that reveal more information for each object in the world rather than simply rendering a scaled-down version of the same 3D graphics used to render the immersive scene.

Level 3: Worlds-in-Wedges

The worlds-in-wedges layer, which serves as inspiration for the name of the overall technique, divides the virtual space surrounding the user into wedge-shaped volumes, where each wedge acts as a volumetric portal, displaying a portion of a world. The wedges provide the first-person, head-tracked, stereoscopic views of virtual worlds, and only the portion of each world that falls into the wedge volume is rendered.

The three layers are visually linked using lines and corresponding colors. Each WIM is linked with a line to a point on the worlds-in-context layer. When the position of the WIM changes, the line's end point moves as well to preserve the visual connection. Each WIM




Figure 3: The horizontal positions of WIMs affect the order of the wedges. From the top screen capture to the bottom, the user has moved the blue WIM to the center, which moves the corresponding wedge to the center.

is linked with the corresponding wedge using a categorical color coding; the color of the WIM outer ring and that of the wedge’s frame is the same. When a new virtual world is activated, a new World-in-Miniature and World-in-Wedge pair is created, and the pair is assigned a new hue to mark it as unique. To help the user maintain spatial awareness, the WIMs are also linked to the wedges by orientation – that is, the WIMs are always rotated so that the mid-line of the miniature wedge drawn on each WIM as an indicator of the user’s current position and orientation is aligned with the mid-line of the larger wedge.

3.2 Interactive Techniques

Worlds-in-Wedges enables exploratory visualization through interaction techniques. These support manipulating the views as well as standard 3D user interface tasks, such as selection and navigation, which need to be adjusted to the new situation of interacting in wedges.

3.2.1 Activating and Deactivating Virtual Worlds

A data analysis session begins with only the worlds-in-context layer visible. To add a new WIM and wedge pair to the visualization, the user selects a point of interest within the worlds-in-context layer using a laser pointer. Depending on the dataset, it may be difficult to select this small target, so we utilize the flashlight technique [17] to select the closest point within a threshold of the ray-plane intersection. When a selection is detected, the ray changes its direction and snaps to the selected point, providing visual feedback to the user. When a wedge is added or removed, the size and location of each wedge is automatically updated to display each virtual world with the largest, non-overlapping view possible within a preset total viewing angle. It is possible for the total angle to be a full 360 degree view, but we often limit the view to 220 degrees for our application, which has a natural preferred orientation (North is forward and the user stands South of the US map looking down on it as if laid on a virtual table). The WIM positions, floating above the worlds-in-context layer, are calculated using a spring system. Each WIM is anchored to its position on the worlds-in-context layer and tries to locate itself directly above this position but is pushed away by other WIMs to avoid overlap.




Figure 4: The angle of each wedge is tied to the vertical position of the corresponding WIM. In the top image, the user has placed the orange WIM higher than others to see the corresponding world in more detail. In the bottom image, the orange WIM has been lowered, causing the angle of the corresponding middle wedge to shrink in proportion.

Deactivating a virtual world is done by selecting and grabbing onto the corresponding WIM using the dominant hand’s laser pointer and trigger button and moving it down toward the worlds-in-context layer. This activates a semi-transparent bounding plane, just above the worlds-in-context layer, and when the WIM is moved below this threshold, the world is deactivated.

3.2.2 Arranging Visuals in the VR Space

During comparison and discussion with colleagues, it is often useful to rearrange the display. For instance, the user may wish to place a specific world in the center to make it a baseline for comparing other worlds to the left and right. Or the user might wish to place closely related worlds next to each other, forming subgroups. This is accomplished by grabbing onto the WIM of interest and moving it horizontally as illustrated in Fig. 3. The relative positions of the WIMs set the ordering of the wedge views. Thus, when a WIM is moved to the left of another WIM, the corresponding wedge will swap places with its neighbor to the left.

It is also useful to resize the wedges based upon the current task. Worlds can be completely deactivated as described above, but this is not always desired because it takes time to find the world again and reactivate it. Instead, if the user wishes to zoom in and focus on one or two of the active worlds temporarily, then this can be done by grabbing onto and moving the corresponding WIM vertically. As shown in Fig. 4, the relative vertical positions of WIMs are used to set the angle of each wedge, like a vertical slider. The angle for each wedge is computed based on its WIM’s position relative to a global maximum and minimum height for all WIMs using the following equation:

$$a_i = \frac{h_i}{\sum_{j=0}^n h_j} a_{total}, \quad (1)$$

where a_i is the angle of $wedge_i$ that corresponds to wim_i , h_i is the normalized height of wim_i relative to the global max and min, n is the number of active virtual worlds, and a_{total} is the total angle of the wedges.

When interacting with either the WIMs described here or the wedges (via the navigation techniques described later), the interface




Figure 5: Ray-casting selection with a virtual laser pointer is used to identify objects of interest in the 3D scene for closer inspection. After pointing to a tree, a pop-up panel shows detailed data on demand.

follows a consistent approach in terms of responding to user input. All active worlds respond to head-tracking input by updating the view and projection matrices used to render the scenes, but only one WIM and wedge pair is considered currently in focus in terms of receiving input from the hand controllers. The wedge in focus is determined based upon a virtual laser pointer that extends out of the dominant hand's controller. Although this pointer can be thought of as "selecting" the active wedge, it is important to note that the wedges are large; so, this does not have the problems typically associated with laser pointer selections in VR. Instead, the technique is designed to feel to the user as if the wedge they are facing, as they naturally turn their body to look, becomes automatically focused.

The one useful exception to the focus rule is when the button held in the non-dominant hand is depressed. We treat this as a modifier, like the Control key on computer keyboards it modifies the current operation so that it applies to all active worlds rather than just the one in focus. When applied to the wedge resizing operation, this modifier means that the angles for all of the wedges are adjusted when one is moved. When the WIM being moved goes up, the other WIMs go down by the same amount. This is useful when the user wants to quickly focus on just one world. Conversely, when the WIM being moved goes down, the other WIMs go up, nearly hiding the world being moved to give the most display resolution to the other worlds.

3.2.3 Accessing Detailed Data on Demand

Pointing with the laser to select an individual object within the 3D wedge scene selects it and this pops up a panel filled with detailed data about the object (see Fig. 5). The panel is placed right next to the dominant stylus and moves along with it. In the case where there are multiple wedges, hence multiple worlds, only the objects in the wedge currently in focus are considered for the ray-intersection.

3.2.4 Navigating in Wedge Space(s)

The 2D analogue input (joystick or touchpad) on the dominant hand's controller is used to perform navigation. We assume the worlds are large, like terrains, somewhere we typically walk. Therefore, we support walking-style and flying-style navigation.

Walking mode is active when the controller is held in a normal, horizontal orientation. Pressing forward on the touchpad or joystick moves in the forward direction within the world that is currently in focus. Pressing back moves backward, and so on. Rotation about the vertical axis is initiated by pressing down on the center nib of the touchpad or another secondary button on the device. The angle of rotation is then set based upon the change in rotation of the controller. Clutching is used to make large rotations. When the




Figure 6: Left: walking mode is active when the controller is held in a horizontal orientation, and touch input is mapped to horizontal movements. Right: pointing the wand upward changes to flying mode, where the ray shoots out from the bottom of the wand, and touch inputs are mapped to vertical movements.

touchpad of the HTC Vive is used to implement this interface, the circular center nib on the touchpad provides a useful haptic aid to the user, making it clear without looking which part of the pad to use for rotation vs. translation.

Flying mode is activated when the wand is pointed up. In this mode, forward and backward touchpad inputs are mapped to flying up/down movement. To make it easy to maintain input focus on the wedge of interest during this operation, the direction of the laser pointer is adjusted so the ray shoots down from the wand instead of forward (see Fig. 6).

When the non-dominant hand's option button is depressed, placing all active worlds in focus, navigation happens simultaneously in all worlds.

3.2.5 Rendering Multiple Worlds

Multi-world rendering is implemented using the stencil buffer and a multi-pass technique. To setup the stencil buffer, a sphere is rendered with a custom fragment shader that computes the wedge shapes and outputs a unique ID number to the stencil buffer for each wedge. Then, in a second pass, each world is rendered using a stencil test that passes only for the corresponding wedge. Users can see the curve of the edges of the wedges along the sphere, but, other than these edges, the sphere used to setup the stencil buffer is not directly visible.

It is useful to think of the user as located within this imaginary sphere. Since, in VR, the virtual camera positions are defined through head tracking, we do not move the virtual cameras directly when users navigate through the worlds. Rather, we translate or rotate the worlds relative to this imaginary sphere that the user occupies.

Other rendering techniques would also be possible, however, the stencil-based approach is sufficient for our application, with frame rates above 30 frames per second reported for all of the examples pictured in this paper and the accompanying video.

4 FORESTRY CASE STUDY AND USER FEEDBACK

The description of the visual layout and interaction techniques in the prior section is intended to be applicable to use with any VR application that can: (1) render head-tracked, stereoscopic views of a data-driven, life-size virtual world, (2) use clipping planes to draw that world in a wedge, and (3) use ray-casting to select specific objects in the scene for displaying details on demand. However, visualization designers will typically want to adjust details of how to represent data within each wedge (e.g., realistic rendering vs.




Figure 7: Tree models used for eight of the tree species. In the data driven forests, the models are non-uniformly scaled to match the heights and trunk diameters recorded in the forest inventory data.

abstract) and within each WIM (e.g., using glyphs or simply a scaled version of the full-size scene) based on the particular application. In this section, we describe our own application-specific design decisions, results, and user feedback, which are useful as a first evaluation of the technique and also, more broadly, as a case study of implementing a complete data VR tool for visualizing an important actively researched dataset in collaboration with domain scientists.

We applied Worlds-in-Wedges to the problem of comparing plots from the US Forest Service’s Forest Inventory and Analysis dataset, which contains more than 20 million trees (20,595,807) across 629,644 forested plots systematically located across the U.S. [19]. Data have been collected via in-the-field measurements on a periodic basis since 1928. Our virtual reconstructions use the tree-level data, which include detailed measurements on diameter, height, species, above and below ground carbon content, damage, growth, and more. The above-ground and below-ground carbon content data are particularly interesting because these factor into estimates that impact the nation’s climate policy.

The data is collected following a regular sampling. Each plot consists of four circular subplots of radius 24 feet. The first subplot is located at the center of the plot, and the other three surrounding it are separated by 120 degrees each. Detailed measurements are made for every tree within the subplots, and additional data on land use and slopes are also collected.

4.1 Implementing the Three Layers

The visualizations created using this data are pictured throughout the paper.

4.1.1 Wedges

Each wedge is a life-size data-driven reconstruction of each plot. Each tree within the plot is rendering using a 3D model selected from a set of models (Fig. 7) based on the tree species data variable.

The model is then non-uniformly scaled based on the height and trunk diameter and positioned within the plot using a combination of three data variables: azimuth, distance, subplot id.

The terrain for each plot is procedurally generated based on the data variables available within each subplot (aspect ratio, slope, subplot id). This calculation is not straightforward because the data were not intended to be used to reconstruct an actual terrain map. Thus, our algorithm reproduces a plausible terrain given the ground slope data, which is only available at the subplot level. For each plot, the data is sufficient for defining just 8 height constraint points, one at the center of each subplot and one at the highest point along the outer radius of the subplot. The center point of the center subplot is the first constraint, we treat this as height 0. Then, using the aspect ratio and slope variables, we compute the height and location for the highest point within the center subplot. This is the second constraint. The algorithm treats this center subplot as special, assuming that its slope, which defines a plane in 3-space, persists to some degree throughout the entire plot. So, the heights for the center points of the three surrounding subplots are determined based on their locations within this plane. This adds three more constrain points. Finally, the highest point within the outer subplots is determined by offsetting from the height at its center using the local slope and aspect ratio. This adds the final 3 constraint points. Using the eight height constraint points as input, the height of each vertex on the terrain grid, $h(i)$, is computed using an inverse distance weighted average,

$$h(i) = \sum_{j=1}^8 \frac{1}{d(i,j)^p}, \quad (2)$$

where $d(i,j)$ is the 2D distance from terrain vertex i to constraint point j and $p = 4$ (an empirically determined constant).

Since the plots do not extend forever, we also include some billboard trees in the style of architectural renderings around the outer region of the plot to give the impression that there is uncertainty about the data beyond these bounds but that it’s likely the trees do not end as abruptly as shown by the 3D models.

4.1.2 WIMs

The most straightforward approach to creating WIMs that are appropriate for use in the worlds-in-miniatures level is to render scaled down versions of the 3D scenes. In a typical graphics engine, this can be accomplished by rendering each scene twice, once at normal scale for the wedge, and a second time with a simple scaling transformation matrix applied for the WIM. However, we construct a more sophisticated data visualization to display on each WIM.

Our approach is to construct an abstract glyph-based visualization for each plot (Fig. 8). Each tree is drawn as a data driven glyph rendered as a billboard. The glyphs are color coded by tree species group, and, as illustrated in Fig. 9, each glyph encodes three important variables: (1) tree height and (2) above ground carbon content, and (3) below ground carbon content. In this way, the WIM visualization adds detail that is not visible in the realistic 3D representation, making it useful not only for wayfinding and spatial context but also as a complementary multivariate data visualization. Within each WIM the billboards are rendered in a non-overlapping way using a mass-spring system to increase legibility.

4.1.3 Context

For this application, the worlds-in-context layer is a US map, and the dots on the map represent the plots in the data. Our current implementation loads data-driven worlds for a subset of 200 of the available plots. The locations of all of the plots are displayed with white dots, and the 200 plots that can be activated as virtual worlds are indicated on the map as green spheres.




Figure 8: Tree glyphs rendered as non-overlapping billboards on a WIM.




Figure 9: Detail of the glyph design. Tree height is mapped to the height of the center line of the glyph. Above ground carbon is mapped to the number of branches on the trunk. Below ground carbon is mapped to the number of underground roots.

4.2 User Feedback from Forest Scientists

Our interdisciplinary team has iteratively co-developed several versions of the software over the course of multiple years. During this time, the visualization researchers have collected and responded to feedback via design review sessions every few months. To provide a first evaluation of the technique, our team treated the most recent collaborative session as a summative design review. The methodology involved two forest scientists working with Worlds-in-Wedges over the course of approximately two hours and using a talk-out-loud protocol while completing an open-ended data exploration task. One forest scientist at a time used the HTC Vive display and controllers while the other watched the interaction on a large TV display setup to mirror the view rendered in the VR headset. The scientists switched roles part way through the session. The visualization researchers began the session with a tour through the features available in the current version of Worlds-in-Wedges and then took detailed notes and prompted discussion to understand points of success and confusion.

The scientists confirmed expected trends and made new discoveries. For example, looking at a forest at a bird-eye view, one scientist pointed out a region where there is a clear boundary. On one side, the trees are much shorter than on the other. This suggests the two areas have different land uses – the side with shorter trees is part of a plantation whereas the other side is not. Specific land usage information is included in the dataset and is something to incorporate into the visualization in the future, but it was interesting to note that, in this instance, users did not need that variable to be explicitly displayed. Given that they were able to immerse themselves in the data, they could see the pattern simply through being there. This example led to a discussion of the potential to use the tool for story telling. For example, telling the story to the general public of where tree plantations are located and how they are managed via a methodology that makes it possible for people to actually see what this means.

As the other scientist was flying around in another forest, he found an example of a few trees that are much taller than the others nearby (see Fig. 10). At first, he thought this looked odd, perhaps a mistake in the data. Then, pointing the ray at the trees to bring up details on demand, he determined that the tall trees are Eastern white and red pines and the short ones are spruces and balsam firs. Thus, the data and visualization actually makes sense given the species and how these would grow within a forest in the northeastern US. The scientists remarked that this is the type of information that is present in the data tables they currently use for analysis, but that the intuitive meaning is lost in those tables. In contrast, with the immersive visualization, the meaning of the data (e.g., the distribution of the tree species, the degree to which some trees differ from the surrounding forest) is immediately clear.

Another visual comparison involved multiple pine forests across the US. During this analysis, the scientists first activated three plots, one in the northeast, one southwest, and one southeast. They observed the three forests in equal-sized wedges to see the overall pattern. Then they raised one WIM to see that plot in more detail. Then, they repeated the same for the other forests, moving back and forth between focusing on one forest or another. In general, when comparing data with juxtaposed views, it is often the case that overall patterns are easily seen but there is a need to zoom in to examine details [13]; this frequent wedge-resizing behavior seems to be one way to support this in VR.

Several suggestions for refinements were also captured. The scientists identified during data analysis that flying to the same height above the ground in all worlds would be a useful way to judge tree heights. This is supported in the current interface. The simultaneous navigation mode supports coordinated movement, but it is coordinated relative movement, not absolute, so if you do not start on the ground in each world, then the heights will not match as you move all views together. Thus, it would also be useful to include a lock feature in order to move the user to the same height across all of the wedges. There was also a suggestion to adopt different gestures for rotating the worlds via WIMs. Since the WIMs look like turn-tables, there is a natural affordance to grab and rotate them via a direct manipulation. It may be that a gesture like a tabletop interface [12] would be a good complement to the current interaction. One scientist reported that the WIMs and the Worlds-in-Context views sometimes got in the way of focusing on the immersive views, suggesting a feature dismiss them on demand or automatically fade them out when the controllers are not being moved.

5 DISCUSSION AND FUTURE WORK

The research raises several interesting implications and future directions.

5.1 New Basic Research in Multi-World Presence

Our current forestry application applies different background-level details (sky maps, ground textures) to the non-data-driven elements in each forest to reinforce the sense that each plot is different. We have done some early experiments with extending this idea to a multi-sensory experience, using sound as an additional cue for distinguishing between the different worlds. In this mode, each wedge is assigned its own “soundscape” (e.g., birds, crickets, wind, water), and the gain applied to each sound is set in proportion to the angle of the corresponding wedge, so as the user adjusts the relative sizes of the wedges, zooming in to focus on one world of interest, the sounds of that world become more prominent. The aim is for the user to feel “more present” in the larger wedge(s).

Although this soundscape feature was added late to the visualization, almost as an after thought, it has sparked some of the most exciting ideas for future research. The topics of embodiment, presence, and useful perceptual illusion (e.g. [3, 22, 26]) show up again and again in the literature and session themes at the major




Figure 10: A plot in the state of Maine (Northeastern U.S.) is being observed from a bird-eye view. The tall trees are Eastern whites and red pines, surrounded by much shorter spruces and balsam firs.

conferences. However, we are not aware of any other VR research that studies these aspects in a “multi-world presence” setting. We believe this concept could open up a new line of basic research in presence, embodiment, and perceptual illusion, and that the potential to use field-of-view, audio, and other cues to adjust levels of presence in multiple simultaneously visible virtual worlds deserves study at a more basic research level. It would be interesting to learn what factors best promote multi-world presence. On the other hand, there is also a possibility that immersion in multiple worlds could break users sense of presence within any one world. Thus, although we believe there is exciting potential, more studies are required to understand this effect.

5.2 Additional Data and Representations

A useful but more straightforward extension of the current forest visualizations is to include additional variables in the data-driven virtual forests. For example, the forest inventory includes realistic contextual information on adjacent water bodies, marshlands, and land use. Incorporating these along with additional procedurally generated elements should increase “perceptual realism” [8], adding to the goal of experiential analysis discussed earlier. In addition, recent work on statistical forest modeling [30] should make it possible to extend the forest reconstructions beyond the limits of the sampled subplots to create continuous virtual forests. Perceptually, it would be easier to read the height of trees on a bar chart as compared to being in the forest, but we have intentionally made a trade-off in that regard, privileging immersion and experience over data legibility. In the future, we envision transitioning between multiple visual representations, some realistic and some abstract.

We are interested in applying the Worlds-in-Wedges technique to other datasets. The diagrams in Fig 2 already demonstrate how small adjustments to the worlds-in-context layer could make the technique applicable to worlds that are linked by time (e.g., comparing structural changes of a historical site at different points in time) or datasets linked in some more abstract way (e.g., parameterized simulation ensembles). We are also interested in exploring the potential of wedges that can be repositioned to focus on looking up (e.g., comparing multiple construction phases of the Sistine Chapel ceiling).

5.3 Insight-Based and Low-Level Evaluations

This research has been motivated directly by an interdisciplinary collaboration. So, it naturally includes a significant data processing

and system building component as well as an initial evaluation via expert user feedback for interdisciplinary collaborators. However, the technique would benefit from additional evaluations in other contexts. On one end of the spectrum, a formal insight-based evaluation with domain experts [23] would be useful. On the other end of the spectrum, it would also be useful to test a synthetic abstract dataset with non-expert participants. In both cases, we are confident based on the current case study and evaluation that data comparison tasks performed in life-size VR data-driven worlds would be faster and more accurate when performed using Worlds-in-Wedges as compared to the status quo approach, immersion in a single life-size virtual world at a time. However, it would be interesting to characterize the degree of the benefit both for scientific analyses and for other uses, such as storytelling.

6 CONCLUSION

Worlds-in-Wedges is a new VR-based visualization and user interaction technique for making simultaneous visual comparisons of multiple virtual worlds at a life-size scale. The approach extends Portal and WIM techniques to create a unifying metaphor for comparative immersive visualization. Within the application to comparing plots from the US Forest Service’s Forest Inventory and Analysis dataset, the technique was successfully used to confirm known and discover new trends in the data. Scientists also suggest the technique shows promise as a public-facing storytelling tool. This application serves as a case study both for evaluating the current technique and, more generally, for how VR can be applied to problems in scientific visualization via interdisciplinary collaborations. Our discussion of lessons learned raises several important directions for future VR research, including study of multisensory cues for evoking different levels of presence in environments that simultaneously display multiple life-size virtual worlds.

ACKNOWLEDGMENTS

The authors thank Ninad Sinha for useful discussion and assistance with implementation. This research was supported in part by the USDA Forest Service (14-CR-11242305-093, 15-JV-11242305-110, 18-JV-11242305-032).

REFERENCES

- [1] D. Acevedo, E. Vote, D. H. Laidlaw, and M. S. Joukowsky. Archaeological data visualization in vr: Analysis of lamp finds at the great

- temple of petra, a case study. In *Proceedings of the conference on Visualization '01*, pp. 493–496. IEEE Computer Society, 2001.
- [2] A. C. Addison. Emerging trends in virtual heritage. *IEEE multimedia*, 7(2):22–25, 2000.
- [3] M. Azmandian, M. Hancock, H. Benko, E. Ofek, and A. D. Wilson. Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences. In *Proceedings of the 2016 chi conference on human factors in computing systems*, pp. 1968–1979. ACM, 2016.
- [4] G. Bruder, F. Steinicke, and K. H. Hinrichs. Arch-explore: A natural user interface for immersive architectural walkthroughs. In *3D User Interfaces, 2009. 3DUI 2009. IEEE Symposium on*, pp. 75–82. IEEE, 2009.
- [5] D. Coffey, F. Korsakov, M. Ewert, H. Hagh-Shenas, L. Thorson, A. Ellingson, D. Nuckley, and D. F. Keefe. Visualizing motion data in virtual reality: Understanding the roles of animation, interaction, and static presentation. In *Computer Graphics Forum*, vol. 31, pp. 1215–1224. Wiley Online Library, 2012.
- [6] D. Coffey, N. Malbraaten, T. B. Le, I. Borazjani, F. Sotiropoulos, A. G. Erdman, and D. F. Keefe. Interactive slice wim: Navigating and interrogating volume data sets using a multisurface, multitouch vr interface. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1614–1626, 2012.
- [7] C. Demiralp, C. D. Jackson, D. B. Karelitz, S. Zhang, and D. H. Laidlaw. Cave and fishtank virtual-reality displays: A qualitative and quantitative comparison. *IEEE transactions on visualization and computer graphics*, 12(3):323–330, 2006.
- [8] K. Devlin, A. Chalmers, D. Brown, et al. Predictive lighting and perception in archaeological representations. In *Conference proceedings of UNESCO World Heritage in the Digital Age: 30th Anniversary Digital Congress (October, 2002)*, [URL: <http://doc.gold.ac.uk/mas01dl/CIS224b/unesco-paper.pdf>]. *Acceso*, vol. 14, p. 2014, 2002.
- [9] A. Forsberg, G. Haley, A. Bragdon, J. Levy, C. I. Fassett, D. Shean, J. Head, S. Milkovich, M. A. Duchaineau, et al. Adviser: immersive field work for planetary geoscientists. *IEEE computer graphics and applications*, 26(4):46–54, 2006.
- [10] M. Gleicher. Considerations for visualizing comparison. *IEEE Transactions on Visualization and Computer Graphics*, 2017.
- [11] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, 2011.
- [12] M. S. Hancock, F. D. Vernier, D. Wigdor, S. Carpendale, and C. Shen. Rotation and translation mechanisms for tabletop interaction. In *Horizontal Interactive Human-Computer Systems, 2006. TableTop 2006. First IEEE International Workshop on*, pp. 8–pp. IEEE, 2006.
- [13] K. Kim, J. V. Carlis, and D. F. Keefe. Comparison techniques utilized in spatial 3d and 4d data visualizations: A survey and future directions. *Computers & Graphics*, 67:138–147, 2017.
- [14] K. Kim, B. Jackson, I. Karamouzas, M. Adeagbo, S. J. Guy, R. Graff, and D. F. Keefe. Bema: A multimodal interface for expert experiential analysis of political assemblies at the pnyx in ancient greece. In *Proceedings of the IEEE Symposium on 3D User Interfaces*, pp. 19–26, 2015.
- [15] A. Kunert, A. Kulik, S. Beck, and B. Froehlich. Photoportals: shared references in space and time. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pp. 1388–1399. ACM, 2014.
- [16] J. J. LaViola Jr, D. A. Feliz, D. F. Keefe, and R. C. Zeleznik. Hands-free multi-scale navigation in virtual environments. In *Proceedings of the 2001 symposium on Interactive 3D graphics*, pp. 9–15. ACM, 2001.
- [17] J. Liang and M. Green. Jdcad: A highly interactive 3d modeling system. *Computers & graphics*, 18(4):499–506, 1994.
- [18] T. Munzner. *Visualization analysis and design*. AK Peters/CRC Press, 2014.
- [19] B. M. O’Connell, B. L. Conkling, A. M. Wilson, E. A. Burrill, J. A. Turner, S. A. Pugh, G. Christiansen, T. Ridley, and J. Menlove. *The Forest Inventory and Analysis Database: Database description and user guide version 7.0 for Phase 2*. U.S. Department of Agriculture, Forest Service. 830 p., 2017.
- [20] R. Olanda, M. Pérez, P. Morillo, M. Fernández, and S. Casas. Entertainment virtual reality system for simulation of spaceflights over the surface of the planet mars. In *Proceedings of the ACM symposium on Virtual reality software and technology*, pp. 123–132. ACM, 2006.
- [21] R. Pausch, T. Burnette, D. Brockway, and M. E. Weiblen. Navigation and locomotion in virtual worlds via flight into hand-held miniatures. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pp. 399–400. ACM, 1995.
- [22] S. Razaque, Z. Kohn, and M. C. Whitton. Redirected walking. In *Proceedings of EUROGRAPHICS*, vol. 9, pp. 105–106. Citeseer, 2001.
- [23] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE transactions on visualization and computer graphics*, 11(4):443–456, 2005.
- [24] F. Steinicke, G. Bruder, K. Hinrichs, M. Lappe, B. Ries, and V. Interrante. Transitional environments enhance distance perception in immersive virtual reality systems. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization*, pp. 19–26. ACM, 2009.
- [25] R. Stoakley, M. J. Conway, and R. Pausch. Virtual reality on a wim: interactive worlds in miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 265–272. ACM Press/Addison-Wesley Publishing Co., 1995.
- [26] E. A. Suma, S. Clark, D. Krum, S. Finkelstein, M. Bolas, and Z. Warte. Leveraging change blindness for redirection in virtual environments. In *Virtual Reality Conference (VR), 2011 IEEE*, pp. 159–166. IEEE, 2011.
- [27] R. Trueba, C. Andujar, and F. Argelaguet. Complexity and occlusion management for the world-in-miniature metaphor. In *International Symposium on Smart Graphics*, pp. 155–166. Springer, 2009.
- [28] J. Viega, M. J. Conway, G. Williams, and R. Pausch. 3d magic lenses. In *Proceedings of the 9th annual ACM symposium on User interface software and technology*, pp. 51–58. ACM, 1996.
- [29] D. Wigdor, C. Shen, C. Forlines, and R. Balakrishnan. Table-centric interactive spaces for real-time collaboration. In *Proceedings of the working conference on Advanced visual interfaces*, pp. 103–107. ACM, 2006.
- [30] B. T. Wilson, A. J. Lister, and R. I. Riemann. A nearest-neighbor imputation approach to mapping tree species over large areas using forest inventory plots and moderate resolution raster data. *Forest Ecology and Management*, 271:182–198, 2012.
- [31] C. A. Wingrave, Y. Haciahmetoglu, and D. A. Bowman. Overcoming world in miniature limitations by a scaled and scrolling wim. In *3D User Interfaces, 2006. 3DUI 2006. IEEE Symposium on*, pp. 11–16. IEEE, 2006.