Motivations, design, and preliminary testing for a 360° vision simulator

Motivations, Design, and Preliminary Testing for a 360° Vision Simulator

Matt Cook
Digital Scholarship Program Manager
Harvard Library
1 Harvard Yard, Cambridge, MA 02138
matt_cook@harvard.edu
(617) 496-2534

John Grime, PhD
Emerging Technologies Developer
University of Oklahoma Libraries
401 West Brooks Street. Norman, OK. 73019.
jgrime@ou.edu
(405) 325-5641

ABSTRACT

Contemporary virtual reality systems enable academics to more efficiently explore and analyze complex three-dimensional (3D) content, but their utility is limited by visual short-term memory. Janus, a geometry agnostic shader script, circumvents this cognitive limitation by automatically rendering complex object meshes to fit entirely within the field-of-view of consumer head mounted displays. The resulting 360° vision experience represents an advantage over existing scientific data visualization tools, which have sought to replicate real-world viewing experiences but have inadvertently replicated associated limitations as well. By presenting data in such a way so as to effectively circumvent cognitive loads associated with body (or object) movement, academics can use the Janus shader to more readily engage in the exploratory analysis of complex 3D data sets, thereby facilitating scientific insight. This paper explores the motivations and design of the Janus shader, and describes preliminary results from user testing conducted under controlled conditions. For the 24 study participants (N=24), statistically significant time-to-completion decreases were observed for spatial analysis tasks taking place in intervention (Janus-enabled) VR scenes of low to moderate complexity.

Keywords: Virtual Reality, Cognitive Load, Scientific Data Visualization, Academic Technologies

Motivations, Design, and Preliminary Testing for a 360° Vision Simulator

Abstract

Contemporary virtual reality systems enable academics to more efficiently explore and analyze complex three-dimensional (3D) content, but their utility is limited by visual short-term memory. Janus, a geometry-agnostic shader script, circumvents this cognitive limitation by automatically rendering complex object meshes to fit entirely within the field-of-view of consumer head mounted displays. The resulting 360° vision experience represents an advantage over existing scientific data visualization tools, which have sought to replicate real-world viewing experiences but have inadvertently replicated associated limitations as well. By presenting data in such a way so as to effectively circumvent cognitive loads associated with body (or object) movement, academics can use the Janus shader to more readily engage in the exploratory analysis of complex 3D data sets, thereby facilitating scientific insight. This paper explores the motivations and design of the Janus shader, and describes preliminary results from user testing conducted under controlled conditions. For the 24 study participants (N=24), statistically significant time-to-completion decreases were observed for spatial analysis tasks taking place in intervention (Janus-enabled) VR scenes of low to moderate complexity.

Keywords: Virtual Reality, Cognitive Load, Scientific Data Visualization, Academic Technologies

1 Introduction

What is the most efficient and effective way to analyze complex surface mesh geometries, especially with low-cost imaging and 3D scanning techniques providing the means to generate detailed models with myriad research applications? Scholarship suggests that a fully optimized virtual reality (VR) viewing system effectively replicates the experience of exploring physical objects in the real world. This is an especially important design consideration for academic applications, as activities that require the analysis of complex 3D data sets – including modeled specimens, numerical data sets, scanned artifacts, etc. – are shown to be improved when models are made available for inspection in VR.

However, the relatively lifelike combination of stereoscopic depth cues and positional tracking (defining characteristics of today's high-end VR systems) represent a viewing paradigm that is unnecessarily beholden to the limitations of human perception. In the course of both real-world and digital exploration of complex 3D objects, for example, users are initially bound to a single viewing angle, and take on cognitive load in moving virtual objects (or themselves) to gain novel perspectives and/or circumvent occlusions. Yet exploring complex 3D objects need not be as cognitively taxing in a digital research environment, since both digital content and visual input can be arbitrarily transformed via customized software.

The Janus geometry shader script, detailed below, represents an advantageous development in the of discipline-agnostic immersive virtual reality viewing systems.[1] By deploying data in a way that effectively minimizes the need for user or model movement, academics can more readily engage with complex 3D objects, especially during the exploratory phase of research, when a high-level mapping of the complete data set might shape the subsequent trajectory of a nascent research program. This human-first system design approach represents a first step towards a "super-fidelity" VR system, which – like existing technologies - preserves real-world stereoscopy and positional tracking but extends human capabilities to 360° vision.

2 Literature Review

Virtual Reality has been employed to support scholarship across a range of academic fields, including: Archaeology, Neuroscience, Engineering, Geology, and Medicine (Acevedo et al. 2001; Van Dam, Laidlaw, and Simpson 2002; Seth, Vance, and Oliver 2011; Donalek et al. 2011; Anderson et al. 2016). Testing these in these instances typically requires special access to costly visualization equipment, such as CAVE-type projection rooms (Cruz-Neira, Sandin and DeFanti 1993). Since the release of the first Oculus Developer Kit (2013), however, VR has hit the academic mainstream, and the benefits documented under controlled conditions have made their way from relatively well-resourced computer science and electrical engineering departments and into the classroom (Milovanovic et al 2017; Thompson 2018; Cook & Lischer-Katz 2019). The relatively low cost of head mounted display hardware, coupled with the ease-of-access to game development software (e.g. Unity3D) has translated to a proliferation of educational VR experiences that all University stakeholders can engage with and benefit from.

2.1 The Benefits of Virtual Reality for Scholarship

Virtual Reality is useful tool for scholarship because it presents the user with digital content in a way that resembles the familiar experience of viewing physical objects in the real world. Such fidelity is enabled by a combination of low-level system affordances including stereoscopy (via appropriately offset displays) and six-degrees-of-freedom positional tracking, the latter a feature that respects the way human beings engage with their surroundings bodily in daily life (i.e. along three axes of locomotion). Given this combination of depth cues and body-centered interfacing capabilities, VR-aided analyses are more efficient than engagement with the same complex 3D objects using traditional (mouse, keyboard, monitor) computer interfaces (Ragan et al. 2013; Laha, Bowman, and Socha 2014).

More specifically, the benefits of stereoscopically induced *depth-of-field* have been studied extensively, and has proven to be an especially valuable feature of contemporary VR systems. Indeed, Ware and Mitchel (2005) describe a full order of magnitude increase in performance for tasks associated with analyzing complex 3D graphs under stereoscopic viewing conditions, while Kersten-Oertel, Chen, and Collins (2014) found that stereoscopic viewing of vascular structures was of practical utility for expert brain surgeons seeking to determine the placement of blood vessels, especially when presented in conjunction with chromatic depth cues.

[1] Janus, the Roman god of transitions, is often depicted as simultaneously seeing forwards and backwards.

Importantly, the embodied human observer is also preserved in virtual reality viewing environments, and, because VR thereby permits, "…participants to use their body as part of the very process of data exploration," novel perspectives on complex 3D objects can be gained efficiently without significant prior training or mental exertion (Prabhat et al. 2008). Positional tracking - including head, upper body, and hand/arm tracking supported by today's higher-end (but still consumer oriented) VR systems - lowers the barrier to entry for learning a new software system by preserving certain intuitive movements and abilities (Cook 2018). In VR, the user can bend down, lean in, crane their neck, perambulate, etc., to take in various aspects of a complex three-dimensional data set; data that may only exist at microscopic or cosmic scales, well beyond the reach of physical human interaction.

Virtual reality systems that combine stereoscopic viewing with positional tracking are among those classified as "high fidelity" in the literature insofar as the combination of these features more accurately reproduce the visual conditions associated with viewing human scale objects in the physical environment. When compared with a 2D illustration of a given object (on a textbook page, perhaps), or even a 3D rendering of the same model viewed on a traditional computer screen, VR exists at the pinnacle of the representational spectrum. This is borne out in the literature by participants demonstrating comparatively superior performance on tasks related to the *search, identification, description,* and *comparison* of complex 3D objects in high fidelity, VR testing environments (Laha, Bowman, and Socha 2014, 7).

2.2 The Limitations of Existing VR Systems

In a controlled comparison of the individual features that determine abovementioned system fidelity, Ragan et al. (2013) summarize an experimental design that makes use of low-level, cross-disciplinary spatial tasks to test the impact of VR platform characteristics: "This task involved three stages: visual search for a potential gap location, navigation to view the potential gap location from an advantageous viewpoint, and judgment of whether a gap was present or not" (pg 5).

Here, one of three basic components of the task – "navigation to view the potential gap location from an advantageous viewpoint" – represents a well-known and tacitly accepted drain on the performance of the VR user: Cognitive load (Anderson et al. 2016). To circumvent the visual occlusion obscuring potentially informative components of a complex 3D object, the VR user must take on cognitive load, in the form of visual short-term memory, while moving between viewpoints, or, alternately, they must manipulate the digital objects' position in virtual space (Lages & Bowman 2018). In either case, the user must recall previously encountered (but subsequently obscured) visual features in order to draw meaningful inferences concerning the overall nature or "global shape" of that object (Silvestri et al. 2010). As in the real world, the VR user must move the virtual model or his or her position to gain a useful perspective. This limitation, inherited real-world viewing conditions, is unnecessary in a digital viewing environment.

Silvestri et al. (2010) explored the impact of this type of cognitive load as it applies to the design disciplines (e.g. Architecture), and determined that the ability to reproduce (draw) complex 3D shapes was improved as the "mental transformations" necessary to compose a comprehensive view of the object were offloaded onto the display capabilities of system itself (pg. 16). Navigation between advantageous viewpoints represents a costly mental transformation insofar as the dimensional richness of the visual scene is offset by the need to manipulate the object (or oneself) to mentally assemble a comprehensive view of all occluded perspectives (Thomas & Seiffert 2010). Importantly, human cognitive systems are quite limited in their ability to keep track of the characteristics of complex 3D objects, which directly impacts the usefulness of contemporary virtual reality systems.

With regards to visual short-term memory, the sort of cognitive load most associated with immersive visualization systems discussed here, Alvarez and Cavanagh (2004) concluded that "[T]he total amount of visual information (the product of the number of objects and the visual information per object) cannot exceed the maximum visual information limit," of 4-5 objects. (pg. 6). Importantly, this maximum varies in relation to target stimuli complexity, and subjects exhibit "…an inverse relation between the information load per object and the number of objects that can be stored". Fortunately, visual information (related to size, color, and presence or absence of a gap, at least), can be conjoined, and Luck & Vogel's (1997) study determined that, within those 4-5 target objects, subjects could reliably track as many as 16 sub-features or "chunks" of visual information. Yet the target stimuli of multiple research teams feature relatively complex objects, comprising various sub-features (abstract architecture, cave structures, vasculature, cortical and biomechanical structures, etc.), and most require some locomotive (or controller) movements to circumvent occlusion and complete experimental tasks associated with determining the connectedness, describing the shape, and/or comparing complex 3D objects (Ragan et al 2013; Evagorou & Heinis 2018; Kersten-Oertel, Chen, and Collins 2014; Laha, Bowman, and Socha 2014). The combination of real-world stimuli and locomotion is cognitively taxing.

In all of these cases, both depth cues and positional tracking aid task completion in virtual reality, but the user runs up against hard cognitive limits that limit performance. In the act of virtual navigation or object manipulation, the human user takes on cognitive load and is required to "keep in mind" various types of visual information during the movement portion of the experimental task. Since the real-world viewing experience is assumed to be the benchmark by which complex 3D objects can be analyzed, even the highest fidelity VR systems available to scholars today are, by definition, constrained by their expectation that users circumvent visual occlusion via object or self movement. Fortunately, we can leverage the digital nature of our immersive visualization systems to effectively circumvent the limitations on human cognitive processing of visual information.

3 System Design

Technically, Janus is a geometry shader script designed to integrate easily with any VR research project that uses Unity as a development platform. Unity is a platform for both 2D and 3D software development that is used extensively in commercial software but also provides free-of-charge licenses for educational uses. The relative simplicity and power of Unity, allied to its extensive documentation and support for cross-platform VR development, make Unity an appealing platform for developing scientific visualization tools. The Janus shader is implemented with the HLSL programming language for GPU calculations (also referred to as "Cg" due to the similarities between the two), and the Janus shader therefore requires GPU hardware that supports DirectX 9 and later (which includes all recent GPU hardware). The use of "standard" programming languages such as HLSL to implement the Janus shader should allow for straightforward porting to other 3D platforms, and portability of the Janus shader to other Unity projects is accomplished via the standard Unity workflow of associating the shader with one or more "materials" that are applied to object(s) in a Unity scene. No custom preprocessing of geometry data is required as Janus warps and renders the position of each of a given mesh's constituent vertices on the GPU in real-time, on a frame-by-frame basis, based on the VR user's head position (See: Figure 1). The resulting experience is akin to viewing the real world through a wide-angle fish eye or catadiopteric lens as demonstrated by a handful of research teams (Ardouin et al 2012; Orlosky et al. 2014; Fan et al. 2014). But, unlike these optical implementations, the virtual player can be deployed at an arbitrary location, including at the *center* of a given data set, theoretically allowing them to view all sides of a complex 3D object simultaneously. The end user experiences layers of overlapping mesh components (whatever their chosen data set) and can thereby make use of myriad stereoscopic depth cues and positional tracking to take in all sides simultaneously.

Beyond providing the means to decrease cognitive load associated with locomotive or object movement in virtual reality, Janus compares favorably to existing VR-based scientific visualization systems via three uniquely beneficial features:

1) Modularity: Janus was designed from the ground up to move between laboratories, classrooms, and other academic environments. By crafting Janus as a single geometry shader it can be easily applied to any surface mesh imported into Unity3D, researchers and instructors can make use of the simulated, stereoscopic 360° viewing condition without abandoning their preferred development environment, models, and associated customizations. Furthermore, this approach enables trivial "selective" application of the Janus effect, allowing GUI elements (and related components of a scene) to remain unmodified. The use of a standard shader language (HLSL) provides straightforward portability to other 3D platforms.

2) Mesh Pre-processing: Janus does not require preprocessing of mesh data, as it effectively warps and compresses every associated 3D mesh on the GPU for every rendered frame. Our preliminary testing indicates that the shader is VR-ready, with performance data indicating frame rates at or above the 75-100 FPM detailed in HMD manufacturer guidelines.

3) Dynamic Compression Characteristics: The Janus system was deliberately designed not to interfere with detail rich center-vision, and to respect the limitations of peripheral vision, which lacks stereopsis in the far periphery and is also relatively poor for discerning and distinguishing some visual phenomena (e.g. color). Moreover, the target field-of-view can be manually updated, in the Unity editor, or by using controls inside the VR experience in real time, to accommodate any existing or future headset specifications.

Further technical details regarding the Janus compression can be found in the *Supporting Information* that accompanies this manuscript.

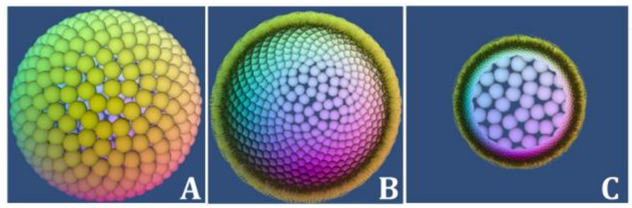

Fig. 1 Simple Janus visualization. (A) Example scene (a spherical arrangement of spheres) as viewed from outside; (B) view "inside" the scene, with camera located at the center of the spheres. Janus compression of scene data "behind" the camera is visible in the dense peripheral ring (unperturbed region of vision set to 20°); (C) as (B), but with "native" field-of-view reduced to 30°.

3.1 Prior Implementations

The dynamic visual field compression characteristics and potential cognitive impact of Janus is informed by the literature concerning the strengths and limitations of peripheral vision and previous, experimental implementations with ultra wide-angle viewing systems.

For example, Orlosky et al. (2014) note the low acuity of peripheral vision in describing their optical ("Fisheye Vision") implementation. However, based on testing that measured the time it took study participants to detect peripherally located objects, the research team determined that, "[P]lacing of virtual objects at more central peripheral angles can improve reaction times", and, hinting at something like the Janus implementation, the authors proceed to suggest that, with system optimization, these detection benefits could, "[P]otentially be expanded past the human visual field" (pg. 7). Because the far periphery is monocular, alternative depth cues – like motion, color, and fog – will be desirable, if the goal is to further differentiate what might otherwise be a confusing mass of visual information (Strasburger, Rentschler, and Jüttner 2011). Human peripheral vision, while lacking the attention driven strength of central vision, is nonetheless adept at detecting movement, which will inform future iterations of the Janus system (See: "Discussion" section).

A multi-institution team based in Japan engineered another wide-angled optical viewing system (Fan et al. 2014). Their testing sought to determine whether users could make use of peripheral camera image data that was blended with a forward-facing image. Intriguingly, researchers concluded that "[H]umans could adapt to having and utilizing eyes on the back of our heads, and could get accustomed to using the back eyes to observe the back instead of turning 180 degrees to use the front eyes" (Fan et al. 2014, 7). Finally, testing of the "FlyViz" system - which combines a helmet-mounted catadioptric sensor with a head mounted display – indicated that users could be trained to both grasp and dodge objects outside of their natural field of view (Ardouin et al. 2012). There are therefore implemented precedents documenting the fact that humans can make sense of, and act upon, 360 degrees of visual information.

4 Preliminary Testing

Preliminary testing of the Janus system comprises pre-test and VR-based test phases, each employing instrumentation designed to determine the effectiveness of the system across various baseline spatial ability levels and to provide data for future improvement of system performance and user interfacing.

Institutional IRB approval was sought and approved. Purposive snowball sampling was employed to recruit a total of 24 test participants (N=24), all of who are affiliated with the University as a student, faculty, or staff member. Testing took place in a closed room, with a total of four (4) Oculus Rift-enabled VR workstations running the Janus software. Each station is separated from neighboring stations and participants by portable divider walls to prevent distraction during testing. The complete testing period takes, on average, between 30 minutes and 1 hour per study participant.

4.1 Pre-Test

Testing begins with a computer-based survey (not in VR), which gathers data concerning prior-experience with 3D visualization tools and technologies. These data are intended to help control for factors that have been positively correlated with mental rotation and spatial perception abilities (Veurink et al. 2009; Maan et al. 2012; Lages & Bowman 2018). Next, the Revised Purdue Spatial Visualization (PSVT:R) test is administered. The PSVT:R has been deployed across hundreds of studies, and has proven reliable as a means to gauge prior spatial ability, especially with regards to the mental rotation of 3D shapes. Moreover, the associated questions, "[C]onsist of spatial problems closely aligned with the tasks routinely required to succeed in the engineering and technology disciplines," which lends itself to the development of a VR system designed specifically for scholarly use (Yoon 2011, 50-51).

Upon completion of the spatial visualization pre-test, study administrators worked with participants to properly seat the virtual reality headset; to familiarize the user with the (Oculus Touch) hand controllers; and, finally, to launch the Janus software. A practice session is initiated, wherein participants are exposed to simplified versions of the testing scenarios (described in detail, below). After 5-10 minutes of pre-test practice, the administrator initializes testing proper.

The testing scenario involves the identification, counting, and input of a total number of idealized arrow–shaped 3D models deployed in 360° virtual reality scene as well as a standard VR viewing environment (i.e. the control condition). Following established experimental design, speed and accuracy of task completion were tracked by the system automatically via the in-app input interface (LaViola et al. 2017). The arrow 3D model, which included a roughly conical "head" and visually chaotic fletching, was chosen as the target object for a number of reasons related to the general applicability of the system for academic purposes.

4.1.1 Target Stimuli Justification

Multiple research teams have sought to balance the potentially confounding impact of prior-knowledge (mostly in the form of subject-specific expertise) with the scientific applicability of virtual reality systems. On the discipline-specific end of this spectrum, some researchers have settled on recognizable but highly simplified representations of scientifically meaningful 3D models for testing, while other teams have deliberately deployed abstract target stimuli, which – while visually unrelated to any discipline (and therefore unrecognizable to an arbitrary subject expert participant) - nonetheless share geometric characteristics and complexities that represent common, spatially relevant, task types. As Laha et al. (2014) conclude, the latter groups' use of abstract model types for testing is doubly useful insofar as the, "[C]ommunity of VR and visualization researchers need to identify and define abstract task categories cutting across various scientific domains". We chose to deploy abstract target stimuli, an arrow.

An arrow is a recognizable 3D object, familiar to all participants, but it is unrelated to any specific scientific discipline. Moreover, an idealized arrow shape can be perceived by the participant as a single, recognizable object, but also as a combination of three primitive sub-components: A cone (arrowhead); tube (arrow shaft); and a sort of explosive feathering (fletching). Critically, the complete arrow model is also elongated, with distinct, asymmetrical end points, which can be deployed so as to span the visual field, thereby necessitating either movement or a compressed field-of-view. To correctly recognize an arrow as an arrow, one will need to accurately perceive a unidirectional (and connected) head, shaft, and fletching, and then successfully combine those visual components via locomotion or object rotation in the case of occlusion.

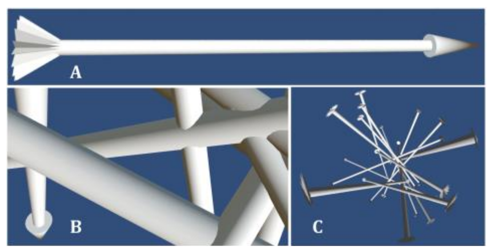

Fig. 2 Target Stimuli ("A", top); Standard field-of-view with typical control-group perspective ("B", bottom-left) and; Janus compression perspective ("C", bottom-right)

4.2 Test

Testing involves the counting of arrow models in space, as presented in simulated 360°, and, in the control, via 100° (Oculus Rift-specific) standard field-of-view (see: Figure 2, above). The speed and accuracy with which the participant completes this counting task represents the performance data that are compared to performance from engagement with repeated (but randomized) stimuli sets in a "traditional" virtual environment with a limited field of view. To minimize variation between control and Janus conditions, only sitting-and-turning-in-place physical movement was permitted, and no rotation or translational control over the target model placement was provided. The virtual scale of the scene and associated stimuli was such that any benefits that might have come from standing (or crouching) would be minimal. Both control and Janus place the user at the exact center of a large-scale virtual scene with varying numbers of complete and incomplete virtual arrow arrangements.

Importantly, there is no way to ascertain for certain either the completeness of a given target object or the total number of target stimuli in a scene from any arbitrary single perspective in the control condition. Rather, the control group has to physically turn their bodies around to take in the full scene and track instances of complete (non-fragmentary) target stimuli. Participants in the Janus condition can theoretically take in the complete scene via a single, global perspective. This is accomplished with the elongated, asymmetric arrow structure as well as "noise" in the form of arrow sub-components distributed throughout the scene on a systematic basis. As per Laha et al., generally applicable VR-based task categories include "Search", "Search, counting", "Spatial Judgement", and "Shape Description", "quantitative estimation", "shape description", and "pattern recognition" (Table 9). While the primary completion task described here requires that a user submit a total arrow count to move to the next scene, the entire process of searching, recognizing a (complete) pattern amidst extraneous data, and then combining (i.e. counting) instances of intact stimuli combines multiple task types.

Rather than verbal reporting, which requires a dedicated human test administrator and risks administrator influence as well as distraction from nearby test takers, counts were submitted with a (Oculus Touch) hand controller. In both the control and Janus conditions, each participant took part in a 5 trial "practice round" for orientation followed by 10 trials each in both the control and Janus viewing conditions ("Perspective 1", and "Perspective 2"). Arrows and arrow parts were procedurally generated, and ranged in number from 3 to 9 per trial.

4.3 Results

Testing was performed under two conditions: *control* (Janus view inactive) and *intervention* (Janus view active). Both phases present the user with identical 3D scenes (albeit in a randomized order).

We examine performance via per-scene changes in both time-to-completion and accuracy; the former is the difference in time required to process a scene under both conditions, the latter is the change in accuracy when counting complete arrows: accuracy = $|N_{inter} - N_{actual}| - |N_{contr} - N_{actual}|$. Here, N_{actual} is the number of complete arrows in the scene, N_{contr} is the user's estimate of N_{actual} under control conditions, and N_{interv} is the user estimate of N_{actual} under intervention conditions. Where users occasionally self-advance to the next scene by accident (i.e., without completing the task), such instances are noted, and those specific data points are removed from any analysis.

First, we consider all users as a single group (N=24), and apply a paired, two-tailed t-test to detect performance differences between the control and intervention phases (H_0 = no difference between control & intervention, $p<0.05$). No statistically significant difference in accuracy was recorded for any scene complexity studied; by contrast, statistically significant improvements were recorded in time-to-completion for scenes of low to moderate complexity (Figure 3).

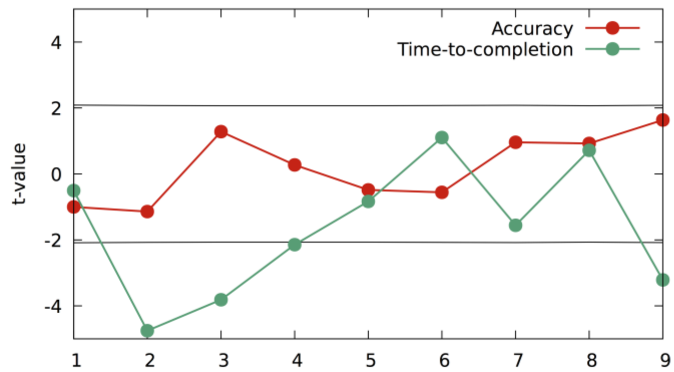

Fig. 3 User performance data as a single group (N=24). Two-tailed, paired t-test (control vs. intervention) was used to examine performance changes for each scene (listed in ascending complexity) changes for each scene. Black lines denote the "bounds" of statistical significance for the relevant degrees of freedom

The user data was then split into two groups via ranked PSVT:R scores to examine potential innate differences between "higher" (N=12) and "lower" (N=12) performers. A paired, two-tailed t-test was use to compare performance differences under control and intervention conditions for the "higher" and "lower" groups (H_0 = no difference between control & intervention, $p<0.05$). Once again, no statistically significant difference in accuracy between control and intervention phases was recorded for either group (Table 2). With respect to time-to-completion, the "lower" PSVT:R cohort displayed a statistically significant increase in performance only for two low-complexity scenes (Table 2); the "higher" PSVT:R group demonstrated a statistically significant performance increase across a wider range of scene complexities (Table 2).

Table 1 User performance for "lower" and "higher" groups, partitioned according to ranked PSVT:R scores (significant changes emboldened)

SceneID	Accuracy: t-value (df_t) "lower" PSVT:R performers	Time-to-completion: t-value (df_t) "lower" PSVT:R performers	Accuracy: t-value (df_t) "higher" PSVT:R performers	Time-to-completion: t-value (df_t) "higher" PSVT:R performers
1	0.00000 (9)	-1.17020 (9)	-1.49071 10	1.16465 10
2	-1.00000 (10)	**-5.48064 (10)**	-0.56061 11	**-2.33103 11**
3	0.00000 (11)	**-2.44386 (11)**	1.77281 11	**-2.84186 11**
4	0.00000 (11)	-1.52866 (11)	0.43179 11	**-2.49642 11**
5	-0.36405 (11)	-1.00685 (11)	-0.35156 11	-0.27396 11
6	-0.21926 (10)	1.27232 (10)	-0.52223 11	0.20338 11
7	1.00000 (10)	-0.46848 (10)	0.45408 10	-1.94146 10
8	1.44887 (11)	1.29458 (11)	-0.00000 11	-0.21638 11
9	1.44463 (9)	-1.48413 (9)	0.80440 11	**-2.94176 11**

In summary, the Janus intervention as not shown to impact accuracy on spatial analysis tasks, but it did allow users work faster, especially in VR scenes of low to moderate complexity. While this speed increase is pronounced in users that perform well on the PSVT:R test, it's also present in the "worse" performers on the PSVT:R, which might indicate the benefits of simulated 360° viewing represent a global effect regardless of existing ability for mental manipulations of complex geometries.

4.4 Limitations

The performance of "lower" and "higher" PSVT:R test groups in both control and intervention phases were further compared by one-way ANOVA, with statistically significant difference ($p<0.05$) only detected for a single scene. This result suggests that PSVT:R test scores may not be a general predictor for performance in our VR tests when used to partition users into roughly equal-sized "lower" and "higher" groups, or that insufficient data prevents the detection of differences.

Moreover, the *ordering* of the scenes may also have impacted performance data. For the "simplest" 3D scene, we do not observe statistically significant improvements in the time-to-completion for Janus (see Table 1). This may be due to the simplicity of this scene not benefiting from Janus, but it could also be because this is the *first* scene shown to the user after Janus is enabled; the initial shock of the Janus view might actually be detrimental to performance while the users adjust. If we exclude those outlier scenes, we have the behavior one might expect: large improvements in time-to-completion for less-complex scenes, with this advantage becoming smaller (and thus, less statistically significant) as the scene complexity increases. Improved experimental design and increase sample sizes in follow-up studies should serve to minimize these outlier discrepancies.

5 DISCUSSION

Again, the basic research question is: What is the most efficient and effective way to analyze complex 3D data sets, especially with low-cost imaging and 3D scanning techniques providing the means to generate detailed models with myriad research applications, including in the humanities (Pfarr-Harfst 2016; Alliez et al 2017)? After all, a scientific data set is fundamentally unlike a macroscopic real-world object, although our familiarity with the latter is a useful baseline when it comes to designing robust enough visualization software to accommodate high levels of visual complexity. When explored at human scale, in three dimensions, and presented with a depth of field, myriad "big" scientific data can be explored efficiently and new discoveries or research directions can be readily facilitated. Numerous examples exist for suitably complex data sets.

Open access meteorological, medical, chemical, and geophysical data sets abound, and each category of data represents viable use cases for the Janus system. In all cases, multi-gigabyte sized data sets are regularly analyzed for patterns, anomalies, outliers, groupings, sub-section connectivity, etc., while the same data are presented to undergraduate and graduate students in the classroom with the expectation that defining characteristics (e.g. pathologies) will be recognizable in a clinical or other practitioner-oriented professional environment. In all cases, the complexity of the 3D objects that emerge from these disciplines test the visual short-term memory limits of the user. After the preliminary testing is complete, and the system is refined, and improved, subject-specific 3D content will be targeted and deployed to support the early-stage, open-ended, and/or diagnostic research efforts of scientists in these disciplines.

Importantly, Janus isn't the catch-all solution for VR-based big data analysis. Rather, the results of this preliminary testing highlights potential time-to-completion benefits assuming search, pattern recognition, and counting activities are goals of the analyses. User performance with regard to other task categories identified in the literature (e.g. "spatial judgement", "shape description") will undoubtedly suffer, given the image warping necessary to compress 360-degrees of visual data into the forward field-of-view. (One might erroneously note the curvature of an anatomical structure, for example, while attempting to analyze a medical data set). So, Janus might best be viewed as a "first pass" tool, allowing the user efficient access to global composition of a moderately complex visual scene.

5.1 Future Directions

Regarding future system design directions, alternative depth cues have been successfully deployed in testing environments to increase performance of analysis on complex 3D data sets. Among these visual cues, color, motion, "fog" (or atmospheric occlusion), and exaggerated scaling have been identified as particularly useful, especially when the limitations of peripheral vision are taken into account and these cues are deployed in tandem with stereoscopic, body-tracking displays (Strasburger, Rentschler, and Jüttner, 2011; Kersten-Oertel, Chen, and Collins 2014; Evagorou & Heinis 2018)

Motion cues represent an especially promising future depth cue insofar as their usefulness isn't confined to center vision, and – in combination with stereopsis – they have been shown to non-linearly aid shape perception under controlled conditions (Johnston, Cumming, and Landy 1994). Indeed, this "most important" combination of depth cues can massively increase the ability of users to identify components of complex 3D structures (Ware, Colin, and Mitchell 2005, 2). In practice, the deployment of ray-traced (eye or controller-based) targeting in a given scene could aid in the connectivity analysis of sub-meshes or mesh components that would otherwise go unnoticed. Evagorou & Heinis (2018) demonstrated the implementation of this connectivity with a highlighting feature that allows the user to trace the path of neocortical structures (pg. 2). Color highlighting won't be as useful when viewed peripherally, however, so in the case of the Janus preliminary target stimuli, this feature might take the form of subtle swaying motion of both arrow end points if the user centers their vision (or points their controller) on the arrow shaft, for example.

Finally, the increased performance of both consumer-oriented head mounted displays and associated graphics processing hardware have provided a sufficiently powerful computing platform for the day-to-day use of virtual reality systems. It appears that all we lack, before we can fully integrate systems like Janus into our current academic workflows, are the highly anticipated "killer apps" that various forecasters point to as the necessary, but heretofore elusive inflection point preventing the widespread adoption of VR across academic (as well as the consumer) markets. Importantly, such software is close at hand.

6 CONCLUSION

In early 2017, researchers from CalTech and JPL announced a Series-A funding round for a startup focused on the development of data visualization software to help users quickly visualize numerical data sets in virtual reality. Virtualitics, as the company is now called, provides the means for researchers to jointly analyze and manipulate complex multi-dimensional data in a way that leverages depth-of-field and positional tracking. Later that same year, at the University of Oklahoma, librarians hosted a multi-campus, expert-guided virtual tour of a remote (and inaccessible) Arizona cave system, which featured archaic-era Native American wall inscriptions. In the case of both abstract/numerical and 3D-scanned scholarly content, Janus is poised to contribute meaningfully. Integration with existing software, which might allow users to toggle between "traditional" and simulated 360° perspectives at will, could provide the means to rapidly comprehend complex 3D objects without succumbing to the limitations associated with cognitively taxing object or self movement.

As Donalek et al. (2014) states, "Massive and complex data sets – no matter how content-rich or how expensively obtained – are of no use if we cannot discover interesting patterns in them" (pg. 1). If seen as falling on a spectrum - with idealized, 2D dimensional illustration (in a source text, perhaps) being the most primitive form of visualization, and contemporary VR systems providing the richest and potentially the most efficient viewing experience - Janus surpasses even the most high-powered immersive visualization software currently in academic use. This is borne out in the preliminary performance increases demonstrated by Janus study participants who were generally able to complete spatial analysis tasks more quickly in the simulated 360° viewing condition than in the standard field-of-view (control) viewing condition. While testing is still in the preliminary phase, Janus nevertheless represents a tool type that both acknowledges and attempts to circumvent cognitive rather than computational limitations. By deploying data in a way that effectively bypasses the need for cognitively taxing (object or self) movement, academics can more readily engage with complex 3D data sets and thereby encourage scientific discovery across disciplines.

REFERENCES

Acevedo, D., Vote, E., Laidlaw, D. H., & Joukowsky, M. S. (2001). 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.

Alliez, P., Bergerot, L., Bernard, J. F., Boust, C., Bruseker, G., Carboni, N., ... & Gautier, H. (2017). Digital 3D Objects in Art and Humanities: challenges of creation, interoperability and preservation. White paper. Digital 3D Objects in Art and Humanities: challenges of creation, interoperability and preservation, 71.

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. Psychological science, 15(2), 106-111.

Andersen, S. A. W., Mikkelsen, P. T., Konge, L., Cayé-Thomasen, P., & Sørensen, M. S. (2016). Cognitive load in mastoidectomy skills training: virtual reality simulation and traditional dissection compared. Journal of surgical education, 73(1), 45-50.

Ardouin, J., Lécuyer, A., Marchal, M., Riant, C., & Marchand, E. (2012). FlyVIZ: a novel display device to provide humans with 360 vision by coupling catadioptric camera with hmd. In Proceedings of the 18th ACM symposium on Virtual reality software and technology (pp. 41-44). ACM.

Cruz-Neira, C., Sandin, D. J., & DeFanti, T. A. (1993). Surround-screen projection-based virtual reality: the design and implementation of the CAVE. In Proceedings of the 20th annual conference on Computer graphics and interactive techniques (pp. 135-142). ACM.

Cook, M. (2018). Virtual serendipity: preserving embodied browsing activity in the 21st century research library. The Journal of Academic Librarianship, 44(1), 145-149.

Cook, M., & Lischer-Katz, Z. (2019). Virtual reality integration in the academic library. To appear in K. Varnum (Ed.), Beyond reality: Augmented, virtual, and mixed reality in the library. ALA.

Donalek, C., Djorgovski, S. G., Davidoff, S., Cioc, A., Wang, A., Longo, G., ... & Mahabal, A. (2014). Immersive and collaborative data visualization using virtual reality platforms. arXiv preprint arXiv:1410.7670.

Evagorou, G., & Heinis, T. (2018). Visual Exploration and Interaction with Scientific Data in Virtual Reality.

Fan, K., Huber, J., Nanayakkara, S., & Inami, M. (2014, March). SpiderVision: extending the human field of view for augmented awareness. In Proceedings of the 5th Augmented Human International Conference (p. 49). ACM.

Laha, B., Bowman, D. A., & Socha, J. J. (2014). Effects of VR system fidelity on analyzing isosurface visualization of volume datasets. IEEE Transactions on Visualization & Computer Graphics, (4), 513-522.

LaViola Jr, J. J., Kruijff, E., McMahan, R. P., Bowman, D., & Poupyrev, I. P. (2017). 3D user interfaces: theory and practice. Addison-Wesley Professional.

Kersten-Oertel, M., Chen, S. J. S., & Collins, D. L. (2014). An evaluation of depth enhancing perceptual cues for vascular volume visualization in neurosurgery. IEEE transactions on visualization and computer graphics, 20(3), 391-403.

Lages, W. S., & Bowman, D. A. (2018). Move the object or move myself? walking vs. manipulation for the examination of 3D scientific data.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. Nature, 390(6657), 279.Mads Sølvsten Sørensen. 2016. “Cognitive load in mastoidectomy skills training: virtual reality simulation and traditional dissection compared.” Journal of Surgical Education 73 (1): 45-50.

Maan, Z. N., Maan, I. N., Darzi, A. W., & Aggarwal, R. (2012). Systematic review of predictors of surgical performance. British Journal of Surgery, 99(12), 1610-1621.

Orlosky, J., Wu, Q., Kiyokawa, K., Takemura, H., & Nitschke, C. (2014). Fisheye vision: peripheral spatial compression for improved field of view in head mounted displays. In Proceedings of the 2nd ACM symposium on Spatial user interaction (pp. 54-61). ACM.

Pfarr-Harfst, M. (2016). Typical workflows, documentation approaches and principles of 3D digital reconstruction of cultural heritage. In 3D Research Challenges in Cultural Heritage II (pp. 32-46). Springer, Cham.

Forsberg, A., Katzourin, M., Wharton, K., & Slater, M. (2008). A comparative study of desktop, fishtank, and cave systems for the exploration of volume rendered confocal data sets. IEEE Transactions on Visualization and Computer Graphics, 14(3), 551-563.

Ragan, E. D., Kopper, R., Schuchardt, P., & Bowman, D. A. (2013). Studying the effects of stereo, head tracking, and field of regard on a small-scale spatial judgment task. IEEE transactions on visualization and computer graphics, 19(5), 886-896.

Seth, A., Vance, J. M., & Oliver, J. H. (2011). Virtual reality for assembly methods prototyping: a review. Virtual reality, 15(1), 5-20.

Silvestri, C., Motro, R., Maurin, B., & Dresp-Langley, B. (2010). Visual spatial learning of complex object morphologies through the interaction with virtual and real-world data. Design Studies, 31(4), 363-381.

Strasburger, H., Rentschler, I., & Jüttner, M. (2011). Peripheral vision and pattern recognition: A review. Journal of vision, 11(5), 13-13.

Thomas, L. E., & Seiffert, A. E. (2010). Self-motion impairs multiple-object tracking. Cognition, 117(1), 80-86.

Thompson, M. (2018). Making Virtual Reality a Reality in Today's Classroom. Retrieved from https://thejournal.com/Articles/2018/01/11/Making-Virtual-Reality-a-Reality-in-Todays-Classrooms.aspx?Page=1

Van Dam, A., Laidlaw, D. H., & Simpson, R. M. (2002). Experiments in immersive virtual reality for scientific visualization. Computers & Graphics, 26(4), 535-555.

Veurink, N. L., Hamlin, A. J., Kampe, J. C. M., Sorby, S. A., Blasko, D. G., Holliday-Darr, K. A., ... & Harris, K. S. (2009). Enhancing visualization skills-improving options and success (EnViSIONS) of engineering and technology students. Engineering Design Graphics Journal, 73(2).

Ware, C., & Mitchell, P. (2005, August). Reevaluating stereo and motion cues for visualizing graphs in three dimensions. In Proceedings of the 2nd symposium on Applied perception in graphics and visualization (pp. 51-58). ACM.

Yoon, S. Y. (2011). Psychometric properties of the revised purdue spatial visualization tests: visualization of rotations (The Revised PSVT: R). Purdue University.