AN EMPIRICAL EVALUATION OF VISUAL CUES FOR 3D FLOW FIELD PERCEPTION

BY

ANDREW HARLAN STEVENS

BA, University of California, Davis, 2007

MS, University of New Hampshire, 2017

DISSERTATION

Submitted to the University of New Hampshire

in Partial Fulfillment of

the Requirements for the Degree of

Doctor of Philosophy

in

Computer Science

May, 2021

This dissertation has been examined and approved in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Computer Science by:

Dissertation Director, Thomas Butkiewicz, Research Assistant Professor of Computer Science

R. Daniel Bergeron, Professor Emeritus of Computer Science

Andrew Kun, Professor of Electrical & Computer Engineering

Larry Mayer, Professor of Earth Sciences

Colin Ware, Professor Emeritus of Computer Science

On February 11, 2021

Original approval signatures are on file with the University of New Hampshire Graduate School.

DEDICATION

to

Jackson Bear

May 9, 2009 – December 16, 2020

ACKNOWLEDGEMENTS

This research would not have been possible without the support received from United States National Oceanic and Atmospheric Administration Grants #NA10NOS4000073 and #NA15NOS4000200. A portion of this work was also funded by United States Department of Energy/Los Alamos National Laboratories Contract DE-AC52-06NA25396, sub. 475412.

I am forever grateful to my advisor, Dr. Thomas Butkiewicz, for being an unwavering source of support throughout this journey. Your confidence in me and my work made all of the difference.

I would also like to sincerely thank Dr. Colin Ware, Dr. R. Daniel Bergeron, Dr. Andrew Kun, and Dr. Larry Mayer for volunteering their time to be members of my doctoral committee and for providing invaluable feedback along the way. It has been a great honor to be guided by such an esteemed group of minds.

Thanks to my colleagues at the CCOM VisLab for the many wonderful and insightful conversations over the years: Roland Arsenault, Briana Sullivan, Erol Aygar, and William Ikedo. I would additionally like to acknowledge the Center for Coastal and Ocean Mapping for fostering such a unique, inspirational community of researchers.

Finally, I would like to thank my friends and family for being my biggest fan club, and my wife, Dallas, for being my partner in all things. Your understanding, patience, and love helped to keep the trek from becoming a trudge.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS

θ — angular error in degrees

θ_z — z-component (depth) of the angular error

θ_{xy} — xy-component (plane projection) of the angular error

$\theta_{corrected}$ — angular error corrected for antiparallel confusion

$\hat{t}$ — unit vector representing the target direction

$\hat{p}$ — unit vector representing the input probe direction

$\hat{t}_z$ — unit vector representing the z-component (depth) of the target direction

$\hat{p}_z$ — unit vector representing the z-component (depth) of the input probe direction

$\hat{t}_{xy}$ — unit vector representing the xy-component (plane projection) of the target direction

$\hat{p}_{xy}$ — unit vector representing the xy-component (plane projection) of the input probe direction

L_a — ambient light intensity from the Blinn-Phong shading model

L_d — diffuse light intensity from the Blinn-Phong shading model

L_s — specular light intensity from the Blinn-Phong shading model

k_a — material ambient reflection constant from the Blinn-Phong shading model

k_d — material diffuse reflection constant from the Blinn-Phong shading model

k_s — material specular reflection constant from the Blinn-Phong shading model

$k_{shininess}(\alpha)$ — material shininess constant from the Blinn-Phong shading model

ABSTRACT

Three-dimensional vector fields are common datasets throughout the sciences. They often represent physical phenomena that are largely invisible to us in the real world, like wind patterns and ocean currents. Computer-aided visualization is a powerful tool that can represent data in any way we choose through digital graphics. Visualizing 3D vector fields is inherently difficult due to issues such as visual clutter, self-occlusion, and the difficulty of providing depth cues that adequately support the perception of flow direction in 3D space. Cutting planes are often used to overcome these issues by presenting slices of data that are more cognitively manageable. The existing literature provides many techniques for visualizing the flow through these cutting planes; however, there is a lack of empirical studies focused on the underlying perceptual cues that make popular techniques successful. The most valuable depth cue for the perception of other kinds of 3D data, notably 3D networks and 3D point clouds, is structure-from-motion (also called the Kinetic Depth Effect); another powerful depth cue is stereoscopic viewing, but none of these cues have been fully examined in the context of flow visualization. This dissertation presents a series of quantitative human factors studies that evaluate depth and direction cues in the context of cutting plane glyph designs for exploring and analyzing 3D flow fields. The results of the studies are distilled into a set of design guidelines to improve the effectiveness of 3D flow field visualizations, and those guidelines are implemented as an immersive, interactive 3D flow visualization proof-of-concept application.

CHAPTER 1
INTRODUCTION

Vector fields are mathematical spaces that assign a direction and magnitude to every point within them. In the natural sciences, vector fields often represent real-world phenomena being observed or simulated, such as fluid flow, magnetic forces, and gravitational fields. They are also widespread in the automotive and aerospace industries for assessing aerodynamic qualities of a vehicle design, and in the medical field, where diffusion tensor imaging is used to capture the magnitude and orientation of water diffusion through the brain over time. These vector fields are often invisible to the human eye in real life, and computer visualization of these phenomena allows insights that are otherwise not possible. Researchers and practitioners usually need to explore and analyze vector fields in the context of fluid flowing through them. When the data is very dense, there can be considerable difficulty interpreting flow direction and magnitude when the data is represented at a scale that is practical for viewing on a computer screen.

In scientific data visualization, photorealism is eschewed in favor of abstracted representations that can be readily interpreted for their salient features and to derive accurate, unbiased insights. The mechanisms of human perception are notoriously complex and not fully understood, so gathering knowledge about the perception of abstract data is imperative for creating visualizations that accurately convey the data they are meant to represent. Flow visualization focuses on visually relaying features of interest such as critical points (saddles, sinks, and sources), maxima and minima, and specific flow patterns. Three-dimensional flow visualization is a natural extension of the two-dimensional case, but is inherently more

challenging because of self-occlusion and visual clutter.

There have been many decades of flow visualization techniques presented in the literature that can be generally grouped into four categories: direct geometric glyphs, which are direct visualizations of the data using objects that encode local information (location, direction, and magnitude); integral geometric glyphs, based on tracing advected particles; texture-based illustration, which uses dense sets of seeds (often a white-noise image) that are "smeared" by the flow field; and feature extraction/topological visualization, which applies computations on a global scale to identify and visualize well-defined characteristics of a flow field. Cutting planes are often used to interactively reduce the dimensionality of 3D vector fields into more cognitively manageable 2D slices. Applying 3D glyphs to these cutting planes is a common choice, as they can present an additional dimension of information as opposed to 2D techniques such as color mapping and texturing.

Many glyph designs have been proposed by the visualization community, and a number of effective styles have been identified. However, there is a lack of empirical studies explicitly examining the relative effectiveness of the different visual elements that comprise these popular glyph designs. Moreover, most related studies are task-based, and do not explicitly address perceptual effectiveness. In the absence of empirical data to guide designs, human perceptual theory can provide a powerful framework to support the creation of novel flow visualizations. Currently, reasoning about the perceptual effectiveness of a 3D flow visualization technique to accomplish a given task is largely a guessing game that relies upon generalizing results from other visualization domains (e.g., node-link relationship graphs) to make perceptual assumptions that may not be valid.

Tasks involving visualization demonstrate marked performance improvements through stereoscopic viewing of the data. On desktop monitors, it simulates human binocular vision by rendering a separate view of the virtual scene from the perspective of each eye. The spatial disparities between these two perspectives are processed by the visual system to provide additional depth information that improves our spatial understanding of the virtual

environment. Though stereoscopic displays have failed to find a place in the typical consumer's home, they continue to serve a valuable role in scientific visualization by supporting the perception of inherently 3D data. The introduction of high-fidelity head-tracking in the early 1990's gave rise to head-mounted displays and so-called "Fishtank Virtual Reality" that updates its virtual projection to match and follow the physical viewpoint of the user, providing an interactive 3D view of the scene on a desktop monitor.

The use of stereo image pairs to gain depth perception also carries with it an increase in geometric distortion of features in the image due to the mismatched perspective between the camera and the actual viewpoint. These can arise from a difference between the viewer's inter-pupillary distance (IPD) and the separation amount between the stereo image pairs, but primarily occur due to uncoupled dislocated viewing, where the virtual and physical viewpoints do not match. Much research has been done to model and quantify these distortions, but relatively little work has explored their perceptual consequences. When visualizing 3D vector field data, it is common practice to map vector magnitude to glyph length. Uncoupled stereoscopic perspective viewing can lead to significant distortion of glyph geometry, potentially causing significant misperceptions of the vector data represented by the glyphs.

Resolving how the principal components of vector fields – direction and magnitude – can be visualized so that they are perceived as accurately as possible will allow us to bridge important gaps in flow visualization research and provide perceptual anchors upon which future 3D flow visualization techniques can be built. Improving our understanding of how these visual cues and viewing conditions impact the perception of three dimensional flow field information can support the development of glyph designs and visualization techniques that are perceptually superior, simplifying the task of interpreting and communicating important details of the flow field.

CHAPTER 2
LITERATURE REVIEW

Three-dimensional vector field visualization may superficially appear to be a trivial extension of the 2D case to an extra dimension, but it presents many formidable visualization challenges.

The primary limitation of glyph-based (and other) approaches is self-occlusion, where glyphs closer in depth to the viewer obstruct the view of other glyphs positioned behind them [1]. Similarly, these approaches can also suffer from extreme visual clutter/complexity as seeding densities increase. When there is too much information for the visual system to deal with effectively, it can lead to an obfuscation of global and local patterns, hindering the tasks they are supposed to support. These issues become immediately apparent when using popular software packages that render vector fields as 3D grids of arrows by default. An example of this effect can be seen in Figure 2.1.

Though there have been efforts to mitigate these particular challenges for 3D glyph-based approaches [2, 3], the overwhelming response by researchers has been to develop alternative visualization metaphors and to strategically sample data to reduce visual density while retaining the communicative power of the visualization. Laramee et al. [4] describe and implement a variety of these approaches, including geometric streamtubes and texture-based techniques, apply them to slices through computational fluid dynamics data, and provide details on their relative strengths and weaknesses.

This dissertation seeks a different approach by understanding the advantages of the various visual cues used in each visualization technique to determine not only why particular

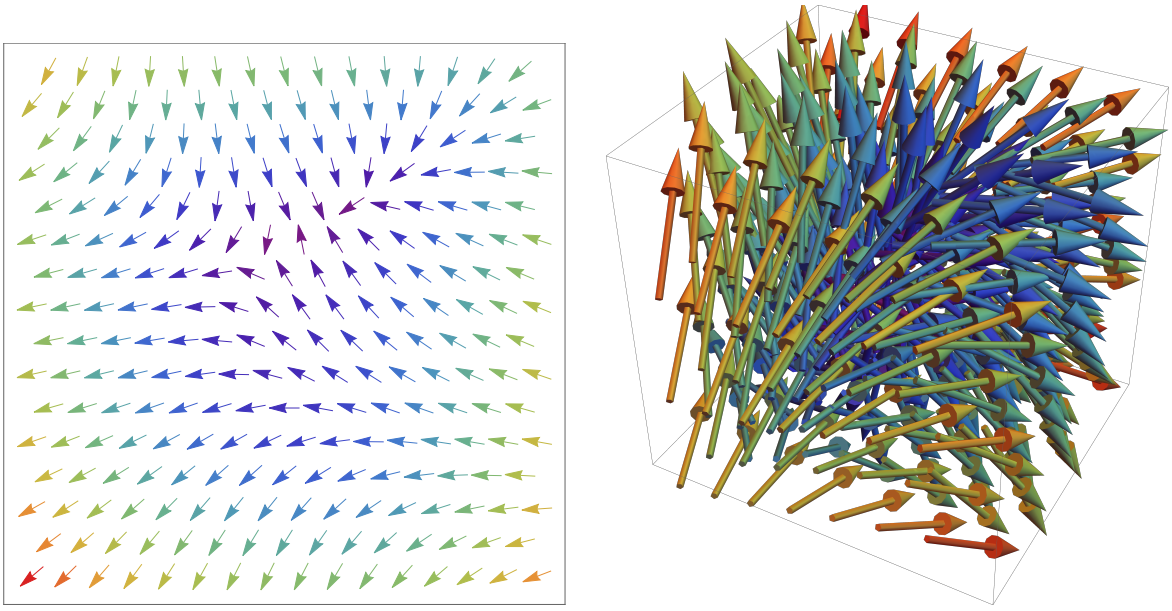

(a) 2D Vector Field with 2D Arrows

(b) 3D Vector Field with 3D Arrows

Figure 2.1: A comparison of 2D vector field visualization and 3D vector field visualization with simple arrow glyphs. Visualization techniques used for 2D vector fields can introduce severe occlusion issues when extended to 3D.

techniques are effective, but to identify sets of compatible visual cues that can be combined to produce demonstrably better visualizations.

2.1 3D Flow Visualization

There are generally considered to be four main categories of 3D flow visualizations: direct flow glyphs, integral flow glyphs, texture flow techniques, and feature-based flow techniques [1, 4]. The first two are very common; the latter two tend to be used in more niche flow visualization domains.

Direct flow glyphs convey local information about the vector field using geometric objects whose substructure is tied directly to the data itself. An example of this is an arrow with a stem length representing vector magnitude, and an orientation indicating flow direction. Figure 2.2 shows a 2D vector field using tapered droplet-style direct glyphs.

Integral flow glyphs give more global information and include methods such as stream-

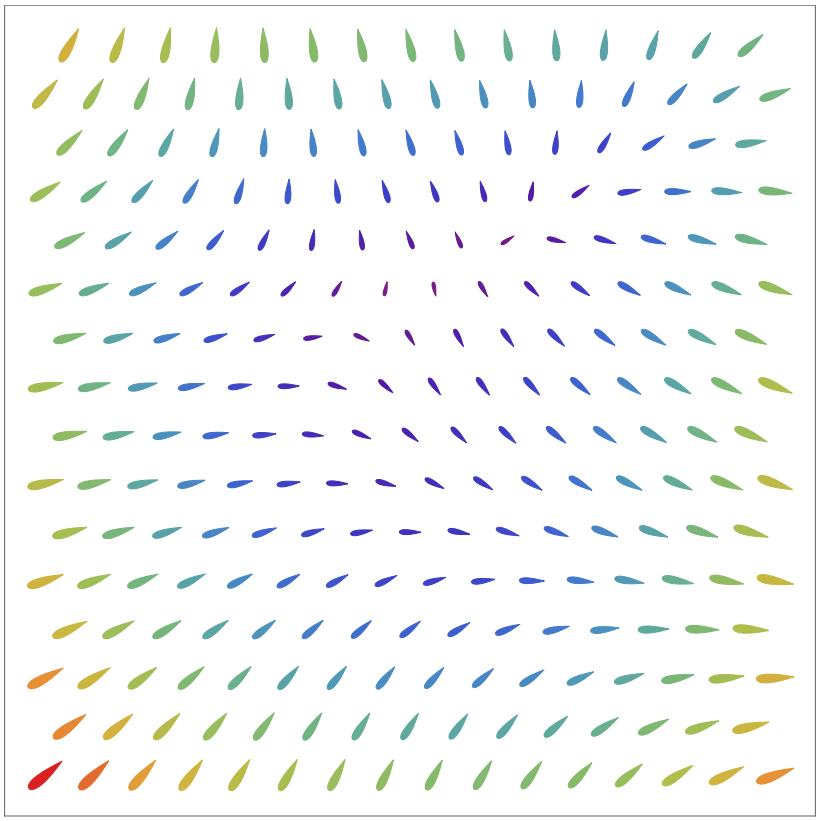

Figure 2.2: A 2D vector field visualized with 2D direct glyphs.

lines, that trace the path of a particle as it advects within a steady flow field; pathlines, the streamline analog within unsteady flow fields; and streamtubes, a three-dimensional extension of streamlines for 3D flow fields. McLoughlin et al. [5] provide a thorough survey on this family of flow visualizations. The same 2D vector field from Figure 2.2 is visualized using similar droplet-style integral glyphs in Figure 2.3 to illustrate the encoding of local versus global information about the vector field.

Texture flow is the use of texture-based methods to indicate flow properties, as in the line integral convolution (LIC) technique introduced by Cabral and Leedom [6]. This algorithm advects pixels from generated white noise images through a vector field, "smearing" the noise image and resulting in a texture depicting the flow field, which densely covers the entire region of interest. While LIC methods applied to 3D scenarios are most often used to show flow on surface boundaries [4], the method has also been extended by Interrante and Grosch to be compatible with 3D volumes [7]. An example can be found in Figure 2.4, where

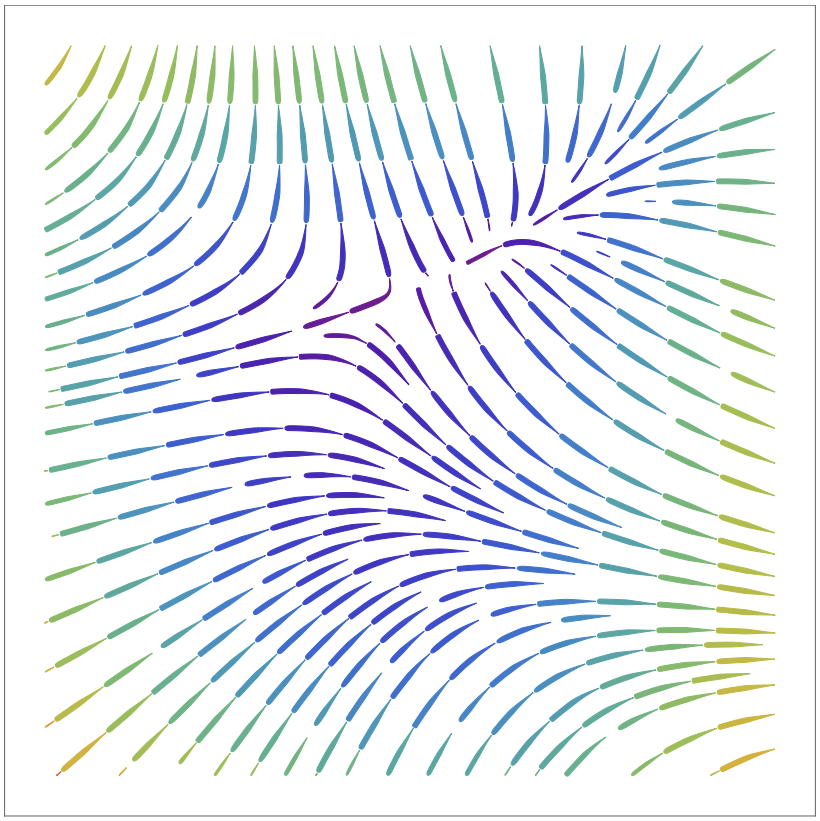

Figure 2.3: A 2D vector field visualized with 2D integral glyphs.

the 2D vector field from Figures 2.2 and 2.3 smears a random noise texture to create the LIC image, which is then colored by vector magnitude, as in the other figures. It provides more coverage than either direct or integral glyphs, but there is no indication of the flow direction (the vector sign, in particular) along the contours.

Lastly, there is feature-based or topological flow, which attempts to computationally extract and visualize the most salient structures of interest, such as critical points and vortices. These techniques are usually employed when the features being sought are known *a priori*, and visualized as discrete, independent objects as opposed to sets or fields of related glyphs. For a detailed exposition on the many solutions presented in the literature, Post et al. [8] provide a useful survey and categorization of the field, along with its challenges.

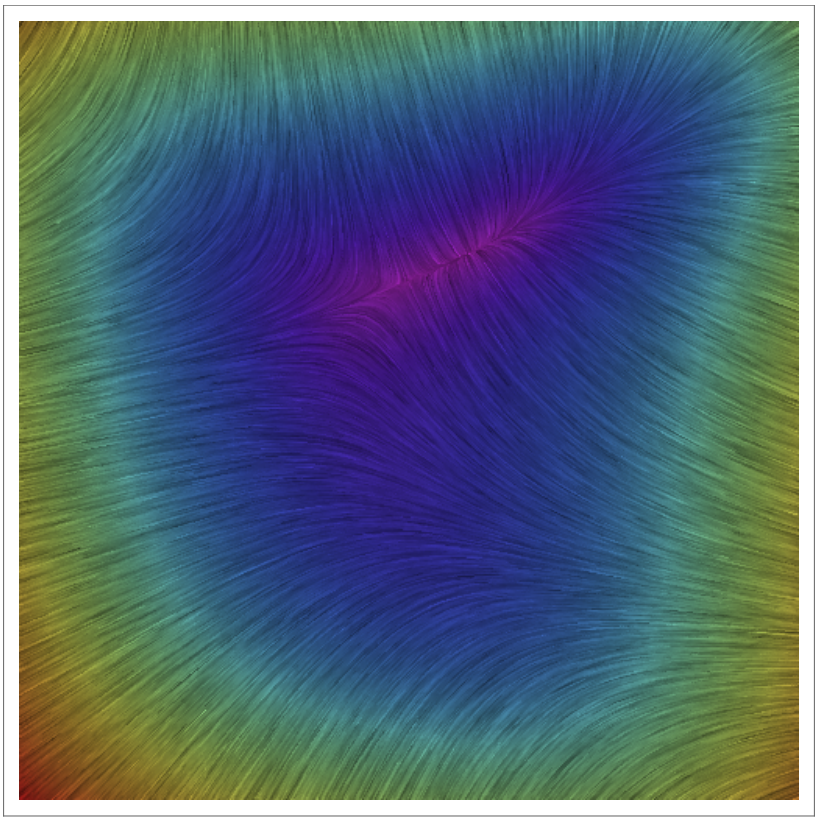

Figure 2.4: A 2D vector field visualized using the Line Integral Convolution (LIC) technique. A noise image is "smeared" by the flow field, and then colored by magnitude.

2.2 Cutting Planes

Cutting planes are tools used to reduce the visual dimensionality of 3D data sets to more manageable 2D slices, which ameliorates many of the issues inherent in volumetric 3D visualizations, such as self-occlusion. Cutting planes are widely implemented in scientific visualization systems [8–11], modern open-source visualization toolkits [12], and turnkey software for industry [13–15]. There have even been educational systems designed to train students in their use and improve performance while employing them as a visualization tool [16]. Figure 2.5 demonstrates how cutting planes slice through the 3D data volume and use volume sampling to drastically reduce the occlusion issues that overwhelm the visualization.

Cutting planes are used in many different ways. They are sometimes used as clipping planes to view structures enclosed within other volumes [4, 17], or as seeding planes for

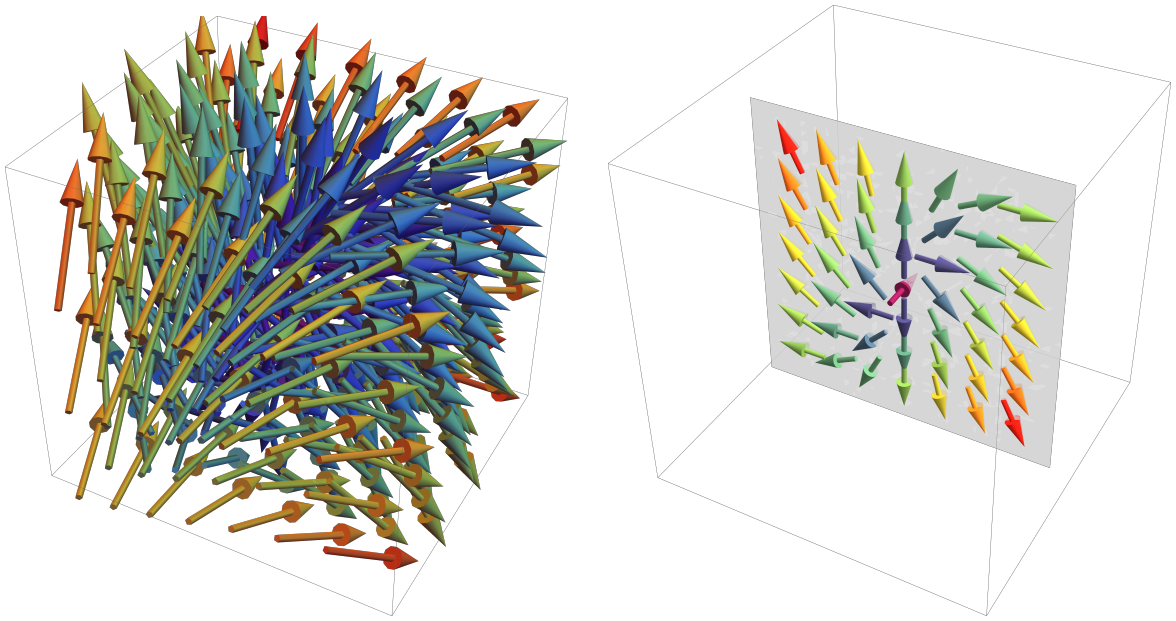

(a) Volume Sampling (b) Cutting Plane Sampling

Figure 2.5: A cutting plane, or slice, through the vector field reduces the dimensionality of the sampled data and greatly improves the clarity and legibility of the visualization, at the cost of reduced coverage of the 3D flow field. The same 3D vector field is visualized in both images.

geometric visualizations like streamlines and streamtubes [4, 9, 18]. Studies have examined the use of physically interactive cutting planes, like Meyer and Globus' study [19] on cutting plane use in virtual reality and Hinckley's [20] use of physical props to interactively orient a cutting plane through a model of the brain. These studies emphasize the value of interactively sweeping the cutting plane through the data field to reveal patterns and structure. Often, cutting planes are color-mapped to show 2D scalar data like temperature and magnitude [18–20] or are used as higher-dimension visual displays that can leverage the effectiveness of 2D visualization methods.

Creating a strong flow visualization on cutting planes is similar to the problem of 2D flow visualization, which has seen some important recent contributions. Laidlaw et al. [21] lay solid quantitative groundwork for assessing the relative merits of commonly used 2D flow methods via empirical task-based evaluation. Liu et al. [22] extend this work to include color mapping and a sophisticated study and analysis design, though with limited evaluation of

glyph styles.

There are many techniques for visualizing 3D vectors through a cutting plane. Konrad-Verse et al. [23] present a deformable cutting plane for surgery planning. Modiano [24] and Schulz et al. [18] place scalar color-mapped cutting planes in series, while Fuhrmann's dashtubes [25] make use of an animated, partially transparent texture applied on streamtubes to strongly disambiguate flow direction and lessen occlusion. However, most techniques employ some form of direct or integral visualization using a glyph or other geometry.

2.3 Glyphs and Integral Geometries for Flow Visualization

Flow glyphs are the geometric representation of local vector field information, which individually present the flow data at their immediate locations. Identifying and understanding global patterns and features typically relies on integral glyph techniques. Munzner [1] gives a good overview of the various types of geometries used for 3D flow visualizations, summarized here.

The most common flow glyphs by far are arrows, as encountered in MATLAB's quiver3 function or Wittenbrink's work on visualizing uncertainty with 3D glyphs [26]. Laramee et al. [4] also include techniques in their computational fluid dynamics (CFD) visualization design study that directly visualize arrows and streamlets (short, tubular glyphs). The FlowVisual 2D teaching tool described by Wang et al. [27] animates arrows along flow paths to reinforce their directional cues.

Integral geometric structures often involve either simple streamlines, or perceptually superior illuminated streamlines [28], that incorporate a shading model. 3D extensions of streamlines, referred to as streamtubes and threads, have been the subject of a number of design studies [4, 29, 30]. There are numerous other variants, e.g. streamribbons [31] which are used to highlight rotation and torsion. When used with cutting planes, the slices typically serve as a seeding plane for these approaches.

Borgo et al. [32] provide a survey of the wide-ranging uses of glyphs, and a design frame-

work by which to compose glyphs. They highlight the importance of simplicity and symmetry in 3D glyph design, and underscore the need to facilitate depth perception for 3D visualizations, chiefly by careful glyph design and illustrative techniques such as halos and depth cueing. Lie et al. [33] also studied glyphs for visualizing 3D data, and while they make the important observation that 3D glyphs are best suited to data that naturally lends itself to a 3D representation, their study focuses on 2D billboarded superellipses and lacks a formal user study.

2.4 Visual Cues for Glyph Perception

Many researchers have sought to increase visual clarity by proposing techniques that add or reinforce perceptual cues. (These include Chen et al.'s framework [34] for 3D vector field visualization, and Brambilla et al.'s survey [35] of flow visualization).

However, a majority of previous work lacks accompanying user studies that examine what perceptual or task-based gains have been made through new techniques. There are few empirical studies in the literature concerned specifically with 3D glyphs, and most focus on representative task performance, not perceptual effectiveness.

Theories of space perception are based on cues that help build our perception of orientation on the plane orthogonal to the line of sight. When combined with 'depth cues', they contribute to our understanding of the layout and orientation of objects in depth [36]. The following sections contain a summary of the perceptual cues that are most relevant to 3D glyph designs.

2.4.1 Orientation and Direction

Two-dimensional vector representations can be conceptualized as a combination of magnitude and direction with respect to orientation and sign. This relation is depicted in Figure 2.6.

Magnitude is often color-coded, but can also be represented by glyph length, size, thick-

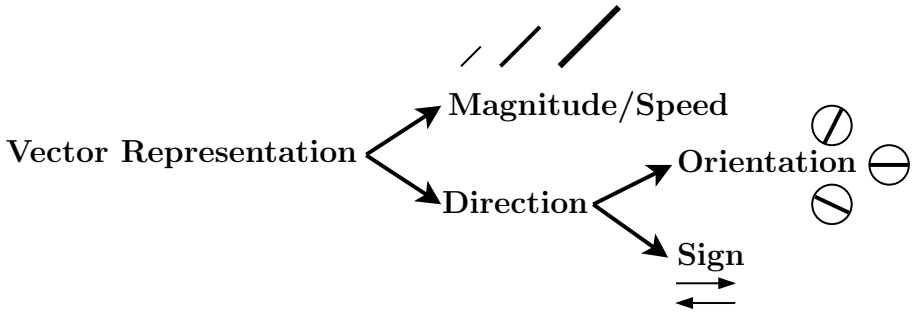

Figure 2.6: The components of a vector representation. In vector visualization, the direction component is usually further separated into an orientation and a binary sign component. Adapted from Ware, 2008.

ness, or other visual weight modifiers.

For the purposes of flow visualization, vector direction is usually separated into an orientation and a binary vector sign. Orientation is a stationary or static alignment, while the sign implies movement of some type: from one position to another over time, or to describe the slope of gradients. For 2D flow, orientation is best conveyed with contours or strong directional glyphs that are tangential to the flow [21].

Vector sign is indicated by adding asymmetry to a path, the perception of which is tied to neural structures known as end-stopped cells [36]. Arrows are a classic way of indicating direction, but more perceptually effective structures exist (e.g. streaklets and streamlets) that can elicit a stronger response out of the end-stopped cells. Arrows also contribute visual noise via the contours of the arrowheads themselves, the directions of which do not correspond to the underlying flow data.

2.4.2 Depth Cues

Information about the depth dimension, defined as distance from the viewpoint, is provided by a set of depth cues which are processed by the brain in various ways [37]. These include pictorial cues, such as linear perspective, occlusion, and size gradients, as well as non-pictorial cues, of which the most important are stereoscopic viewing and structure-from-motion (SfM). However, research has also shown that depth cues are combined in ways that differ according

to the perceptual task [38]. In other words, there is no single optimal way of providing depth information.

Depth cues arguably play the most important role in a 3D flow visualization. We rely on them to help us make spatial distinctions and judgements about objects in our world, and they come in a variety of forms. In Figure 2.7, the same 3D vector field is sampled and visualized with line-based arrows and 3D tube-based arrows. It is obvious from this juxtaposition that the contours and shading of the 3D arrow shape convey much more depth information, and that the design of a 3D glyph has a very large impact on the effectiveness of the 3D flow visualization.

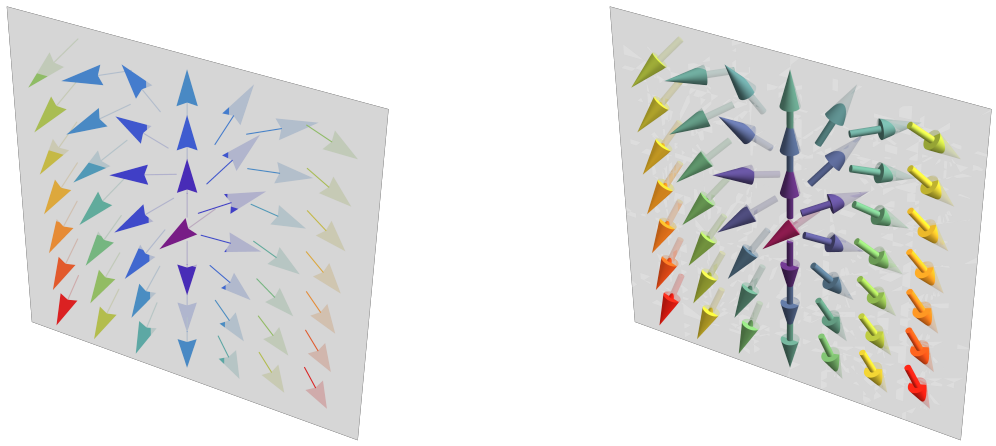

(a) Line-based 3D Arrow Glyphs

(b) Tube-based 3D Arrow Glyphs

Figure 2.7: A comparison of line-based arrow glyphs and 3D tube-based arrow glyphs. The tube-based design provides many more depth cues to the visual system than the line-based arrow.

Occlusion is not just an annoyance for 3D visualization researchers; it acts as an indispensable cue for distinguishing the relative depths of objects. When one object occludes another, we understand the occluding object to be closer than the occluded object, but we do not have any extra information to determine the distance between them.

Shading-based cues are also of key import to depth perception. Shape-from-shading is the cue that helps us to understand the contours and ridges of 3D surfaces through the

shading and texture applied to it. In particular, regularly structured linear surface textures such as simple grids have been found to be helpful in the perception of surface shape [39].

Cast shadows can also serve as a strong cue, when used correctly. They mainly provide information about the height of an object, but can also give an indirect depth cue by conveying the position of an object within the environment relative to a surface and light source. As visual scenes become more complex, the ability to tie shadows to the objects that cast them diminishes, and the cue quickly goes from useful to useless or, worse, a hindrance to understanding the visualization.

Stereoscopic depth cues stem from small differences between the image received by the right eye and that received by the left eye. Though often thought of as *the* 3D depth cue, many other depth cues exist, and experiments have shown that stereoscopic depth is often not the most useful cue for depth perception in certain tasks. Nonetheless, its prominent role in successful 3D flow visualizations has been reinforced by a number of recent quantitative studies [34, 40, 41].

Another depth cue that has been the subject of recent empirical studies is structure-from-motion (SfM). The two main flavors of SfM are kinetic depth, the phenomenon of being able to understand the 3D shape of a rotating bent wire structure solely by viewing its projection on a screen, and motion parallax, where objects closer to us appear to move more quickly than those closer to the horizon when we move laterally. Like stereoscopic depth, SfM has emerged as an important cue for depth perception in controlled evaluations [40, 41].

A number of applied studies have shown that for certain visualization tasks, both stereoscopic viewing and SfM can greatly improve performance. For example, both SfM and stereo greatly improve performance on tracing paths in 3D networks, but SfM is the more powerful cue [42–44]. The same appears to be true for perceiving structure in point clouds [45]. These non-pictorial cues are invaluable for visualizing 3D networks and point clouds, data types for which pictorial cues offer little help. However, for other tasks stereo and SfM are less powerful. For example, when perceiving the shape of surfaces, shape-from-shading and

surface texture may be more useful [46].

Also, for visually guided reaching, stereoscopic depth perception is more important than SfM [47]. However, Chen et al. [34] found that for tasks involving the 3D tracing of fiber bundles in the brain represented by streamtubes, performance was hindered rather than helped by stereoscopic viewing. In a related task involving 3D node-link diagrams or "information nets", Ware and Mitchell [48] showed that rotating the diagram at 36 seconds per cycle provided depth cues to experienced users which dominated any stereo viewing benefits for understanding large graph layouts. Though error rates were lower with SfM, task completion times were significantly higher in both conditions that included motion. Arthur et al. [49] produced similar results in a 3D graph-tracing task where SfM was provided through head-coupled perspective viewing instead of rotating the graph itself.

In considering the best methods for representing 3D vector fields it is important to consider which depth cues are likely to contribute most to the perception of flow patterns. Perception of patterns orthogonal to the line of sight is relatively unproblematic, since the patterns are directly expressed in the 2D images on the retinas. Research suggests that streamline-based methods are the most effective for 2D flow [21], and that animation of streamlets (short streamline segments) along their paths can enhance pattern perception [50].

The value of SfM and stereo may depend greatly on the kind of rendering method. For example, with simple streamlines pictorial cues are weak. However, if the data are represented with streamtubes having a significant diameter, and bands around them, the shape of the bands provides orientation cues. There is also the interesting question of whether animation of streamtubes in the manner of Fuhrman and Groller [25] would hinder any benefits from SfM cues. It is possible that multiple types of animation, one to achieve the kinetic depth effect, and a second animation along the streamtubes (to show flow direction) might cause perception interference when combined simultaneously, meaning that both 3D orientation perception and flow direction perception could be hindered.

Nevertheless, viewing 3D vector fields is a challenging perceptual task where both stereo-

scopic viewing and SfM can be expected to provide substantial benefits, especially for line-based designs, where the value of pictorial cues (such as perspective) are minimal. However, rendering streamlines as shaded streamtubes can provide more depth information to help with the perception of their 3D direction.

2.4.3 Texture

Texture is an important component of glyph design that can encode additional information or reinforce existing cues. To quote Ware directly, "even if we texture all objects in exactly the same way, this can help us perceive the orientation, shape, and spatial layout of a surface." [36] In fact, textures can be conceptualized as dense fields of glyphs, mapping certain properties of the object to which they are applied.

The human visual system is particularly sensitive to high-contrast grating-type patterns [36]. Textures that employ patterns over certain orientations and spatial frequencies elicit particularly strong responses, as do the boundaries between regions of textures with different orientations.

2.4.4 Evaluations

Ware [40] and Forsberg et al. [41] obtained conflicting results when comparing line-based flow visualization techniques to tube-based techniques. While Ware found that tube-based rendering techniques had a significant advantage over those that are line-based, Forsberg et al. did not observe the same difference, and even found line-based techniques to be better in some circumstances and preferred by the majority of users. It is hard to identify the precise cause of this disagreement, because the experimental designs differed on key aspects. Ware evaluated lit lines along with shaded and ring-textured tubes on a high-resolution display, while Forsberg et al. compared unlit but colored lines along with larger shaded tubes augmented with auxiliary directional glyphs, a dense direction-encoded texture, and halos to help disambiguate the occluding tubes displayed on a standard-resolution monitor. An

interesting result from Ware's study is that tube-based visualizations performed remarkably well even under monoscopic conditions. These differences helped to form the initial questions which lead to this dissertation topic.

2.5 Correct Stereoscopic Viewing

As Figure 2.8 illustrates, there is a correct pair of viewpoints to generate a stereo image pair, and this includes an assumed eye separation. When a stereo image is viewed from the wrong distance or the wrong viewing angle, distortions should result and the magnitude of these distortions can be calculated using the geometry illustrated in Figure 2.8. For example, the diagram in Figure 2.8b shows that because of the incorrect viewing angle, the bar should appear shorter.

Fishtank VR is a technique for 3D visualization that uses head tracking and a stereoscopic desktop display [51,52]. The effect is as though the user were looking into a virtual 3D volume that extends into the monitor itself, like a window into a fishtank. Because Fishtank VR generates images that are always correct from a particular viewpoint, the distortions from off-axis viewing should not exist, and the bar in Figure 2.8c should appear to have the correct length.

A geometric distortion model for stereoscopic image acquisition and representation was developed by Woods et al. [53] and further expanded by Held et al. [54]. For a full technical assessment and description of stereoscopic errors and distortions, see the report prepared by Boev et al. [55]. An analysis by Son et al. [56] claims that the geometric distortion introduced by turning one's head is negligible in comparison to the distortions caused by the wrong viewing distance or viewing angle.

Banks et al. examined the perception of surface orientation for viewers of monoscopic [57] and stereoscopic [58] images at different viewpoints, both correct and where there was a mismatch between virtual and actual viewpoints. Essentially, their results showed that viewers of monoscopic images were able to compensate for incorrect viewpoints, but viewers

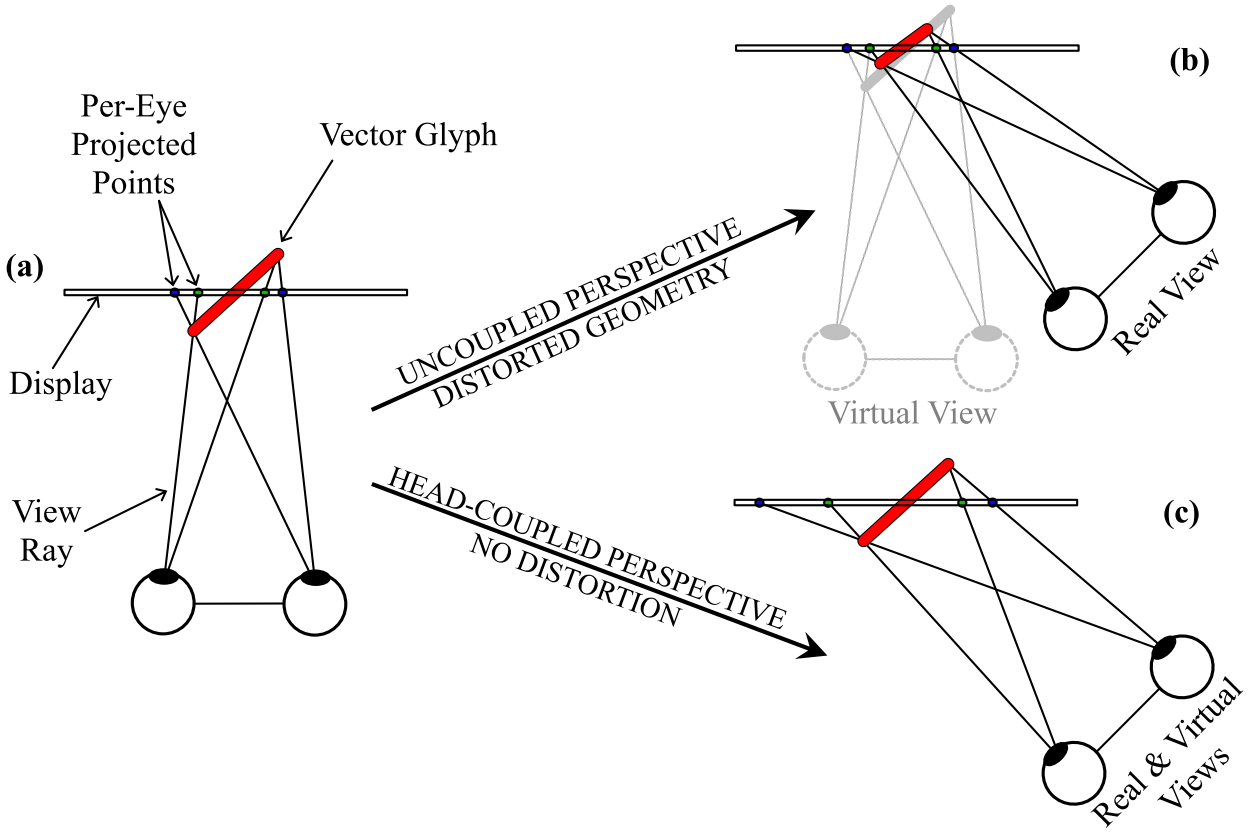

Figure 2.8: Correct stereoscopic viewing on a desktop monitor requires the user's view and the virtual projection to match (a), (c); otherwise the glyph length (red) will appear distorted (b) because the visual system receives incorrect information about the 3D scene.

of stereoscopic images did not compensate for mismatched viewpoints.

Fishtank VR has been shown to improve the accuracy and speed with which some 3D tasks can be accomplished. Arthur et al. [49] conducted a series of experiments to evaluate Fishtank VR performance alongside a traditional graphics workstation, and found that a 3D graph tracing task saw significant performance improvements from head-coupled perspective and stereopsis, though the head-coupled perspective was responsible for the majority of the reduction in errors. Participants in that study also reported a strong preference for head-coupled stereo over uncoupled or non-stereo conditions. However, a number of other studies have shown that kinetic depth from motion (not necessarily coupled to head movements) is responsible for reductions in errors in 3D path tracing tasks [42, 43].

Understanding 3D flow models is an extremely challenging visualization task, and stereo

viewing seems likely to make more depth information available. But since stereo images on desktop monitors are rarely viewed from the correct positions, there may be significant distortions in the judged angles and lengths of 3D glyphs representing the underlying 3D flow field. It seems plausible that using Fishtank VR will help with this. A thorough examination of the literature uncovered no prior studies that test this hypothesis.

CHAPTER 3
PROBLEM STATEMENT

Flow visualizations help researchers and practitioners understand complex patterns and features in their data. Visualizing three-dimensional flow fields is particularly challenging due to occlusion and visual noise; the past few decades have seen many techniques proposed to improve these visualizations, often using clever seeding strategies and illustration methods to increase legibility and interpretability of the final product. A common tactic is to reduce the visualization space by sampling the 3D flow field on a surface or 2D plane. Yet, there are few empirical studies which can be drawn upon to build perceptually effective flow visualizations. This research investigates three lines of inquiry to begin addressing the lack of flow visualization evaluations motivated by human perceptual theory.

The primary goal is to empirically determine which depth cues are most important for perceiving three dimensional flow information. Many techniques have been presented and studied by researchers, and a number of effective 3D glyph designs have been identified. Most related studies are task-based, and do not explicitly address perceptual effectiveness. Understanding and identifying the visual cues which make visualization techniques successful can help us develop guidelines that can inform future design decisions. This research seeks to determine the principal cues involved and how they contribute to the effectiveness of 3D glyphs in conveying flow information.

A related secondary focus is finding which sets of visual cues can be combined to produce demonstrably better visualizations. Depth cues are presented as an ensemble in data visualizations, so it is essential to also investigate which combinations of these cues reinforce one

another, and which combinations have the potential to clash and undermine one another.

The final goal is to gain a better understanding of how viewing conditions may affect perception of flow information. Correct stereoscopic viewing depends upon a variety of conditions that are often difficult to wholly satisfy, particularly the assumption that the physical viewpoint and the virtual viewpoint used to render the scene are congruent. To that end, the perception of glyph length are examined when images are viewed from the correct position, or from incorrect positions where length judgement errors can be calculated and predicted using the geometrical distortion model, and then compared to the empirically-derived results to assess how well the model predicts perceptual performance.

3.1 Hypotheses

This research addresses the following hypotheses regarding human perception of 3D flow glyphs seeded on a cutting plane:

H1: Glyph designs incorporating the most depth and orientation cues will be the most successful for estimating 3D flow direction.

H2: Tube-type designs should outperform line-type designs because of their stronger pictorial cues.

H3: The cutting plane seeding density significantly affects perceptual performance for a given glyph design.

H4: The effectiveness of tube-type designs depends upon their diameter; diameters that are too small will not be able to adequately show lighting and texture cues, and those that are too big waste space and add unnecessary visual noise.

H5: Adding lighting cues to line-type designs (i.e., illuminated lines) can enhance accuracy.

H6: Cast shadows can enhance accuracy.

H7: Stereoscopic viewing can enhance accuracy.

H8: Structure-from-motion can enhance accuracy.

H9: The benefit of stereo and SfM cues should be greater in the case of line-type designs than for tube-type designs, because lines otherwise lack depth information.

H10: Animating glyph textures can improve the perception of flow direction.

H11: Animation will hinder SfM. In other words, when glyphs are animated to show flow direction, the benefit from SfM should be reduced.

H12: Perceptual performance of glyph length judgements follows the predictions made by the geometric distortion model.

H13: Head-coupled viewing reduces errors in perceived glyph length compared to uncoupled, dislocated viewing.

H14: Perceptual performance is consistent across typical glyph length ranges.

CHAPTER 4
RESEARCH APPROACH

The hypotheses of this dissertation were explored through three experiments that build upon one another to provide a basis for understanding the perception of 3D flow glyphs seeded on a cutting plane. These experiments were approved by the University of New Hampshire Institutional Review Board for the use of human subjects (see Appendix A).

Past studies [21,22,41,59] have concentrated on measuring performance for a wide variety of tasks. They evaluate visualizations in the context of tasks such as estimating advection trajectories; identifying important patterns in the flow structure; and perceiving magnitude, vorticity, turbulence, or other features of a flow field. These studies provide valuable practical insights for common tasks involving flow fields. However, Ware's evaluation [40] of 3D contour orientation is the only qualitative study that focuses specifically on the perceptual performance of these visualization techniques, and it would be beneficial to have more studies that tie perceptual theory to visualization performance.

The following tasks were identified by Ware as key goals of flow visualization [36]:

- Judging the speed, orientation, and direction at an arbitrary point.
- Identifying the location and nature of critical points.
- Judging an advection trajectory.
- Perceiving patterns of high and low speed (or magnitude).
- Perceiving patterns of high and low vorticity/curl.

- Perceiving patterns of high and low turbulence.

Of the set of representative tasks, this dissertation focuses on the first, as it is a strongly spatial task that supports the objective of perceptual evaluation. Judging the 3D direction of flow at an arbitrary point on a plane is one of the most basic and fundamental tasks in understanding a flow field, and serves as an ideal springboard for continuing to piece together the perceptual framework of 3D flow visualization.

4.1 Empirical Evaluations

Three experiments were designed to test the hypotheses proposed in the previous chapter. Each experiment deals with a different subset of the hypotheses. The first experiment addresses the more basic hypotheses, and the two following experiments build on the observations of the first to answer the remaining hypotheses. The subset of hypotheses tested through each of the three experiments are listed in Table 4.1.

Table 4.1: Proposed experiments and the hypotheses they address.

Experiment	Hypotheses Tested
1	H1 - H6
2	H1, H7 - H11
3	H7, H12 - H14

The first experiment [60] focuses on the effects of non-stereoscopic, pictorial depth cues on the perception of simple 3D glyph directions. It seeks to understand how the more basic visual depth cues incorporated into 3D glyphs help with the perception of direction. It tests hypotheses H1 through H6.

The second experiment [61] extends the first by looking at how more complex 3D integral glyphs are affected by stereoscopic, motion-based depth in terms of perceiving their directions. Its design closely follows the first experiment, so that the results may be interpreted as an ensemble and more generalized observations may be made. It looks at unanswered aspects of hypothesis H1 from this first experiment along with hypotheses H7 through H11.

The final experiment takes a slightly different approach and examines how stereoscopic viewing irregularities may affect the perception of 3D direct glyph length. Glyph length was chosen because it is a very common mapping for flow magnitude, and can undergo extreme visual distortions if certain stereoscopic viewing conditions are not met. The expected distortion can be described and modeled geometrically so that potential perceptual errors in glyph length judgement can be predicted using this model and compared to observed errors from a human factors experiment. This experiment investigates the stereoscopic viewing benefit hypothesis, H7, along with the remaining three hypotheses, H12 through H14.

4.2 Discussion and Design Guidelines

Each of the hypotheses are evaluated with respect to the experimental results and discussed on an individual basis. Acceptance or rejection of the hypothesis is indicated based on the available empirical evidence.

Finding answers to the proposed hypotheses makes it possible to infer best practices for designing an effective 3D flow visualization. Distilling the results into a set of actionable guidelines increases the practical value of the research, and makes explicitly clear what is most important to take into account while developing a perceptually strong 3D flow visualization. It also connects perceptual theory to practice by deriving prescriptive design guidelines from the descriptive empirical results.

4.3 Application

In data visualization, the insights gained from evaluations should be further validated by implementing them, bridging the gap between perceptual theory used to make predictions about the effectiveness of a visualization and its practical application to flow visualization. The proposed design guidelines are the layer connecting the theory to practice, so it follows that an application should be developed to confirm these guidelines in a real-world setting.

The results of these evaluations have been implemented as an interactive flow visual-

ization tool that can be used to inspect 3D flow fields and create visualizations following the guidelines established through this perceptual research. The final product combines the perceptual research carried out in this dissertation with an immersive virtual environment featuring natural interaction metaphors to facilitate exploring, analyzing, and creating visualizations of any three-dimensional flow field.

CHAPTER 5
EMPIRICAL EVALUATIONS

5.1 Experiment 1: Non-Stereoscopic, Pictorial Cues

The primary purpose of this study is to understand how different visual cues incorporated into direct glyphs help with the perception of 3D vector direction under static, non-stereoscopic viewing conditions.

5.1.1 Glyph Designs

Five glyph designs were selected from the literature for evaluation. Each of these approaches represents a class of common glyph designs, and each design provides a different set of visual cues that support the estimation of flow direction. The chosen designs are detailed in the following sections.

In all cases, a flow glyph is anchored to a seeding point on a cutting plane within the experimental flow field, and extends in the direction of the flow according to the magnitude at that point. Each glyph has a spherical head attached to its terminal end to disambiguate the overall direction. One unit of magnitude is mapped to one millimeter of glyph length on the display, except for the Shadowed Lines design, which requires magnitude normalization as explained in its section.

These glyphs encode only local information about the flow field directly at the seeding point: flow direction and magnitude. Transfer functions map these values to visual elements of a glyph. Typically, direction is mapped to the principal axis of the glyph so that its geometry visually aligns with the flow, while magnitude can be mapped to length, size,

thickness, visual weight, color, and other parameters of a specific glyph design. Because direct flow glyphs are almost always designed to visually align with flow direction, a better understanding of how a design contributes to the accurate judgement of flow direction should yield insights that improve the perceptual effectiveness of 3D flow visualizations at a basic level.

Plain Lines

The plain lines design serves as a control or baseline condition with which to compare the other designs. As shown in Figure 5.1, the glyph is a plain white unlit line, which provides no depth cues except for the point at which it intersects its spherical head.

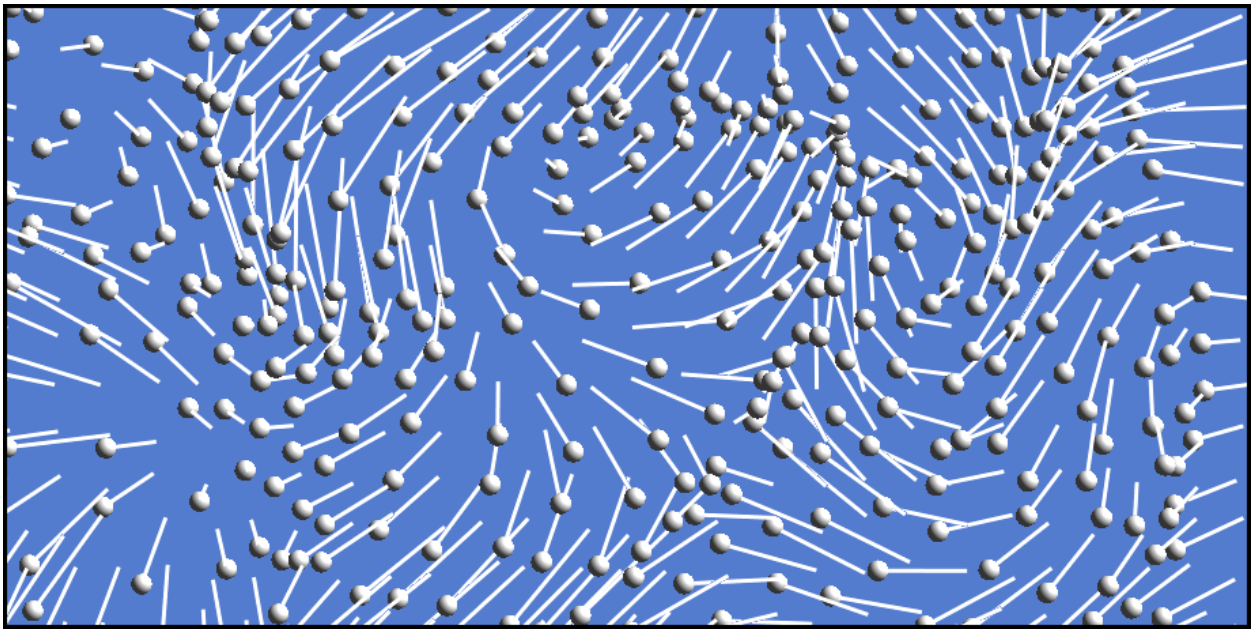

Figure 5.1: Example of the "Plain Lines" glyph design.

Illuminated Lines

The Illuminated Lines library [62] was used to augment the plain line design by adding shading as an additional visual cue. The original method, developed by Zöckler et al. [28], was extended by determining the shading for the line through a model that approximates the

total amount of light falling upon a cylinder with an infinitesimally-small diameter, which is then applied as the final color for the line. See Figure 5.2 for an example of this method. This condition will address hypothesis H5.

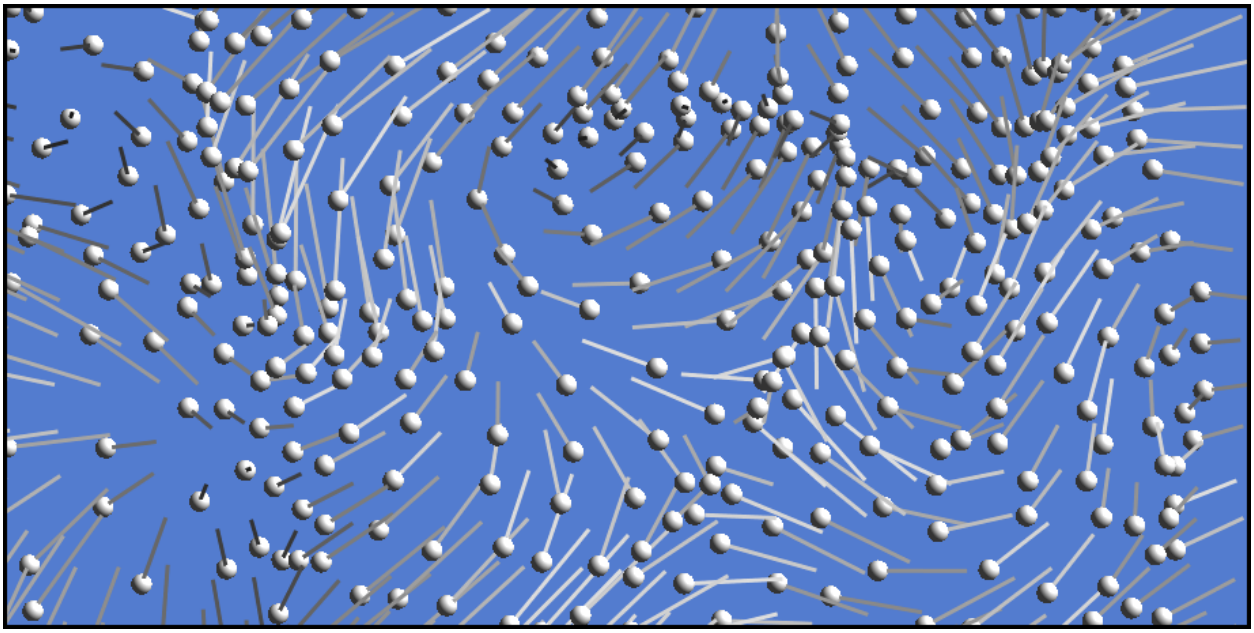

Figure 5.2: Example of the "Illuminated Lines" glyph design.

Shadowed Lines

This design is based on the "shadowed hedgehogs" technique developed by Klassen and Harrington [63], is a line-based design that removes all shading and instead uses shadows to provide direction and depth information; an ideal design for studying cast shadows in relation to the other depth cues. The glyphs, shown in Figure 5.3, are still rendered in a perspective 3D projection but are colored black with no shading, and are anchored to the cutting plane by either their heads or tails, depending on whether the flow is directed into or out of the cutting plane. This modification facilitates the association of a glyph with its shadow; the original design displaced the shadow plane behind the glyph plane, but this modification yielded a better design when applied to a cutting plane.

Additionally, pilot testing found that the blue background made it very difficult to see the

shadows. Since the glyphs themselves are entirely black, conditions using this glyph design were displayed on a white background to increase shadow contrast as much as possible. The enhancement that this modification offers outweighs the potential confounding effects of having a different background color from the other conditions.

To ensure regularity and eliminate glyph-glyph occlusions, grid jittering was not performed, and glyph lengths were normalized to $\frac{1}{2}$ glyph spacing. Each 'hedgehog' glyph casts a light grey shadow onto the cutting plane based on the light source, giving a cue that, together with the glyph itself, theoretically provides all the information necessary to judge its direction. This condition addresses hypothesis H6.

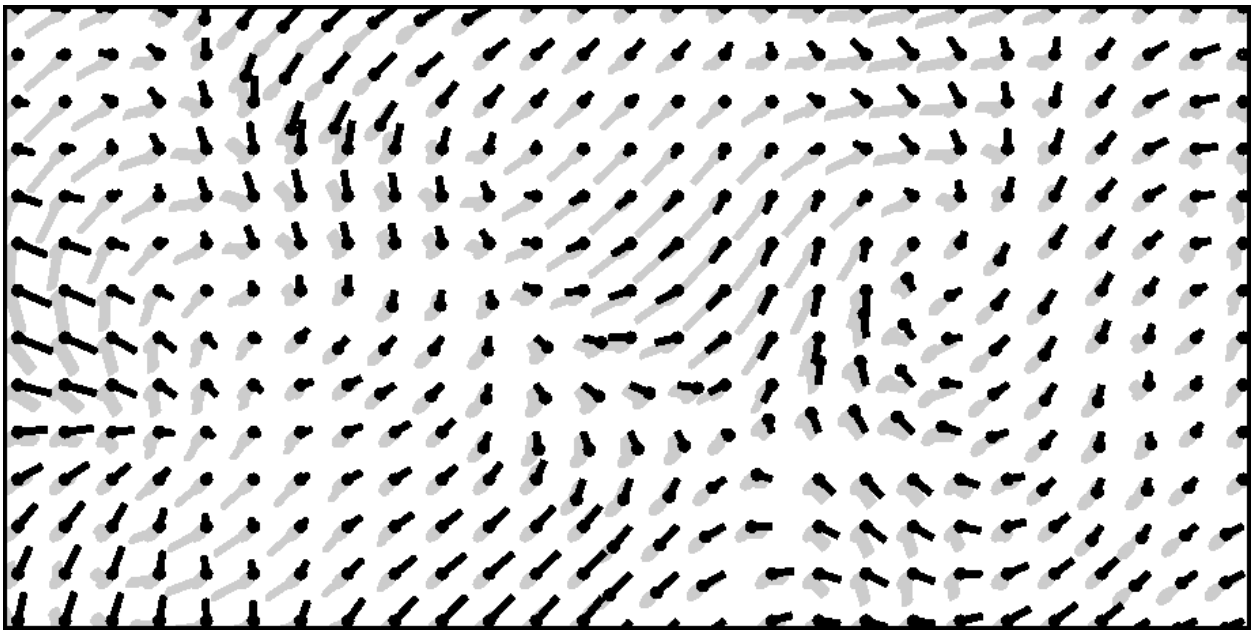

Figure 5.3: Example of the "Shadowed Lines" glyph design.

Plain Tubes

Tube-based glyphs have the advantage of encoding multiple visual cues in their design, including shading, occlusion, and perspective. When one tube occludes another, we understand the occluded tube to be behind the occluding one (relative to ourselves). Perspective in this case results in foreshortening, where the projection of the tube diameter decreases

with distance.

Another noteworthy theoretical benefit of tubes, which can be seen in Figure 5.4, is the appearance of the circular shape of the tail end. As a tube's orientation becomes more oblique with respect to the line of sight, the circular endcap deforms into an ellipse. These additional cues have the potential to increase the effectiveness of tube-based glyphs over plain lines. This condition is the baseline condition to help answer hypothesis H2.

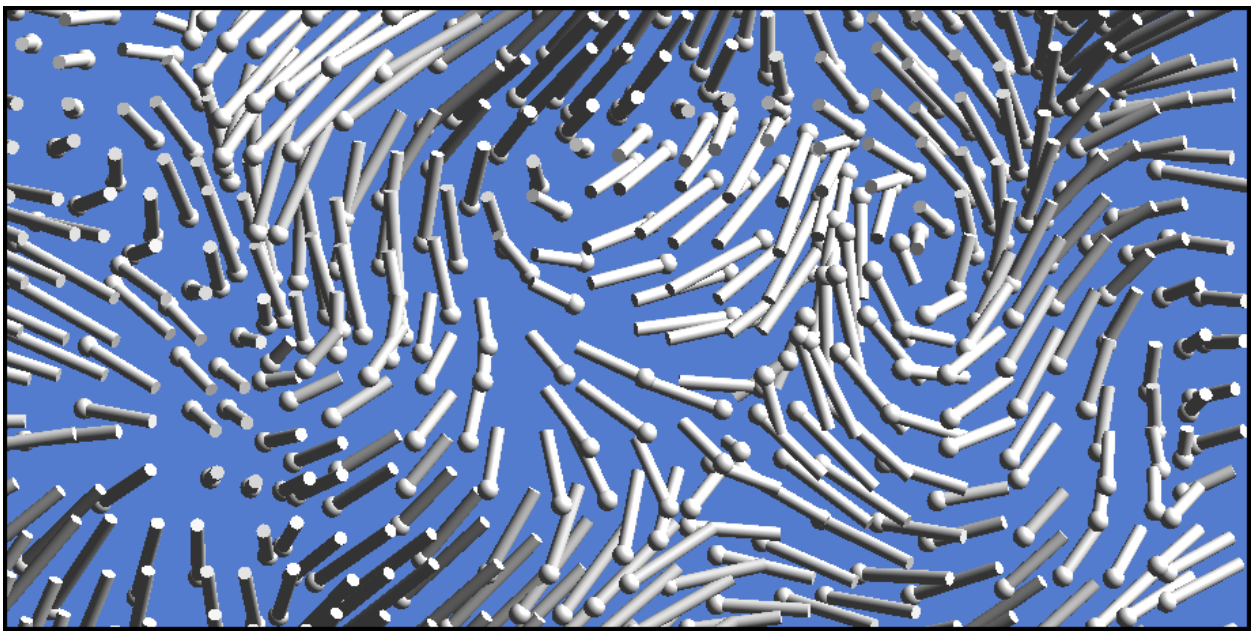

Figure 5.4: Example of the "Plain Tubes" glyph design.

Ringed Tubes

Ringed tubes (pictured in Figure 5.5) are plain tubes with the addition of a regular ringed texture. Each of these individual rings help with perceiving orientation by reinforcing perspective cues in the same way as the singular circular tail end of a plain tube. However, the addition of this texture does increase visual noise, which has the potential to hinder orientation perception. The rings also serve as a redundant encoding of magnitude – longer glyphs will contain more rings. This condition will help to answer hypotheses H1 and H2.

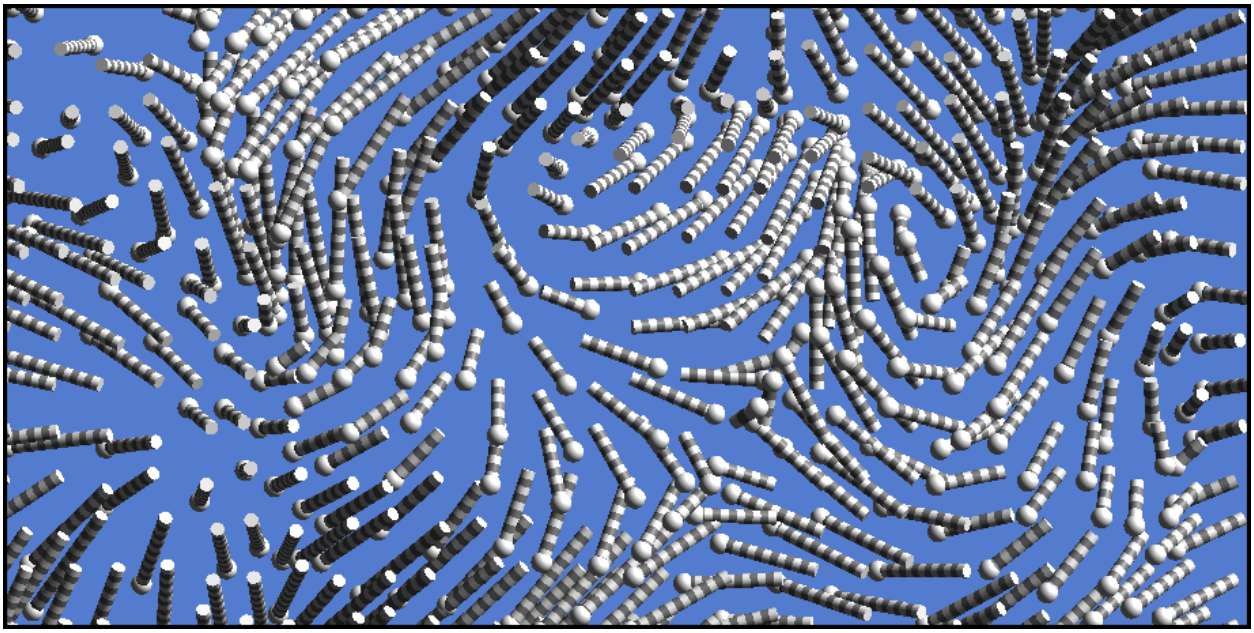

Figure 5.5: Example of the "Ringed Tubes" glyph design.

5.1.2 3D Glyph Diameter

For the tube-based designs, an additional parameter of interest is glyph diameter. Thinner diameters make shading and circular-distortion orientation cues more difficult to discern, but allow for more densely-seeded cutting planes before occluding effects become an issue. On the other hand, thicker diameters have a higher chance of occluding one another at sparser seeding densities, but enhance the shading and orientation cues, making them easier to distinguish. This experiment tests tube glyph diameters of 0.5mm, 1mm, and 2mm, which were determined through pilot study evaluations and can be seen in Figure 5.6. This parameter addresses hypothesis H4.

5.1.3 Seeding Density

The density of seeds on the cutting plane is also varied to understand its effects for each glyph design and find points of diminishing returns. Too few seeds risks missing finer flow patterns. Too many seeds leads very quickly to problems with occlusion and overwhelming visual noise.

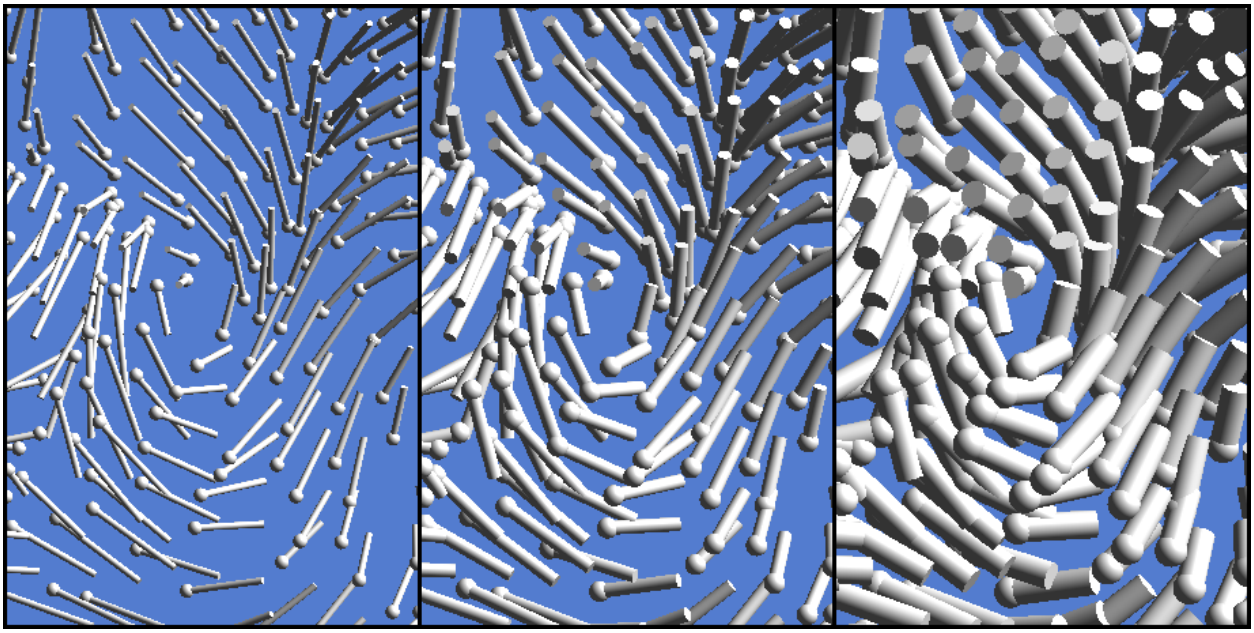

Figure 5.6: An example of the three different tube diameters (0.5mm, 1mm, and 2mm, respectively), rendered with "Plain Tubes" glyphs at a density of 3 seeds/cm.

The three seeding densities evaluated in this experiment use intervals of 2, 3, and 4 seeds per centimeter, which correspond to 5mm, 3.33mm, and 2.5mm Cartesian grids, respectively. These conditions were determined from the pilot study to be adequate parameters for representing the usable ranges of possible seeding densities and are shown in Figure 5.7. This parameter addresses hypothesis H3.

5.1.4 Cutting Plane Design

There does not seem to be a consensus on the best lighting style to use; Forsberg et al. [41] and Mallo et al. [62] suggest an additional headlight to further disambiguate lit geometry, but Penney et al. [59] found no performance difference for static scenes between ray-traced and standard OpenGL multi-lighting scenarios for depth perception tasks.

The stimuli renderings in this experiment use a single, directional light source, keeping lighting parameters consistent throughout all experimental conditions. The lighting direction originates from above the right shoulder of the viewer at infinity ((-1, -1, -1), assuming right-handed coordinate system, xy-plane on the screen, and z-axis towards viewer). Participants

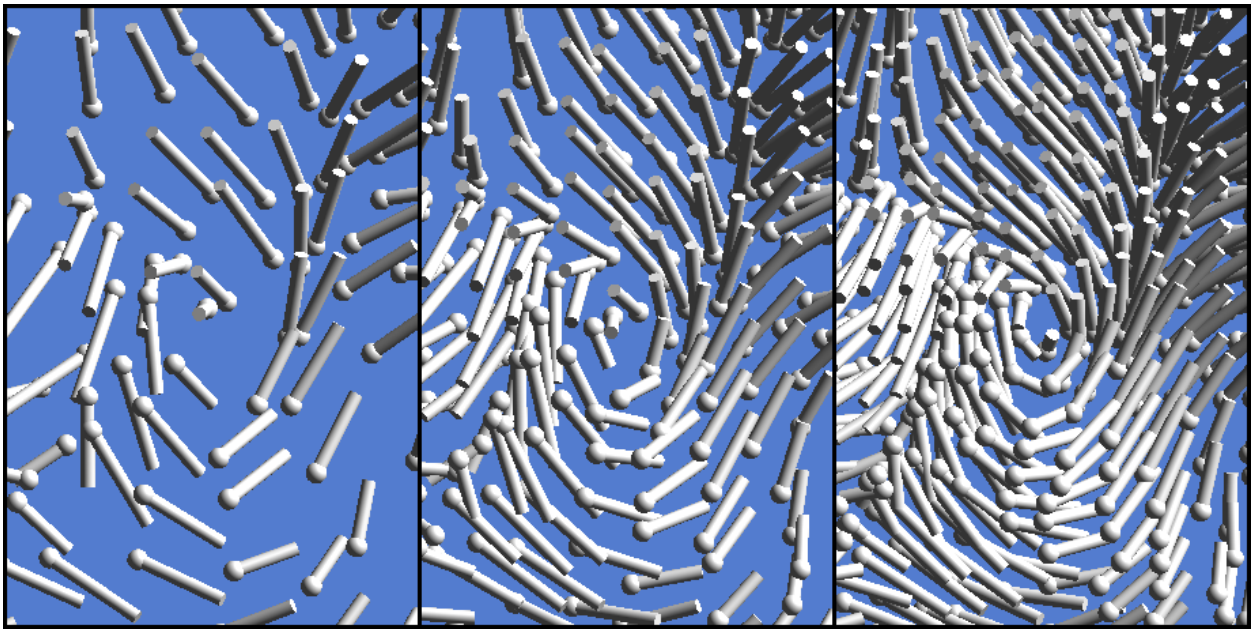

Figure 5.7: An example of the three different seeding density conditions (2, 3, and 4 seeds/cm, respectively), rendered using "Plain Tubes" glyphs with a 1mm diameter.

are made explicitly aware of the lighting direction during the training phase to help with their judgements.

Glyphs are displayed atop a light blue background to maintain an effective level of contrast for the stimuli, except for the Shadowed Lines design, which uses a white background as described in Section 5.1.1. This background color was tuned during a pilot study; other colors were found to be too distracting or fatiguing, while white/grayscale backgrounds made it more difficult to use shading cues. A standard Blinn-Phong shading model is used to shade the spherical glyph heads and tube-based glyphs with the following lighting (L) and material (k) parameters: $L_a = 0.2$, $L_d = 1.0$, $L_s = 0.5$; $k_a = 0.5 * k_d$, $k_d = 1.0$, $k_s = 0.5$, $k_{shininess} = 10$. Though an orthographic projection of the stimuli could be used, the stimuli are rendered with a perspective projection from a frustum modelled after the physical viewing conditions of the study to match reality as closely as possible.

Seeds are evenly distributed across the cutting planes, spaced on a regular grid based on the desired density value. Seeds are then slightly jittered by moving them randomly within $\frac{1}{4}$ of the spacing value before generating the glyph. This jittering is done as a precaution,

because regular spacing can potentially induce meaningless patterns with some glyph designs [32,36]. Furthermore, adding jitter is unlikely to negatively affect results; Laidlaw et al. [21] found no performance difference between regular and jittered grids for 2D flow glyphs.

A random location on the cutting plane is then selected as the target where participants must estimate the direction of the flow. Error due to perspective projection distortion is mitigated by ensuring targets are selected from within the center half of the slice, where this error is much less pronounced.

5.1.5 Sampled Data

The cutting plane slices are sampled from randomly-generated trivariate vector fields. These artificial vector fields are created by summing multiple random Gabor functions oriented along the three planes formed by each principal axis pair. After sampling, the slice is rendered using the current experimental condition and is displayed orthogonal to the participant's line of sight, filling the display screen.

In the experimental flow fields, target vector magnitudes averaged 6.78 and ranged from 0.27 to 17.7 with a standard deviation of 2.56.

5.1.6 Task

Participants are first screened with a vision test to verify at least 20/20 binocular visual acuity prior to participation. Additionally, each participant completes a training session to ensure they understand the experimental task and how to properly use the physical probe device to input perceived orientations.

At the beginning of training sessions, a virtual scale model of the probe is displayed, anchored at the center of the display and with its direction mirroring that of the physical probe. A second virtual target probe is displayed in a similar fashion to serve as the goal. Participants are guided through four different levels of visual aids to align the virtual probe to the goal and become accustomed to using the probe. The first level of visual aid provides

color cues on the goal to indicate alignment between the probe and the goal; the second level removes the color cues; the third level removes the virtual probe but restores the color cues; and the fourth removes both the virtual probe and color cues, displaying only the goal probe. Participants are provided with feedback on the amount of error for each practice trial, and at least three practice trials using each level of visual aid must be completed with an error of approximately 10° or less. Participants are then shown examples of each stimulus condition and asked to orient the probe with the flow direction at the target marker while being given verbal feedback on their angular error with the target.

During the experiment, participants are shown flow visualizations (as described in the next section) and asked to estimate the direction of the flow through the cutting plane at a point indicated by a target marker (shown in red in Figure 5.8). Participants indicate this by orienting the physical probe in front of them, near their line of sight with the display (so that there is minimal perspective distortion between the stimuli and the probe, similar to Figure 5.10). When participants are satisfied with the probe's direction, they press a key to record the direction and advance to the next trial. The angle of the input probe and the angle of the target glyph stimulus are both recorded, along with elapsed decision time.

5.1.7 Apparatus and Display

The experimental apparatus was constructed using a high-definition digital monitor to display the stimuli and a three-dimensional direction tracking system to record participant responses. These elements were respectively inset and mounted to a rigid piece of black foamboard in order to provide the participant with a clean and distraction-free interface, shown in Figure 5.9.

Despite the increasing availability of high-quality consumer 3D displays, they still have not become a commonplace item in visualization workstations. In light of this, and because of the numerous studies already confirming the usefulness of stereoscopy in 3D visualization perception, this study was limited to a single non-stereoscopic viewing condition. This also

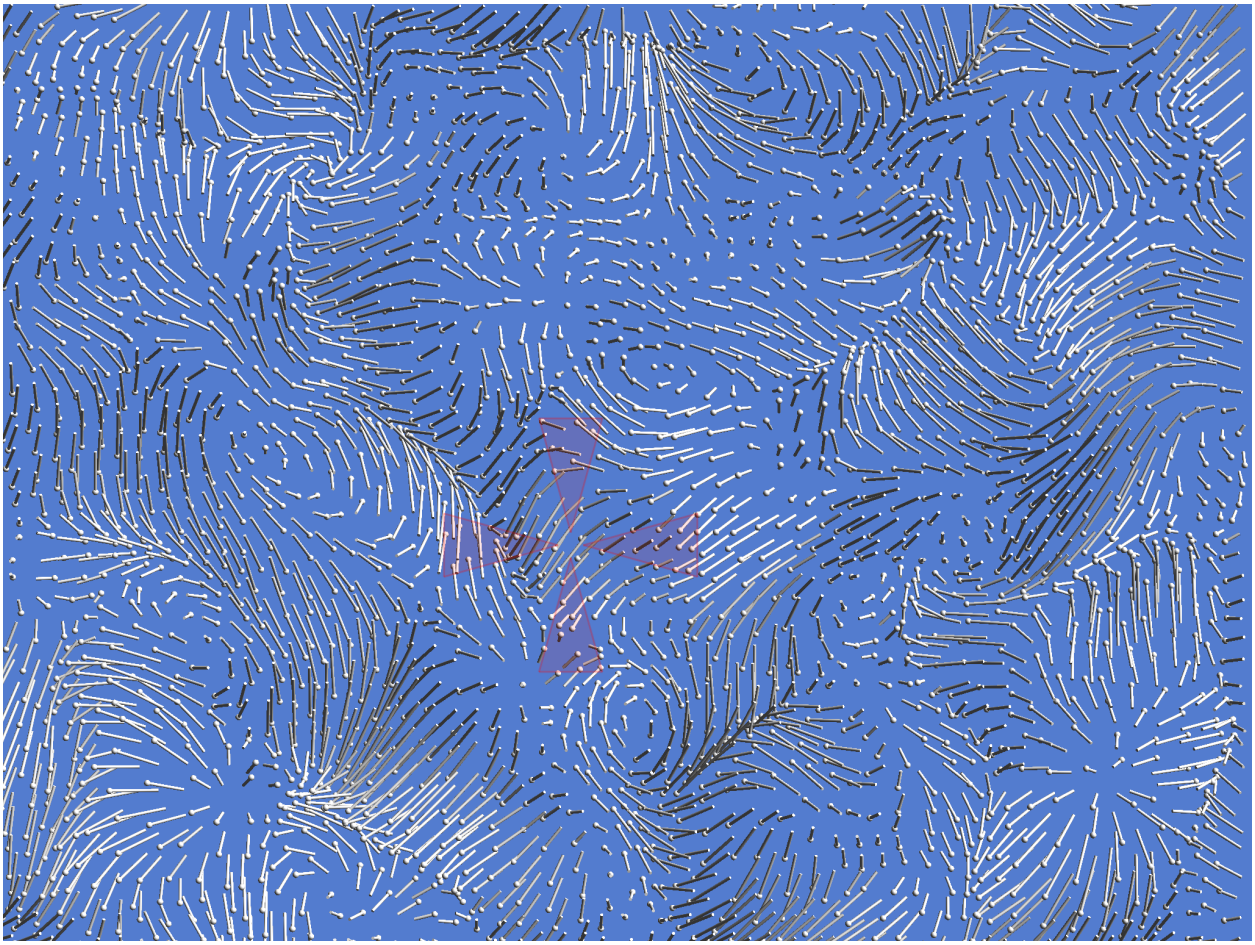

Figure 5.8: A screen capture showing the red target marker.

allows the experiment to be expanded to cover additional glyph designs, which is the primary focus of this initial experiment.

An LG LP097QX1 9.7" 60Hz display panel is used due to its high pixel density relative to its size; the 2048 x 1536 pixel across its 197.1mm x 147.8mm screen yields a pixel size of approximately 0.096mm and pixel density of 10.4px/mm. With participants seated approximately 57cm from the display, each pixel subtends a visual angle of 34.7 arc seconds, providing about 104 pixels per degree of visual angle. This approximates the size of foveal receptors in the eye [64], allowing for an optimal display resolution congruent with that of the human visual system as well as the finest resolvable non-stereoscopic point and grating acuities (approximately 60 seconds of arc) [36].

Figure 5.9: The experimental apparatus, consisting of a 25cm diagonal high-resolution "Retina" display and an electromagnetically tracked hand-held probe for inputting perceived direction.

Direction Tracking

Perceived direction is input using an electromagnetic spatial tracker attached to a physical hand-held probe (shown in Figure 5.10). The system collects data at 60Hz, with a resolution of 0.1°, and a root mean square error of 0.75°.

The electromagnetic tracker is attached to the tip of a physical probe, consisting of a 2cm x 2cm x 10cm block of rigid foam, textured with a $1cm^2$ grid pattern. This physical shape and texture combination is designed to provide strong linear perspective depth cues, while being visually dissimilar to any of the rendering methods, thereby eliminating the possibility of the experimental task devolving into one-to-one matching between physical probe and on-screen stimuli.

Figure 5.10: The physical hand-held probe device that participants used to indicate perceived direction.

The tracker system is physically calibrated to the experimental setup, and electromagnetic field interference from the display and other nearby equipment was found to contribute no more than approximately 1° error on average.

5.1.8 Experimental Design

The experimental model is a fully-crossed 9×3 within-participants randomized complete block design that evaluates three primary cutting plane parameters: nine different glyph conditions comprised of five different glyph designs and three diameter values for the two 3D tube-based glyph designs; and three seeding densities of these glyph conditions on the cutting plane. The experimental conditions are listed in Table 5.1. The order of conditions

is randomized per-block and per-participant to help control for ordering effects.

Table 5.1: Experiment 1 Conditions

Parameter	Values
Glyph Design	Plain Lines, Illuminated Lines, Shadowed Lines, Plain Tubes*, Ringed Tubes*
*Tube Glyph Diameters (mm)	0.5, 1, 2
Glyph Density (seeds/cm)	2, 3, 4
Total	**27 unique conditions**

Every participant saw each of the 27 rendering conditions a total of 10 times, resulting in 270 trials per participant. For each participant, trials were presented in two randomized blocks, with each block containing five sequential replicates of the randomized conditions.

Analysis of the data is performed in two stages. The first stage examines overall effects of the nine glyph conditions and three seeding densities, plus the possible interaction between these on total angular errors, decomposed angular errors, and decision time. Post-hoc analysis with Tukey's Honestly Significant Difference (HSD) test at the $p < 0.05$ confidence level conservatively separates the glyph conditions and seeding densities into statistically-distinct performance groups. The second analysis stage examines only the tube-based conditions for significant main effects of glyph design, seeding density, and glyph thickness on total angular error, decomposed angular error, and decision time. The two-way interactions between these terms is also analysed, and a Student's t-test reveals any significant difference between the plain and ringed tube glyphs.

Total Angular Error

The total angular error is the angle subtended by the probe direction and the flow vector at the target point, as given in Equation 5.1. A perfectly aligned probe and target yields a total angular error of 0°, and a probe and target aligned but facing in opposite directions result in a 180° total angular error.

$$\theta = \frac{180}{\pi} \arccos(\hat{t} \cdot \hat{p}) \tag{5.1}$$

In this equation, θ is angular error, $\hat{t}$ is the normalized target vector, and $\hat{p}$ is the normalized direction probe vector input by the participant.

Absolute Depth Error

Depth error is the absolute difference between the angle from the cutting plane to the target vector and the angle from the cutting plane to the probe vector. This depth error provides a measure of accuracy for depth perception. To assess the amount of error due to misperception of depth, the total error is decomposed to find the depth component of the error according to Equation 5.2:

$$\theta_D = \frac{180}{\pi} |\arccos(\hat{t}_z) - \arccos(\hat{p}_z)| \tag{5.2}$$

In the above equation, θ_D is absolute depth error, $\hat{t}_z$ is the depth component (z) of the normalized target vector, and $\hat{p}_z$ is the depth component (z) of the normalized direction probe vector input by the participant.

Weighted Projection Error

To assess non-depth-cue-related perception of vector direction, the weighted projection component of the error was also examined, i.e., the angle subtended by the probe and target vectors when projected onto the cutting plane. However, when target vectors are nearly parallel with the viewing direction, even small overall angular errors can generate very big projection errors. The solution is to apply a weighting factor, so that parallel target vectors (which come straight at or away from viewers) receive the lowest weight, while vectors orthogonal to the viewing direction (i.e. along the cutting plane) are assigned the highest weight. Equation 5.3 describes the weighted projection error:

$$\theta_P = \frac{180}{\pi} \arccos\left(\frac{t_{xy} \cdot p_{xy}}{||t_{xy}|| \cdot ||p_{xy}||}\right) \cdot (1 - |t_z|) \tag{5.3}$$

In the above equation, θ_P is weighted projection error, $\hat{t}_x y$ is the projection of the normalized target vector onto the orthogonal (xy) plane, and $\hat{p}_x y$ is the projection of the normalized direction probe vector input by the participant onto the orthogonal (xy) plane.

5.1.9 Analysis

Analysis of the data was performed using JMP statistical software in two stages. The first stage looked at the overall effects of the nine glyph conditions and three seeding densities, plus the possible interaction between these on total angular errors, decomposed angular errors, and decision time. Post-hoc analysis with Tukey's Honestly Significant Difference (HSD) test at the $p < 0.05$ confidence level was used to more conservatively group the glyph conditions and seeding densities.

The second analysis stage examined only the tube-based conditions for significant main effects of glyph design, seeding density, and glyph thickness on total angular error, decomposed angular error, and decision time. Two-way interactions between these terms were also analysed, and a Student's t-test was used to compare the plain and ringed tube glyphs for any significant differences.

Before running the statistical analyses, the raw data were transformed using the Johnson SL distribution function [65] to normalize the distributions of the variables in order to satisfy the prerequisites of the chosen analytical methods.

5.1.10 Results

A total of 14 adult participants participated in the study. Of these, 6 were female and 8 were male. Each participant was paid $27 for their participation, and each session lasted approximately 75 minutes: about 15 minutes for training and an average of 43 minutes viewing stimuli, the remaining 17 minutes representing between-condition and between-block

breaks to help control for participant fatigue.

Trials where participants took an inordinate amount of time to make a decision were filtered using Leys et al.'s approach [65] to finding outlier decision times. A conservative cutoff value of 26s over the mean time of 6.88s removed 2% of trials, leaving 3792 trials total.

It should be noted that F-tests for significant effects use the Kenward-Rogers approximation [66], which reduces bias from small sample sizes and responds well to removed outlier samples. Because participants' overall performance varied widely, the 14 participants were evenly divided into two performance groups: one containing the top 7 performers and one containing the bottom 7 performers, based on their overall mean total angular error (28.9 deg and 36.6 deg, respectively). A highly significant main effect of group membership ($F(1, 377.1) = 39.65$, $p < 0.0001$) was found on total angular error (see 5.1.8); the following results are thus presented both for the entire participant pool ($n = 14$), as well as split by each performance group ($n = 7$) for cases where the split reveals more nuanced effects.

Total Angular Error

As shown in Figure 5.11, for total angular error, glyph condition was a highly significant factor, $F(8, 351.9) = 26.34$, $p < 0.0001$. The Tukey's HSD test split the glyph designs into five significantly different but overlapping ranks at the $p < 0.05$ confidence level: 2mm plain tubes came out on top, followed by the other tube-based conditions, except for 0.5mm ringed tubes, whose performance was not significantly different from the shadowed lines group. The two worst performing glyph conditions were illuminated lines and plain lines.

While the low-performance participant group demonstrates a stratification and overlap of glyph condition rankings similar to the overall data set, the Tukey HSD analysis found only three distinct and completely separate groups within the high-performance group: all of the tube-based conditions in the top rank, followed by shadowed and illuminated lines, then plain lines.

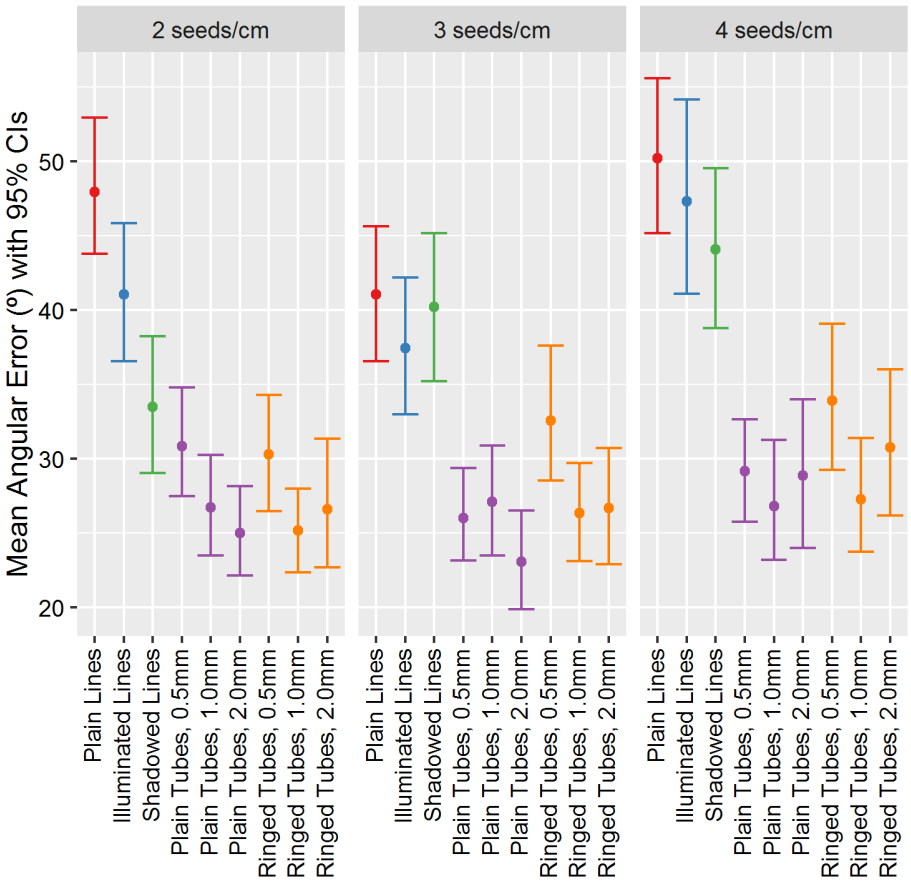

Figure 5.11: Mean total angular error with bootstrapped 95% confidence intervals for each glyph condition and seeding density.

Overall, seeding density was a significant factor $F(2, 351.9) = 3.59$, $p = 0.0285$. The Tukey HSD analysis identified 3 seeds/cm as being significantly better than 4 seeds/cm at the $p < 0.05$ confidence level, but no significant differences between the 2 seeds/cm level.

There was no significant interaction found between glyph condition and density for the overall participant pool, nor for the individual performance groups.

For tube based glyphs, tube diameter had a highly significant effect ($F(2, 239) = 6.85$, $p = 0.0013$) on total angular error, and the Tukey HSD test found that 1 mm and 2 mm diameters are significantly better than 0.5 mm diameters at the $p < 0.05$ confidence level. Interestingly, breaking this result down by performance group reveals that the middle-sized 1 mm tube diameter was only significantly better than the 0.5 mm diameter for high-performing

participants; i.e. low-performing participants did not perform as well with the 1 mm tube diameter conditions. There was no significant main effect of seeding density nor glyph technique, and none of the interaction terms reached significance.

There was some concern that the experimental design may provide the two tube-based glyph techniques with an advantage by being seen 3 times as frequently as the other line-based glyph techniques, due to the three different diameter conditions that only apply to tube-based techniques. To assess this potential bias, the first 15 trials from each glyph technique (regardless of tube diameter) for each participant were extracted to enable a balanced and unbiased analysis. An ANOVA of the total angular error on this data subset revealed a highly significant effect of glyph condition ($F(4, 194) = 12.10$, $p < 0.0001$) but not seeding density. The glyph techniques were grouped into four overlapping but significantly different groups at the $p < 0.05$ confidence level: plain tubes first, followed by ringed tubes, shadowed lines, illuminated lines, and plain lines in the worst group. As these rankings follow the same pattern as those from the larger data set, it can be concluded that any biasing effect introduced by the relative frequency of tube-based techniques is overpowered by the larger effects of technique difference.

Absolute Depth Error

Absolute depth errors for the different glyph conditions and seeding densities are shown in Figure 5.12. Again, there was a highly significant effect ($F(8, 350.2) = 31.71$, $p < 0.0001$) of glyph technique, with performance of the different techniques being almost identical to the observations of total angular error in Section 5.1.10. The Tukey HSD test grouped the conditions into four main groups at the $p < 0.05$ confidence level: 2 mm tubes are better than 0.5 mm tubes, but 1 mm tubes are not significantly different from either; the next group contains shadowed lines, and illuminated lines lies between that group and the worst group which contains only plain lines.

Seeding density also had a significant ($F(2, 350.2) = 3.66$, $p = 0.0267$) effect on depth

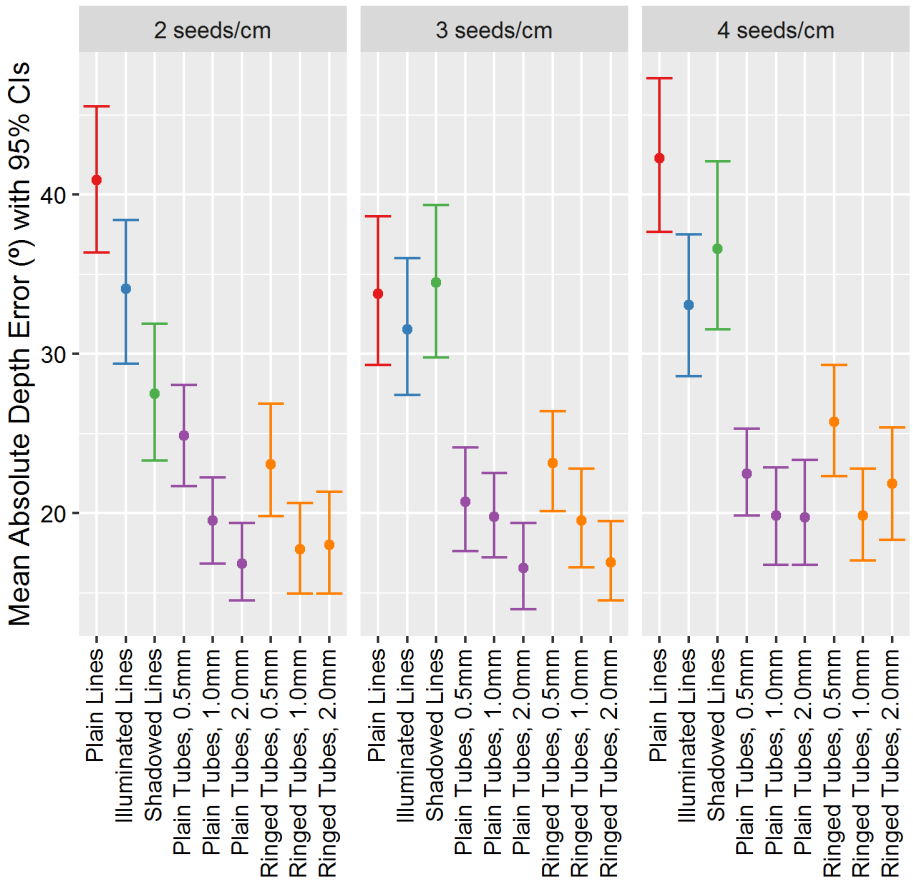

Figure 5.12: Absolute depth error with bootstrapped 95% confidence intervals for each glyph condition and seeding density.

errors, and these were grouped in the same manner as they were for total angular error in Section 5.1.10 by the Tukey HSD test at a $p < 0.05$ confidence level.

For tube based glyphs, again only the diameter had a significant effect ($F(2, 237.5) = 13.99$, $p < 0.0001$) on depth error. The diameter groupings found by a Tukey HSD analysis at the $p < 0.05$ confidence level are the same as those found for total angular error in Section 5.1.10. An interesting finding is that, while the high-performance participant group showed the same significant effects and groupings as the combined data, the low-performance group had an additional significant main effect of seeding density, $F(2, 111.8) = 3.47$, $p = 0.0344$. A Tukey HSD analysis found that 3 seeds/cm are significantly better than 4 seeds/cm at the $p < 0.05$ confidence level, just like what was found for total angular error.

Weighted Projection Error

As shown in Figure 5.13, for projection errors, the performance differences between different glyph conditions and seeding densities is much less apparent. Considering all participants, glyph technique had a significant effect ($F(8, 350.9) = 2.56$, $p = 0.0101$), but only in the case of 0.5 mm plain tubes performing better than shadowed lines. When broken down by performance group, this advantage of 0.5 mm plain tubes over shadowed lines was only observed in high-performing participants; low-performing participants showed no significant difference between any glyph techniques.

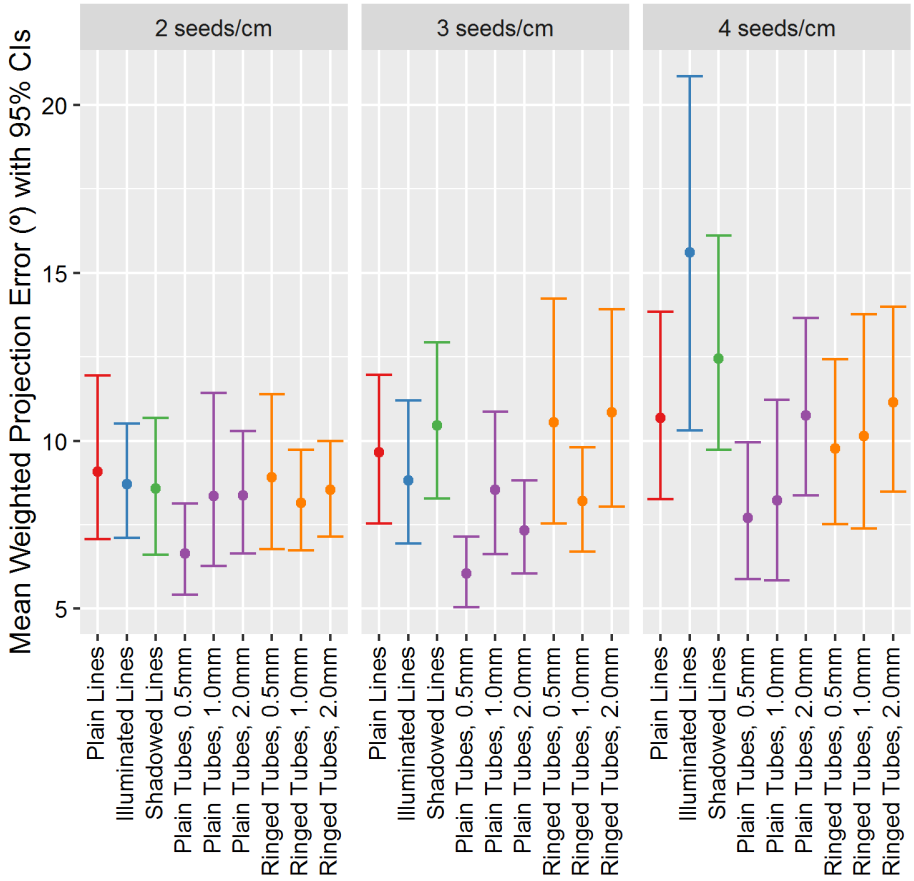

Figure 5.13: Mean Weighted projection error with bootstrapped 95% confidence intervals for each glyph condition and seeding density.

Seeding density had no significant main effect on projection error, nor did the interaction

between seeding density and glyph.

For tube-based glyphs, neither tube diameter nor seeding density had a significant effect on projection error, but glyph technique was highly significant ($F(1, 238) = 7.27$, $p = 0.0075$). The Student's t-test found plain tubes to be significantly better than ringed tubes at the $p < 0.05$ confidence level.

Response Times

The analysis of the data did not show any significant main effects or interactions on response times for either the overall data set or within the two performance groups.

Feedback

A post-study questionnaire was given to each participant to collect impressions about the different glyph techniques. Participants were asked to rate each technique on a scale from one (worst) to five (best) in response to four questions. A reproduction of the questionnaire can be found in Appendix C.

When asked whether the glyph technique was intuitive and easy to understand, participants rated plain tubes highest ($avg = 4.1$, $SD = 0.5$), followed by ringed tubes ($avg = 3.8$, $SD = 0.8$), illuminated lines ($avg = 3.4$, $SD = 0.5$), plain lines ($avg = 2.4$, $SD = 0.8$), and shadowed lines ($avg = 2.3$, $SD = 1.2$).

When asked whether the participant thought they got better and faster using a glyph technique over the course of the study, participants rated plain tubes highest ($avg = 4.2$, $SD = 0.4$), followed by ringed tubes ($avg = 4.0$, $SD = 0.6$), illuminated lines ($avg = 3.9$, $SD = 0.7$), plain lines ($avg = 3.6$, $SD = 0.6$), and shadowed lines ($avg = 3.1$, $SD = 1.3$).

When asked whether the glyph technique was aesthetically pleasing and enjoyable to look at, participants rated plain tubes highest ($avg = 3.8$, $SD = 1.0$), followed by ringed tubes ($avg = 3.6$, $SD = 1.1$), illuminated lines ($avg = 3.4$, $SD = 0.9$), plain lines ($avg = 3.0$, $SD = 1.0$), and shadowed lines ($avg = 2.4$, $SD = 1.2$).

When asked whether it took a long time to make a decision using the glyph technique, participants rated plain tubes and ringed tubes fastest ($avg = 2.6$, $SD = 0.9$; $avg = 2.6$, $SD = 1.1$; respectively), illuminated lines ($avg = 3.0$, $SD = 0.6$), plain lines ($avg = 3.4$, $SD = 0.7$), and shadowed lines ($avg = 3.9$, $SD = 1.4$) slowest.

5.2 Experiment 2: Stereoscopic & Motion-Based Cues

In general, the method and experimental design of this evaluation is very similar to Experiment 1. This experiment examines the value of stereo and structure-from-motion cues for visualizing 3D flow through cutting planes. The prior experiment investigated direct, straight glyphs, which only reveal the direction of flow at the point of plane intersection. This experiment features the more frequently used integral geometric glyphs: curved, variable length streamlines that can reveal flow patterns both on and in the vicinity of the plane, as well as depict the magnitude of the flow. The prior experiment used static, pictorial images; this experiment examines the use of stereoscopic viewing, structure-from-motion, and animation to convey the direction of flow, and evaluate whether this animation interferes with depth perception.

5.2.1 Glyph Designs

Five different integral glyph designs were selected for investigation based on the depth and directional cues they incorporate.

Static Streamlets - The line-based streamlet heads are white and fade away to medium gray. See Figure 5.14. This design addresses hypothesis H9.

Animated Streamlets - This is an animated variant of the static streamlet design. A white streamlet head is spawned at the beginning of a streamline, and travels along its path for 1s. Its fading tail length is half the streamline length. Streamlet spawning was staggered to avoid the visual distraction of the entire streamlet field appearing or fading at the same time. See Figure 5.15. This design addresses hypotheses H9, H10, and H11.

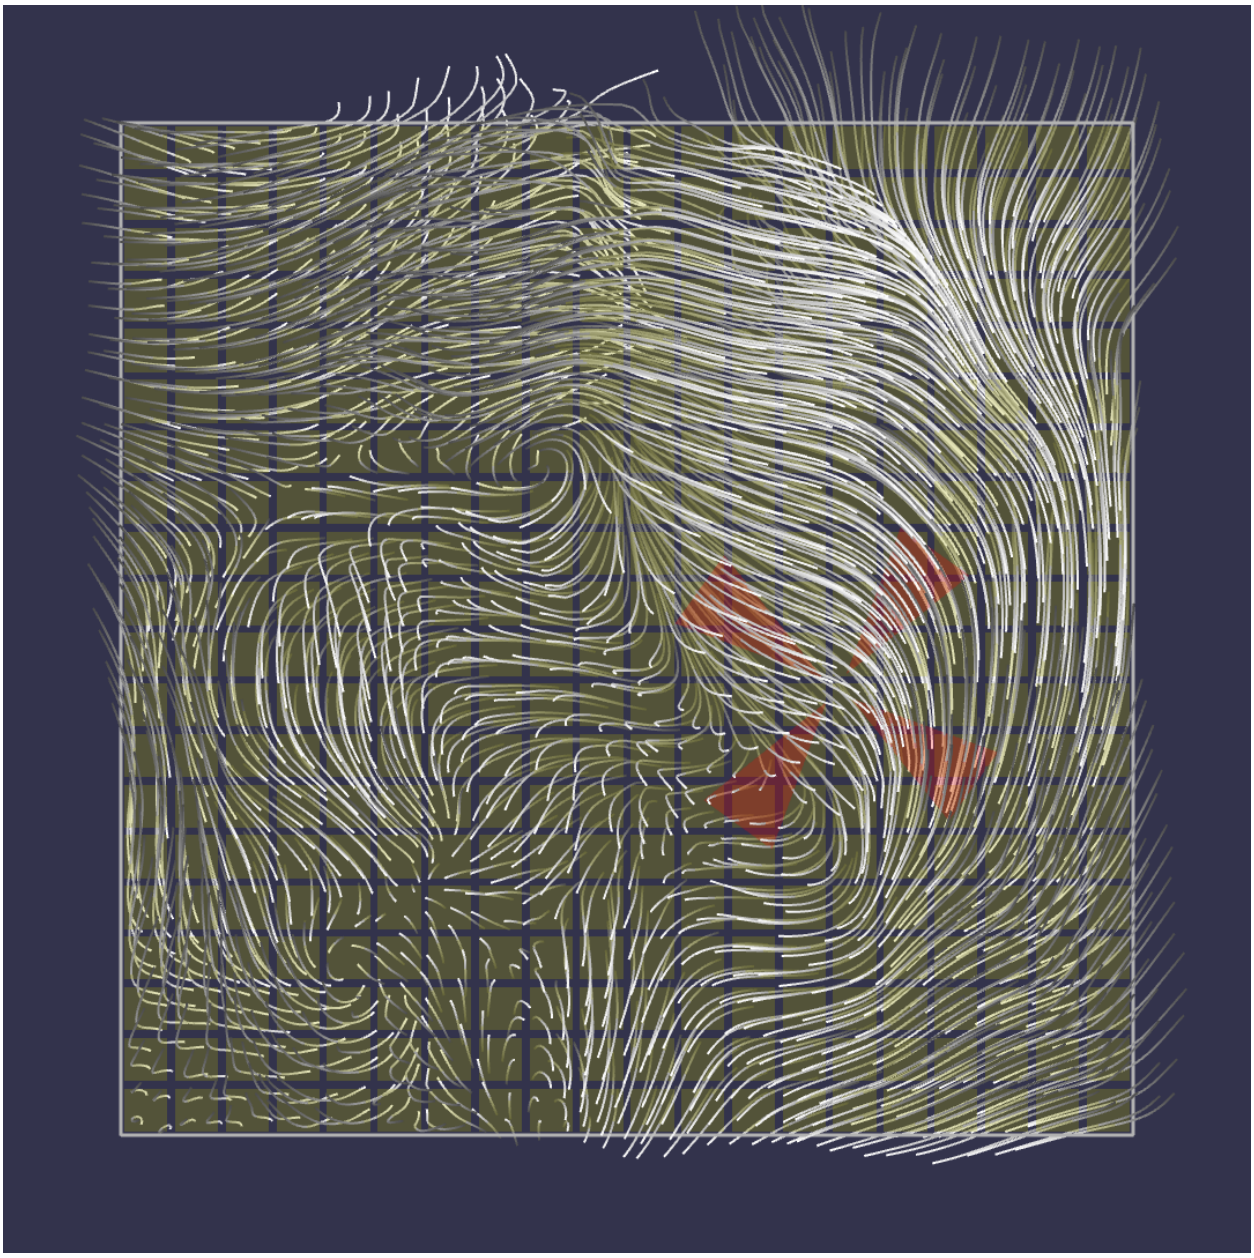

Figure 5.14: The static streamlets glyph design.

Static Streamtubes - The radius of each streamtube is approximately 2.3mm. Each streamtube is textured with five gradient bands, which go from dark to light to indicate direction. See Figure 5.16. This design addresses hypothesis H2.

Animated Streamtubes - These share the same features as the streamtubes, but the texture gradient bands are animated at a rate of 0.5s per gradient. Both the texture gradient and animation convey direction. This design addresses hypotheses H2 and H10.

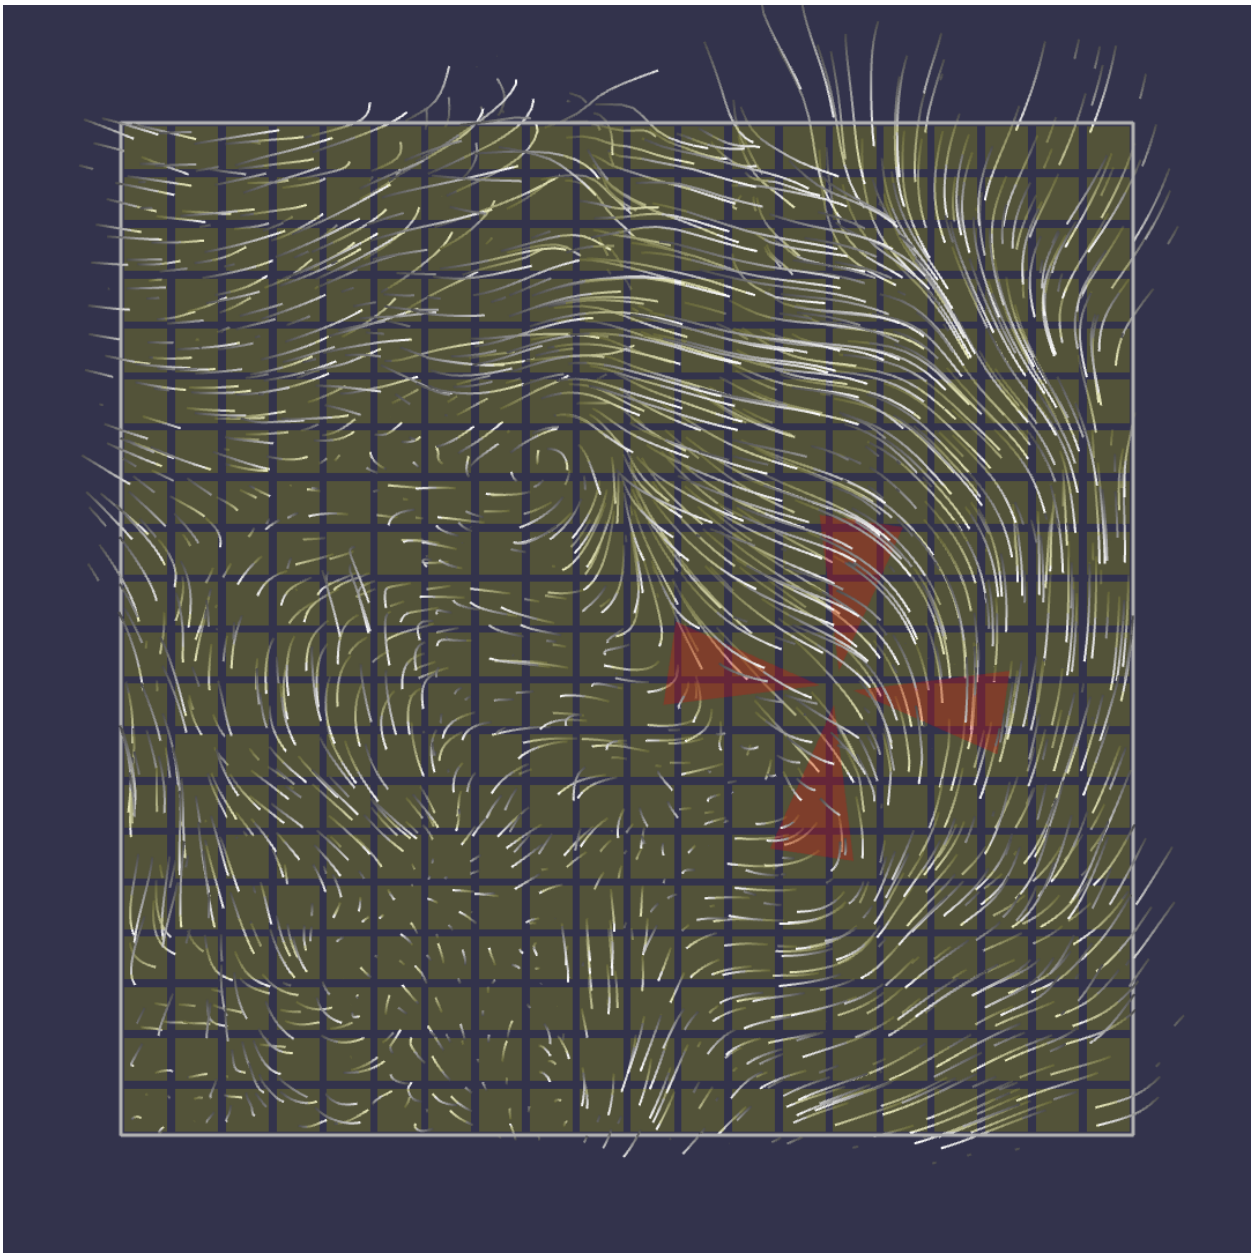

Figure 5.15: The animated streamlets glyph design.

Streamcones - Base radius is approximately 3.1mm. They are textured with five pairs of dark and light alternating bands. The direction is implied by shape, going from base to tip like a 3D arrowhead. See Figure 5.17. This design addresses hypotheses H2 and H4.

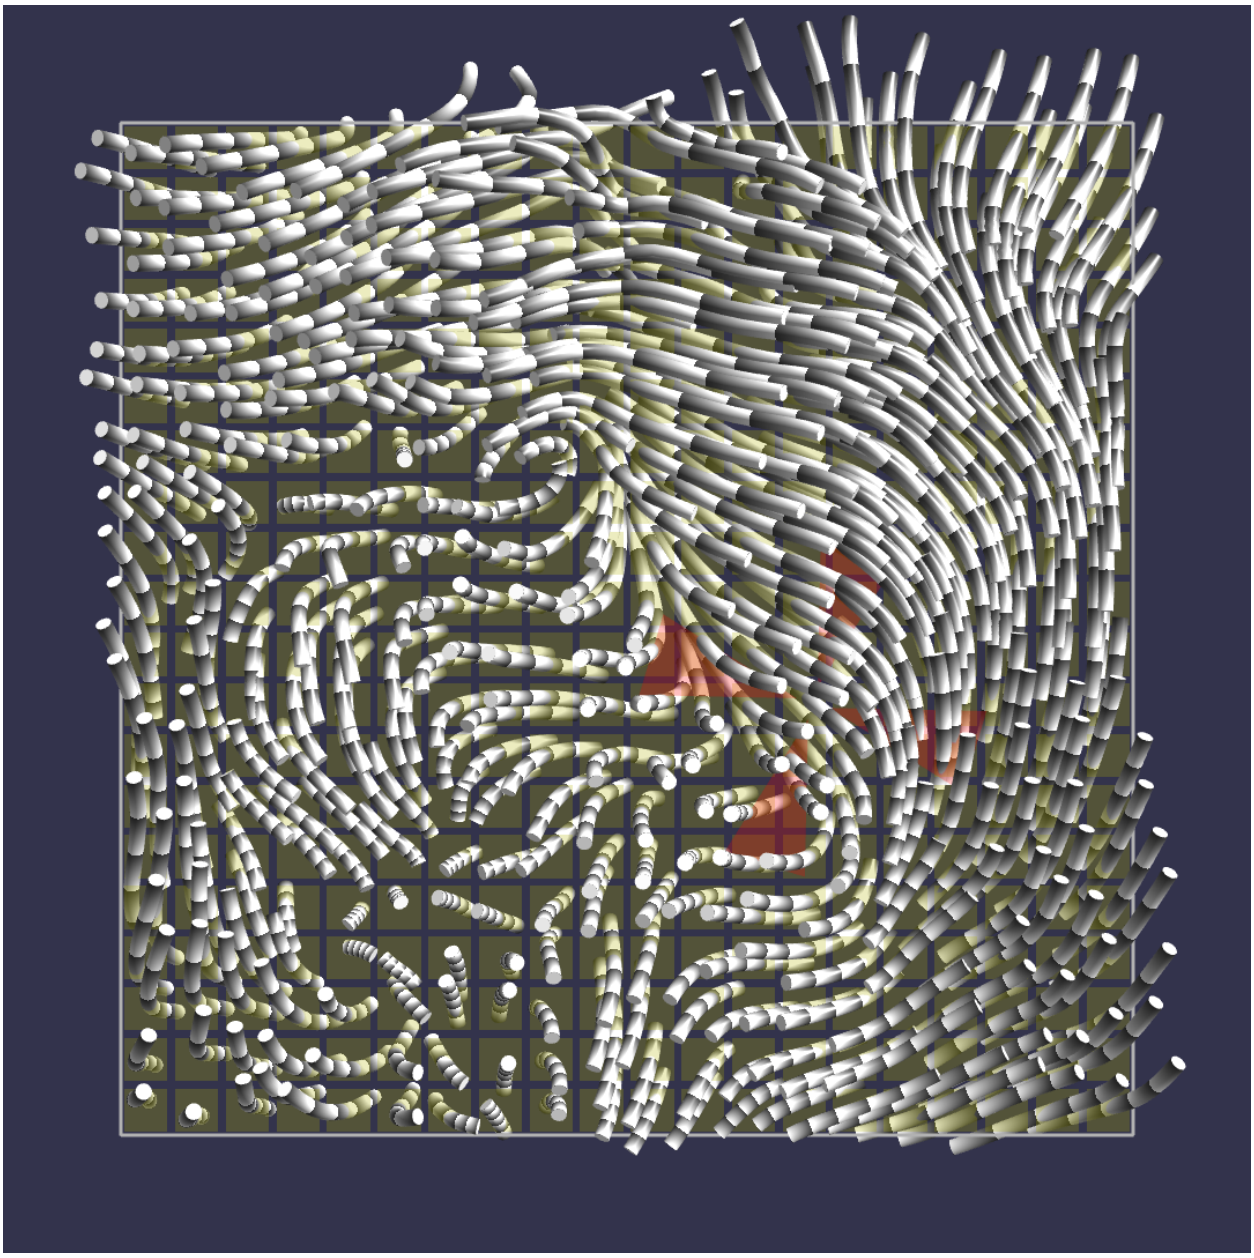

Figure 5.16: The streamtube glyph design with a gradient ring texture. Flow direction is from dark to light.

5.2.2 Stereoscopic Viewing

The integral glyph designs are presented both with and without stereoscopy using the 3D monitor and active shutter glasses described in Section 5.2.5. Each participant's interpupillary distance (IPD) is measured and the separation of the virtual cameras rendering the stereoscopic content is set to match the IPD. This helps to account for any perceptual

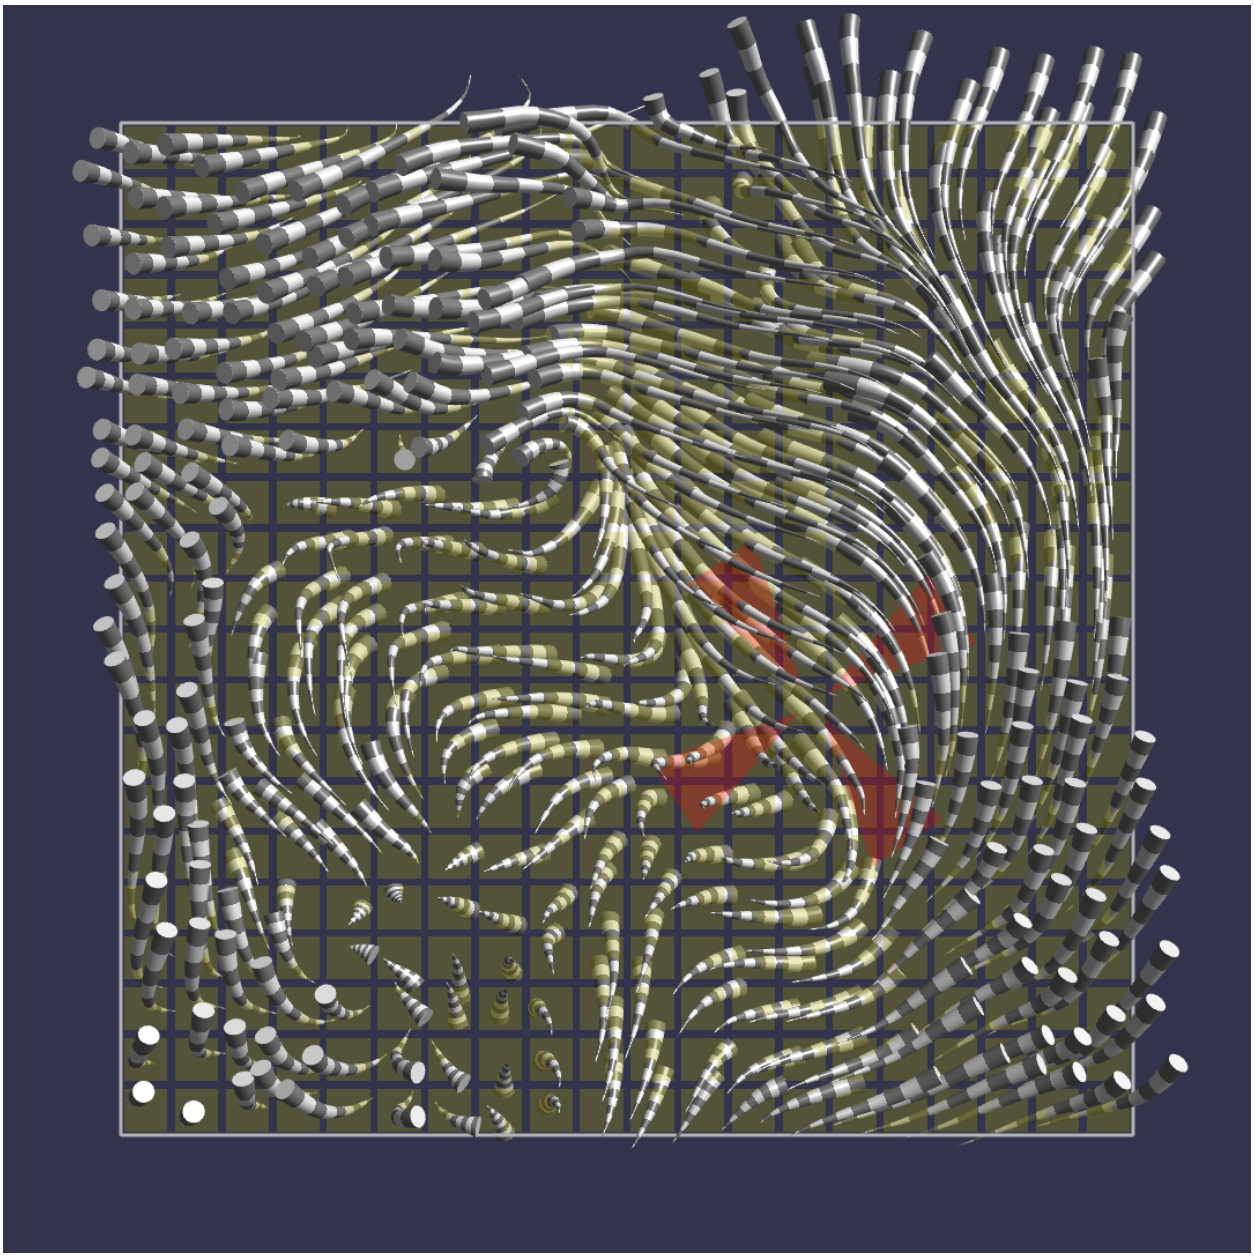

Figure 5.17: The streamcone glyph design. The direction of flow is from the base to the tip.

distortions that may arise through incorrect stereo disparity. This condition will address hypothesis H7 and H9.

5.2.3 Structure-from-Motion

The cutting planes were presented both with and without structure-from-motion by oscillating the cutting plane about its vertical bisecting axis a total of 20° over a 2.5s period

using a sinusoidal easing function. These parameters are based on the SfM parameters used successfully in previous research [45]. The oscillations from SfM help to reveal more of the 3D shape of the streamline geometry. This condition addresses hypotheses H8, H9, and H11.

5.2.4 Cutting Plane Design

The cutting plane is positioned at neutral parallax (at the depth of the screen) and was sized to fill most of the display at 30cm square (a visual angle of approximately 29.5°). Streamline seeding is accomplished by dividing the plane into 20x20 cells and jittering their midpoints by up to $\frac{1}{4}$ of the cell width and height. Pilot testing indicated that the two streamlet conditions would benefit from additional seeds, and the seeding density was increased to 40x40 (a four-times increase), as can be seen by comparing Figures 5.14 to 5.15.

The cutting plane is rendered as semitransparent yellow tiles to help disambiguate streamlines in front of or behind the slice. Streamlines are forward- and reverse-propagated from the seed points using fourth-order Runge-Kutta integration, with 10 steps forward and backward each. With forward propagation, the streamline seed is advected based on the flow direction and magnitude at the current seed point. In reverse propagation, the flow direction is reversed while the magnitude is preserved, which provides a way to simulate "tracing back" the flow from a given sampling point.

5.2.5 Apparatus and Display

The experimental stimuli are displayed on an Asus PG278Q 27" 2560x1440 px 3D monitor with Nvidia 3D Vision active shutter glasses, which are worn in both stereoscopic and non-stereoscopic viewing conditions to control for brightness and contrast levels consistency. Active stereo glasses use frame-sequential 3D, ensuring the same resolution whether or not 3D content is being viewed on the display.

With users seated approximately 57cm from the display, the per-pixel visual angle of the display is 84.4 arc seconds, or approximately 43 pixels per degree of visual angle. The

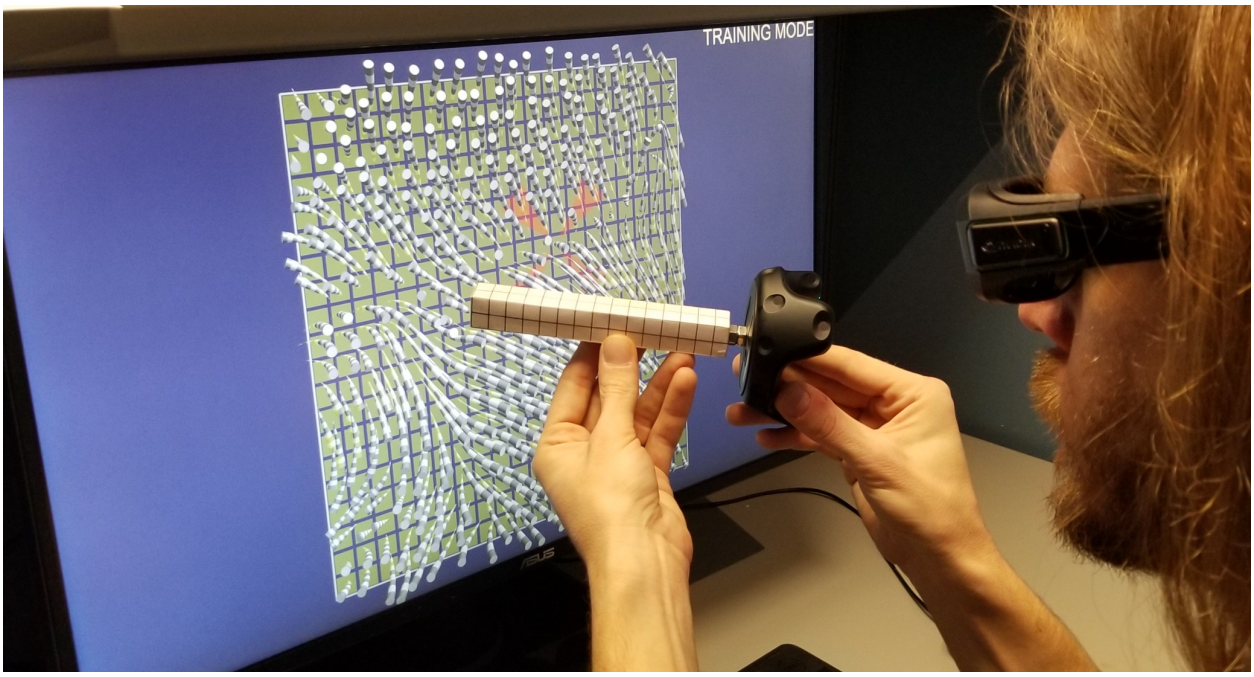

Figure 5.18: The experimental setup showing the wireless direction probe, stereoscopic 3D monitor, and active shutter glasses during a training session.

display sits in an enclosed desk space free of visual distractions, and the lights in the room are dimmed during the experiment. An example of the experimental setup is pictured in Figure 5.18

Direction Tracking and Input

The physical probe pictured in Figure 5.19 allows the user to input perceived glyph directions by aligning the probe in front of the display and pressing a key. An HTC Vive Tracker and two Lighthouse tracking cubes provide 6-degree-of-freedom (6DOF) tracking, though only 3DOF were used to record direction. The tracker was affixed to the end of a 2cm x 2cm x 10cm rigid block of foam with a $1cm^2$ grid pattern, identical to the design from Experiment 1.

During pilot studies, the accuracy and jitter in the spatial tracking system was noted to be improved over the electromagnetic tracking system used in Experiment 1.

Before each participant begins the experiment, the tracking system is calibrated to the

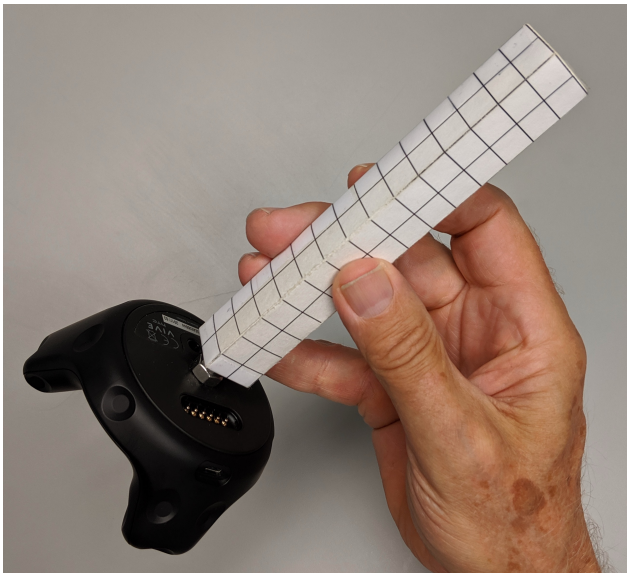

Figure 5.19: Perceived directions are input using the probe, which points away from the black HTC Vive wireless tracker attached to its base.

display in order to co-register the virtual scene's coordinate system with the real world.

5.2.6 Sampled Data

The test data set comes from a supercomputer simulation of the birth of the universe. It was selected because it contains a complex swirling vector field with many interesting flow patterns. On each trial a randomly oriented slice is chosen from a random position within the data set. The random orientation ensures that all vector directions are equally represented.

The data volume is 400^3 units, and the cutting plane is scaled to 100x100 units. For each trial, the center of the cutting plane is randomly positioned within the 200^3-unit center of the data volume, and its direction is randomly selected from a uniform distribution over a unit sphere. This provides a large amount of variation in the flow features and patterns, because the simulation results had many local and global characteristics at different scales throughout its domain.

While other approaches were considered, including random vector field generation using radial basis functions, the simulation output was chosen due to its rich flow characteristics and because it represents a real data product in need of improved visualization strategies.

Since this experiment focuses primarily on the perception of vector directions, the chosen data set provides a wide array of flow features which generalize to other flow visualization problems.

5.2.7 Task

After a vision test and stereoblindness test, participants are guided through an interactive training session to become familiar with the input probe and the experimental task.

Each training session begins by displaying a to-scale virtual model of the input probe at the center of the screen and matched to the direction of the physical probe. An additional target probe also appears at the center of the screen. Subjects are shown how to align the virtual probe to the target goal, with the assistance of visual feedback that decreases over the course of the training session. The first level of visual feedback changes the color of the target goal depending on the angle between the input probe and the target goal; the second level eliminates the color cues from the first level; the third level does not render the virtual probe on the screen, but the color cues from level one are applied to the target goal; and the fourth level shows only the target goal, with no input probe rendered nor color applied to the target goal.

At any point during the training, the misalignment error in degrees can be displayed using a toggle key. Each subject needs to achieve an error of 10° or less on at least three practice trials for each level of visual aid. Examples of each stimulus condition are displayed to the participant, and they are shown how to orient the probe with the flow direction at a target marker on the cutting plane. Feedback in this portion of the training is provided verbally.

During the study session, participants see a slice through a 3D vector field. The flow direction at the point specified by the red target reticle is indicated by aligning the wireless probe pictured in Figure 5.19 with the perceived flow direction. Participants are instructed to hold the probe as close to their line-of-sight as possible, so that there is minimal perspective

distortion between the displayed stimulus and the direction probe. The experimental setup in Figure 5.18 shows the probe in use during a practice trial at the end of a training session.

Each trial is based on a randomly oriented slice taken from a random position within the data set, as described in the previous section. Target points are randomly chosen within the center $\frac{2}{3}$ of each slice. A red target reticle is centered over each target point, and is slowly rotated for added visibility in visually dense/obscured configurations (e.g., Figure 5.16). All visuals (reticle, cutting plane, flow geometries) persist until the participant presses the key to record their direction input and advance to the next trial.

5.2.8 Experimental Design

This experiment followed a fully-crossed $5 \times 2 \times 2$ within-subjects design for a total of 20 total conditions, listed in Table 5.2. Each participant completed 15 trials per condition, yielding 300 trials per participant. All trials for a condition were presented in the same block, and blocks were randomly ordered.

Table 5.2: Experiment 2 Conditions

Parameter	Values
Glyph Design	Static Streamlets, Animated Streamlets, Static Streamtubes, Animated Streamtubes, Streamcones
Stereoscopic Viewing	On, Off
Structure-from-Motion	On, Off
Total	**20 unique conditions**

The data was analyzed using a procedure similar to Experiment 1. The total angular errors (Equation 5.1) were decomposed into their constituent absolute depth errors (Equation 5.2) and weighted projection errors (Equation 5.3) for further statistical analyses.

5.2.9 Results

A summary of the raw results can be found in Figure 5.20.

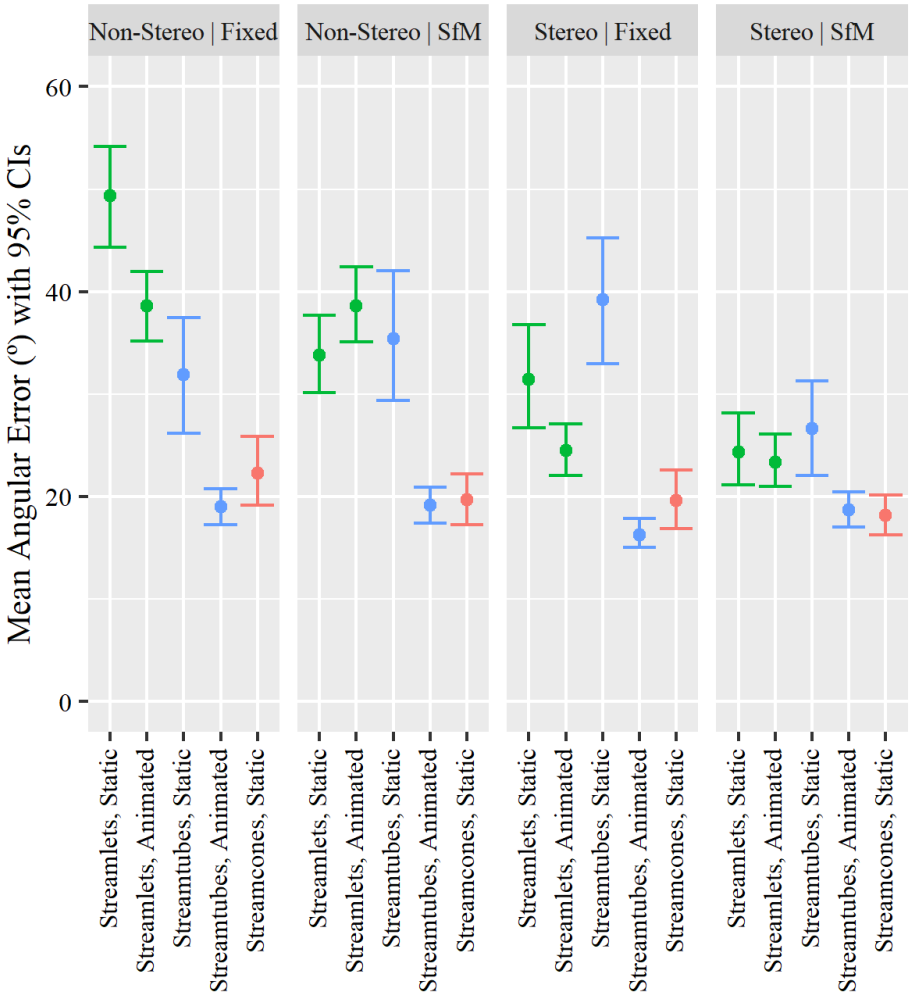

Figure 5.20: Uncorrected Mean Total Angular Error with 95% confidence intervals. The results are colored by the geometry condition (streamline, streamtube, and streamcone).

Antiparallel Correction

A common challenge in flow visualization using streamlines is that the orientation can be clear, but the sign (forward vs. backward) along the streamline can be ambiguous.

The mapping of cone geometry to direction is an arbitrary choice (i.e. the tip of the cone could point either forward or backward); there are no established conventions about which way a cone points. The direction chosen for the cone glyphs in this study follow the MATLAB "coneplot" [13] alignment, where the taper of the cone follows the direction of the flow, i.e. the smaller, tip end points forward in the direction of the flow.

Similarly, the texture gradient used in the static streamlet and streamtube designs is also directionally ambiguous. Thus, it is inevitable that without extensive training, a portion of the participant pool will have difficulty consistently adhering to the chosen directionality conventions for the aforementioned parameters if it conflicts with their natural intuition of them. This represents an *antiparallel confusion* by the participant.

In three-dimensional Euclidean space, two directed lines are antiparallel when they are parallel but point in opposite directions. Antiparallel confusion by the participant yields the largest angular error (180°) when the target and the probe are completely aligned but pointed in opposite directions. This means that participants who are very spatially accurate may still produce large errors if they are inconsistent in using the conventions chosen for the experiment, which confounds the results in terms of angular accuracy.

Participants were explicitly trained to follow our chosen directional conventions for the cone geometry (from base to tip) and the texture gradient (from dark to light). They were observed during the pre-study practice session to ensure adherence to the chosen conventions, and an additional visual reference sheet was provided during the study in case the participant forgot or became confused about the convention. (This was observed to occur in some participants whose natural intuitions were opposite the chosen conventions – the participant's intuition would begin to override the training over time, resulting in some trials where the participant accidentally followed their own intuition about the directionality, rather than the chosen convention for the study.)

Evidence of antiparallel confusion is demonstrated through the histogram of raw errors shown in Figure 5.21. There is a clear second, smaller distribution of errors near the far end of the positive tail. The histogram has been colored by whether or not the condition included a static or animated texture to indicate direction. This reveals that antiparallel confusion occurs overwhelmingly during trials with static textures, where the participant must follow our chosen directional convention which may conflict with their intuition. While this error could be dealt with *a priori* by allowing the participant to choose the directional convention

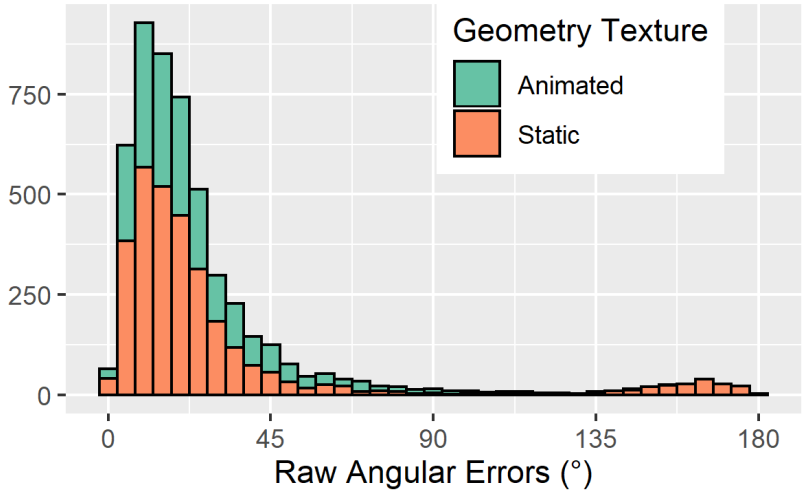

Figure 5.21: A histogram displaying the distribution of uncorrected, raw angular errors. The histogram is binned in 5° increments on the range $[0 - 180]$. The bimodal distribution is due to antiparallel confusion by some participants, creating a second (and much smaller) distribution in the far right tail of the main distribution.

they prefer, as opposed to forcing an arbitrary convention onto those with strong intuitions, it is better to keep all experimental variables consistent between participants, and correct for this error in the data.

To detect antiparallel confusion, individual trials are filtered where total angular errors over 130°, recorded for further evaluation of antiparallel metric, and corrected for further analysis using the equation:

$$\theta_{corrected} = 180° - \theta \tag{5.4}$$

In the above equation, $\theta_{corrected}$ is the corrected error, and θ is the raw result that needs to be transformed (i.e., is $\geq 130°$).

This correction yields a more accurate evaluation of perceived vector directions for participants with very strong directional intuitions that ran counter to the chosen cone direction conventions.

Figure 5.22 shows the total number of times that participants' individual trial results were corrected for antiparallel confusion according to Equation 5.4.

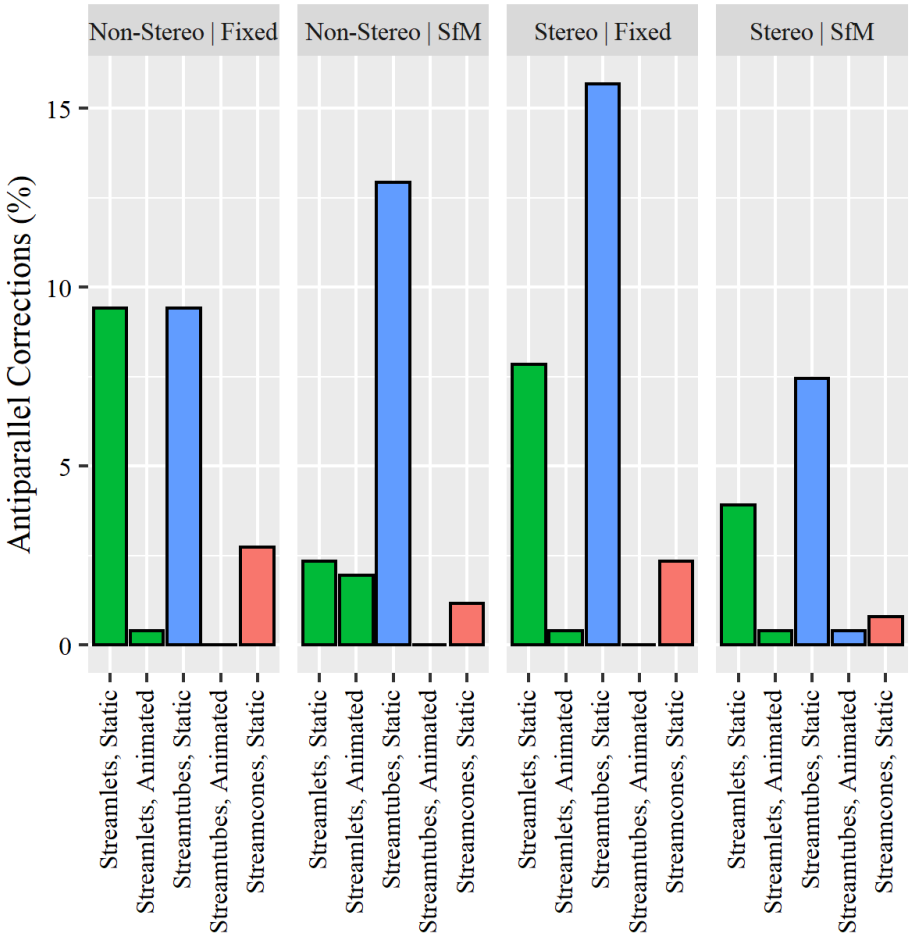

Figure 5.22: A count of trials where the antiparallel correction (Eq. 5.4) was applied. The results are colored by the geometry condition (streamline, streamtube, and streamcone). Note that corrections overwhelmingly affected the static streamlets and static streamtubes conditions.

Analysis

Since angular errors have a performance floor at 0°, the distribution of errors is heavily skewed to the right (see Figure 5.21), with a long tail ending at 180°. Transformations of the data were performed to try to correct the distribution of residual errors after running an analysis of variance (ANOVA) parametric test, but did not find a transformation that satisfied the assumption of residual error normality enough to be confident in the appropriateness of a parametric test. Therefore, the non-parametric Kruskal–Wallis one-way analysis of variance

was used; this compares the median values of similarly-distributed groups to test the null hypothesis that the two groups have the same median value [67]. A pairwise Wilcoxon test with a Benjamini-Hochberg multiple comparisons correction (intended to reduce the rate of false positives) was then used to find the relative performance amongst the group members at a $p = 0.05$ confidence level [68].

Corrected mean angular error rates can be found in Figure 5.23 and Figure 5.24, providing a much clearer picture of the relative performances between conditions. Overall, glyph design

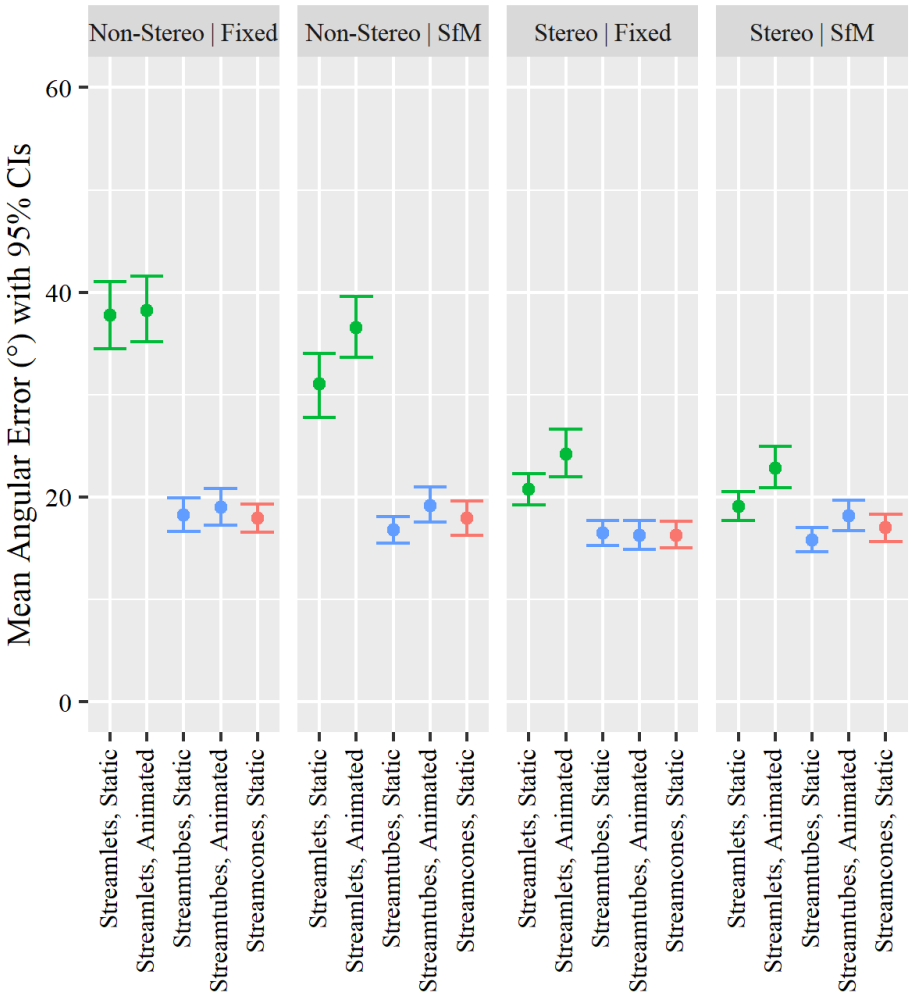

Figure 5.23: Corrected Mean Total Angular Error with 95% confidence intervals, grouped by stereoscopic viewing and SfM conditions. The results are colored by the geometry condition (streamline, streamtube, and streamcone).

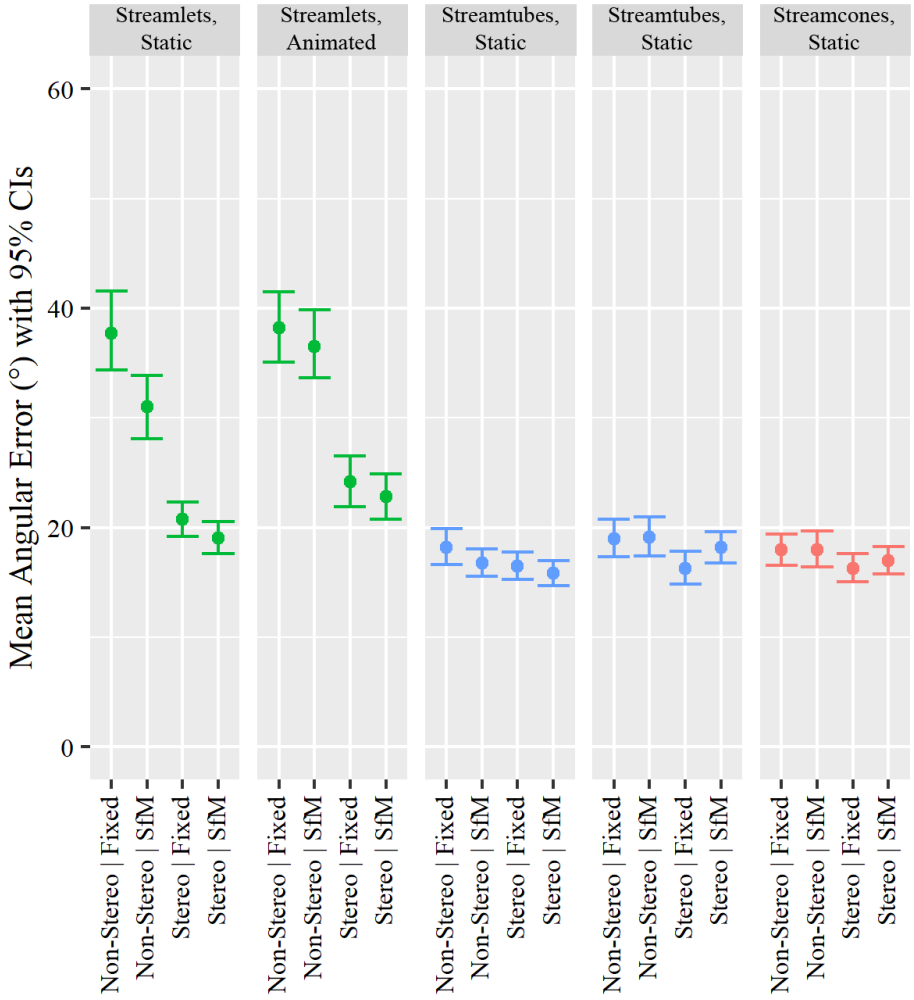

Figure 5.24: Corrected Mean Total Angular Error with 95% confidence intervals, grouped by glyph design. The results are colored by the geometry condition (streamline, streamtube, and streamcone).

($\chi^2(4) = 418.62$, $p < 0.00001$) and stereo viewing ($\chi^2(1) = 99.524$, $p < 0.00001$) have a highly significant effect on error rates, but the effect of motion parallax ($\chi^2(1) = 1.69$, $p = 0.1931$) is not significant.

Among the glyph designs, the animated streamlets produced significantly higher errors than all other designs. Static streamlets performed significantly worse than either streamtube design or the streamcone design, but the two streamtube designs (static and animated) along with the streamcone design did not perform significantly differently from one another. Stereo viewing yielded significantly lower errors than the non-stereo viewing condition. Figure 5.23

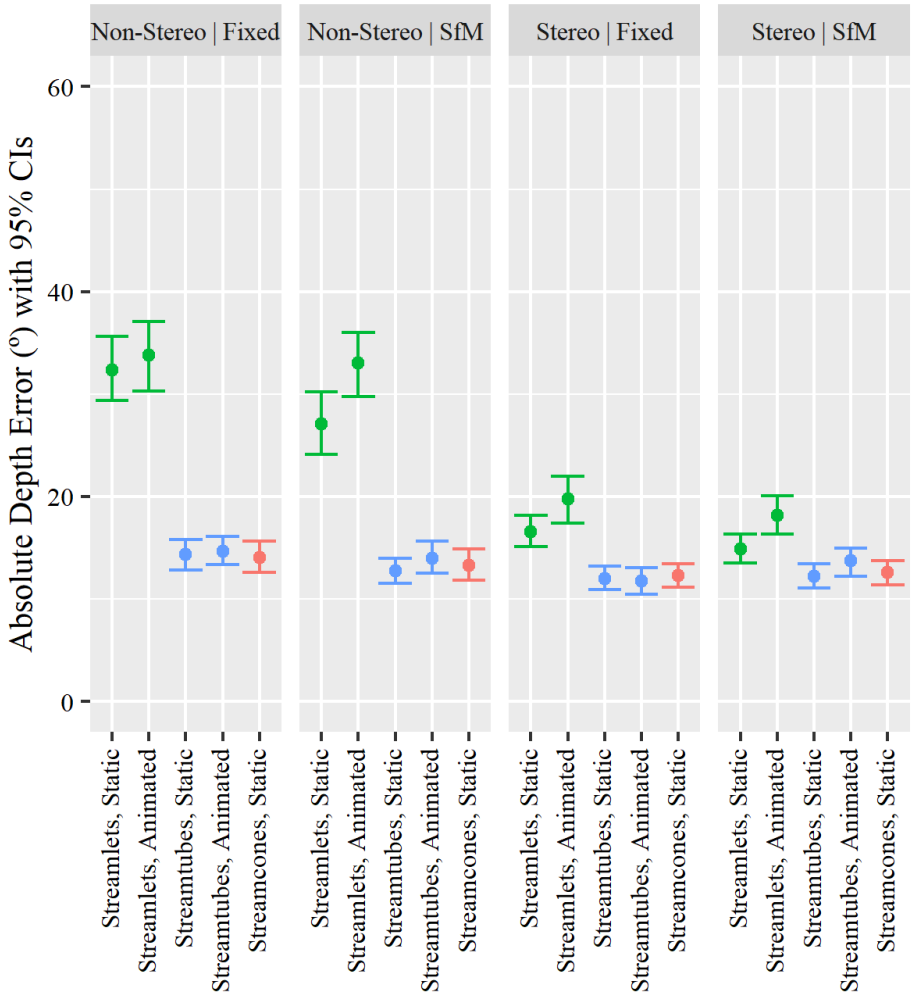

Figure 5.25: Mean Absolute Depth Error. The results are colored by the glyph geometry condition (streamline, streamtube, and streamcone).

helps to illustrate these results.

Digging deeper into the glyph designs, we find that motion parallax only helps in the case of streamlets ($\chi^2(1) = 6.989$, $p = 0.0082$), significantly reducing the error rate when SfM is available. Error rates with streamlets and streamtubes are significantly reduced with stereo viewing ($\chi^2(1) = 114.68$, $p < 0.00001$), but streamcone performance was not significantly affected by motion parallax or stereo.

Animating the glyph textures had a significant effect for streamlets ($\chi^2(1) = 8.603$, $p = 0.00335$), where animated streamlets performed worse than static streamlets. Streamtube

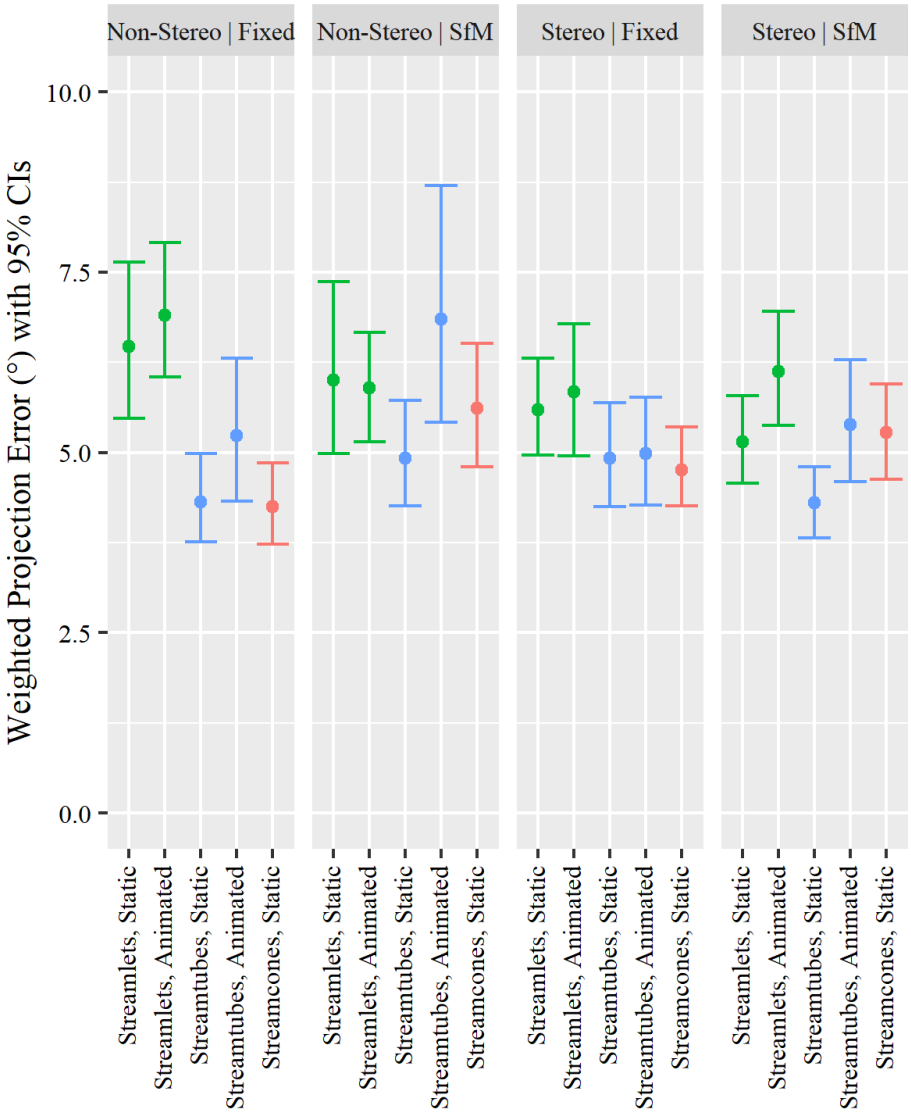

Figure 5.26: Mean Weighted Projection Error with 95% confidence intervals. Note that the severity and range of errors is much lower than in Figures 5.20, 5.23, and 5.25. The results are colored by the geometry condition (streamline, streamtube, and streamcone).

error rates were not significantly different between static and animated conditions.

A plot of the absolute depth error is shown in Figure 5.25. We can see that this component makes up the majority of the angular error, and follows almost exactly the same pattern as the corrected angular error. Glyph design was highly significant ($\chi^2(4) = 383.5$, $p < 0.00001$), as was stereo viewing ($\chi^2(1) = 104.05$, $p < 0.00001$). Motion parallax was not a significant effect on depth errors. The relative performance between conditions of the groups

were the same as described above for the corrected mean total angular error.

The weighted projection error plot can be found in Figure 5.26. Projection errors were, on average, much lower than depth errors. Glyph design had a highly significant effect on weighted projection errors ($\chi^2(4) = 35.704$, $p < 0.00001$), but neither stereoscopic viewing nor SfM had a significant effect. For the glyph designs, animated streamlets performed significantly worse than every other glyph design except static streamlets, and static streamtubes and streamcones produced significantly lower weighted projection errors than static streamlets.

5.3 Experiment 3: Stereoscopic Misperceptions of Vector Magnitude

Mapping vector magnitude to glyph length is a well-established practice in 3D flow visualization. When a visualization is viewed from a perspective other than that from which it was rendered, significant distortions of glyph geometry can arise, causing misperceptions of the vector data represented by the glyphs, which could ultimately lead to a misunderstanding of the true data. This experiment examines the perception of glyph length when images are viewed from the correct position, or from increasingly oblique angles, under different viewing conditions and glyph parameters to test hypotheses H12, H13, and H14.

Banks et al. [57, 58] previously showed that viewers of monoscopic images were able to compensate for incorrect viewpoints, but viewers of stereoscopic images did not compensate for mismatched viewpoints. Whether this pattern carries over to glyph length judgement is not known.

The geometric distortion model [53,54] can be used to predict the distortion of a 3D glyph when viewed from an incorrect viewpoint. Of particular value is the ability to calculate the true distorted length of the glyph, which can be used as a metric to assess the perceived distortion. These models were used in the design and analysis of this experiment.

5.3.1 Stimuli

A single idealized vector represented by a glyph was chosen as the primary stimulus for the experiment. The glyph design employs shading, 3D cylinder geometry, and a ringed Perlin-noise texture to improve the accuracy of orientation perception; an example of the stimulus can be seen in Figure 5.27. The glyph is 5 mm in diameter and either 10 cm or 20 cm in length to probe for consistency in performance between moderate and large vector magnitudes (hypothesis H14). Stretching the ringed texture instead of repeating it prevented participants from estimating length by simply counting rings. Figure 5.27a contains an example of the two glyph lengths.

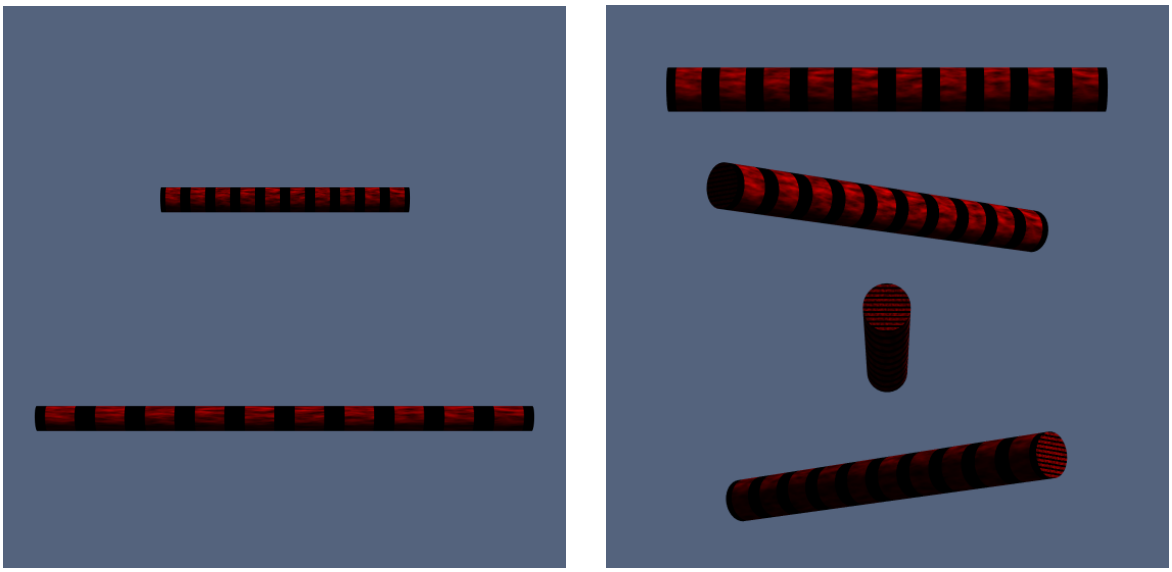

(a) Glyph Length (b) Glyph Angle

Figure 5.27: The stimulus and its experimental parameters: (a) Glyph Length and (b) Glyph Angle.

The glyph's midpoint is anchored at screen depth (neutral parallax) and displaced vertically 8.4 cm above the center point of the screen to accommodate the virtual measuring tool's vertical range (described below) and to provide perspective information for glyphs aligned with the view direction, as glyphs in this configuration visually degenerate into a simple 2D circle if kept at the center of the projection. Since these stereoscopic distortions

only occur at non-neutral parallax values, the glyph orientation is restricted to the xz-plane, i.e., the glyph can rotate around the vertical y-axis, by which it is bisected and to which it is orthogonal. Four glyph angles were chosen to probe perceptual performance at different glyph orientations and amounts of parallax: 0°, 45°, 90° and 135°. See Figure 5.27b for a visual example.

The stimuli are viewed in both head-coupled and uncoupled stereoscopic perspective projection environments (Figure 5.28) to evaluate the effects of distortions introduced by uncoupled dislocated viewing, and to assess the potential mitigation of those effects by head-coupled viewing. They are also viewed in stereoscopic and non-stereoscopic viewing conditions, which provides additional insight into how the non-stereoscopic geometric distortion model may differ from the stereoscopic model in terms of glyph length judgements.

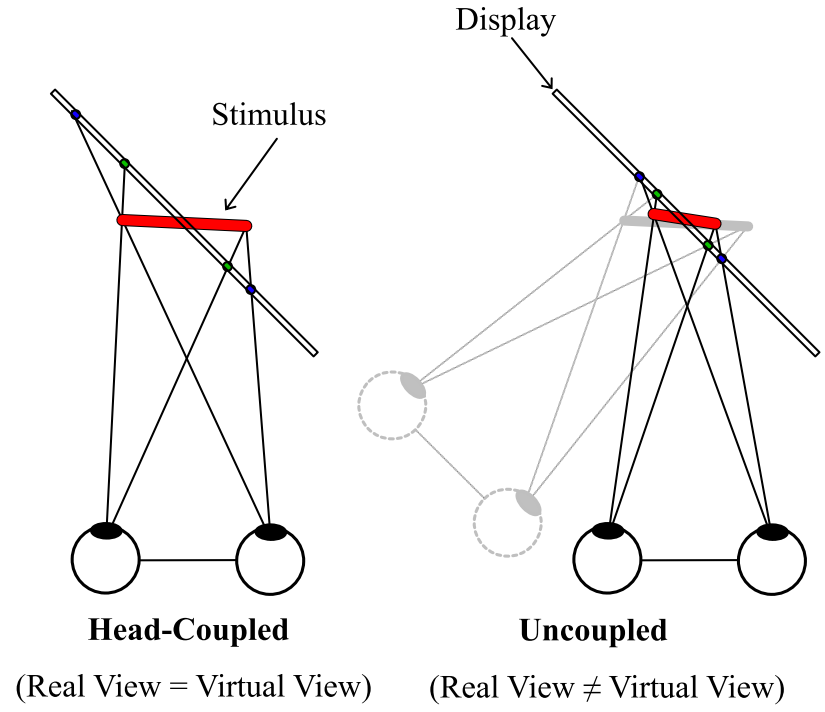

Figure 5.28: The two stereoscopic perspective viewing mode conditions. The head-coupled condition renders a "correct" view of the stimulus based on the offset viewpoint; the uncoupled condition does not account for viewpoint, and the stereoscopic visual representation is distorted.

Since distortions are induced by dislocated viewing, it is necessary to add a condition to vary the viewing angle to the screen. Along with the orthogonal, always-coupled 0° setting, moderate and extreme settings of 15° and 30°, respectively, were chosen as representative viewing angles for a typical usage scenario. A depiction of these viewing angles can be found in Figure 5.29.

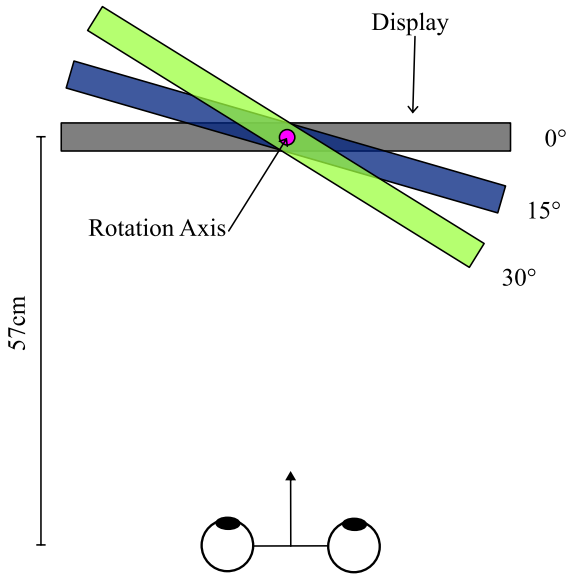

Figure 5.29: The three viewing angle conditions, viewed top-down. The viewpoint is held fixed while the display swivels about its vertical center to offset the viewing angle by 0°, 15°, or 30°.

A vertical white shaded virtual rod serves as a measuring tool to input perceived glyph lengths. Virtual measuring tools have been compared with physical and haptic variations for stereoscopic depth judgment tasks by Hartle et al. [69], but no significant difference between those input techniques were found. One end of the measuring rod is anchored at neutral parallax 5.6 cm above the center point of the screen and extends downward along the vertical axis as the participant interactively changes its length. A vertical orientation for the measuring tool was chosen because its geometry is minimally distorted by incorrect stereoscopic viewing as long as the tool is placed at zero parallax and aligned with the physical rotation axis depicted in Figure 5.29. This design was chosen because it allows both

the stimulus glyph and measurement rod to coexist in the same virtual space.

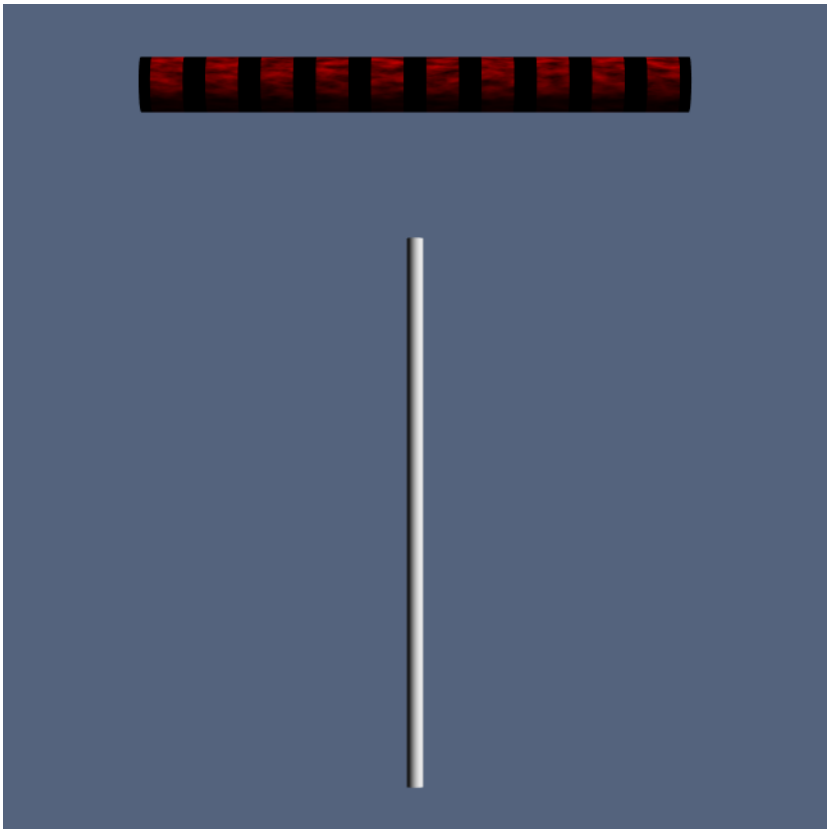

Figure 5.30: An image of the experimental interface showing the target glyph in red and the measuring rod in white.

The diameter of the measuring rod is 1.5 mm, significantly smaller than the target glyph diameter. This choice was made to minimize any skewing distortion for the uncoupled, non-orthogonal conditions, which could cause the vertical cylinder to appear more like an elliptic cylinder. An example of the experimental interface is shown in Figure 5.30.

5.3.2 Apparatus and Display

The apparatus consists of an Asus PG278Q 27" 3D display running at 120 Hz (60 Hz per eye using active-shutter glasses) with a resolution of 2560 × 1440 pixels, affixed to a vertical mount whose rotational angle can be controlled through software. The mount has a rotation range of 180° (±90°) about the vertical axis, and an angular resolution of approximately 1°. The apparatus is pictured in Figure 5.31.

Figure 5.31: The stereoscopic monitor is affixed to a swiveling base that can be programmatically set to different viewing angles.

A chin rest is positioned 57 cm from the screen and adjusted so that the view direction is orthogonal to the center of the screen. At this distance, a 1 cm stimulus on the display subtends a visual angle of 1°, and the per-pixel visual angle of the display is approximately 82 seconds of arc.

Stereo image pairs are generated by an interactive OpenGL application and an nVidia Quadro M4000 graphics card with quad-buffered stereo enabled, viewed through nVidia 3D Vision active shutter glasses. User input is obtained through a standard keyboard connected

to the system.

The design of the apparatus allows the observer to remain at a stationary viewpoint (the chinrest) while the display itself is rotated about the vertical axis at the center of the screen. This is isomorphic to real-world usage, where the position and orientation of the display is fixed and the viewer moves to different positions around it. The two arrangements are physically identical and only differ in their frame of reference. See e.g. Fig. 1 in [70] for a visual representation and explanation.

5.3.3 Task

The experiment is conducted in a laboratory station set up to minimize distractions and with dimmed lights. Participants begin by reading through and signing a consent form, and are then administered an informal visual acuity and stereoacuity test to verify 20/20 vision and the ability to properly fuse stereo image pairs. Each participant's IPD is measured with a specialized ruler, and entered into the study software to account for varying IPDs amongst participants.

Each participant is seated at the station and instructed to adjust the chair so that the chin rest is comfortable to use. An oral description and guided example of the procedure are followed by free interaction until participants indicate their readiness to proceed with the experiment. Participants are allowed to ask questions during the training and are informed that they can take short breaks between trials, if desired.

At the beginning of each trial, a splash screen is displayed indicating the number of remaining trials and prompting the participant to press the space bar to initiate a trial. Participants place their head onto the chin rest and press the space bar on the keyboard to indicate their readiness to proceed. The monitor mechanically rotates to the angle indicated by the current trial condition (if necessary) over 2.5 s, and the stimulus is displayed for 10 s. While the stimulus is displayed, participants use the up and down arrows on the keyboard to adjust the length of the measuring rod by 1 mm increments to match the perceived length

of the target glyph. At the end of the 10 s stimulus display time, the length of the measuring rod is recorded and the monitor rotates back to the orthogonal position and returns to the splash screen for the next trial. No feedback is provided to the participant regarding their performance at any point during the experiment.

5.3.4 Experimental Design

This experiment follows a fully-crossed $2 \times 2 \times 3 \times 4 \times 2$ mixed-model design for a total of 96 experimental conditions. Stereoscopic Viewing, Perspective Viewing Mode, Viewing Angle, Glyph Angle, and Glyph Length are all treated as within-subjects factors and can be found listed in Table 5.3. A rejection level of $p < 0.05$ is used for all analyses.

Table 5.3: Experiment 3 Conditions

Parameter	Values
Stereoscopic Viewing	On, Off
Perspective Viewing Mode	Uncoupled, Head-coupled
Viewing Angle	0° (orthogonal to screen), 15°, 30°
Glyph Angle	0°, 45°, 90°, 135°
Glyph Length	10 cm, 20 cm
Total	**96 unique conditions**

5.3.5 Results

The relative percent error for glyph length judgment of each condition is plotted in Figure 5.32. The head-coupled glyph angles are plotted according to their apparent angles: 180° - Relative Glyph Angle + Viewing Angle = Effective Glyph Angle. For example, a head-coupled condition with a relative glyph angle of 135° and viewing angle of 30° yields an effective glyph angle of 15°. Conditions with a 0° Viewing Angle are redundantly plotted at 180° as well to reflect the fact that they are identical conditions, and to better reveal any symmetric patterns in the data.

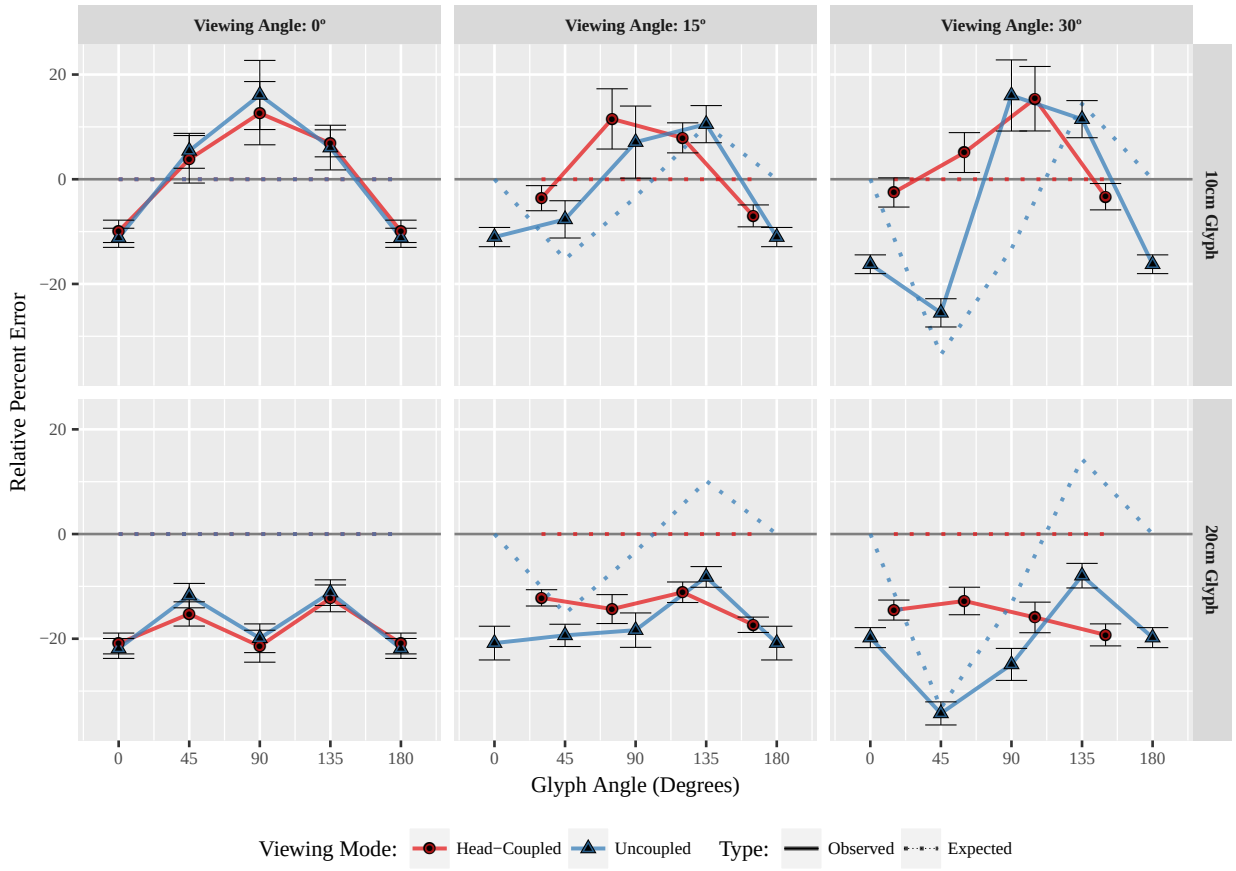

Figure 5.32: Relative percent error of the perceived glyph length under stereoscopic viewing plotted with standard error bars. The dotted lines represent the predicted error according to the geometric distortion model. A glyph with its angle set to 0° is orthogonal to the view direction at the 0° viewing angle; a glyph at 90° is parallel to a 0° view direction.

A four-way ANOVA found a highly significant effect of Glyph Angle ($F(3, 983) = 45.1$, $p < 0.0001$), Perspective Viewing Mode ($F(1, 983) = 7.95$, $p < 0.005$), and Glyph Length ($F(1, 983) = 350.9$, $p < 0.0001$). Viewing Angle failed to reach the $p < 0.05$ threshold for significance.

The two-way interactions between Glyph Angle and each other factor were found to be significant: Perspective Viewing Mode ($F(3, 983) = 4.59$, $p < 0.005$), Viewing Angle ($F(6, 983) = 5.79$, $p < 0.0001$), and Glyph Length ($F(3, 983) = 20.44$, $p < 0.0001$). Additionally, the interaction between Perspective Viewing Mode and Viewing Angle was found to be significant ($F(2, 983) = 5.07$, $p < 0.01$).

CHAPTER 6
DISCUSSION

6.1 Assessment of Hypotheses

The following section discusses the results in the context of each hypothesis.

6.1.1 Hypothesis 1

> H1: Glyph designs incorporating the most depth and orientation cues will be the most successful for estimating 3D flow direction.

Overall, the results appear to support this hypothesis. In Experiment 1, tube-based glyph designs proved to be significantly better for direction perception than the illuminated lines and shadowed lines designs, which were better than plain lines. This makes sense, as the tubes provide more visual cues for orientation than the other glyph designs, and are easier to distinguish from one another in denser seeding conditions than line-based methods.

Further examination of the high-performing split group in Experiment 1 reveals a clear hierarchy of glyph designs, ordered by available visual cues. The lack of significant effects on weighted projection errors supports this grouping, since no depth information should be needed to accurately determine the projection of a vector onto a viewing plane.

In Experiment 2, stereoscopic viewing cues yielded better performance in every visualization that used them, in addition to confirming the same hierarchy of tubes over lines when it comes to 3D contour perception. Line-based streamlet designs saw much more benefit from stereoscopic viewing than did tube-based designs, likely because tubes already incorporate many more redundant cues so that the effect of stereoscopic viewing is not as strong. The

results also supported the addition of texture animation for strong, unambiguous directional cues.

The results of the non-stereo, fixed conditions of Experiment 2 can be compared to the results of Experiment 1, which used neither stereo nor structure-from-motion. The seeding density in Experiment 2 is closest to the most sparse seeding density in Experiment 1 (2 seeds/cm). Overall, errors in the static streamlet (line-based geometry) and static streamtube conditions are lower than those observed for plain lines and ringed tubes. The biggest difference can be found in the absolute depth error, where static streamlets yielded on average 10° less error than plain lines. This could be in part due to the occlusion cues available in the streamlet design, which was not present in the direct-glyph plain lines design of the previous study. The depth error for static streamtubes was similar to the ringed tubes; occlusion cues did not provide additional perceptual benefit like they did for static streamlets. Weighted projection errors between the two studies were comparable and very low.

6.1.2 Hypothesis 2

> H2: Tube-type designs should outperform line-type designs because of their stronger pictorial cues.

This hypothesis had strong evidence to support it between the first two experiments. Tube-based designs were consistently more performant than line-based methods in every comparison, even illuminated lines whose lighting is modeled by an infinitesimally small cylinder. The only difference between the tube and line designs is the visual space they provide for cues. Thus, it can be reasoned that the additional pictorial cues of the tube designs are responsible for the improved perceptual performance over line designs.

Experiment 2 provided further evidence that depth cues from tube designs are much stronger than those from line designs. Both streamtubes and streamcones yielded significantly lower errors than line-based streamlets. When stereoscopic viewing was enabled, streamlets saw a much more significant performance increase than did the tube designs.

This is because stereoscopic viewing is mostly reinforcing the already-strong cues of the tube design, while the line designs are visually impoverished and the addition of stereoscopic viewing helps convey much more information about their direction. Since streamlets did receive considerable benefit from both stereo and structure-from-motion, this hypothesis is supported by the results.

6.1.3 Hypothesis 3

> H3: The cutting plane seeding density significantly affects perceptual performance for a given glyph design.

The seeding density condition from Experiment 1 was found to have a significant effect on task performance. The medium density and sparse densities both outperformed the densest seeding, likely because of the detrimental amount of visual noise that dense seeding adds. An interesting result was that tube-based techniques exhibited resistance to the performance decrease observed in other techniques as seeding density increased, evidenced by the significant interaction between glyph technique and seeding density within the high-performing split group.

No interaction effect was found between the set of tube-based glyph designs and the seeding density, so the diameter of the tube does not appear to be tied to the seeding density in terms of task performance (assuming reasonable diameters and seeding densities are chosen). This suggests that streamtubes are less sensitive to overly-dense seeding conditions when compared to their line-based counterparts.

Increasing the seeding density of the streamlet conditions in Experiment 2 was carefully considered before being implemented, because it could give an unfair advantage by sampling four times the amount of information as the other conditions. Ultimately, pilot testing confirmed that occlusion and structure-from-motion cues degrade quickly when the seeding density gets sparse, as the thin line geometry of the streamlets fails to provide enough local depth information to reasonably judge flow direction. The results clearly show that any

bias in favor of the streamlet conditions was dominated by the objectively better perceptual cues in the streamtube- and streamcone-based conditions. It is possible that the animated streamlets performed worse than the static streamlets because the seeding density was too high, and the substantial number of small streamlets moving through the cutting plane became visual noise rather than a helpful cue.

This hypothesis can be considered partially supported, and the results tell us that an ideal density will be somewhere towards the medium range, as opposed to very dense seeding that introduces undesirable visual noise or the very sparse range that may not provide as many ancillary cues like occlusion when used with integral glyphs. Algorithms or heuristics that take pixel coverage or screen-space glyph occlusions into account could be deployed to automatically adapt and tune the seeding density and distribution as the viewpoint and the features being visualized change.

6.1.4 Hypothesis 4

> H4: The effectiveness of tube-type designs depends upon their diameter; diameters that are too small will not be able to adequately show lighting and texture cues, and those that are too big waste space and add unnecessary visual noise.

In Experiment 1, tube diameter had a highly significant effect on total angular error, with 1mm and 2mm diameters performing significantly better than 0.5mm diameter. Breaking down the results into performance groups revealed that the lower-performing participants only had a significant performance increase with the largest tube diameters, while the higher-performing participants were able to derive significant performance benefits from the middle tube diameter.

This supports the hypothesis that tube-style glyph effectiveness does depend upon its diameter, but with the caveat that the most effective diameter depends not only on the flow features being visualized, but also on the viewer. It is therefore unlikely that a single number or formula exists to model the ideal tube diameter for good visual cues. A qualitative

suggestion for finding this balance is to parameterize the tube diameter for adjustments without allowing the values to become too extreme for the given data.

6.1.5 Hypothesis 5

> H5: Adding lighting cues to line-type designs (i.e., illuminated lines) can enhance accuracy.

In Experiment 1, the performance of the plain lines glyph design was not found to be significantly different from that of the illuminated lines design. However, when the data were split by performance group, a significant performance benefit with illuminated lines over plain lines appeared. There is no surprise in this, as high-performers are able to take better advantage of the available visual cues to make direction judgements.

Thus, while there is some evidence to support this hypothesis, that evidence is not strong, and indicated that lines are a bad design decision in general, since the few additional cues that can be added to them – lighting and cast shadows – do not improve the effectiveness of the visualization anywhere near an amount that would justify their use.

6.1.6 Hypothesis 6

> H6: Cast shadows can enhance accuracy.

The cast shadows cue did appear to give the shadowed lines technique an advantage over the other line-based techniques, but it was also highly sensitive to seeding density. At the sparsest density (5mm spacing), the performance of the shadowed lines technique approached that of the tube-based techniques. This is noteworthy, as the tube-based techniques utilize form-revealing shading that directly visualizes an orientation, whereas the shadowed lines technique requires the integration of two individual components to understand the whole glyph. These results suggest that cast shadows may be a strong cue for perceiving the orientation of objects lacking strong shading or for partially occluded objects whose shading is obscured.

Despite the results supporting the use of this cue, the shadowed lines design was rejected by the participants as being visually displeasing and did not instill confidence in their glyph direction judgements. Nonetheless, the relative performance success of the shadowed lines design compared to other line-based techniques hints at the potential for shadows to be a useful additive depth cue to be used alongside sparsely-seeded 3D glyphs, as long as their inclusion is not unpleasant for the user.

6.1.7 Hypothesis 7

H7: Stereoscopic viewing can enhance accuracy.

The results of Experiment 2 support this hypothesis. Every condition with stereoscopic viewing enabled lead to improved performance with every glyph design evaluated. Stereoscopy is known to be a very powerful cue, and this confirms its strength with perceiving 3D flow glyph visualizations.

It is important to also note the results from Experiment 3 here. Though this experiment focused on the effects of head-coupled perspective viewing, the results showed that judging glyph lengths with any type of stereoscopic viewing on a 3D monitor is wildly inaccurate due to many potential complicating factors, such as the horizontal-vertical illusion and the shape of the monitor itself.

While accuracy of glyph *direction* judgements improved under stereoscopic viewing, doing so via a 3D desktop monitor made glyph *length* judgements very difficult to accurately estimate. So, the support for this hypothesis comes with caution in terms of stereoscopic display use, and it is suggested that a modern stereoscopic viewing solution be used, such as a virtual reality head-mounted display (HMD), to avoid the potential perceptual side-effects that a 3D monitor may introduce.

6.1.8 Hypothesis 8

H8: Structure-from-motion (SfM) can enhance accuracy.

Structure-from-motion provides little overall benefit for judging flow glyph direction. Only in the case of simple streamlets was there any benefit, and even there it was quite small. This is surprising, and in marked contrast to prior studies of different visualization tasks, such as path tracing in 3D networks [43] and perception of patterns in point cloud data [45] where SfM provided a greater benefit than stereoscopic viewing.

One explanation for the lack of benefit from SfM is that the motion of the cutting plane and glyphs introduced more visual noise than helpful depth cues for these glyph designs. Another explanation is that the motion is so distracting that it degrades perceptual performance. But that does not help to clarify why SfM ends up being useful in other domains, by comparison.

The most likely reason for the poor performance of SfM is that it is better suited for revealing global structures and relationships between objects, as opposed to localized depth information that would be needed to make orientation judgements about individual glyphs. Tasks like network path tracing and point cloud pattern perception rely more heavily upon relative depths of the constituent objects; the task used in this research needs absolute depth cues. So, SfM does not provide absolute depth information to help with the task of judging local glyph direction. Therefore, this hypothesis is rejected.

The SfM tested in Experiment 2 was external motion, in that the cutting plane object moved while the viewpoint remained fixed. The complement to this is SfM through head-coupled perspective viewing. It is possible that the subtle motion from natural head movements controlled by the viewer would strengthen SfM cues to the point where they provide an appreciable benefit. This remains to be explored.

6.1.9 Hypothesis 9

> H9: The benefit of stereo and SfM cues should be greater in the case of line-type designs than for tube-type designs, because lines otherwise lack depth information.

The results very clearly supported this hypothesis. Experiment 2 showed that stereo has a very slight effect on tube-type designs (streamtubes and streamcones), but is much stronger and more pronounced when used with line-type designs (see Figure 5.11).

Line-based static streamlets were the only condition to benefit from structure-from-motion cues. Examining Figure 5.24, a large drop in error occurs for the non-stereo static streamlets when structure-from-motion cues were available. Stereocopic viewing drops the error rates for static streamlets down to approximately 20°, and another small drop in error is seen with structure-from-motion, but it is much smaller and does not significantly affect error rates when stereoscopic viewing is available.

6.1.10 Hypothesis 10

H10: Animating glyph textures can improve the perception of flow direction.

This hypothesis was very strongly supported by the results of Experiment 2. In fact, it can be said that animation is the single best disambiguator for direction. The animated streamtubes performed remarkably well (only one antiparallel correction across conditions over 1020 trials; $< 0.1\%$ correction rate, see Figure 5.22).

Animated streamlets with structure-from-motion were the exception; without stereo, the extra motion from texture animation degraded performance (see Figure 5.20). The absence of additional depth cues to fall back on made it difficult to disambiguate rotational movement from texture animation in the animated streamlet condition. Since structure-from-motion did not have a significant effect on performance, texture animation should be prioritized, but used with caution in cases where other motion is also present, as in time-varing 3D flow data.

For static visualizations where animation is not feasible, streamcones convey direction almost as well as animated streamtubes. In fact, cones are arguably the best general-purpose streamline glyph design to show flow direction; the depth cues they provide are strong enough that they do not receive any significant additional benefits from stereo or structure-from-

motion cues, as they have effectively reached a performance plateau. The caveat, however, is that there is no standard convention for a streamcone's direction, and may be perceived as going in the opposite direction than intended if the viewer's understanding of the convention is not clear.

6.1.11 Hypothesis 11

> H11: Animation will hinder SfM. In other words, when glyphs are animated to show flow direction, the benefit from SfM should be reduced.

Structure-from-motion did not provide much benefit overall, and although animated streamtubes in the stereo condition did perform slightly worse with SfM, it was not significant enough to draw a strong conclusion about this hypothesis. Some participants required more training than others to understand how to align the probe (as though the oscillating plane were fixed orthogonally to the viewer), but their performance after being trained was no different from other participants. It may be possible that this is a systemic error in the experimental design confounding the results, but the extensive training each participant received should have minimized this as much as possible.

6.1.12 Hypothesis 12

> H12: Perceptual performance of glyph length judgements follows the predictions made by the geometric distortion model.

No support was found for the hypothesis that perceptual performance would follow the predictions of the geometric distortion model. The general trends of the 20 cm Head-Coupled conditions are reasonably close to constant, but with an average underestimation of the glyph length by 15.6%. The 30° Viewing Angle conditions performed somewhat closely to the predicted error, and it could be weakly argued that the geometric distortion model has moderate prediction power at more oblique viewing angles.

The large errors for the 0° Viewing Angle and 0° Glyph Angle conditions were initially surprising, as there was no distortion of the geometry and no immediately obvious reason why average errors of -15.4% would be recorded. One possibility is that the configuration of the horizontal glyph at the top of the screen and the vertical measuring rod induced the Horizontal-Vertical Illusion and the Bisection Illusion (HVI/BI) [71, 72], resulting in significant underestimates in the measuring rod magnitude. The use of an anisotropic widescreen display for the stimulus may have exacerbated the HVI/BI as well.

Perhaps this could be avoided by masking the display so that its physical properties are hidden, e.g., a circular viewport mask. These perceptual illusions have been notoriously difficult to model [73], but more recent efforts have sought to further characterize the phenomena empirically [74] and produce simple models, such as the one proposed by Mamassian and Montalembert [75]. Ultimately, these are issues that arise when using desktop monitors, which are largely unavoidable due to the dominance of such displays in the workplace.

Another solution is to employ an immersive head-mounted display (HMD), where the virtual displays are stretched to fill (most of) the entire field of view. There would be no anisotropic display to bias length judgements, and the effect should mimic real-world egocentric distance/length judgements.

6.1.13 Hypothesis 13

> H13: Head-coupled viewing reduces errors in perceived glyph length compared to uncoupled, dislocated viewing.

There was significant evidence to support the hypothesis that head-coupled perspective viewing would reduce glyph length judgment errors over uncoupled dislocated viewing. It is especially apparent after examining the 30° Viewing Angle column in Figure 5.32 that head-coupled perspective drastically reduced errors in almost every instance. The reversal of performance at the 135° Glyph Angle setting can be explained by the similarity of the Head-Coupled Viewing Mode at a 30° Viewing Angle to the 90° Glyph Angle setting at a

0° Viewing Angle: i.e., the view direction and the glyph orientation are parallel (see Figure 5.27b), which makes the length of the glyph difficult to estimate. The uncoupled view of the 135° glyph angle provides a better perspective on the glyph for judging length in both a colocated and dislocated position, even with distortions.

6.1.14 Hypothesis 14

H14: Perceptual performance is consistent across typical glyph length ranges.

The most unexpected result from Experiment 3 was the strong effect of glyph length on judgment errors. Since 10 cm conditions tended to be overestimated and 20 cm conditions underestimated, perhaps there is some optimal range of visual angle a glyph should subtend in order to maximize length perception. The reason(s) for this are not clear; the significant interaction between glyph length and viewing angle tells us that something is occurring that is more complex than the geometric distortion model is able to capture, and thus a more refined model could be employed to explore this particular effect.

This issue may only have an impact on direct glyphs, since integral glyphs do not represent discrete magnitudes with their geometry in the same way a direct glyph does. If using direct glyphs in a design where magnitude information is important, the magnitude should be encoded by more than just the glyph length to strengthen its perception.

6.2 Depth Cues for 3D Vector Field Visualization

An effective 3D visualization should incorporate as many useful depth cues as can be reasonably accommodated into its design. Some depth cues were found to have a much stronger impact on judging glyph direction than others, and should be preferred if they are available.

The strongest such depth cues come from the shape-from-shading, linear perspective, and occlusion offered by tube-style glyph designs. Line-based glyphs are ubiquitous in computer-aided visualization because they are simple to implement and cheap to render. But, the results of the experiments carried out here prove that the perceptual benefits of a 3D tube

shape far outweigh the convenience of a line-style glyph for 3D flow visualizations. Therefore, line-based glyphs should be avoided for representing 3D flow information if it is to be conveyed accurately.

If a line-style glyph *must* be used, stereoscopic viewing should be very strongly considered. Errors in directional judgements of the two streamlet designs in Experiment 2 were almost halved when stereoscopic viewing was enabled. This means that it is a very powerful depth cue on its own, and may act to speed up judgements as well, but task speed was not evaluated here.

The ringed textures applied to tubes in Experiment 1 were added to enhance depth, but the ringed tubes did not perform any better than the plain tubes. The ringed texture may have contributed to glyphs visually fusing with one another. Since the ring color contrasts with the lighter tube color, it may diminish the ability of the visual system to use the subtler shading cues and high-contrast contour edges to disambiguate individual glyphs. For helping to distinguish occluding 3D elements from one another, the halo technique has been used extensively [7, 30, 76]. This modification could provide the necessary visual separation between glyphs to allow textures to provide the strong additive benefits that they have exhibited in other studies.

6.3 Perceiving 3D Glyph Direction

Glyph direction is best conveyed by animating texture elements along the glyph geometry. This provides the strongest and most unambiguous cue for conveying the vector sign of an oriented 3D flow glyph. Animation speed was not tested, but it is likely that very slow animation rates could be subtle enough to clearly disambiguate direction without becoming visually distracting. Since the animation is providing a binary cue – forward or backward – it only needs to be slightly above its perceptual threshold to be effective. This threshold could be analyzed by future studies.

Texture cues themselves did not appear to play a significant role in perceiving direction or

orientation, as ringed tubes did not perform any differently than plain tubes in Experiment 1. An explanation for this may be that the addition of texture introduces extra visual cognitive load instead of helpful cues, harming performance instead of helping it. This merits further exploration in the role of textures that directly encode information, as in Stoll et al.'s texture arrows [29] or Fuhrmann and Gröller's animated dashtubes [25] to indicate flow direction.

When animation is not an option for a visualization, a textured cone glyph geometry does a fine job of conveying direction, though not as well as animation does. It is important, however, to establish and make clear what convention is being used for the alignment of the cone, as the tip can be understood to point towards or against the direction of flow.

One potential issue with the use of a cone is that it could be confused with a tube foreshortened by perspective projection: the viewer perceives a tube retreating away from them, but it is actually the diameter of the cone getting smaller. This would only occur when other strong depth cues like occlusion, stereoscopic viewing, and animation are unavailable to disambiguate the true 3D shape. If neither animation nor stereo viewing cues are available, improving glyph-glyph occlusion cues could help to correct any misinterpreted foreshortening cues caused by the cone geometry. Applying a haloing technique to each glyph would help to reinforce the relative depth order of each glyph, so when the tapered end of a haloed cone occludes another glyph, any illusion of a foreshortened tube would be immediately shattered.

6.4 Perceiving 3D Glyph Length

Accurately judging glyph lengths appears to be a much more formidable task compared to perceiving their directions. Similar research has found equally complicated effects on distance and length judgements using 3D displays in object-grasping and estimation tasks, and that length judgements are very prone to perceptual biases and distortions that depend on many factors [77]. Those results also suggest that additional aids like haptic feedback and visual feedback from grasping the object being estimated would improve length judgements.

Viewing 3D flow glyphs on stereoscopic monitors may lead to some very inaccurate per-

ceptions of the flow glyph lengths, whether or not a head-tracked perspective is rendered. This research did not explore the potential differences between a stereoscopic 3D monitor and a stereoscopic head-mounted display, but the improved coupling between the head position and the display in the headset may improve spatial judgements and improve glyph length perceptions. It is a complex subject that should be investigated further.

Beyond its mapping to the length of a direct glyph design, the perception of vector speed/magnitude on integral glyphs merits further study, as this is an important characteristic of flow fields, and integral glyphs are just as common as direct glyphs for visualizing 3D flow fields.

6.5 Practical Considerations

Feedback is a critical facet of flow visualization design. When users are put off by a visualization, they will get little from it. Of course, cues that are objectively helpful in theory are not always practical. User feedback on visualizations is just as important as a grounding in perceptual theory; a visualization is ultimately worthless if users find it obnoxious or disagreeable. This was evidenced by the relatively good performance, but terrible user rating, of shadowed lines design in Experiment 1.

User feedback was overwhelmingly negative for the shadowed lines design. Although it achieved respectable performance results when compared to the other line-based techniques (see Figure 5.11), participants rated shadowed lines as the most negative in every category: aesthetics, ease of interpretation, speed of decision, and how well they thought they performed using the technique. These results underscore the importance of real-world evaluations to strike a balance between what works in the laboratory and what works in practice.

A separate practical consideration is the type of stereoscopic display that is feasible for use with a 3D vector field visualization. At the beginning of this research project, stereoscopic 3D computer monitors and television displays were mass-produced and widely available. In

a few short years, they have all but disappeared. However, there are now head-mounted stereoscopic displays that can largely take the place of the stereoscopic desktop displays used for this research. The performance impacts of this method of stereoscopic viewing over others are unknown in this context, but should be explored in future studies.

6.6 3D Flow Visualization Design Guidelines

The results of the three evaluations have provided a number of useful insights concerning the visualization of 3D flow data. Taken together, the following guidelines have been created to aid in the design of 3D flow field visualizations to improve their ability to accurately convey information by making them perceptually stronger and more robust.

G1: Use cones instead of tubes or lines for better orientation and direction perception. The streamcone provides stronger depth cues than either tube- or line-based glyph designs, and are a viable general-purpose glyph design because of their superb performance across many conditions.

G2: Use tubes instead of lines for better orientation perception. Their thicker 3D form presents stronger shape-from-shading cues, a clear advantage over line-based techniques.

G3: Avoid overly dense seeding of glyphs on cutting planes. Tightly-packed glyphs can visually fuse together. Tube-based glyphs are less sensitive to fusion issues. However, occlusion is an obstacle with any glyph design.

G4: Use caution when applying texture to glyphs. Unless the texture adds specific information, the additional visual noise can be detrimental.

G5: Cast shadows can improve direct glyph designs. They provide redundant orientation cues, as long as they can be unambiguously associated with the objects casting them.

If they are to be used, they must be applied with care, as they are tied to the lighting conditions of the visualization and may only be appropriate in limited situations.

G6: Stereoscopic viewing should be used whenever possible. It helps to reduce depth errors and disambiguate occlusion conflicts.

G7: When stereoscopic viewing is used, it should either be head-coupled through viewpoint tracking, or the non-tracked viewpoint should be explicitly indicated to the user (e.g. "best if viewed at 60cm away from center of screen").

G8: Animation is the best disambiguator for flow direction. Care should be taken when there are other sources of movement, such as the viewpoint changing, data set moving, or data changing over time; the multiple sources of motion can be difficult to separate and degrade perceptual performance.

G9: Use streamcones or a similar glyph design incorporating directional cues into its geometry when animation is not feasible.

G10: User feedback is just as important as perceptual theory for guiding a final design. Even if a design is demonstrably more effective than another, it must also appeal to the user to have practical value.

CHAPTER 7
APPLICATION

This chapter details the implementation of an interactive 3D flow visualization application using the design guidelines developed through the research reported in this dissertation.

7.1 Background and Motivation

To verify and validate the perceptual guidelines established in the prior chapter, a flow visualization tool was developed which conforms to those guidelines. This completes the bridge between perceptual theory, empirical evaluation of the predictions supported by the theory, and practical application of those empirical results. The application presented here is intended as a proof-of-concept interactive tool demonstrating a perceptually-superior visualization design.

Visualization, interaction, and presentation comprise the core features of the application. The visualization aspects of the application have been described and justified in the preceding sections, and applied in accordance with the design guidelines proposed in the previous section. Interacting with a visualization allows users to explore and gain insights into their data, and may even help improve spatial judgements [77]. The level of interaction should support the tasks they are meant to accomplish in a way that is most unobtrusive to the user; ideally, interactions would be intuitive so that little training is required, and demand little to no active attention from the user so that focus can remain on the task.

Recalling, capturing, and sharing data visualizations for presentation are not often given the same attention as visualization design and interaction. However, the results of the

evaluations and the goal of providing visualizations with the most useful visual cues to accurately perceive them suggests taking a more nuanced approach. Stereoscopic viewing improves accurate flow field perception so much that it should always be used when available, and modern virtual reality headsets put stereoscopic computer graphics within reach of more users than ever before. But what is the best approach to take if a user who finds an interesting visualization using stereoscopic viewing wants to include it in printed materials, or share it with a colleague who does not have (or may not be able to use) stereoscopic viewing equipment? Depth cues can combine in ways that can enhance or degrade the quality of a 3D visualization, so by taking into account a target viewing environment, the application or user can adjust the visual parameters to accommodate the most effective display of the selected flow visualization.

Applying the 3D glyph design lessons gleaned through this research in an interactive flow visualization tool is a method for substantiating the proposed design guidelines. This chapter describes (and justifies, where needed) the design choices and implementation of a 3D flow visualization application that combines intuitive data interactions, perceptually superlative visual representations, and an adaptable snapshot system for storing and sharing high-quality 3D flow visualizations.

7.1.1 Flow Visualization Software

Software tools are typically built upon other existing libraries and frameworks. The degree of reliance upon external code can depend on factors like security, licensing, resources, and, ultimately, the availability and robustness of these components. Computer graphics engines provide a framework that abstracts away much of the lower-level programming and hardware knowledge often required with computer graphics application programming interfaces. They are meant to facilitate a workflow, and do this by providing off-the-shelf components that help the user accomplish their task without needing to know the details of the underlying implementation.

For 3D flow visualization, the two most compatible engines to be considered are general-purpose scientific visualization engines and computer game engines. Scientific visualization engines have analytical functions as a primary focus, and visualizations do not tend to incorporate the most up-to-date techniques or graphical capabilities. On the other hand, computer game engines are focused on state-of-the-art graphical rendering and interactions for entertainment purposes.

Flow visualization software tools are not a novel concept. ParaView is one of the most widely-used turnkey applications for 3D scientific and industrial visualization [78]. It offers a variety of analytical tools in addition to its visualization capabilities, and promotes itself as a general-purpose scientific visualization toolkit. A visualization created with ParaView is included in Figure 7.1, and shows a "rake" of streamtubes advecting through a computational fluid dynamics (CFD) simulation of a motorbike in a wind tunnel. Depth cueing is used to blur the parts of the image to enhance depth and draw attention to the region of interest along the front tire. Shadows and lighting are also added to improve depth perception. In this particular image, it is only obvious which way the flow is directed because of the motorcycle giving context to the streamtubes. On their own, the streamtubes here do not indicate the vector sign of the direction whatsoever.

MATLAB [13] is another popular software tool with visualization capabilities for 3D flow visualization, though they are basic. The vector field images that appear in Chapter 2 of this dissertation were created using yet another commercial software tool, Mathematica by Wolfram Research, Inc. [14]. All of these software packages handle basic visualization and analysis, but 3D data is also difficult to interact with efficiently using a standard keyboard and mouse. The visualization capabilities of these software packages vary considerably, but all of them leave something to be desired when it comes to top-tier perceptual visualizations. For these reasons, a computer graphics and interaction platform will serve as the basis for the flow visualization application implemented here.

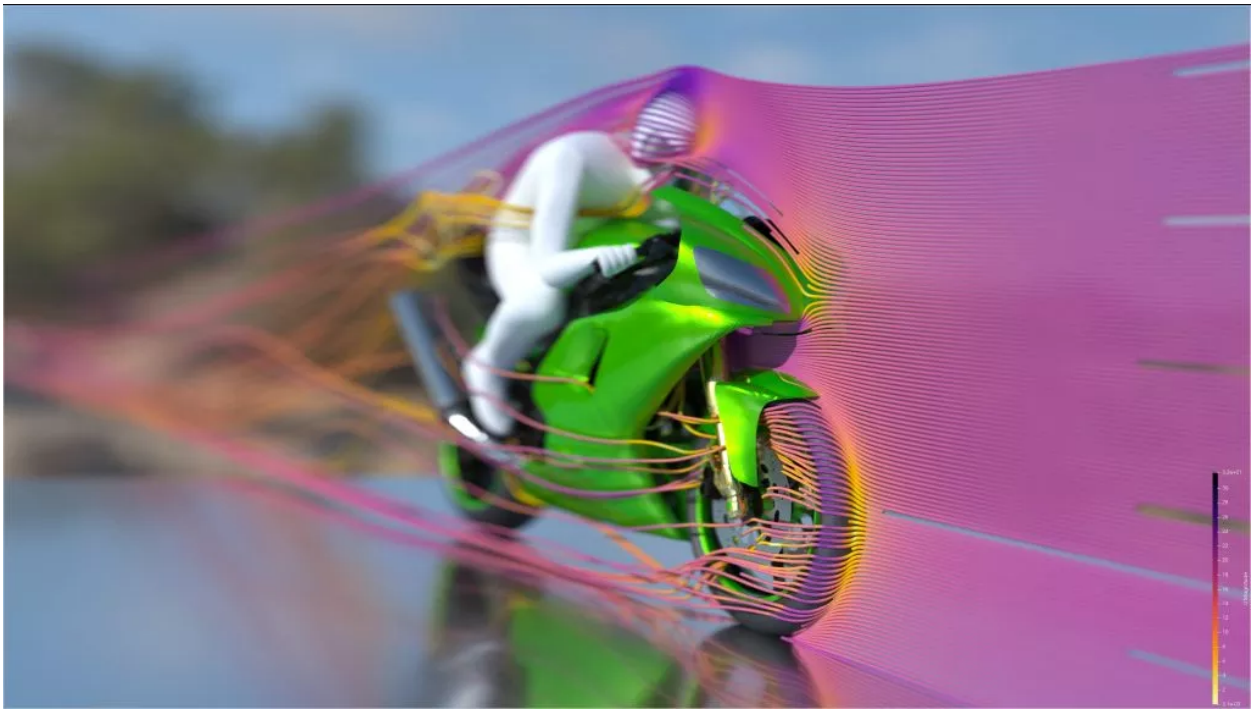

Figure 7.1: A sample visualization from the ParaView online image gallery showing depth of field and shadow effects to increase the sense of depth and draw the viewer's eye to the area of interest.

7.2 Approach

The goal of this proof-of-concept application is to demonstrate improved exploration and explanation of flow field features using guidelines derived from the results of the experiments discussed in the preceding chapter.

7.2.1 Visualization

The visual aspects of the tool are the most important, as they serve to corroborate the results of the experiments conducted in this dissertation. For that reason, the default style uses a streamcone geometry with ringed textures. The streamcones will taper along the direction of the flow, and the rings are parameterized to allow for flexibility in their design – including having no ring at all. The ring texture can be animated along the streamcone to redundantly and unambiguously emphasize the direction of the flow. The streamcones can cast shadows

on cutting planes, and have shadows cast upon themselves as well, including self-shadowing. Each type of shadow can be enabled or disabled as appropriate for the data being visualized. Since lighting and shading are important aspects of 3D form and contour perception, the lighting parameters of the visualization should be flexible as well.

In addition to providing strong depth cues, stereoscopic viewing makes 3D interactions like pointing and grabbing faster and more accurate. Stereoscopic desktop monitors no longer enjoy the popularity they once did, and affordably-priced, consumer-grade displays are becoming exceedingly rare to source. Fortunately, the recent proliferation of low-cost consumer-grade virtual reality headsets puts high-quality, head-tracked stereoscopic computer graphics within reach of more people than ever before. Therefore, stereoscopic head-mounted displays were targeted as the viewing hardware for the tool.

7.2.2 Data Interaction

Bruckner et al. [79] model the directness of spatial interaction in 3D visualizations using six related coordinate system spaces grouped into three categories. One-to-one direct interaction is desirable because it collapses all of those spaces into a single unified space, so there is no mental transformation required between the input space, the visualization space, and the display space, as is the case with mouse and keyboard input on a desktop monitor. This supports the use of spatially tracked six-degree-of-freedom (6DOF) input devices and head-mounted displays (HMDs) to minimize the extra cognitive burden imposed by mental transformations between coordinate system spaces.

Physical controllers are preferred for user input, because current gesture and hand detection systems are not yet dependable enough to truly facilitate tasks. It may be acceptable to miss every n^{th} gesture in casual applications like gaming and entertainment, but for serious work it does not lead to a dependable system when compared to the near 100% input reliability of a mouse and keyboard. A mouse click always registers, a keyboard press always registers, but the same cannot be said for detecting a pinch versus point gesture, or an open

versus closed hand. It is also much more fatiguing than a mouse and keyboard, because the wrists are not supported and must be held aloft in particular poses to be optimally recognized.

A mouse and keyboard is a very precise input system, relying on hands supported by a desk or other surface to steady them. Spatial interactions using handheld controllers are less precise because they do not have the same physical support, but we can help to overcome some of these limitations by providing strong ancillary cues. There are two main feedback mechanisms for error correction: visual and somatosensory (touch, force, temperature) feedback. These mechanisms have many subtleties and interact in complicated ways that are not yet fully understood [80], but there is strong evidence that combining the two increases the speed and precision of our hand movements [81].

7.3 Design and Function

The application was implemented within the Unity 3D development platform by Unity Technologies [82]. The Unity engine has a tight coupling between its editor interface used to construct a virtual scene and its player interface that acts as a live, debuggable preview of the virtual scene. In addition to significantly easing implementation effort of the core application, this feature enables rapid prototyping of new visualization designs thanks to its ability to reflect changes to the code and editor immediately without reloading the virtual scene.

The Unity Engine offers OpenXR support, which allows for hardware-agnostic virtual reality, mixed reality, and augmented reality application development. It supports a wide array of modern HMDs, and 6DOF input from myriad hardware. It is now possible to buy a fully-fledged, high-fidelity, spatially-tracked virtual reality kit for less than the cost of a new computer, making quality scientific visualization and interaction tools available to practically anyone.

An example of the Unity interface running a fully-interactive 3D flow visualization scene

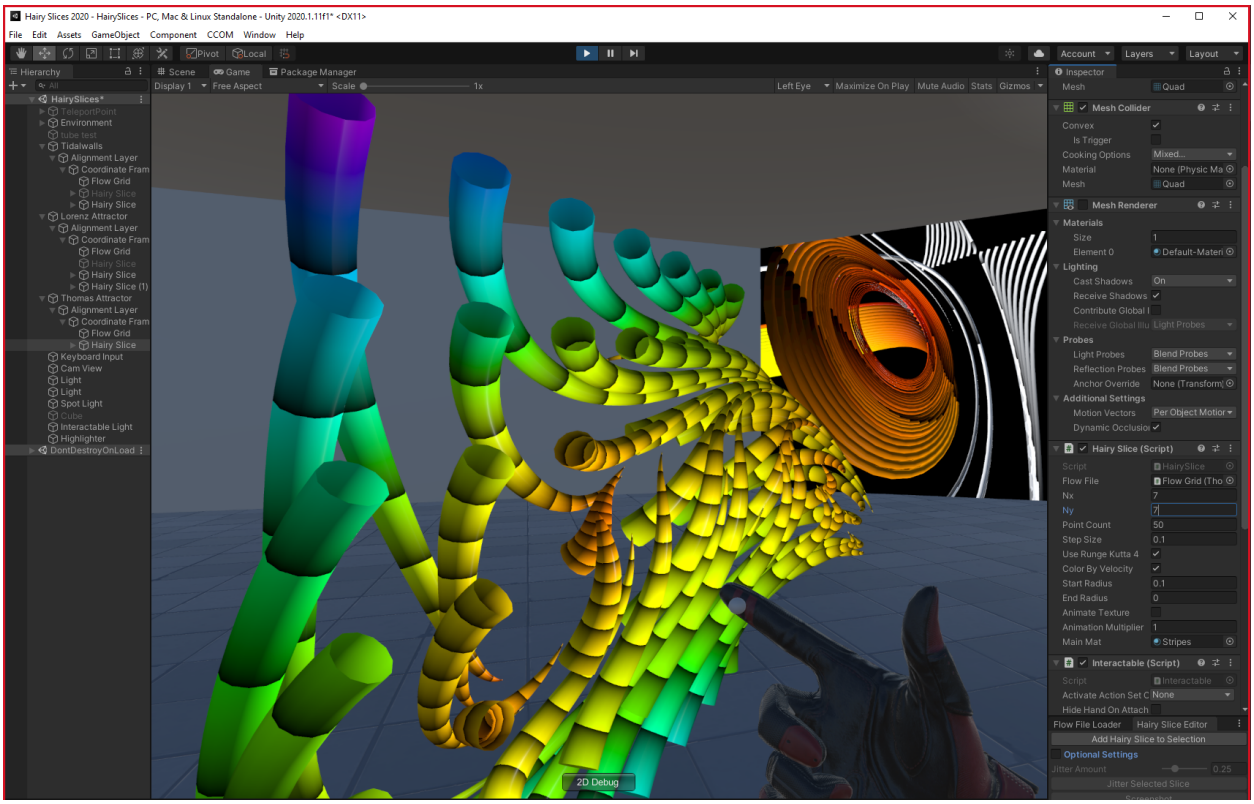

Figure 7.2: A screen capture of the Unity editor interface showing the custom extensions implemented to facilitate the workflow.

is shown in Figure 7.2. Adjustments to the flow visualization can be made interactively by the user wearing the headset and spatial controllers, or at the desktop computer using the controls of the Unity editor interface. Flow field data can be pre-compiled into a scene or loaded at runtime. Time-varying 3D (4D) flow data is also supported.

7.3.1 Code Organization

The main implementation was written as a "Hairy Slices" plugin that can be loaded into a Unity project to add 3D vector field data ingestion, interaction, and representation capabilities. Three main classes comprise the plugin: the Flow Source class, the Hairy Slice class, and the Hair class. Their structures and relationships are briefly described in the following sections. Additionally, the Unity editor itself was also extended to speed up the workflow needed to create an interactive visualization.

Flow Source Class

A Flow Source is the base class that takes care of data processing and defines pure virtual loading and sampling functions along with metadata accessors. Including compatibility for a certain data specification is accomplished by extending this base class and implementing its loading and sampling member functions. Analytical flow fields may be trivially implemented through sampling functions alone, since their solutions are exact and can be calculated as needed. In the case of numerical flow fields with sampled data points, the interpolation function (usually trilinear interpolation) is also implemented in this class to calculate flow values that lie between sampled points in space. In addition to steady-state flow fields, time-varying flow fields are also supported through simple linear interpolation between time steps.

The Flow Source converts the flow data to a local coordinate system in order to avoid precision issues that occur when performing mathematical operations on large floating point values – a very common occurrence with flow models that cover large geographical areas. The Flow Source class stores metadata and properties such as the original domain boundaries, the local domain boundaries, the range of flow velocities, and the amount of vertical exaggeration to apply to the depth dimension (needed for ocean models, which can be orders of magnitude larger in longitude and latitude than they are in depth).

There are visual components to the Flow Source class that can be enabled to aid user interaction or improve the clarity of the visualization. A 3D bounding box shows the extents of the Flow Source domain in the virtual environment, and provides visual cues and feedback to improve interactions like repositioning and rescaling the Flow Source. To further indicate the extents of the Flow Source and to increase visual clarity in certain situations, a dynamic backdrop can be enabled to make opaque the inner faces of the flow volume that are farthest from the viewer, so that the flow visualization within the Flow Source volume is always contrasted with the backdrop behind it.

Hairy Slice Class

The Hairy Slice class is essentially a bookkeeping layer between the Flow Source entities in a visualization and the individual "hair" entities that comprise the visualization. The class also provides functions to spatially distribute the hairs throughout the slice, like aligning seeds to a grid or randomly jittering their positions. Though a hairy slice implies a 2D plane, the class can accommodate 2D surfaces, 3D volumes, and paths; the hairy slice defines the spatial boundaries that limit the distribution of the hairs.

A Hairy Slice may be associated with any number of Flow Sources, so that a slice may extend through multiple Flow Sources, or none at all. Hairs, however, are only associated with a single Hairy Slice, though a Hairy Slice can have many hairs that it manages. The Hairy Slice class provides an optional bounding box for visualizing its extents, so that it is as unobtrusive as possible.

Hair Class

The Hair class is responsible for the visual components of the system. It stores the advection path, creates the mesh that is rendered from that path, and keeps track of the appropriate graphics shader and texture references. It is the seeding point from which flow field sampling initiates, and keeps the set of sampled points so that different mesh designs can be applied to the same streamline. This class controls all aspects of the hair's propagation through the flow field. This enables each Hair in a Hairy Slice to have its own visual parameters, if required.

7.3.2 Exploring and Visualizing 3D Flow Fields

The virtual environment is fully immersive, allowing the user to walk around as though it is a real, physical space. For seated use and smaller physical spaces, locomotion in the virtual world is accomplished through teleportation. The teleportation target is indicated by pointing the controller at the desired movement location, a visual marker is placed on

the virtual floor where the pointer intersects it, and a button press confirms the action. To control motion sickness effects, the scene is quickly faded to black and back in a fraction of a second with the user at their new location. This technique – called "blinking" – can sometimes be slightly disorienting, but should be preferred over potential nausea induced by motion sickness when the technique isn't used.

A single controller is used to directly reposition and reorient interactable objects (Flow Sources, Hairy Slices, light sources, snapshot camera, etc.) with a natural pinching or gripping gesture. This is similar to the the "one-handed grab" interaction in the Responsive Workbench [83]. When a second controller grabs the same object, the size of the object can be adjusted as the two controllers are brought together or apart in the "symmetric scale" fashion [83]. In addition to the visual feedback of "locking" the object to the grabbing hand at the grabbing point during the action, the controller emits a short haptic vibration pulse to reinforce the success of the grabbing action. This also makes it feasible to grab virtual objects without any visual feedback at all.

The visualization's construction and properties can be set *a priori* through the desktop editor interface, or interactively during runtime. This facilitates quickly setting up a rough flow visualization scene in the editor, and then making further adjustments in the virtual environment by leveraging the better viewing conditions of the HMD and the easier 3D interactions of the spatial controllers. Examples of 3D flow fields visualized using the virtual environment are shown in Figure 7.3.

A rectangular Hairy Slices may be interactively spawned by the user through a three-step single-handed technique. First, the user indicates an anchor point for one corner of the Hairy Slice with input from the spatial controller. The opposite corner of the slice is then anchored to the controller until the user is satisfied with the size and shape of the slice and ends the actions with an additional input from the controller. The Hairy Slice object is then instantiated, and can then be fine-tuned with scaling and grabbing interactions described earlier in addition to further adjusting its properties such as hair distribution and length. To

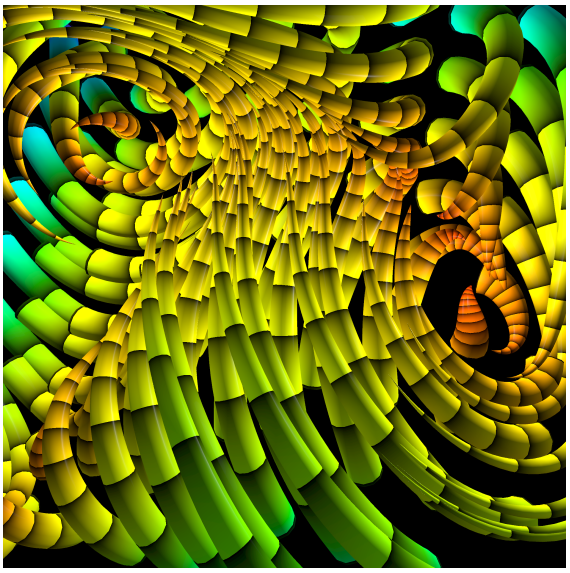

(a) Detail of a hairy slice through a flow field, visualized with long streamcones and a gradient ring to enhance direction and depth cues.

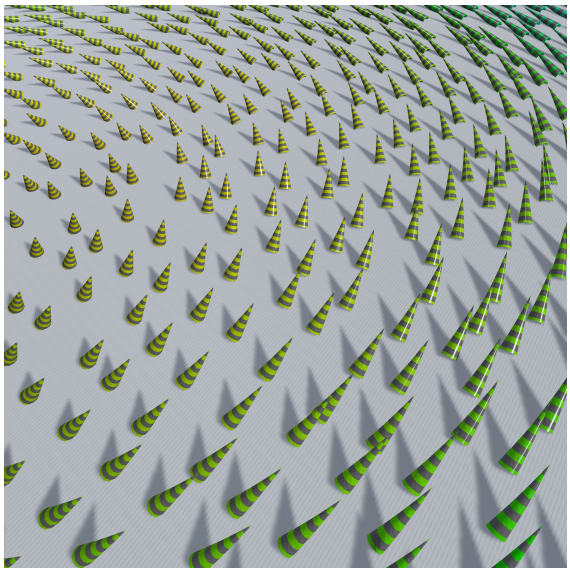

(b) A hairy slice using shorter streamcones with a ringed texture and cast shadows to improve local glyph perception in a close-up shot. The effectiveness of the shadows depends entirely upon the lighting, and is not always helpful.

Figure 7.3: Streamcone visualizations created with the tool incorporate strong depth and directional cues through shape, lighting, texture, and cast shadows.

destroy a Hairy Slice, it is grabbed with a controller and a button is held down for a short duration to avoid accidental activations.

Good lighting for a visualization plays a large part in its effectiveness and will depend upon the features of the flow field. Point lights can be added on-the-fly by pressing a button and grabbing the small grey spherical representation to position it. Using the same two-handed grab-and-scale interaction technique described earlier on the point light representation adjusts the point light intensity instead of the object scale, a logical extension of the metaphor that combines multiple actions into one interaction metaphor. A point light can have its position and orientation set to stay relative to the virtual environment or to a particular object (e.g., a Flow Source, a Hairy Slice, or even a single Hair itself). This configurable spatial relationship means that lighting parameters can be locked in for specific regions while other parts of the visualization are modified or moved for a better view.

Streamcone textures can be animated along the flow at any rate desired. When a ring

texture is set to transparent, the result looks like small tubelets racing through the flow field. It is a surprisingly powerful effect, with the added benefit that the visual density of the visualization is reduced significantly by the unrendered ring sections. It is, however, a technique that only works well with animation, and may be misleading if the flow field varies over time. Further investigation would help to clarify its strengths and weaknesses as a visualization technique.

7.3.3 Communication and Collaboration

Part of the power of this visualization tool comes from the facilitated interactions with 3D vector field data in a fully immersive virtual environment for enriched and perceptually stronger exploration of complex flow field data. The other novel component that this tool offers is a straight-forward, intuitive way to capture and share the high-quality visualizations put together and refined within the virtual environment. It relies upon a simple, but compelling virtual camera metaphor.

A virtual camera object can be instantiated via controller input. It has the appearance of a real-world digital camera with a telephoto lens to make the metaphor as clear as possible. It is an interactive object that can be grabbed and posed within the virtual environment. Unlike its real-world analog, it is not affected by gravity, so when the virtual camera is released, it stays in place. Camera properties like image aspect ratio and focal length are fully adjustable. The virtual camera can be set to capture the flow visualization objects (Hairs) with the rest of the virtual environment excluded and replaced with a solid color of the user's choosing in order to emphasize only the flow field. The back of the camera model has a virtual viewfinder that renders a live preview of the virtual camera's perspective. Additional live previews may be added to any surface. Figure 7.4 shows one such preview applied to a wall-sized surface in the virtual environment.

Having these additional views provides feedback to speed up the time it takes to get the camera into the desired pose. Another feature of the tool, which is more of a side-effect from

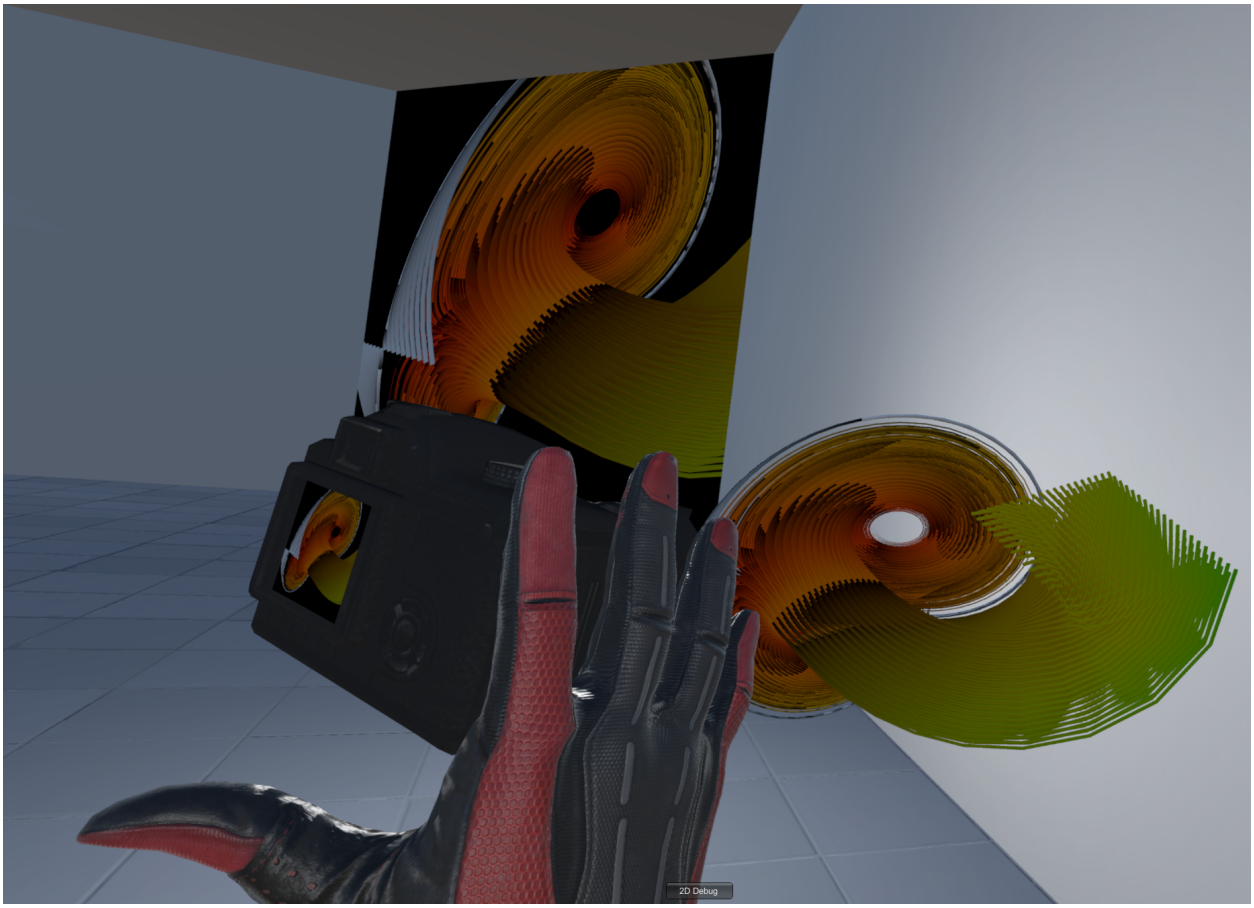

Figure 7.4: A screen capture demonstrating the interactive virtual camera that saves snapshots of flow visualizations. The back of the camera object shows a preview of the image it will capture, and another wall-sized image preview is displayed in a corner of the virtual environment to give multiple references while setting up a shot.

robust two-handed interactions, is that the virtual camera and the Flow Source or Hairy Slice may be grabbed and posed simultaneously to take advantage of the improved precision of bimanual interactions. When the object that needs the most precision in placement (the virtual camera) is held in the dominant hand, and the other object it is being posed to capture is held in the non-dominant hand, it forms what is referred to as a kinematic chain [84]. The dominant and non-dominant hands work together in an asymmetric fashion to apply error corrections along the kinematic chain from shoulder to each hand that, overall, increase the precision of the dominant hand's movements with regard to the non-dominant hand.

Once the camera is posed, there are a number of options available to save and share what

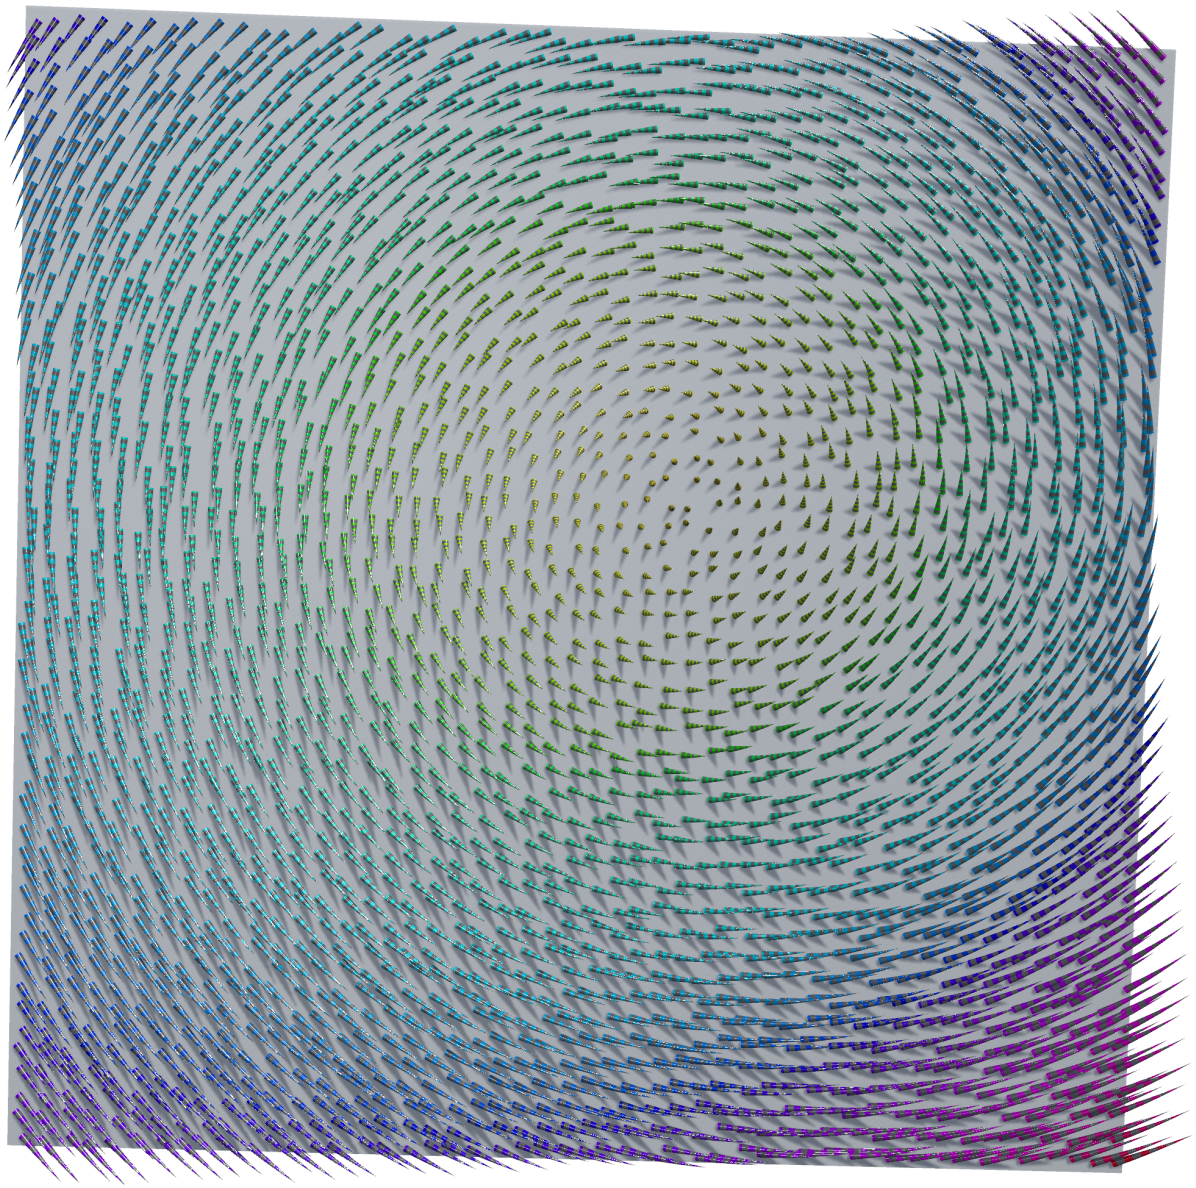

Figure 7.5: A snapshot made with the interactive virtual camera of a hairy slice through a flow field visualized using streamcones colored by magnitude and enhanced with a dark gray ring texture and cast shadows.

the camera sees. First, and most simply, leaving a camera instance pointed at a feature of interest would draw the attention of another user in the virtual environment. Secondly, a high-resolution snapshot can be saved to disk as fully-rendered images, or as a data snapshot of the related system states that can be recalled within the virtual environment. The saved

images can optionally be written as a stereo pair with a given separation from the virtual camera's perspective. In addition to still captures, sequences of images can be captured as the camera or visualization components are moved about the scene, and combined into an animation. Figure 7.5 is an example of a still image captured by the virtual camera system using a white background to isolate the flow visualization and increase the visual contrast when placed on the page.

Aside from the virtual camera, sharing a visualization scene can be accomplished by building a curated scene as its own self-contained virtual experience that can be run independently of the Unity editor. This is the ideal way to share a visualization produced with this system, because it allows for the full level of visual and interaction fidelity as the tool used to create it.

One useful feature of the Unity editor interface is that it can be accessed at the same time the virtual environment is being rendered to the headset. This opens up the possibility for one user to control the Unity scene editor while a separate user interacts within the virtual world via the HMD and 6DOF controllers. Unity also includes off-the-shelf networking plugins to handle the synchronization of many clients for multiplayer gaming, allowing remote collaborations to take place within the same virtual environment.

7.4 Outlook

This chapter describes an immersive 3D flow visualization tool built with the goal of facilitating the exploration and illustration of 3D vector fields. The application follows best-practice design guidelines enumerated in the previous chapter by implementing streamcone glyphs with animated textures using a head-tracked stereoscopic headset display. Interaction metaphors improve the speed and precision of 3D data manipulation in the virtual environment. A virtual camera system extends the visualization tool by providing an intuitive method for capturing detailed images of flow field features that can be saved and shared in multiple flexible ways.

A human factors user study with similar tasks to the three experiments in this research would help to assess the relative strengths and weaknesses of the application. This is especially true in the context of the 3D interactions this tool employs, as research has shown that spatial interactions can change our perception of the shape and layout of objects, and this might modulate the strength of a glyph design that was successful in the non-interactive stimulus used for the three earlier experiments. This tool could also serve as a platform from which to conduct 3D flow field interaction and visualization experiments.

One possibility is to explore the effects of SfM cues from a tracked HMD on a similar glyph direction judgement task within the virtual environment. There are additional proprioceptive cues available through even subtle head movements that can improve the spatial perception of an immersive virtual environment, but it is not known how these additional cues might affect a more localized spatial judgement task.

Evaluating glyph length judgement using a similar setup to Experiment 3 but swapping the stereoscopic desktop monitor for an HMD is a straightforward extension that would help us understand whether the patterns observed were a systemic effect of the monitor and viewing setup, or if there is something more complex occurring. Conducting this experiment with an HMD frees up some of the experimental constraints imposed by the use of a desktop display in Experiment 3. In particular, the input measuring rod would no longer need to stay oriented vertically to avoid distortion, so the glyph stimulus and input rod could be swapped in their horizontal and vertical placements in order to probe any Horizontal-Vertical Illusion effects in a virtual environment viewed through a stereoscopic HMD.

Seeding strategies are another area ripe for further investigation within this tool. The results of Experiment 1 suggest that seeding can be moderately dense, but the design of the glyph itself will have an impact upon how dense the seeds can be before visual noise begins to degrade the visualization's effectiveness. One potential strategy for reducing visual noise is to simply count the number of pixels belonging to glyphs in a rendered visualization, and re-seed and re-render the visualization until the percentage of glyph pixels is within a

desired range. Taking this idea further, an additional data buffer could be used to count how many times a glyph object was rendered to each pixel, so that areas with many glyphs being rendered on top of one another could be re-seeded and re-evaluated until glyph-glyph occlusions are reduced below a certain threshold.

Another line of research could look at further improving directional cues through animation and texture. Animation likely needs to just barely be noticeable in order to unambiguously understand the direction of flow, and an experiment could be designed to find this threshold rate of animation. There are also additional textures to evaluate that unambiguously encode direction into them, like faded chevrons or drawing arrows/droplets along a glyph, so that antiparallel confusion is kept to a minimum.

There are many other tasks involved in flow visualization that could be studied as well. There are some objectives that are not easily answered by computer algorithms because we do not yet know what we are seeking. Exploratory pattern-finding, judging relative high and low areas of magnitude/vorticity/turbulence related to these patterns, and identifying regions with swirling and tumbling flow characteristics are examples of such goals. These tasks rely on a solid holistic understanding of the visualization when compared to the more localized judgement tasks that were studied in this research, and gaining insight into the performance of glyph designs for those ends will further improve our ability to produce effective 3D flow visualizations.

The clarity with which flow features are visualized and the engaging, immersive environment suggest some potential as an educational tool, similar to the FlowVisual tool by Wang et al. for 2D flow visualization [27]. Analytical capabilities could be added so that critical points and types of flow patterns can be highlighted and visually explained, as well as the differences between flow visualization techniques like streamlines, pathlines, streaklines, and timelines. Other 3D mathematical objects (e.g., a Gaussian Radial Basis Function or other 3D kernel functions) could also be explored and better understood with this type of interactive visualization tool.

Another promising use is with next-generation augmented reality devices, leaving behind the virtual environment and depicting the flow visualizations in the real world. Combining a 3D flow simulation visualization with the real, physical subject being modeled can enable powerful insights into a model's design or potential improvements. This type of multi-modal or mixed reality visualization is not novel from a research standpoint, and has seen decades of research into the topic, but has much potential from a practical standpoint now that high-resolution, lightweight augmented reality headsets can let us apply these visualization techniques outside of a purely virtual visualization environment.

CHAPTER 8
CONCLUSION

The visualization of 3D flow fields is a mature topic, but still presents many inherent challenges, due to issues such as occlusion and the difficultly of providing effective depth cues on small glyphs. Myriad visualization techniques have been developed, and many successful designs have been identified. Despite widespread use of these visualizations in the sciences, there have been few empirical studies of these designs to ascertain the underlying perceptual cues that make them effective. This work expanded upon previous perceptual research into effective glyph design by conducting new human factors experiments that examined the benefits of the most salient depth and directional cues as applied to a variety of 3D glyphs. This research directly addresses the lack of objective, quantitative formal studies that explore the perceptual effectiveness of three-dimensional flow visualizations by conducting a series of experiments investigating the visual cues integral to the use of 3D glyphs on cutting planes.

A set of human factors experiments confirmed the hypotheses that increasing the dimensionality of lines into tubular structures enhances their ability to convey orientation through shading, and that increasing their diameter intensifies this effect. It also confirmed that adding shadows to lines increases depth perception. Interestingly, while it was hypothesized that adding a ringed texture to tubes would increase depth and orientation perception, no significant effect was found, possibly due to increased visual noise. Nonetheless, ringed textures remain useful for indicating the vector sign component of the flow direction when animated. A better alternative may be to use a texture which unambiguously encodes direction, so that animation is not needed to remove all doubt about the true flow direction.

The results also strongly support the use of 3D glyphs such as streamtubes and streamcones in rendering flow patterns through cutting planes. Although they cannot display data at the same density as thinner, line-based glyphs, they more effectively convey 3D direction information and are less prone to antiparallel misperceptions. These glyph designs can be so effective that their performance reaches a plateau where no additional accuracy benefits accrue from the addition of stereoscopic viewing or structure-from-motion.

Human performance at judging the length of a simple tube-like glyph with stereoscopic viewing from a correct and incorrect viewpoint did not follow the predictions of the geometric distortion model that can be used to describe them. The findings have nonetheless indicated that head-coupled perspective viewing of stereoscopic images generally leads to a more genuine perception of the geometries depicted, and should be used whenever possible to provide more faithful viewing of 3D scientific data. Participants were able to more accurately judge glyph lengths with correct stereoscopic viewing. The benefits of head-coupled stereoscopic viewing become increasingly important as viewing angles become more oblique to the display. There was also evidence that the size of the glyph representation itself has a significant effect on the human visual system's ability to discern its length; shorter glyph lengths are more accurately estimated, especially at more oblique viewing angles. Although new observations have been gathered about how humans perceive stereoscopic distortions, it is still not clear how this phenomenon should be modeled, and it will require further investigation.

From these experimental results, a set of guidelines were created for designing effective glyph-based 3D vector field visualizations. These guidelines are prescriptive, in that they steer a designer in what to choose, rather than being descriptive observations of the visualization techniques. They discuss a series of best-practices for improving a 3D flow visualization with superior depth and directional cues, and should be expanded by conducting additional research on 3D flow glyph effectiveness.

An immersive, interactive 3D flow visualization tool was developed following the guide-

lines enumerated in this dissertation in order to corroborate them in a real-world setting outside of the laboratory. It was built to facilitate exploring and sharing 3D vector fields by implementing streamcone glyphs with animated textures using a head-tracked stereoscopic headset display and six degree of freedom controllers. Intuitive interaction metaphors were implemented to improve the speed and precision of 3D data manipulation in the virtual environment. A flexible virtual camera system provides a straightforward method of capturing high-resolution flow field images that can be saved and shared. The tool also has promise as a 3D flow interaction and visualization platform for future human factors studies, as well as potential educational applications.

This research has begun resolving some of the key perceptual issues surrounding the accurate visualization of the principal components of 3D vector fields. It has established a series of strong, informative results upon which future 3D flow visualization techniques can be built. Our understanding of the visual cues and viewing conditions that impact the perception of three dimensional flow field information has expanded, but still remains a largely unexplored and complex topic. Only through sustained investigation into the many flow visualization tasks and design choices can we work to reveal the secrets and surprises of human perception and apply those lessons to continue improving 3D flow field visualizations into the future.

LIST OF REFERENCES

[1] Tamara Munzner. *Visualization analysis and design*. CRC press, 2014.

[2] Ed Boring and Alex Pang. Directional flow visualization of vector fields. In *Proceedings of Seventh Annual IEEE Visualization'96*, pages 389–392. IEEE, 1996.

[3] Niklas Elmqvist and Philippas Tsigas. A taxonomy of 3d occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1095–1109, 2008.

[4] Robert S Laramee, Daniel Weiskopf, Jürgen Schneider, and Helwig Hauser. Investigating swirl and tumble flow with a comparison of visualization techniques. In *IEEE Visualization 2004*, pages 51–58. IEEE, 2004.

[5] Tony McLoughlin, Robert S Laramee, Ronald Peikert, Frits H Post, and Min Chen. Over two decades of integration-based, geometric flow visualization. *Computer Graphics Forum*, 29(6):1807–1829, 2010.

[6] Brian Cabral and Leith Casey Leedom. Imaging vector fields using line integral convolution. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 263–270, 1993.

[7] Victoria Interrante and Chester Grosch. Visualizing 3d flow. *IEEE Computer Graphics and Applications*, 18(4):49–53, 1998.

[8] Frits H Post, Benjamin Vrolijk, Helwig Hauser, Robert S Laramee, and Helmut Doleisch. Feature extraction and visualisation of flow fields. *Eurographics (STARs)*, 2, 2002.

[9] Gordon V Bancroft, Fergus J Merritt, Todd C Plessel, Paul G Kelaita, R Kevin McCabe, and Al Globus. Fast: a multi-processed environment for visualization of computational fluid dynamics. In *Proceedings of the First IEEE Conference on Visualization: Visualization90*, pages 14–27. IEEE, 1990.

[10] Mike Giles and Robert Haimes. Advanced interactive visualization for cfd. *Computing Systems in Engineering*, 1(1):51–62, 1990.

[11] Craig Upson, TA Faulhaber, David Kamins, David Laidlaw, David Schlegel, Jefrey Vroom, Robert Gurwitz, and Andries Van Dam. The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42, 1989.

[12] Will J Schroeder, Bill Lorensen, and Ken Martin. *The visualization toolkit: an object-oriented approach to 3D graphics*. Kitware, 2006.

[13] MATLAB. *version 9.0.0 (R2016a)*. The MathWorks Inc., Natick, Massachusetts, 2016.

[14] Wolfram Research, Inc. Mathematica, Version 12.2. Champaign, IL, 2020.

[15] PowerFLOW. *version 5.4*. Exa Corporation, Burlington, Massachusetts, 2017.

[16] James R Osborn and Alice M Agogino. An interface for interactive spatial reasoning and visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 75–82, 1992.

[17] Joachim Diepstraten, Daniel Weiskopf, and Thomas Ertl. Interactive cutaway illustrations. *Computer Graphics Forum*, 22(3):523–532, 2003.

[18] Martin Schulz, Frank Reck, W Bertelheimer, and Thomas Ertl. Interactive visualization of fluid dynamics simulations in locally refined cartesian grids. In *Proceedings Visualization'99 (Cat. No. 99CB37067)*, pages 413–553. IEEE, 1999.

[19] Tom Meyer and Al Globus. Direct manipulation of isosurfaces and cutting planes in virtual environments. *Department of Computer Science, Brown University*, 1993.

[20] Ken Hinckley, Randy Pausch, John C Goble, and Neal F Kassell. Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 452–458, 1994.

[21] David H Laidlaw, Robert M Kirby, Cullen D Jackson, J Scott Davidson, Timothy S Miller, Marco Da Silva, William H Warren, and Michael J Tarr. Comparing 2d vector field visualization methods: A user study. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):59–70, 2005.

[22] Zhanping Liu, Shangshu Cai, J Edward Swan, Robert J Moorhead, Joel P Martin, and TJ Jankun-Kelly. A 2d flow visualization user study using explicit flow synthesis and implicit task design. *IEEE Transactions on Visualization and Computer Graphics*, 18(5):783–796, 2011.

[23] Olaf Konrad-Verse, Arne Littmann, and Bernhard Preim. Virtual resection with a deformable cutting plane. In *SimVis*, pages 203–214, 2004.

[24] David Modiano. *Visualization of three dimensional CFD results*. PhD thesis, Massachusetts Institute of Technology, 1989.

[25] Anton Fuhrmann and Eduard Groller. Real-time techniques for 3d flow visualization. In *Proceedings Visualization'98 (Cat. No. 98CB36276)*, pages 305–312. IEEE, 1998.

[26] Craig M Wittenbrink, Alex T Pang, and Suresh K Lodha. Glyphs for visualizing uncertainty in vector fields. *IEEE transactions on Visualization and Computer Graphics*, 2(3):266–279, 1996.

[27] Man Wang, Jun Tao, Chaoli Wang, Ching-Kuang Shene, and Seung Hyun Kim. Flowvisual: Design and evaluation of a visualization tool for teaching 2d flow field concepts. In *2013 ASEE Annual Conference & Exposition*, pages 23–609, 2013.

[28] Malte Zockler, Detlev Stalling, and H-C Hege. Interactive visualization of 3d-vector fields using illuminated stream lines. In *Proceedings of Seventh Annual IEEE Visualization'96*, pages 107–113. IEEE, 1996.

[29] Carsten Stoll, Stefan Gumhold, and H-P Seidel. Visualization with stylized line primitives. In *VIS 05. IEEE Visualization, 2005.*, pages 695–702. IEEE, 2005.

[30] Andreas Wenger, Daniel F Keefe, Song Zhang, and David H Laidlaw. Interactive volume rendering of thin thread structures within multivalued scientific data sets. *IEEE Transactions on Visualization and Computer Graphics*, 10(6):664–672, 2004.

[31] Dave Darmofal and Robert Haimes. Visualization of 3-d vector fields-variations on a stream. In *30th Aerospace Sciences Meeting and Exhibit*, page 74, 1992.

[32] Rita Borgo, Johannes Kehrer, David HS Chung, Eamonn Maguire, Robert S Laramee, Helwig Hauser, Matthew Ward, and Min Chen. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics (STARs)*, pages 39–63, 2013.

[33] Andreas E Lie, Johannes Kehrer, and Helwig Hauser. Critical design and realization aspects of glyph-based 3d data visualization. In *Proceedings of the 25th Spring Conference on Computer Graphics*, pages 19–26, 2009.

[34] Jian Chen, Haipeng Cai, Alexander P Auchus, and David H Laidlaw. Effects of stereo and screen size on the legibility of three-dimensional streamtube visualization. *IEEE transactions on Visualization and Computer Graphics*, 18(12):2130–2139, 2012.

[35] Andrea Brambilla, Robert Carnecky, Robert Peikert, Ivan Viola, and Helwig Hauser. Illustrative flow visualization: State of the art, trends and challenges. *Visibility-oriented Visualization Design for Flow Illustration*, 2012.

[36] Colin Ware. *Information visualization: perception for design*. Morgan Kaufmann, 2019.

[37] Jeremy M Wolfe, Keith R Kluender, Dennis M Levi, Linda M Bartoshuk, Rachel S Herz, Roberta L Klatzky, Susan J Lederman, and Daniel M Merfeld. *Sensation and Perception*. Sinauer Associates, 2015.

[38] Mark F Bradshaw, Andrew D Parton, and Andrew Glennerster. The task-dependent use of binocular disparity and motion parallax information. *Vision Research*, 40(27):3725–3734, 2000.

[39] Graeme Sweet and Colin Ware. View direction, surface orientation and texture orientation for perception of surface shape. In *Proceedings of Graphics Interface 2004*, GI '04, page 97–106, Waterloo, CAN, 2004. Canadian Human-Computer Communications Society.

[40] Colin Ware. 3d contour perception for flow visualization. In *Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, pages 101–106, 2006.

[41] Andrew Forsberg, Jian Chen, and David Laidlaw. Comparing 3d vector field visualization methods: A user study. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1219–1226, 2009.

[42] R. L. Sollenberger and P. Milgram. Effects of stereoscopic and rotational displays in a three-dimensional path-tracing task. *Human Factors*, 35(3):483–499, 1993.

[43] Colin Ware and Glenn Franck. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15(2):121–140, 1996.

[44] Maurice HPH van Beurden, Andre Kuijsters, and Wijnand A IJsselsteijn. Performance of a path tracing task using stereoscopic and motion based depth cues. In *2010 Second International Workshop on Quality of Multimedia Experience (QoMEX)*, pages 176–181. IEEE, 2010.

[45] Erol Aygar, Colin Ware, and David Rogers. The contribution of stereoscopic and motion depth cues to the perception of structures in 3d point clouds. *ACM Transactions on Applied Perception (TAP)*, 15(2):1–13, 2018.

[46] Vilayanur S Ramachandran. Perception of shape from shading. *Nature*, 331(6152):163, 1988.

[47] Roland Arsenault and Colin Ware. The importance of stereo and eye-coupled perspective for eye-hand coordination in fish tank vr. *Presence: Teleoperators & Virtual Environments*, 13(5):549–559, 2004.

[48] Colin Ware and Peter Mitchell. Reevaluating stereo and motion cues for visualizing graphs in three dimensions. In *Proceedings of the 2nd Symposium on Applied Perception in Graphics and Visualization*, pages 51–58, 2005.

[49] Kevin W. Arthur, Kellogg S. Booth, and Colin Ware. Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems*, 11(3):239–265, 1993.

[50] Colin Ware, Daniel Bolan, Ricky Miller, David H Rogers, and James P Ahrens. Animated versus static views of steady flow patterns. In *Proceedings of the ACM Symposium on Applied Perception*, pages 77–84, 2016.

[51] Michael Deering. High resolution virtual reality. *ACM SIGGRAPH Computer Graphics*, 26(2):195–202, 1992.

[52] Colin Ware, Kevin Arthur, and Kellogg S Booth. Fish Tank Virtual Reality. *Proc. InterCHI*, pages 37–42, 1993.

[53] Andrew J. Woods, Tom Docherty, and Rolf Koch. Image distortions in stereoscopic video systems. *Stereoscopic Displays and Applications IV*, pages 36–48, September 1993.

[54] Robert T. Held and Martin S. Banks. Misperceptions in stereoscopic displays. *Proceedings of the 5th symposium on Applied perception in graphics and visualization - APGV '08*, 1(212):23, 2008.

[55] Atanas Boev, Danilo Hollosi, and Atanas Gotchev. Classification of stereoscopic artefacts. *Mobile3DTV Project report*, 2008.

[56] Jung-Young Son, Yuri Gruts, Joohwan Chun, Yong Jin Choi, Ji-Eun Bahn, and Vladimir I. Bobrinev. Distortion analysis in stereoscopic images. *Optical Engineering*, 41(3):680, 2002.

[57] Martin S. Banks, Emily A. Cooper, and Elise A. Piazza. Camera Focal Length and the Perception of Pictures. *Ecological Psychology*, 2014.

[58] Martin S. Banks, Robert T. Held, and Ahna R. Girshick. Perception of 3-D layout in stereo displays, 2009.

[59] Devon Penney, Jian Chen, and David H Laidlaw. Effects of illumination, texture, and motion on task performance in 3d tensor-field streamtube visualizations. In *2012 IEEE Pacific Visualization Symposium*, pages 97–104. IEEE, 2012.

[60] Andrew H Stevens, Thomas Butkiewicz, and Colin Ware. Hairy slices: Evaluating the perceptual effectiveness of cutting plane glyphs for 3d vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):990–999, 2016.

[61] Andrew H. Stevens, Colin Ware, Thomas Butkiewicz, David Rogers, and Greg Abram. Hairy Slices II: Depth Cues for Visualizing 3D Streamlines Through Cutting Planes. *Computer Graphics Forum*, 39(3):25 – 35, 2020.

[62] Ovidio Mallo, Ronald Peikert, Christian Sigg, and Filip Sadlo. Illuminated lines revisited. In *VIS 05. IEEE Visualization, 2005.*, pages 19–26. IEEE, 2005.

[63] R Victor Klassen and Steven J Harrington. Shadowed hedgehogs: A technique for visualizing 2d slices of 3d vector fields. In *Proceeding Visualization'91*, pages 148–153. IEEE, 1991.

[64] FW Campbell and DG Green. Optical and retinal factors affecting visual resolution. *The Journal of physiology*, 181(3):576, 1965.

[65] Christophe Leys, Christophe Ley, Olivier Klein, Philippe Bernard, and Laurent Licata. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, 49(4):764–766, 2013.

[66] Michael G Kenward and James H Roger. Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*, pages 983–997, 1997.

[67] William H Kruskal and W Allen Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.

[68] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.

[69] Brittney Hartle and Laurie M. Wilcox. Depth magnitude from stereopsis: Assessment techniques and the role of experience. *Vision Research*, 125:64–75, 2016.

[70] Dong Wook Kim, Kwang Hoon Lee, Seon Kyu Yoon, Yong Moo Kwon, Eun Young Chang, and Sung Kyu Kim. Analysis and compensation of distorted stereoscopic images due to the observer's viewing orientation. *Optics Communications*, 285(19):3909–3916, 2012.

[71] Pom Charras and Juan Lupiáñez. The relevance of symmetry in line length perception. *Perception*, 38(10):1428–1438, 2009.

[72] Pom Charras and Juan Lupiáñez. Length perception of horizontal and vertical bisected lines. *Psychological Research PRPF*, 74(2):196–206, 2010.

[73] Uta Wolfe, Laurence T Maloney, and Mimi Tam. Distortions of perceived length in the frontoparallel plane: Tests of perspective theories. *Perception & psychophysics*, 67(6):967–979, 2005.

[74] Kyriaki Mikellidou and Peter Thompson. The vertical-horizontal illusion: Assessing the contributions of anisotropy, abutting, and crossing to the misperception of simple line stimuli. *Journal of vision*, 13(8):7–7, 2013.

[75] Pascal Mamassian and Marie de Montalembert. A simple model of the vertical–horizontal illusion. *Vision Research*, 50(10):956–962, 2010.

[76] Maarten H Everts, Henk Bekker, Jos BTM Roerdink, and Tobias Isenberg. Depth-dependent halos: Illustrative rendering of dense line data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, 2009.

[77] Carlo Campagnoli, Sholei Croom, and Fulvio Domini. Stereovision for action reflects our perceptual experience of distance and depth. *Journal of Vision*, 17(9):21–21, 2017.

[78] James Ahrens, Berk Geveci, and Charles Law. Paraview: An end-user tool for large data visualization. *The visualization handbook*, 717(8), 2005.

[79] Stefan Bruckner, Tobias Isenberg, Timo Ropinski, and Alexander Wiebel. A model of spatial directness in interactive visualization. *IEEE transactions on visualization and computer graphics*, 25(8):2514–2528, 2018.

[80] Liliana Rincon-Gonzalez, Jay P Warren, David M Meller, and Stephen Helms Tillery. Haptic interaction of touch and proprioception: implications for neuroprosthetics. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(5):490–500, 2011.

[81] Katherine J Kuchenbecker, Netta Gurari, and Allison M Okamura. Effects of visual and proprioceptive motion feedback on human control of targeted movement. In *2007 IEEE 10th International Conference on Rehabilitation Robotics*, pages 513–524. IEEE, 2007.

[82] Unity Engine. *version 2020.1.11f1*. Unity Technologies, Inc., San Francisco, California, 2020.

[83] Lawrence D Cutler, Bernd Frohlich, and Pat Hanrahan. Two-handed direct manipulation on the responsive workbench. In *Proceedings of the 1997 Symp. on Interactive 3D graphics*, 1997.

[84] Yves Guiard. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of motor behavior*, 19(4):486–517, 1987.

APPENDIX A

The Institutional Review Board for the Protection of Human Subjects in Research (IRB) Approval

University *of* New Hampshire

Research Integrity Services, Service Building
51 College Road, Durham, NH 03824-3585
Fax: 603-862-3564

14-Jun-2017

Stevens, Andrew H
Center for Coastal & Ocean Mapping,
24 Colovos Rd
Durham, NH 03824

IRB #: 6719
Study: Virtual and Augmented Realities to Support the Acquisition, Processing, Analysis and Display of Oceanographic Data
Approval Date: 11-Jun-2017

The Institutional Review Board for the Protection of Human Subjects in Research (IRB) has reviewed and approved the protocol for your study as Expedited as described in Title 45, Code of Federal Regulations (CFR), Part 46, Subsection 110.

Approval is granted to conduct your study as described in your protocol for one year from the approval date above. At the end of the approval period, you will be asked to submit a report with regard to the involvement of human subjects in this study. If your study is still active, you may request an extension of IRB approval.

Researchers who conduct studies involving human subjects have responsibilities as outlined in the document, *Responsibilities of Directors of Research Studies Involving Human Subjects.* This document is available at http://unh.edu/research/irb-application-resources. Please read this document carefully before commencing your work involving human subjects.

If you have questions or concerns about your study or this approval, please feel free to contact me at 603-862-2003 or Julie.simpson@unh.edu. Please refer to the IRB # above in all correspondence related to this study. The IRB wishes you success with your research.

For the IRB,

Julie F. Simpson
Director

cc: File
Moore, Cheryl
Butkiewicz, Thomas

University *of* New Hampshire

Research Integrity Services, Service Building
51 College Road, Durham, NH 03824-3585
Fax: 603-862-3564

22-May-2018

Stevens, Andrew H
Center for Coastal & Ocean Mapping,
24 Colovos Rd
Durham, NH 03824

IRB #: 6719
Study: Virtual and Augmented Realities to Support the Acquisition, Processing, Analysis and Display of Oceanographic Data
Review Level: Expedited
Approval Expiration Date: 11-Jun-2019

The Institutional Review Board for the Protection of Human Subjects in Research (IRB) has reviewed and approved your request for time extension for this study. Approval for this study expires on the date indicated above. At the end of the approval period you will be asked to submit a report with regard to the involvement of human subjects. If your study is still active, you may apply for extension of IRB approval through this office.

Researchers who conduct studies involving human subjects have responsibilities as outlined in the document, *Responsibilities of Directors of Research Studies Involving Human Subjects.* This document is available at http://unh.edu/research/irb-application-resources or from me.

If you have questions or concerns about your study or this approval, please feel free to contact me at 603-862-2003 or Julie.simpson@unh.edu. Please refer to the IRB # above in all correspondence related to this study. The IRB wishes you success with your research.

For the IRB,

Julie F. Simpson
Director

cc: File
Butkiewicz, Thomas
,

University *of* New Hampshire

Research Integrity Services, Service Building
51 College Road, Durham, NH 03824-3585
Fax: 603-862-3564

16-May-2019

Stevens, Andrew H
Center for Coastal & Ocean Mapping,
24 Colovos Rd
Durham, NH 03824

IRB #: 6719
Study: Virtual and Augmented Realities to Support the Acquisition, Processing, Analysis and Display of Oceanographic Data
Review Level: Expedited
Approval Expiration Date: 11-Jun-2020

The Institutional Review Board for the Protection of Human Subjects in Research (IRB) has reviewed and approved your request for time extension for this study. Approval for this study expires on the date indicated above. At the end of the approval period you will be asked to submit a report with regard to the involvement of human subjects. If your study is still active, you may apply for extension of IRB approval through this office.

Researchers who conduct studies involving human subjects have responsibilities as outlined in the document, *Responsibilities of Directors of Research Studies Involving Human Subjects.* This document is available at http://unh.edu/research/irb-application-resources or from me.

If you have questions or concerns about your study or this approval, please feel free to contact me at 603-862-2003 or Julie.simpson@unh.edu. Please refer to the IRB # above in all correspondence related to this study. The IRB wishes you success with your research.

For the IRB,

Julie F. Simpson
Director

cc: File
Butkiewicz, Thomas

University *of* New Hampshire

Research Integrity Services, Service Building
51 College Road, Durham, NH 03824-3585
Fax: 603-862-3564

14-May-2020

Stevens, Andrew H
Center for Coastal & Ocean Mapping,
24 Colovos Rd
Durham, NH 03824

IRB #: 6719
Study: Virtual and Augmented Realities to Support the Acquisition, Processing, Analysis and Display of Oceanographic Data
Review Level: Expedited
Approval Expiration Date: 11-Jun-2021

The Institutional Review Board for the Protection of Human Subjects in Research (IRB) has reviewed and approved your request for time extension for this study. Approval for this study expires on the date indicated above. At the end of the approval period you will be asked to submit a report with regard to the involvement of human subjects. If your study is still active, you may apply for extension of IRB approval through this office.

Researchers who conduct studies involving human subjects have responsibilities as outlined in the document, *Responsibilities of Directors of Research Studies Involving Human Subjects.* This document is available at http://unh.edu/research/irb-application-resources or from me.

If you have questions or concerns about your study or this approval, please feel free to contact me at 603-862-2003 or Julie.simpson@unh.edu. Please refer to the IRB # above in all correspondence related to this study. The IRB wishes you success with your research.

For the IRB,

Julie F. Simpson
Director

cc: File
Butkiewicz, Thomas

APPENDIX B

Research Study Participant Consent Form

CONSENT FORM FOR PARTICIPATION IN A RESEARCH STUDY

PERCEPTUAL OPTIMIZATION FOR DATA VISUALIZATION
This is a research study carried out by Drew Stevens, a Ph.D. student in computer science, for the Data Visualization Research Lab at the University of New Hampshire.

WHAT IS THE PURPOSE OF THIS FORM?
This consent form describes the research study and helps you to decide if you want to participate. It provides important information about what you will be asked to do in the study, about the risks and benefits of participating in the study, and about your rights as a research participant. You should:

- Read the information in this document carefully.
- Ask the research personnel any questions, particularly if you do not understand something.
- Not agree to participate until all your questions have been answered, or until you are sure that you want to.

WHAT IS THE PURPOSE OF THIS STUDY?
The purpose of this study is to find more effective ways of displaying and interacting with numerical data. There will be approximately 30 total participants in this study.

WHAT DOES YOUR PARTICIPATION IN THIS STUDY INVOLVE?
To participate, you must be aged 18 or older and have 20/20 vision (corrected or uncorrected). You will wear active-shutter 3D glasses to view data in a virtual environment on a stereoscopic computer monitor. You will then perform a set of tasks to complete the study objectives, which may be repeated using different configurations.

There will be an interactive training session prior to the study to familiarize you with the tasks, tools, and techniques involved. Your participation is expected to take approximately one hour.

WHAT ARE THE POSSIBLE RISKS OF PARTICIPATING IN THIS STUDY?
There are no known risks associated with this study. You are completely free to halt your participation in the study at any time for any reason.

WHAT ARE THE POSSIBLE BENEFITS OF PARTICIPATING IN THIS STUDY?
The potential benefit of this research is that data displays and data interactions will be better designed in the future.

WILL YOU RECEIVE ANY COMPENSATION FOR PARTICIPATING IN THIS STUDY?
You will be given $20 in cash for your time to participate in this study.

DO YOU HAVE TO TAKE PART IN THIS STUDY?
Taking part in this study is completely voluntary. You may choose not to take part at all. If you agree to participate, you may refuse to answer any question. If you decide not to participate, you will not be penalized or lose any benefits for which you would otherwise qualify.

CAN YOU WITHDRAW FROM THIS STUDY?
If you agree to participate in this study and you then change your mind, you may stop participating at any time. Any data collected as part of your participation will remain part of the study records. If you decide to stop participating at any time, you will not lose any benefits for which you would otherwise qualify, however, the $20 compensation is contingent upon the full completion of the study tasks.

HOW WILL THE CONFIDENTIALITY OF YOUR RECORDS BE PROTECTED?

I plan to maintain the confidentiality of all data and records associated with your participation in this research. There are, however, rare instances when I may be required to share personally-identifiable information with the following:

- Officials at the University of New Hampshire,
- The sponsor(s), or
- Regulatory and oversight government agencies.

To help protect the confidentiality of your information, the data will be stored on the UNH Box platform in accordance with UNH research data retention policies. Your raw data will only be accessible to me, my faculty advisor Dr. Thomas Butkiewicz, and lab director Dr. Colin Ware. Your data will be reported in aggregate and, if the dataset is to be shared with other researchers, any personally-identifiable information will be anonymized. The results may be used in reports, presentations, and publications.

WHOM TO CONTACT IF YOU HAVE QUESTIONS ABOUT THIS STUDY

If you have any questions pertaining to the research you can contact Drew Stevens or Dr. Thomas Butkiewicz (astevens@ccom.unh.edu, tbutkiewicz@ccom.unh.edu) to discuss them.

If you have questions about your rights as a research subject you can contact Dr. Julie Simpson in UNH Research Integrity Services, (603) 862-2003 or Julie.simpson@unh.edu to discuss them.

Yes, I, _____consent/agree to participate in this research project.

_____ _____

Signature Date

APPENDIX C

Experiment 1 Post-Study Questionnaire

Participant: _____ Vision checked for 20/20 acuity: ____

Hairy Slices Post-Study Questionnaire

Please indicate to what degree you agree or disagree with the following statements as they apply to the 5 different rendering methods.

1. Overall, the rendering method was intuitive and easy to understand.

	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
Plain Lines	1	2	3	4	5
Illuminated Lines	1	2	3	4	5
Plain Tubes	1	2	3	4	5
Ringed Tubes	1	2	3	4	5
Shadowed Hedgehogs	1	2	3	4	5

2. I got better and faster at interpreting the rendering method over the course of the study.

	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
Plain Lines	1	2	3	4	5
Illuminated Lines	1	2	3	4	5
Plain Tubes	1	2	3	4	5
Ringed Tubes	1	2	3	4	5
Shadowed Hedgehogs	1	2	3	4	5

3. The rendering method was aesthetically pleasing and enjoyable to look at.

	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
Plain Lines	1	2	3	4	5
Illuminated Lines	1	2	3	4	5
Plain Tubes	1	2	3	4	5
Ringed Tubes	1	2	3	4	5
Shadowed Hedgehogs	1	2	3	4	5

4. It took me a long time to make a decision when using this rendering method.

	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
Plain Lines	1	2	3	4	5
Illuminated Lines	1	2	3	4	5
Plain Tubes	1	2	3	4	5
Ringed Tubes	1	2	3	4	5
Shadowed Hedgehogs	1	2	3	4	5

5. Please enter any additional comments below and/or on the back of this paper.