

Hairy Slices II: Depth Cues for Visualizing 3D Streamlines Through Cutting Planes

Andrew H. Stevens¹ Colin Ware¹ Thomas Butkiewicz¹ David Rogers² and Greg Abram³

¹Center for Coastal and Ocean Mapping, University of New Hampshire, NH, USA

²Los Alamos National Laboratory, NM, USA

³Texas Advanced Computing Center, University of Texas at Austin, TX, United States

Abstract

Visualizing 3D vector fields is challenging because of occlusion problems and the difficulty of providing depth cues that adequately support the perception of direction of flow lines in 3D space. One of the depth cues that has proven most valuable 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. We carried out an experiment of the perception of direction for short streamlines passing through a cutting plane. The conditions included viewing with and without structure-from-motion and with and without stereoscopic depth. Conditions also include comparing streamtubes to lines. The results show that for this particular task, stereo provided an effective depth cue, but structure-from-motion did not. Ringed streamtubes and streamcones provided good 3D direction information, even without stereoscopic viewing. We conclude with guidelines for viewing slices through vector fields.

CCS Concepts

• **Human-centered computing** → *Empirical studies in visualization; Scientific visualization;* • **Computing methodologies** → *Perception;*

1. Introduction

Occlusion problems make it impossible to perceive all regions of a 3D vector field simultaneously. This is the case for 3D arrows, streamlines, and most other rendering methods. Because of occlusion, parts of the vector field further from the viewpoint become hidden from view, and for this reason it is necessary to simplify the visualization. One method is to use cutting planes, which are a tried and true technique commonly used to show scalar values on a slice through a volume [LWSH04, BMP*90, FG98, MG93]. When visualizing 3D flow patterns with a cutting plane, it is usually important to show the angle at which vectors pass through the plane (unless we are only interested in the vector component tangential to the plane, e.g. [LWSH04]) and this can be done using 3D glyphs. Because a dense set of glyphs passing through a plane can resemble fur, we call the method “Hairy Slices” [SBW16].

In our first investigation of this problem [SBW16], we examined a number of methods for using local 3D glyphs to show 3D vector direction at the point of intersection with a cutting plane. The most substantial effect found in that study was that tubes were considerably better than lines for showing 3D vector direction, even when the lines were shaded. Another finding was that thicker tubes were better than thinner tubes, especially when the spatial density of glyphs was low. There is an obvious trade-off between glyph diam-

eter and glyph density; smaller glyphs can be placed with a higher density, and thereby potentially reveal more detail in flow patterns. In that study, only pictorial depth cues were evaluated. These included using shading of the tubes and perspective. But two of the most important depth cues were not included: namely structure-from-motion (SfM), and stereoscopic viewing. The study in this paper extends this prior work by including these critical variables. The hope was that by providing better depth information it would be possible to decrease the size of glyphs and increase their density, thereby enabling more detailed information to be perceived. In addition, we evaluated both the use of animation to show flow direction, and the use of curved glyphs conforming to streamlines through the cutting plane.

2. Perceptual Issues

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 [WKL*15]. 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 [BPG00]. In other words, there is no single optimal way of providing depth information.

A thorough background on many of these cues and an exhaustive discussion of related works can be found in our previous report [SBW16]. To provide some of that background, it is summarized below.

There are, in general, four families of 3D flow visualization methods: flow glyphs, geometric flow, texture flow, and feature flow [LWSH04, Mun14]. While our previous study examined flow glyphs, which convey local information about the vector field using geometric objects whose substructure is tied directly to the data itself, this study focuses on geometric flow. Geometric flow includes integral methods like streamlines, where a particle is traced as it advects through a steady flow field; pathlines, which are the unsteady flow field equivalent to streamlines; and streamtubes, which apply 3D geometric shapes to streamlines in 3D flow fields. This family of flow visualizations is thoroughly detailed by McLoughlin et al. [MLP*10], and Post et al. provide a more general survey and categorization of flow visualizations [PVH*02].

Cutting planes serve to reduce the visual dimensionality of multidimensional datasets by displaying 2D slices, which help avoid self-occlusion, which is one of the many issues inherent in visualizing 3D data. Cutting planes have a variety of uses in scientific visualization: as clipping planes to view structures enclosed within or obscured by other volumes [LWSH04, DWE03], and as seeding planes for streamlines and streamtube visualizations [LWSH04, BMP*90, SRBE99]. Many techniques have been presented to visualize 3D vectors through a cutting plane, but none of them provide empirical data about the perceptual success of these techniques. One of our goals in this study is to gather empirical data regarding the perceptual effectiveness of various integral geometries on cutting planes visualizations.

Integral geometric structures are most often rendered as simple streamlines, or sometimes illuminated streamlines [ZSH96], which apply a shading model to the streamline to provide superior perceptual cues over the plain variant. Streamtubes and threads are three-dimensional extensions of streamlines [LWSH04, SGS05, WKZL04]. A cutting plane used with integral geometries is typically treated as a seeding plane from which the streamline structures “grow”.

A number of applied studies have shown that for certain visualization tasks, both stereoscopic viewing and SfM can greatly improve performance. For example, it has been shown that both SfM and stereo greatly improve performance on tracing paths in 3D networks, but SfM is the more powerful cue [SM93, WF96, vBKI10]. The same appears to be true for perceiving structure in point clouds [AWR18]. For both 3D networks and point clouds, pictorial cues offer little help. However, for other tasks these cues are less powerful. For example, when perceiving the shape of surfaces, shape-from-shading and surface texture may be more useful [Ram88]. Also, for visually guided reaching, stereoscopic depth perception is more important than SfM [AW04]. However, Chen et al. [CCAL12] 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 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 [LKJ*05], and that animation of streamlets (short streamline segments) along their paths can enhance pattern perception [WBM*16].

We surmised at the start of this study that the value of SfM and stereo would depend greatly on the kind of rendering method. For example, with simple streamlines pictorial cues will be weak. However, if the data are represented with streamtubes having a significant diameter, and bands around them, then the shape of the bands provide orientation cues. There is also the interesting question of whether animation of streamtubes in the manner of Fuhrman and Groller [FG98] 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 stereoscopic viewing and SfM can be expected to provide substantial benefits, especially for line-based designs, where the value of pictorial cues (such as perspective) will be minimal. However, rendering streamlines as shaded streamtubes can provide more depth information to help with the perception of their 3D direction.

In this study, we extend our prior work in a number of important ways:

1. **Stereo and motion depth cues** - In this study we evaluate the value of stereo and structure-from-motion cues for visualizing 3D flow through cutting planes.
2. **Advanced streamlines** - The prior study investigated straight glyphs, which only reveal the direction of flow at the point of plane intersection, and do not show speed/magnitude. Here we evaluate 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.
3. **Motion for direction along streamlines** - The prior study used static images, while this study examines the use of animated streamlets to show the direction of flow, and evaluates if this animation interferes with depth perception.

In accordance with many prior studies, we developed the following hypotheses:

- H1.** Stereoscopic viewing will enhance accuracy.
- H2.** Structure-from-motion will enhance accuracy.
- H3.** 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.
- H4.** Tube-type designs should outperform line-type designs because of their stronger pictorial cues.

H5. Animation will hinder SfM. In other words, when glyphs are animated to show flow direction, the benefit from SfM should be reduced.

3. Method

In general, the method and experimental design follows those described in the previous paper [SBW16], with some improvements.

3.1. Apparatus and Display

The experimental stimuli were presented on an Asus PG278Q 27" 2560x1440 px 3D monitor with Nvidia 3D Vision active shutter glasses, which were worn in both stereoscopic and non-stereoscopic conditions to keep brightness and contrast levels consistent amongst conditions. 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 display sat upon an enclosed desk space which was free of visual distractions, and the lights in the room were dimmed during the study.

3.2. Orientation Tracking and Input

The physical probe pictured in Figure 1 allowed the user to input perceived orientations by aligning the probe in front of the display and pressing a key. An HTC Vive Tracker and two Lighthouse tracking cubes provided 6-degree-of-freedom (6DOF) tracking, though only 3DOF were used to record orientation. The tracker was affixed to the end of a 2cm x 2cm x 10cm rigid block of foam with a 1cm² grid pattern. This design provides strong linear perspective depth cues, and does not resemble any of the geometric designs used in the study so that the task does not become a visual matching exercise.

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

Before each study participant began, the tracking system was calibrated to the display in order to co-register the virtual scene's coordinate system with the real world.

3.3. Stimuli

3.3.1. Glyph Designs

Five different integral glyph designs were implemented:

1. **Streamlet Lines** - The streamlet heads are white and fade away to medium gray. See Figure 2a.
2. **Animated Streamlets** - Born at the beginning of a streamline, travel along it for 1s, tail length is half the streamline length, dies at end of streamline, with 0.5s to allow the tail to flow to the end and complete. Streamlet spawning was staggered to avoid the visual distraction of the entire streamlet population being born or dying at the same time. See Figure 2b.



Figure 1: Orientations are input using the probe, which points away from the black HTC Vive wireless tracker attached to its base.

3. **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 2c.
4. **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.
5. **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 2d.

3.3.2. Stereoscopic Viewing

The integral glyph designs were presented both with and without stereoscopy using the 3D monitor and active shutter glasses. Each participant's interpupillary distance (IPD) was measured and the binocular disparity of the stereoscopic content was set to match the IPD. This helps to account for any perceptual distortions that may arise through incorrect stereo disparity.

3.3.3. 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 were based on the SfM parameters used successfully in previous research [AWR18]. The oscillations from SfM help to reveal more of the 3D shape of the streamline geometry.





(a) *Static streamlets.*(b) *Animated streamlets.*(c) *Streamtubes with gradient texture. Flow direction is from dark to light.*(d) *Streamcones. The direction of flow is from the base to the tip.*

Figure 2: The glyph designs evaluated in the study. Animated streamtubes are not explicitly shown here because they are animated and visually equivalent to static streamtubes in a still picture. The animated streamtube design mirrored the design of static streamtubes, but the texture moved along the streamtube in the direction of the flow.

3.3.4. Cutting Plane Design

The cutting plane was 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 was accomplished by dividing the plane into 20×20 cells and jittering their midpoints by up to $1/4$ 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 40×40 (a four-times increase), as can be seen in Figures 2a and 2b. This increase of seed density is supported by the results of our prior work [SBW16], which found that thin glyph geometries could be successfully seeded at higher densities than thicker geometries without a reduction in perceptual performance.

The cutting plane was rendered as semitransparent yellow with transparent grid lines to help disambiguate streamlines in front of or behind the slice. Streamlines were 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 “trace back” the flow from a given sampling point.

3.4. Test Data Sets

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

The data volume was 400^3 units, and the cutting plane was scaled to 100×100 units. For each trial, the center of the cutting plane was randomly positioned within the 200^3 -unit center of the data volume, and its orientation was randomly selected from a uniform distribution over a unit sphere. This provided a large amount of variation in the flow features and patterns, because the single dataset 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, we chose the simulation output due to its rich flow characteristics and because it represents a real data product in need of improved visualization strategies. Since our study focuses primarily on the perception of vector orientations, we believe the chosen data set provides a wide array of flow features which generalize to other flow visualization problems. The data set has been made available for download to the reader [AHSA20].

3.5. Task

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

Each training session began by displaying a to-scale virtual




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

model of the input probe at the center of the screen and matched to the orientation of the physical probe. An addition target probe also appeared at the center of the screen. Subjects were shown how to align the virtual probe to the orientation of the target goal, with the assistance of visual feedback which decreased over the course of the training session. The first level of visual feedback changed the color of the target goal depending on the angle between the input probe and the target goal; the second level eliminated the color cues from the first level; the third level did not render the virtual probe on the screen, but the color cues from level one were applied to the target goal; and the fourth level showed 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 could be displayed using a toggle key. Each subject needed 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 were shown to the participant, and they were asked to orient the probe with the flow direction at a target marker on the cutting plane. Feedback in this portion of the training was provided verbally, including instructions on the conventions used for the direction of the cone geometry and texture gradients.

During the study session, participants saw a slice through a 3D vector field, like the ones shown in Figure 2. They were required to estimate the flow direction at the point specified by the red target reticle by aligning the wireless probe pictured in Figure 1 with the perceived flow direction. Participants were instructed to hold the probe as close to their line-of-sight as possible, so that there would be minimal perspective distortion between the current stimulus and the orientation probe. The experimental setup is depicted in Figure 3, and shows the probe in use during a practice trial at the end of a training session.

Each trial was based on a randomly oriented slice taken from a random position within the data set, as described in the previous section. Target points were randomly chosen within the center $2/3$ of each slice. A red target reticle was centered over each target point, and it slowly rotated for added visibility in visually dense/obscured configurations (e.g., Figure 2c). All visuals (reticle, cutting plane, flow geometries) remained throughout each trial

until the participant hit the key to record their orientation input and the next trial was loaded.

3.6. Participants

A total of 17 adults (4 female, 13 male) participated in the study. The participants ranged in age from 22 to 79 ($\bar{x} = 35.2$), were trained in the task before participation, and received compensation for their time. This study was approved by the university's Institutional Review Board. Participants were allowed breaks between experimental blocks.

Two of the authors were part of the participant pool, and collected two sets of data each to gauge any difference in performance patterns between experts in flow visualization and non-experts. (There were no notable differences.)

All participants were screened for 20/20 vision and stereoblindness. One potential participant was disqualified due to stereoblindness.

Due to significant performance issues, two participants' data sets were discarded. The cause for this is not known, as these two data sets were markedly different than the rest, so we suspect a calibration error before or during the trial lead to these results. We do not believe it would have been an issue with understanding the task, given the extensive training each participant received where they were observed by the experimenter to perform adequately in each condition before being able to begin the study.

4. Experimental Design

We used a fully-crossed 5x2x2 within-subjects design for a total of 20 total conditions, listed in Table 1. 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 1: *Experimental 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

4.1. Total Angular Error

The total angular error is the angle subtended by the probe orientation and the flow vector at the target point, as given in Equation 1. A perfectly aligned probe and target will yield a total angular error of 0° , and a probe and target aligned but facing in opposite directions result in a 180° total angular error.

$$\theta = \frac{180}{\pi} \arccos(\hat{t} \cdot \hat{p}) \quad (1)$$

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

4.2. Absolute Depth Error

To assess the amount of error due to misperception of depth, we decompose the total error to find the depth component of the error according to Equation 2:

$$\theta_D = \frac{180}{\pi} |\arccos(\hat{t}_z) - \arccos(\hat{p}_z)| \quad (2)$$

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 orientation probe vector input by the participant.

4.3. Weighted Projection Error

To assess non-depth-cue-related perception of vector direction, we examined the weighted projection component of the error, 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. See the previous paper [SBW16] for an in-depth discussion of this challenge and our solution, which 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 3 describes the weighted projection error:

$$\theta_P = \frac{180}{\pi} \arccos\left(\frac{t_{xy} \cdot p_{xy}}{\|t_{xy}\| \cdot \|p_{xy}\|}\right) \cdot (1 - |t_z|) \quad (3)$$

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

5. Results

A summary of the raw results can be found in Figure 4. All results and analyses performed on them are available to the reader to download [AHSA20].

5.1. Antiparallel Correction

A common challenge in flow visualization using streamlines is that the orientation can be clear, but the direction (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), as no conventions exist. Similarly, the texture gradient used in the




Figure 4: *Uncorrected Mean Total Angular Error with 95% confidence intervals. The results are colored by the geometry condition (streamline, streamtube, and streamcone).*

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 tip to base) 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.)




Figure 5: *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.*

Evidence of antiparallel confusion is demonstrated through the histogram of raw errors shown in Figure 5. 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 they prefer, as opposed to forcing an arbitrary convention onto those with strong intuitions, we think it is better to keep all experimental variables consistent between participants, and correct for this error in the data.

To detect antiparallel confusion, we filter for individual trials with total angular errors over 130° , record them for further evaluation of antiparallel metric, and correct them for further analysis using the equation:

$$\theta_{corrected} = 180^\circ - \theta \quad (4)$$

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

We believe this correction yields a more accurate evaluation of perceived vector directions for participants with very strong directional intuitions which ran counter to our chosen conventions.

Figure 6 shows the total number of times that participants’ individual trial results were corrected for antiparallel confusion according to Equation 4.

5.2. Analysis

Since angular errors have a performance floor at 0° , the distribution of errors is heavily skewed to the right (see Figure 5), 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




Figure 6: A count of trials where the antiparallel correction (Eq. 4) was applied. The results are colored by the geometry condition (streamline, streamtube, and streamcone). Note that corrections overwhelmingly affected both streamlet conditions (green) and static streamtubes.

find a transformation that satisfied the assumption of residual error normality enough to be confident in the appropriateness of a parametric test. Therefore, we used the non-parametric Kruskal–Wallis one-way analysis of variance, which compares the median values of similarly-distributed groups to test the null hypothesis that the two groups have the same median value [KW52]. We then use a pairwise Wilcoxon test with a Benjamini-Hochberg multiple comparisons correction (intended to reduce the rate of false positives) to find the relative performance amongst the group members at a $p = 0.05$ confidence level [Wil45].

Corrected mean angular error rates can be found in Figure 7 and Figure 8, providing a much clearer picture of the relative performances between conditions. We find that, overall, glyph design ($\chi^2(4) = 418.62$, $p < 0.00001$) 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 ($\chi^2(1) = 1.69$, $p = 0.1931$) was not significant.

Amongst 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 7 helps to illustrate these results.




Figure 7: 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).

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

A plot of the absolute depth error is shown in Figure 9. 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 10. 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, an-




Figure 8: 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).

imated 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.

6. Discussion

As can be seen, stereoscopic viewing provided a considerable benefit when simple streamlets were drawn, greatly reducing depth errors and supporting hypothesis H1. When streamtubes were used, surprisingly, there was only a small benefit to stereoscopic viewing.

In contrast, SfM provides little overall benefit for this task. Only in the case of simple streamlets was there any benefit, and even there it was small. This is surprising, and in marked contrast to prior studies of different visualization tasks, such as path tracing in 3D networks [WF96] and perception of patterns in point cloud data [AWR18] where SfM provided a greater benefit than stereoscopic viewing. Therefore, hypothesis H2 is rejected, but since streamlets did receive considerable benefit from both stereo and SfM, hypothesis H3 is supported by our results.

Animation is the best disambiguator for direction. The animated streamtubes performed remarkably well (only one antiparallel correction across conditions over 1020 trials; < 0.1% correction rate). Animated streamlets with SfM were the exception; without stereo, the extra motion from texture animation degraded performance (see




Figure 9: Mean Absolute Depth Error. The results are colored by the geometry condition (streamline, streamtube, and streamcone).

Figure 4). The absence of additional depth cues to fall back on made it difficult to disambiguate SfM from texture animation in the animated streamlet condition.

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 it provides are strong enough that it does not receive any significant additional benefits from stereo or SfM cues, as it has effectively reached a performance plateau. Both streamtubes and streamcones yielded significantly lower errors than streamlets, supporting hypothesis H4.

Our results did not support hypothesis H5. SfM 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 than 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.

When designing the experiment, we were hesitant to increase the seeding density of the streamlet conditions, as we did not want to give them an unfair advantage by sampling four times the amount of information as the other conditions. Ultimately, pilot testing confirmed that occlusion and SfM cues degrade quickly when the seed-




Figure 10: Mean Weighted Projection Error with 95% confidence intervals. Note that the severity and range of errors is much lower than in Figures 4, 7, and 9. The results are colored by the geometry condition (streamline, streamtube, and streamcone).

ing density gets sparse, as the thin line geometry of the streamlets fails to provide enough local depth information to reasonably judge flow direction. Our results clearly show us 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.

We can compare the results of the non-stereo, fixed conditions to our previous study [SBW16], which used neither stereo nor SfM. Our seeding density in this study is closest to the most sparse seeding density in the prior study (2 seeds/cm). We find that, 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.

7. Conclusions

Visualization of 3D flow fields is a mature topic, but still presents many inherent challenges, due to issues such as occlusion and the difficulty of providing effective depth cues on small glyphs. This work expands upon previous perceptual research into effective glyph design by conducting a new human factors experiment that examines the benefits of the depth cues provided by stereoscopic viewing and structure-from-motion as applied to a variety of advanced streamline visualization techniques.

Overall, the results 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 stereo or SfM.

While this research only studied the perception of flow direction (as did our prior study), it should be further expanded in future work by examining the perception of vector speed/magnitude, as this is also an important characteristic of flow fields, and improving its perception can be combined with this research to further optimize the visualization of complex 3D flow fields.

Acknowledgment

This research was made possible through the support of NOAA Grant NA15NOS4000200 and DOE/LANL DE-AC52-06NA25396, sub. 475412.

References

- [AHS20] ANDREW H. STEVENS COLIN WARE T. B. D. R., ABRAM G.: *Hairy Slices II Data, Results, and Analyses*, 2020. URL: <https://github.com/ahstevens/Hairy-Slices-II-Data>. 5, 6
- [AW04] ARSENAULT R., WARE C.: The importance of stereo and eye-coupled perspective for eye-hand coordination in fish tank vr. *Presence: Teleoperators & Virtual Environments* 13, 5 (2004), 549–559. 2
- [AWR18] AYGAR E., WARE C., ROGERS D.: 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 (2018), 9. 2, 3, 9
- [BMP*90] BANCROFT G. V., MERRITT F. J., PLESSEL T. C., KELAITA P. G., MCCABE R. K., GLOBUS A.: Fast: a multi-processed environment for visualization of computational fluid dynamics. In *Proceedings of the First IEEE Conference on Visualization: Visualization90* (1990), IEEE, pp. 14–27. 1, 2
- [BPG00] BRADSHAW M. F., PARTON A. D., GLENNERSTER A.: The task-dependent use of binocular disparity and motion parallax information. *Vision Research* 40, 27 (2000), 3725–3734. 2
- [CCAL12] CHEN J., CAI H., AUCHUS A. P., LAIDLAW D. H.: Effects of stereo and screen size on the legibility of three-dimensional streamtube visualization. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2130. 2

- [DWE03] DIEPSTRATEN J., WEISKOPF D., ERTL T.: Interactive cut-away illustrations. In *Computer Graphics Forum* (2003), vol. 22, Wiley Online Library, pp. 523–532. 2
- [FG98] FUHRMANN A., GROLLER E.: Real-time techniques for 3d flow visualization. In *Proceedings Visualization'98 (Cat. No. 98CB36276)* (1998), IEEE, pp. 305–312. 1, 2
- [KW52] KRUSKAL W. H., WALLIS W. A.: Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* 47, 260 (1952), 583–621. 8
- [LKJ*05] LAIDLAW D. H., KIRBY R. M., JACKSON C. D., DAVIDSON J. S., MILLER T. S., DA SILVA M., WARREN W. H., TARR M. J.: Comparing 2d vector field visualization methods: A user study. *IEEE Transactions on Visualization and Computer Graphics* 11, 1 (2005), 59–70. 2
- [LWSH04] LARAMEE R. S., WEISKOPF D., SCHNEIDER J., HAUSER H.: Investigating swirl and tumble flow with a comparison of visualization techniques. In *IEEE Visualization 2004* (2004), IEEE, pp. 51–58. 1, 2
- [MG93] MEYER T., GLOBUS A.: Direct manipulation of isosurfaces and cutting planes in virtual environments. In *Technical Report CS-93-54*. (1993), Department of Computer Science, Brown University. 1
- [MLP*10] MCLOUGHLIN T., LARAMEE R. S., PEIKERT R., POST F. H., CHEN M.: Over two decades of integration-based, geometric flow visualization. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 1807–1829. 2
- [Mun14] MUNZNER T.: *Visualization Analysis and Design*. AK Peters Visualization Series. CRC Press, 2014. URL: <https://books.google.com/books?id=dznSBQAAQBAJ>. 2
- [PVH*02] POST F. H., VROLIJK B., HAUSER H., LARAMEE R. S., DOLEISCH H.: Feature extraction and visualisation of flow fields. In *Eurographics (STARs)* (2002). 2
- [Ram88] RAMACHANDRAN V. S.: Perception of shape from shading. *Nature* 331, 6152 (1988), 163. 2
- [SBW16] STEVENS A. H., BUTKIEWICZ T., WARE C.: Hairy slices: Evaluating the perceptual effectiveness of cutting plane glyphs for 3d vector fields. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 990–999. 1, 2, 3, 5, 6, 10
- [SGS05] STOLL C., GUMHOLD S., SEIDEL H.-P.: Visualization with stylized line primitives. In *VIS 05. IEEE Visualization, 2005.* (2005), IEEE, pp. 695–702. 2
- [SM93] SOLLENBERGER R. L., MILGRAM P.: Effects of stereoscopic and rotational displays in a three-dimensional path-tracing task. *Human Factors* 35, 3 (1993), 483–499. 2
- [SRBE99] SCHULZ M., RECK F., BERTELHEIMER W., ERTL T.: Interactive visualization of fluid dynamics simulations in locally refined cartesian grids. In *Proceedings Visualization'99 (Cat. No. 99CB37067)* (1999), IEEE, pp. 413–553. 2
- [vBK10] VAN BEURDEN M. H., KUIJSTERS A., IJSSELSTEIJN W. A.: Performance of a path tracing task using stereoscopic and motion based depth cues. In *2010 Second International Workshop on Quality of Multimedia Experience (QoMEX)* (2010), IEEE, pp. 176–181. 2
- [WBM*16] WARE C., BOLAN D., MILLER R., ROGERS D. H., AHRENS J. P.: Animated versus static views of steady flow patterns. In *Proceedings of the ACM Symposium on Applied Perception* (2016), ACM, pp. 77–84. 2
- [WF96] WARE C., FRANCK G.: Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics (TOG)* 15, 2 (1996), 121–140. 2, 9
- [Wil45] WILCOXON F.: Individual comparisons by ranking methods. *Biometrics Bulletin* 1, 6 (1945), 80–83. URL: <http://www.jstor.org/stable/3001968>. 8
- [WKL*15] WOLFE J. M., KLUENDER K. R., LEVI D. M., BARTOSHUK L. M., HERZ R. S., KLATZKY R. L., LEDERMAN S. J., MERFELD D. M.: Sensation and perception. 1
- [WKZL04] WENGER A., KEEFE D. F., ZHANG S., LAIDLAW D. H.: Interactive volume rendering of thin thread structures within multivalued scientific data sets. *IEEE Transactions on Visualization and Computer Graphics* 10, 6 (2004), 664–672. 2
- [ZSH96] ZÖCKLER M., STALLING D., HEGE H.-C.: Interactive visualization of 3d-vector fields using illuminated streamlines. In *IEEE Visualization* (1996), vol. 96, pp. 107–113. 2