3D Contour Perception for Flow Visualization

Colin Ware*
University of New Hampshire

Abstract

One of the most challenging problems in data visualization is the perception of 3D flow fields because judging the orientation of 3D paths is perceptually difficult. It is hypothesized that perception of the orientations of streamlines in space can greatly benefit from stereoscopic depth cues and motion parallax. In addition, because stereoscopic depth perception is a super-acuity and relies on such factors as small-scale stereoscopic disparity gradient based on texture, stereoscopic depth judgments will be exceptionally sensitive to display quality. In conventional displays, the aliasing of pixels gives completely spurious texture information to the mechanisms of stereoscopic depth perception. We carried out a study to evaluate the importance of 3D cues in perceiving the orientation of curved contours. The result showed that stereo and motion cues are essential to perceiving the orientation of 3D lines. If, however, the contours are rendered as shaded tubes, good orientation information is available even without stereo or motion depth cues.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces – evaluation/methodology, theory and methods.

Keywords: stereo viewing, flow visualization, depth cues.

1 Introduction

The perception of the orientation of streamlines and pathlines in space is important for understanding 3D flow fields. In perceptual terms this can be thought of as the problem of perceiving oriented curved contours in 3D-space. This presents a considerably greater challenge to the human visual system than, say, perceiving the shape of curved surfaces. The shape from shading cues and conforming textures that enable surface shape to be perceived in 2D perspective pictures are typically unavailable in most flow visualizations where simple unshaded lines are used to represent the flow direction.

Stereoscopic depth and structure-from-motion (kinetic depth) are probably the most important cues in perceiving the orientation of 3D lines in space but perceptual studies of stereopsis suggest that current displays are, for the most part, woefully inadequate to display 3D contours in space because they lack sufficient resolution.

Here, we report an empirical investigation of the perceptual requirements for accurately perceiving the orientation of curved contours in 3D space making use of an ultra-high resolution digital stereoscopic display with 9.2 million pixels for each eye. This enables the presentation of visual information at, or close to, the resolution of the human eye. The problem of effective visualization of 3D streamlines has two components. The rendering issue – how best to draw the contours so that 3D paths can be perceived, and the perceptual issue – should stereoscopic viewing and motion depth cues be provided? In the following sections we summarize prior work that relates to these matters.

2 Rendering Issues

Although pathlines and streamlines are often presented as simple unshaded lines a number of methods have been developed that are more effective. These include the use of solid 3D arrows or solid tubes [Geiben and Rumpf, 1991; Schroeder et al, 1991]. Interrante, and Grosch, [1998] showed how line integral convolution can also be used to produce an effect similar to streamtubes by using texture elements integrated through a volume and then rendering the external surfaces of the extrusions that resulted. If this method is used then the perceptual issues are essentially the same as for stream tubes.

One of the ways of improving the ability to see fine streamlines in 3D is to shade them [Zockler et al, 1996] using the dot product of the streamline tangent at each point and the lighting vector. The result cannot give as much information as, for example, a thicker shaded tube, but does reduce the uncertainty in 3D flow direction. It is also possible to use shaded ribbons to visualize a stream surface [Hultquist, 1992].

A common solution to the 3D flow problem is to visually simplify the flow and only show those portions that lie on some curved manifold. This can be done by displaying grids of 3D arrows "hedgehogs" on 2D slices through 3D flow fields [Klassen and Harrington, 1999] or by draping a line-integral convolution on a shaded representation of a curved surface [Laramee et al 2004]).

In all of the above examples, perceiving the 3D orientation of contours would seem to be the core perceptual problem.

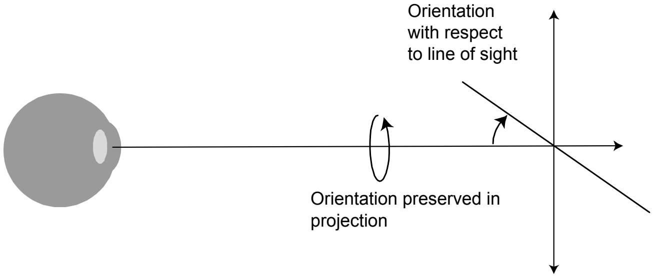

Figure 1. Angular error in the perception of the orientation of a 3D contour can be decomposed into components that are preserved in the picture plane projection and an angular difference from the view vector.

---

*cware@ccom.unh.edu

APGV 2006, Boston, Massachusetts, July 28–29, 2006.
© 2006 ACM 1-59593-429-4/06/0007 $5.00

3 Perceptual Issues

The problem of judging the orientation of a 3D contour can be decomposed into the determination of orientation as it is projected on to the picture plane (a plane orthogonal to the line of sight), and the angle that contour has with respect to the line of sight (See Figure 1). Depth cues only apply to the latter judgment, since they are not necessary to judge orientation of a contour projected on the picture plane.

If a streamline is rendered simply as a constant width, unshaded line only the stereoscopic depth and motion parallax depth cues are available for the slant judgment. If however, the streamline is rendered as a textured line or a shaded tube, other cues, such as texture gradients, shading and conforming contours may provide useful information. The following sections discuss these factors in more detail, particularly stereoscopic depth, kinetic depth and their combination.

Stereoscopic viewing

Stereoscopic depth perception relies on the detection of relative differences, called disparities between pairs of features imaged in the two eyes. This gives rise to the so-called correspondence problem. In order to find disparities the brain must match parts of the images seen by the two eyes. In the case of long straight lines, with no textural features along their lengths, this may be impossible (van Ee and Schor, 2000). However, if lines are textured, the texture elements can provide a disparity gradient, a series of corresponding points of decreasing or increasing disparity. Frisby et al [1966] carried out a series of experiments that used a set of physical textured twig-like objects in order to investigate how accurately judgments of length could be made under real-world viewing conditions. They used monocular and stereoscopic viewing and showed that the judgment of the lengths of differently oriented highly texture real sticks was much better than the reported judgment of the length of simplified lines using artificial computer generated stimuli (e.g. Tittle et al 1995). In addition, performance for their twig stimuli was also remarkably good even without stereoscopic viewing and motion parallax cues.

Our ability to see stereoscopic depth allows for extra-ordinarily fine judgments. For example, Tyler [1975] found that acuity for discriminating a wavy line varying in depth was better than 1 arc second. (It was best for a wave period of 1 cycle per degree). This is much better than could be predicted from receptor size and requires the visual system to integrate the signal from multiple receptors. Tyler also measured discrimination of sinusoidal wavy lines having different spatial frequencies. He found that stereoscopic depth was not perceived for spatial frequencies above 3 cycles /degree.

Other patterns, however, yield different stereoscopic disparity thresholds. For example, the threshold for detecting line disparity is about 12 arc seconds [Howard and Rogers, 1995]. This is still remarkably precise but much larger than the 1 sec that Tyler measured. The extreme sensitivity of the human eye to disparities is the reason for using the very high resolution display chosen for this study. Showing complex 3D flow patterns, necessarily requires the perception of the orientation of fine contours. Conventional screens only have a resolution of about 0.3 mm per pixel and at normal viewing distances this translates to 2 minutes of arc. This is a much coarser resolution than is optimal for stereoscopic depth perception.

Structure from motion cues

It has been long known that the projected image of a rotating 3D wire objects appears strongly three dimensional. This is called the kinetic depth effect [Wallach, 1959]. A number of studies have compared the relative value of stereoscopic depth and kinetic depth for the task is tracing paths in 3D node-link structures. These have generally shown that motion parallax is a more effective cue than stereoscopic depth for this particular task [Sollenberger and Milgram, 1993, Ware and Franck, 1996].

For the task of surface shape perception, stereo and motion cues appear to be roughly equivalent [Norman et al, 1996] although this may depend on the shape of the objects being observed and for how long. For cylindrical objects under stereoscopic viewing it is easier to resolve curvature differences for horizontal cylinders than for vertical cylinders [Rogers and Gagnello, 1989]. In the case of kinetic depth, such asymmetries may be expected to depend on the axis of rotation. Viewing time is also a factor. A study by Uumori and Nishida [1994] showed that for random dot surfaces, motion parallax was the most important cue at the start but after a few seconds stereoscopic depth became dominant.

The study using real twig-like objects [Frisby et al, 1966], discussed in the previous section, also contained a comparison of motion and stereoscopic depth cues. The results showed that the stereoscopic depth cue was more important than the kinetic depth cue for this task. This may have been at least partially due to the presence of fine visual textures on the surfaces of the twigs.

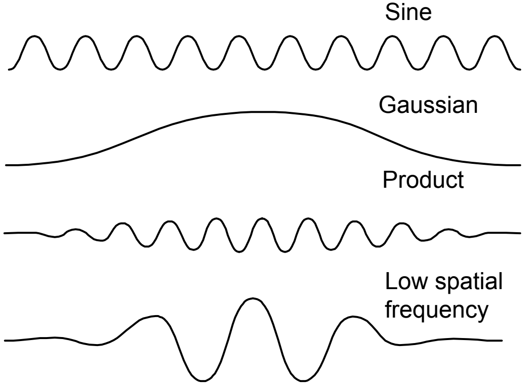

Figure 2. The stimuli for this study were planar Gabor functions. The spatial frequency and amplitude were scaled proportionately while the Gaussian envelope was constant.

The above are only a few of the many studies that have been carried out to compare motion and stereoscopic cues for different tasks. What has become clear is that the relative value of these two cues is highly task dependent. It is also dependent on the exact stimulus configuration. Hence it is not possible to predict, based on prior studies, the relative values of different cues for the perception of curved contours in space. Nevertheless, it is possible to propose a number of plausible hypotheses.

1) Both stereo and motion cues will be important in judgments concerning the orientation in depth of 3D contours.
2) Rendering the paths as tubes will add additional shape-from-shading depth cues and this should improve orientation judgments.
3) Adding surface attributes, such as rings around the tubes should also help in providing a disparity gradient.

4) The resolution of the display is likely to be important, especially for stereoscopic viewing.

The primary purpose of this study was to examine the importance of stereoscopic and motion parallax cues in perceiving the orientation of 3D paths rendered either with lines or with tubes.

A second goal was to determine *amount of detail* that could be seen in 3D paths under different conditions. In order to accomplish this, the pattern used for the experiments was a Gabor function. This function consists of the product of a sinusoid with a Gaussian as illustrated in Figure 2. The amount of detail was varied by changing the spatial frequency of the sinusoidal pattern.

4 Method

Apparatus:

The apparatus consisted of two Viewsonic VP 2290b displays arranged in a Wheatstone stereoscope arrangement as shown in Figures 3 and 4 [Wheatstone, 1838]. Front surface mirrors were use to eliminate ghosting. The frame rate achieved was 20 frames per second. Each of these displays has 3840x2400 pixels with a display area 47.7 x 29.7 cm giving an individual pixel size of 0.0125 cm.

The screens were set at a viewing distance of 109 cm. This yielded a visual angle per pixel of approx 23 seconds of arc. This is comparable to the size of receptors in the fovea and is sufficient to display the finest grating pattern that can be resolved by the human eye – about 60 cycles per degree corresponding to 1 cycle per minute of visual angle. [Campbell and Green, 1965].

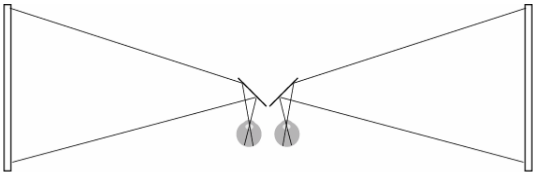

Figure 3. The configuration of a Wheatstone stereoscope.

Figure 4. The Wheatstone stereoscope display. The image on the screen has been added using Photoshop.

Task

The test pattern consisted of a Gabor function arranged vertically as illustrated in Figures 2 and 5. The subject's task was to set the paddle (shown in the lower portion of the screen) to match the plane in which the sinusoidal pattern varied.

On each trial, a new Gabor contour was presented with a different random orientation. The orientation was varied about a vertical axis between +/- 45 degrees from the plane orthogonal to the line of sight.

Conditions:

Stereoscopic depth: (2) In the stereo viewing condition different views were generated for each eye assuming an eye separation of 6.4 cm. For the no-stereo condition, both views were identical.

Motion parallax: (2) To generate motion parallax the entire scene (the Gabor pattern and the paddle) was oscillated sinusoidally about a vertical axis with an amplitude of +/- 15 deg and a period of 10 seconds. For the no-motion condition the scene did not move.

Tube renderings: (4) There were four rendering styles for the gabor path. Tubes: 1mm and 2mm in diameter, 2 pixel width lines, and 2 pixel width shaded lines. The tubes had alternating grey and white bands around them. The lines were shaded using the cosine of the angle between the illumination vector and a plane orthogonal the path at each point.

Spatial frequencies: (6) Cycle/degree: {0.18, 0.37, 0.75, 1.5, 3.0, 6.0}. It is possible that the amplitude of the 2D projection of the pattern could be used as a cue to its orientation. To avoid this the amplitude was varied randomly by a factor of 2. The net result of the generation algorithm was that the amplitude varied between 0.2 and 0.4 times the wavelength of the sinusoid.

The total number of conditions was the product of the above 2x2x4x6 = 96.

There were 24 trial for each condition and it was possible to complete a trial in about 5 seconds on average. Settings were done in blocks of 12 for each condition. All 96 conditions were carried out in a random order, then the process was repeated with a different random ordering. A different random order was used for each subject. The whole process took about 6 hours of observation per subject distributed over several sessions. To prevent fatigue sessions were no longer than one hour.

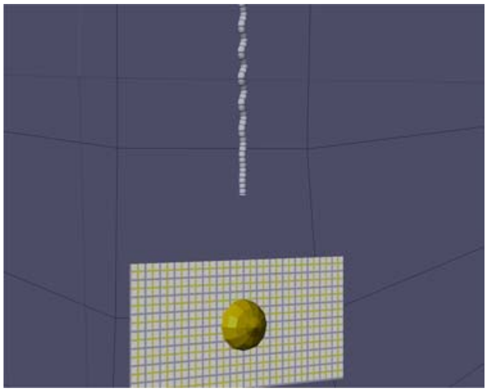

Figure 5. A part of the display showing the thick tube pattern, and the paddle used to match the perceived orientation of the tube.

Subjects:

Because of the difficult nature of these observations we used trained observers. One was the author of this paper. A second was a postdoctoral fellow in the psychology department who had employed psychophysical techniques for his PhD. The third was a research assistant in the lab who had participated in a number of previous experiments.

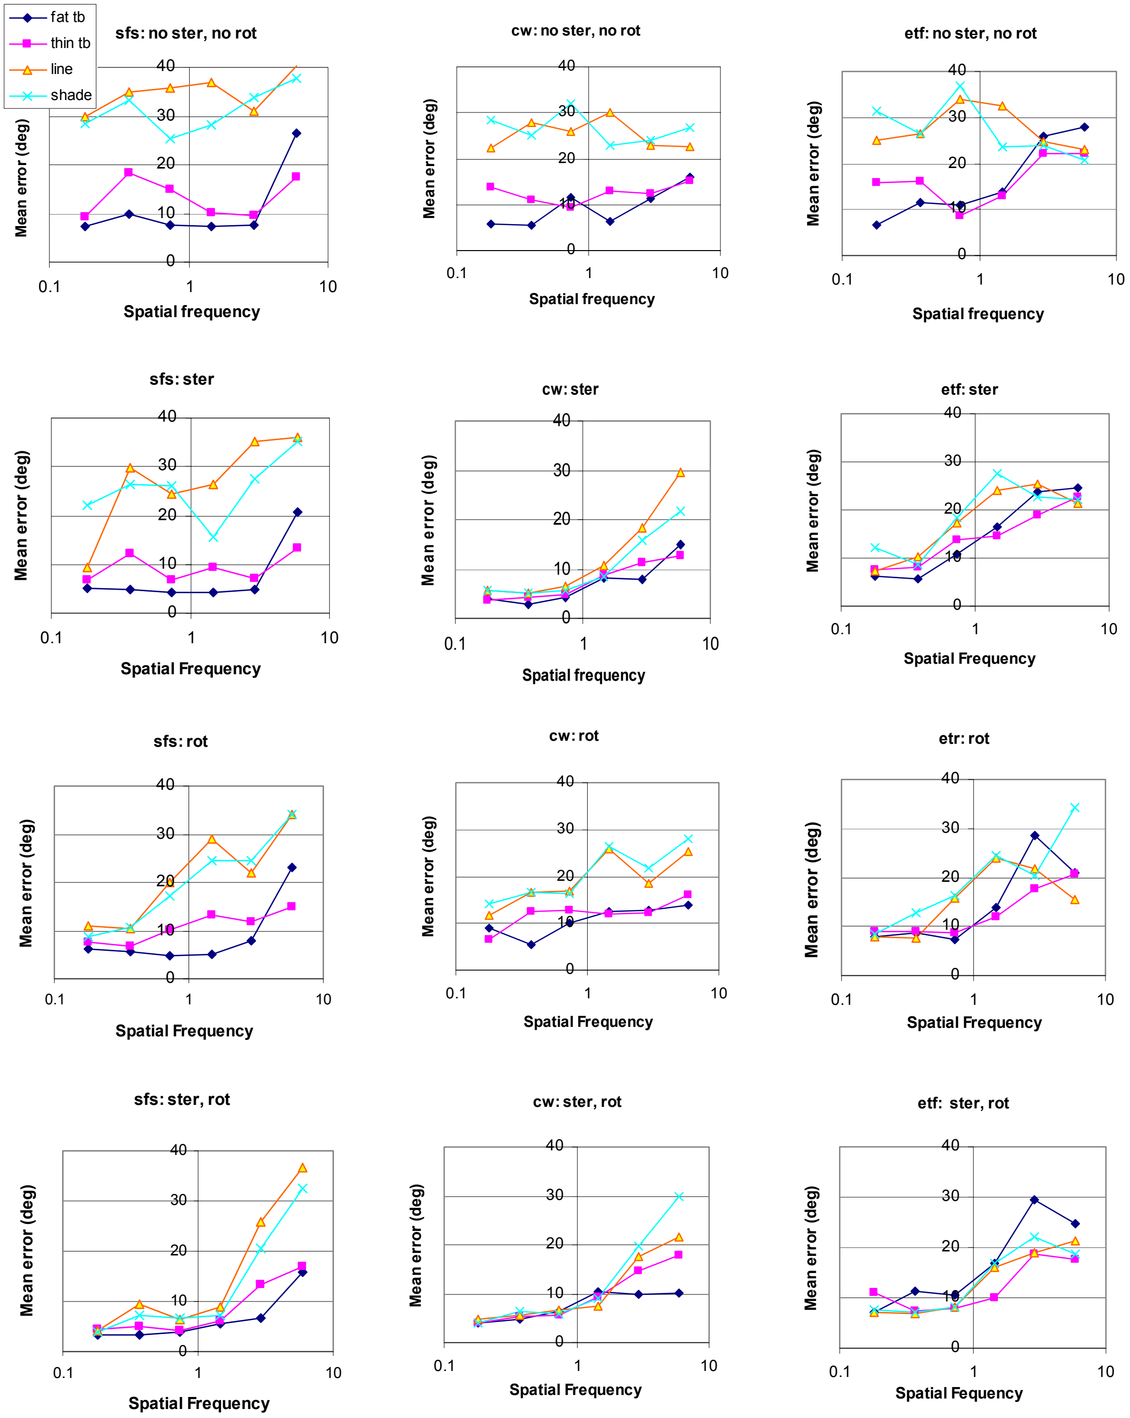

Figure 6. Results for each of the three subjects. Each row shows a different viewing condition. Each column shows results from a different subject. Data points represent mean absolute error from 24 observations.

5 Results

The results are summarized in Figure 6. Each curve shows how accuracy varied with spatial frequency under a particular viewing condition. We describe them row by row according to the stereo and motion cues available.

No stereo, no motion

The top row of graphs shows the results obtained where only pictorial cues were available. In the case of the unshaded lines this meant that no depth cues were present at all, unless there had been something we had overlooked. As can be seen the mean error for the unshaded lines was approximately. 30 degrees. A 30 degree error on average is what would be expected from random setting between +/- 45 deg. (Note that a smaller mean error would result if subjects just left the paddle orthogonal to the line of site).

All subjects did well with the shaded tube, with mean errors of around 10 degrees. These errors are somewhat smaller than those obtained with the larger diameter tubes for two of the three subjects. The errors increase somewhat for the highest spatial frequencies.

Stereo only

There were marked individual differences in the relative benefits of stereo and motion cues. For subject *sfs*, the results in the stereo only condition and line renderings were not much better than the no stereo, no rotation conditions This subject did, however, show improvement with both the thin and thick tubes. Subjects *cw* and *etf* benefited considerably form the stereoscopic depth cue across all conditions excepting for spatial frequencies at and above 3 cycles per degree.

Rotation only

All three subjects benefited to a roughly equal extent from the kinetic depth cue. As with *stereo only* the benefit was smaller at and above 3 cycles per degree.

Rotation and stereo

As expected, the most accurate settings were obtained with the combination of stereo and kinetic depth cues. The advantage of tubes over lines was much reduced for subjects *sfs* and *etf* and eliminated at lower spatial frequencies for subject *cw*. Low spatial frequency errors were reduced two around 5 degrees for *sfs* and *cw* and around 8 degrees for *etf*.

The results reveal individual differences in the relative benefits of thick and thin tubes. Subject *sfs* was considerable more accurate with thick tubes than thin tubes. Subject *cw* was somewhat more accurate with thick tubes. For subject *etr* thick tubes were not systematically better.

Over the entire set of results shaded lines appeared to confer no systematic benefits over simple lines.

6 Discussion

Even though there are clearly large individual differences in the results there are some striking consistencies. For all subjects the largest factor influencing accuracy was whether paths were rendered using tubes as opposed to lines. Even without stereo and motion depth cues, tubes allowed for surprisingly accurate judgments. Thus the strongest recommendation that comes from this study is that tubes should be used to render 3D pathlines or streamlines.

It is worth noting that the thin (1mm) tubes we used would be equivalent to 0.5 mm diameter tubes at more normal viewing distances, and this should allow quite dense patterns of flowlines to be rendered. We cannot say to what extent the high resolution rendering was a factor in this since we made no comparison to low resolution rendering. But with a more conventional screen, a 0.5 mm tube would only be 2-3 pixels wide, and it seems unlikely that the same accuracy could be obtained.

The results of this study agree with the prior findings of Tyler [1975] that stereoscopic depth resolution falls off for detail finer than 3 cycles per degree. Thus it would appear that for the finer details rendering using fine tubes is the best solution since for two of the three subjects orientation judgments were still good for fine tubes having a 3 cycle per degree spatial frequency.

We were surprised by the lack of benefit conferred by shaded lines. One possible reason for has to do with the particular test pattern we used. Our paths were mainly vertical and the line shading method, mostly conveys information about the orientation with respect to the vertical axis in this case. But this was not measured in this experiment. Therefore it is likely that with more complex paths line shading would convey a greater benefit.

Overall, although our result showed that the biggest win came from rendering with tubes, the most consistently good judgments were obtained when stereoscopic and kinetic depth cues were available.

Finally, we recognize that this experiment is highly artificial. Gabor function paths were used to enable systematic variation of the effect of level of detail, but they are not representative of typical 3D flow patterns. For more realistic flows, rendered with multiple pathlines, the benefits of stereoscopic viewing and motion parallax might increase. Only further research will tell.

Acknowledgements

Thanks are due to Roland Arsenault for his contribution to the stereoscope software. This work was supported by the National Science Foundation ITR's 0324899, and the NOAA-funded Center for Coastal and Ocean Mapping.

References

CABRAL, B. AND LEEDOM, C. 1993. Imaging vector fields using line integral convolution. Proceedings of SIGGRAPH, '93, 27, 263-272.

CAMPBELL, F.W. AND GREEN, D.G. 1965. Monocular versus binocular visual acuity. Nature 209: 191-192.

FRISBY, J.P., BUCKLEY, D., AND DUKE, P.A. 1996. Evidence for recovery of lengths of real objects seen with natural stereo viewing. Perception, 25, 129-154.

GEIBEN, M., RUMPF, M. 1993. Visualization of finite elements and tools for numerical analysis. In Advances in Scientific Visualization. F.H. Post and AJS Hin (eds) Springer Verlag.

HOWARD, I.P. AND ROGERS, B.J. 1995. Binocular Vision and Stereopsis. Oxford Psychology Series No 29. Oxford University Press. Oxford.

HULTQUIST, J.P.M. 1992. Constructing Stream Surfaces in Steady 3D Vector Fields. Proceedings of Visualization '92. 171-198

INTERRANTE, V., AND GROSCH, C. 1998. Visualizing 3D flow. IEEE Computer Graphics and Applications. 18(4):49-53.

LAIDLAW, D.H. KIRBY, R.M., DAVIDSON, J.S., MILLER, T.S. DASILVA, M., WARREN, W.H. AND TARR, M. (2001) Quantitative comparative evaluation of 2D vector visualization methods. In Proceedings of IEEE Visualization '01, October.

LARAMEE, R.S., WEISKOPF, D., SCHNEIDER, J., AND HAUSER, H. 2004 Investigating swirl and tumble with a comparison of visualization techniques. Proceedings of IEEE Visualization (Vis 2004), 51-58, October 15-19, Austin,

VAN EE, R., AND SCHOR, C.M. 2000. Unconstrained stereoscopic matching of lines. Vision Research, 40, 151-162.

NORMAN, J.F. TODD, J.T. PEROTTI, V.I AND TITTLE, J.S. 1996. The visual perception of 3-D length. Journal of Experimental Psychology: Human Perception and Performance. 22, 173-186.

REZK-SALAMA, C. HASTREITER, P., TEITZEL, C., AND ERTL, T. 1999 Interactive exploration of volume line integral convolution based on 3D texture mapping. Proceedings IEEE Visualization '99, 233-240.

ROGERS, B. AND CAGENELLO, R. 1989. Disparity curvatures and the perception of three-dimensional surfaces. Nature 339: 137-139.

SOLLENBERGER, R.L., AND MILGRAM, P. 1993. The effects of stereoscopic and rotational displays in a three dimensional path tracing task. Human Factors, 35(3) 483-500.

SCHROEDER, W., VOLPE, C.R. AND LORENSEN, W.E. 1991. The stream polygon: a technique of 3D vector field visualization. Proceedings IEEE Visualization '99, 126-132.

TYLER, C. 1975. Spatial organization of binocular disparity sensitivity. Vision Research 15, 583-590.

UUMORI, K. AND NISHIDA, S. 1994. The dynamics of the visual system in combining conflicting KDE and binocular stereopsis cues. Perception and Psychophysics, 55(5), 526-536.

WALLACH, H. 1959. The perception of motion. Scientific American. July, 56-60.

WARE, C. AND FRANCK, G. 1996. Evaluating stereo and motion cues for visualizing information nets in three dimensions. ACM Transactions on Graphics. 15(2) 121-140.

WHEATSTONE, C. 1838. Contributions to the physiology of vision. Pert the first. On some remarkable and hitherto unobserved phenomena of binocular vision. Philosophical Transactions of the Royal Society. 128, 371-394.

ZOCKLER, M., STALLING, D., AND HEGE, H.C. 1996. Interactive Visualization of 3D Vector fields using illuminated streamlines. In Proceedings of IEEE Visualization '96 San Francisco. 107-113.