Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

On the Benefits of Using Constant Visual Angle Glyphs in Interactive Exploration of 3D Scatterplots

Stenholt, Rasmus

Published in:
A C M Transactions on Applied Perception

DOI (link to publication from Publisher):
[10.1145/2677971](https://doi.org/10.1145/2677971)

Publication date:
2014

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Stenholt, R. (2014). On the Benefits of Using Constant Visual Angle Glyphs in Interactive Exploration of 3D Scatterplots. *A C M Transactions on Applied Perception*, 11(4), 1-23. [19]. <https://doi.org/10.1145/2677971>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

On the Benefits of Using Constant Visual Angle Glyphs in Interactive Exploration of 3D Scatterplots

RASMUS STENHOLT, Aalborg University

Visual exploration of clouds of data points is an important application of virtual environments. The common goal of this activity is to use the strengths of human perception to identify interesting structures in data, which are often not detected using traditional, computational analysis methods. In this article, we seek to identify some of the parameters that affect how well structures in visualized data clouds can be identified by a human observer. Two of the primary factors tested are the volumetric densities of the visualized structures and the presence/absence of clutter around the displayed structures. Furthermore, we introduce a new approach to glyph visualization—constant visual angle (CVA) glyphs—which has the potential to mitigate the effect of clutter at the cost of dispensing with the common real-world depth cue of relative size. In a controlled experiment where test subjects had to locate and select visualized structures in an immersive virtual environment, we identified several significant results. One result is that CVA glyphs ease perception of structures in cluttered environments while not deteriorating it when clutter is absent. Another is the existence of threshold densities, above which perception of structures becomes easier and more precise.

Categories and Subject Descriptors: H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Evaluation/Methodology*; H.1.2 [Models and Principles]: User/Machine Systems—*Human information processing*

General Terms: Human Factors

Additional Key Words and Phrases: 3D shape perception, glyphs, scatterplots, visualization

ACM Reference Format:

Rasmus Stenolt. 2014. On the benefits of using constant visual angle glyphs in interactive exploration of 3D scatterplots. *ACM Trans. Appl. Percept.* 11, 4, Article 19 (December 2014), 23 pages.
DOI: <http://dx.doi.org/10.1145/2677971>

1. INTRODUCTION

Visualization of abstract data to make sense of it to humans is an old idea. In Tufte [2001], the discipline is regarded as having originated in the late 18th century. Much more recently, the use of computers and advances in computer graphics and virtual reality technology has allowed visualization forms that were completely infeasible in the past. Today, it is possible to view data visualizations from any conceivable angle, using any desired projection of the data, all in real time. The real-time nature of

The author wishes to thank Fonden Vision Nord for generously donating the funds needed to acquire some of the equipment used in this study.

Author's address: R. Stenolt, Aalborg University, Dept. of Architecture, Design & Media Technology, Rendsburggade 14, DK-9000 Aalborg, Denmark; emails: rs@create.aau.dk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 1544-3558/2014/12-ART19 \$15.00

DOI: <http://dx.doi.org/10.1145/2677971>

Fig. 1. Two versions of the same 3D scatterplot viewed from the same position and orientation: using the proposed CVA glyphs (left) and using traditional perspective glyphs (right). With CVA glyphs enabled, the blue, box-shaped structure appears through the yellow clutter while remaining hidden and imperceptible using traditional glyph rendering. The screenshots are taken at a distance (approximately 1m) from the structure, where the use of CVA glyphs is beneficial. At greater distances, traditionally rendered glyphs will feature less occlusion than their CVA counterparts.

computer visualizations also opens the possibility of interacting with the data. Examples of possible interactions are marking regions of interest, annotating data, or masking out undesired parts of the visualization. Including a human in the process of data analysis is based on the idea that the strong pattern recognition skills of human visual perception can be combined with expert knowledge of the domain from which the data originated to discover interesting patterns and structures otherwise missed by conventional computer analysis. This process is often referred to as exploratory visual data mining.

One common approach to data visualization is to map a subset of the variables of the records in a database to the visual attributes of small objects called *glyphs*. Thus, each observation in the dataset becomes a glyph in the visualization. The visual attributes of glyphs can be 3D position, color, size, shape, texture, and so forth. Glyph-based visualization of arbitrary data in real-time virtual environments is addressed in Nagel [2005], Feng et al. [2009], and Ropinski et al. [2011]. Although many other types of data visualization exist, such as volumetric rendering based on voxels [Kratz et al. 2006] or surface visualizations based on polygons or splats [Zhang and Kaufman 2006], the issues addressed in this article are aimed at 3D glyph-based visualizations. In a comprehensive survey of perceptual issues related to visualizations, Ware [2004] notes that little investigation has been made on the human perception of points and structures in 3D scatterplots. A search of literature on the subject reveals that although some work has been made since Ware's survey, there still seems to be many unexplored areas. Since a 3D scatterplot is typically the type of plot produced by a 3D glyph visualization, it seems relevant to study some of the parameters that may govern how well a human can perceive structures in such visualizations. In this article, a structure specifically refers to a contiguous 3D structure made up from a set of discrete glyphs by their placement in 3D space. This interpretation of 3D glyph placement is in line with Ropinski et al. [2011].

The presence of undesired glyphs, or *clutter*, around a structure is a condition that must be expected in 3D scatterplots created from most real-world datasets. Thus, occlusion caused by clutter becomes a frequently encountered problem, as demonstrated in the right part of Figure 1. In an attempt to

diminish this problem, we introduce and test a new way of rendering glyphs, referred to as constant visual angle (CVA) glyphs. Briefly explained, a CVA glyph is a 3D glyph that, when fully visible, is rendered such that it always covers the same amount of pixels on the screen and hence approximately the same visual angle, regardless of the distance from the viewer to the glyph. The rationale and technical details of CVA glyphs are explained in Section 3. The three parameters studied in this work are therefore (1) the volumetric density of the glyphs in the scatterplot, (2) the presence/absence of clutter that obscures the view of a structure, and (3) the use of CVA glyphs to diminish the adverse effects of clutter. The study will attempt to characterize the perception of the visualizations both from an objective and a subjective perspective. This approach allows us to understand if any subjective, perceptual differences reported by subjects can also be objectively measured.

This article is structured into six main sections. Following the Introduction, a short review of related works is presented in Section 2. Then, an in-depth walkthrough of the concepts and theory of CVA glyphs is given in Section 3. In Section 4, the hypotheses and design of the performed experiment is explained in detail. The results of the experiment are given in Section 5. Finally, a discussion and conclusion is given in Section 6.

2. RELATED WORKS

2.1 Visual Data Mining

Visual data mining (VDM) is a frequently studied area of scientific visualization. The collection of papers by Simoff et al. [2008] gives a good overview of recent work in this field. The concept of *serendipity*—that is, the human ability to discover something useful even when not looking specifically for it—plays an important role in exploratory VDM. In Granum and Musaeus [2002], the concept is explained in a VDM context. Interestingly, it is noted that the concept of *perceptual grouping* is important in VDM. Perceptual grouping is the tendency of human perception to group similar items into larger groups or structures and to focus on the structure rather than its individual, constituent objects. Perceptual grouping is treated in-depth in Treisman [1982]. The ability of perceptual grouping is closely related to the well-established Gestalt laws introduced in the early 20th century [Koffka 1935]. A more recent review of Gestalt laws is found in Ware [2004]. If the overall shape of a structure is sufficiently simple, it is also known as a geon [Irani and Ware 2003; Biederman and Gerhardstein 1993]. Geon structures are readily and immediately recognized by the human perceptual system, independent of viewpoint, as demonstrated in Biederman and Gerhardstein [1993], which makes them attractive for experiments with glyph-based visualizations. In this article, a box-shaped geon is used as the target structure in all experiments. Thus, a glyph visualization method that actively supports and reinforces perceptual grouping should have an advantage over one that does not do so. This observation is influential in the design of CVA glyphs.

In Granum and Musaeus [2002], some guidelines for visual parameters that can be associated with glyphs are also given. A distinction between static and dynamic attributes is made, with static attributes being those that do not vary with time. Dynamic attributes, on the other hand, are those that do vary with time, thus constituting different modes of motion, or animation, of the glyphs. In Nagel et al. [2008] and Nagel [2005], the authors describe the implementation of a framework for exploring 3D scatterplots in immersive virtual environments using both static and dynamic glyph attributes. However, they do not seem to have considered the possibility of improving on the perceptual grouping aspect by changing the rendering method of the glyphs. The main design consideration with respect to rendering of glyphs in these works seem to be performance related, in that only tetrahedral glyphs are used—that is, the simplest possible, closed 3D polygon surfaces. Another method for rendering glyphs is known as splatting [Krüger et al. 2005; Feng et al. 2009]. Using splatting, glyphs and other objects

are rendered using small images, called *splats*, instead of meshes of triangles. In Feng et al. [2009], splats are used in 3D scatterplots instead of mesh-based glyphs to present multivariate datasets using a method referred to as scaled data-driven spheres. Rendering splats instead of mesh-based glyphs is closely related to the CVA method presented in this article.

Other approaches to glyph-based visualization than those based on small 3D objects of fixed shape in a scatterplot have been studied in Forsell et al. [2006]. In that paper, a method for rendering glyphs as a continuous 2D surface is presented. With this type of glyph rendering, each glyph becomes a tile on a 2D grid, where the properties of the glyphs are represented as surface curvature, color, and shape.

Mazeika et al. [2008] present a method of turning the glyphs in a 3D scatterplot into a plot of continuous isodensity surfaces. The rationale of doing this is to increase the perception of large-scale structures in the data at the cost of removing the individual observations. In VDM terms, this means that the computer makes the perceptual grouping analysis that normally would be the human observer's task. As will be explained in Section 3, CVA glyphs have the potential to offer a view of both overall structures and individual observations, depending on the user's movements.

In a comparative study of user performance in different types of virtual environments [Qi et al. 2006], the results indicate that typical VDM tasks such as recognizing different glyph shapes and regions of high glyph density are more difficult using head-mounted displays (HMDs) than in fish tank VR-type displays. The authors speculate that this is caused by a lack of overview and a consequently higher mental load when using the HMD. This indicates that the tasks performed by our participants, who also use a HMD, may be difficult if no support through overview is given. However, as mentioned earlier, the fact that CVA has the potential to provide both overall structure and details may mitigate this effect. However, no comparison with other display systems will be made.

In Wang et al. [2010], a user study based on glyphs rendered using point sprites is presented. This is a technique that can also be used to render CVA glyphs. The authors, however, seem to have completely overlooked the potential perceptual benefits of such glyphs. Instead, they remove the CVA aspect of the glyph rendering to make a more realistic visualization. The main focus of the study is on the ability of users to identify and point out maxima and minima in a volumetric visualization of regularly grid-sampled temperature distributions, depending on parameters such as glyph size and density, and background environment. Their motivations for using point sprites is not stated, but it is probably a pure computer performance consideration, such that more glyphs can be rendered at the same computational cost. In the experiment described in Wang et al. [2010], the users viewed the visualization in 3D stereo on a projection screen without the ability to change their viewpoint much. For instance, it was impossible for the user to move into or away from the cloud of glyphs, which is exactly the mode of navigation that reveals the benefits of CVA glyphs. For these reasons, we claim that the VDM potential of CVA glyphs has been unexplored prior to the study presented in this article.

2.2 Managing Clutter in 3D Data Visualizations

Currently, there are several strategies for managing glyphs that obscure the view of a structure of interest in a 3D visualization. The first one is to remove unwanted glyphs from the visualization. This can be done interactively, such as by letting the user place additional clipping planes using his or her hands. The clipping plane approach is used in Wang et al. [2010] and Zudilova and Sloot [2003]. If the user-controlled clipping planes form a closed volume in space, a 3D region of interest is created, which is also known as a 3D magic lens [Viega et al. 1996]. User-controlled clipping planes are frequently used with medical datasets to see an arbitrary slice from a 3D dataset, such as in an MRI scan. Another possible approach to occlusion management through clipping planes would be to move the virtual camera's near clipping plane suitably far away from the viewer, such that any clutter in between the viewer and the near clipping plane is culled from the visualization. Compared to CVA

glyphs, this approach shares the common attribute that it provides a viewpoint controlled clutter management method. However, moving the near clipping plane away from the user would not add any of the reinforced perceptual grouping benefits that CVA glyphs provide. Instead, it would likely provide a more realistic experience than CVA glyphs, because it does not affect or remove any of the common, natural depth cues. The idea of modifying camera parameters to deal with occlusion from clutter in 3D has been studied in Elmqvist and Tsigas [2007]. There, they propose to let the user switch between using a realistic perspective and a parallel (orthographic) mode of projection to manage clutter. This idea is similar to CVA glyphs. However, CVA glyphs make simultaneous use of the two projection methods, instead of having two separate projection modes that the viewer must switch between. It is exactly this simultaneous mixture of projection modes that may reinforce perceptual grouping when using CVA glyphs.

Another approach to dealing with clutter is to make the glyphs semitransparent. This approach has the advantage of showing the clutter and the structure together. However, it diminishes the ability to assign visual attributes to the surface of the glyphs, such as texture and shape. Furthermore, if a perceived structure is related to the colors of the glyphs, the colors experienced by the user will be a mixture of the clutter's and structure's colors, which may possibly be confusing. Splatting relies heavily on the use of semitransparency to render discrete point sets as continuous-looking surfaces [Westover 1991; Zwicker et al. 2001]. The rendering method used for the CVA glyphs in this article, as stated earlier, is closely related to splatting.

A third approach called *distortion mapping* [Carpendale et al. 1996] makes a user-controlled geometric distortion of the 3D space, which pushes away glyphs far from the user's hand/cursor while condensing those nearby. The effects of distortion mapping are similar to CVA glyphs. However, with distortion mapping, nearby glyphs may still occlude the entire field of view (FoV). Furthermore, it requires the inclusion of a cursor or hand tracking into the virtual environment. This is not needed with CVA glyphs, as they are viewpoint controlled.

Yet another approach is simply to avoid 3D altogether and map all data variables to the properties of 2D objects or brush strokes [Forsell et al. 2006; Healey et al. 2004]. The use of 2D brush strokes is presented in Healey et al. [2004], where a nonphotorealistic visualization method is used for producing artistically inspired 2D visualizations from multivariate datasets. Removing the third dimension from the visualization can potentially eliminate all clutter. However, at the same time, it excludes the use of the third dimension as a useful visualization parameter. Excluding the use of the third dimension in data visualizations can be misleading to humans, as demonstrated in Shah and Carpenter [1995]. The methods presented in this article are specifically aimed at 3D visualization contexts. As will be explained in Section 3, CVA glyphs offer a different way of mitigating the adverse effects of clutter compared to the methods presented previously.

2.3 Interactions with 3D Scatterplots

The common interactions with a 3D scatterplot, like the interactions in most other applications of virtual environments, fit into the categories of navigation, selection, manipulation, or system control [Bowman et al. 2005]. There are, however, a couple of specialized versions of selection tasks that are frequently made in VDM applications: (1) brushing and linking and (2) multiple object selection (MOS). A prerequisite for both of these important interactions in VDM contexts is the ability to recognize structures. Brushing and linking is the activity of selecting and highlighting (called *brushing*) a subset of the glyphs in the current projection of the data to see what this highlighted subset looks like under different projections of the data, thereby *linking* an area of interest across several plots of the same dataset. The technique was introduced in Buja et al. [1991]. A more recent example of application of brushing and linking is presented in Isenberg and Fisher [2009].

Whereas selection tasks traditionally deal with indicating a single object of interest in a virtual world, MOS, or volumetric selection, involves the same task applied to a set of objects. This is an important task, which enables interactions such as brushing and linking as described earlier. The task given to the subjects in the experiment documented in this article amounts to a MOS-type task. In a study by Ulinski et al. [2007] with experimental factors similar to the ones investigated in this article, semitransparent splats were used for rendering the glyphs. Ulinski's study showed that the MOS performance of the participants was adversely affected by the presence of clutter. Furthermore, the results showed that correct selection was more difficult when the target splats were sparsely spread out in space instead of being densely clustered. No quantitative interpretation of the terms *sparse* and *dense* was given, nor were any intermediate density levels tested.

In a volumetric selection task, where users had to select an entire 3D object with all of its associated polygons [Zhai 1995], demonstrated that using a box-shaped, semitransparent selection volume called a *silk cursor* significantly increased selection accuracy compared to just having a wireframe version of the same selection box. For this reason, semitransparency will be used for the selection volume used in the experiments of this article.

3. THEORY OF CONSTANT VISUAL ANGLE GLYPHS

3.1 Definition

A central contribution of this work is the introduction of the CVA glyph. A CVA glyph is a 3D glyph like any other, except that its projected size does not change as the viewpoint moves closer or farther away. This implies that the perspective foreshortening experienced in real life does not affect the projected size of CVA glyphs, which is much like the effect gained by rendering objects using orthographic projections. The immediate consequence of this can be expressed in two ways:

- (1) CVA glyphs have a constant screen space size measured in pixels, independent of the viewing distance.
- (2) CVA glyphs cover an approximately constant solid angle of the viewer's FoV, independent of the viewing distance.

Both of these immediate consequences are only true if a glyph is not occluded or outside of the FoV, in which case the glyph will be partially, or entirely, invisible, thus reducing its pixel coverage/visual angle. In addition, the visual angle is only approximately constant if viewed at a reasonable FoV. For instance, if viewed using a large-perspective FoV, then the visual angle covered by a glyph will be significantly smaller at the periphery than at the center of the FoV. Objects with the CVA property have also been rendered prior to this work, such as the text labels presented in Polys et al. [2011]. Indeed, any visualization based on point sprites without perspective size correction has the CVA property. However, making active use of CVA as a beneficial property assigned to 3D objects (i.e., not just billboards of text) in a VDM context has not been the subject of previous study. The definition of CVA glyphs given previously has several implications. Some of these are related to the human perception of a visualization of CVA glyphs. Other implications concern the technical implementation of CVA glyph visualizations. These two areas of consequences are treated in the two following subsections.

3.2 Perceptual Implications of Constant Visual Angles

When perceiving the real world, we are used to the effect of the visual angle of objects increasing as we move in closer and vice versa as we move farther away. This effect is physically caused by the perspective transformation produced by the lens and retina of the human eye. However, since real-world objects generally do not change their physical size as we change our distance to them, the

Fig. 2. The consequence of CVA objects. The two objects, A and B, cover the same visual angle of the viewer. However, A and B are placed at different distances from the viewer. If the viewer is able to perceive this depth difference due to other depth cues than relative size, the logical consequence must be that A is smaller than B, as shown in the figure. Furthermore, if the viewer moves closer to A and the visual angle of A remains constant, then the physical size of A must be decreasing. Conversely, A's physical size will be perceived to be increasing if the viewer moves away from A, and the visual angle remains constant.

perceptual interpretation of this phenomenon is that an object's perceived size is constant, although its visual angle varies depending on the viewpoint. This fact is known as the size-distance invariance hypothesis [Blessing et al. 1967], a relation also expressed in Emmert's law [Emmert 1881]. The fact that perceived size is generally constant means that the visual angle covered by an object can be used to judge the distance to the object. This is often known as the relative size, or size gradient, depth cue.

With CVA glyphs, the visual angle remains constant, and hence the relative size depth cue will not work. Instead, the most likely inference about a CVA glyph's perceived size is that it must be changing as the viewer moves to accommodate its visual angle remaining constant. For instance, when moving closer to a CVA glyph, it will appear to become smaller and vice versa when moving farther away. This is shown in Figure 2. As such, CVA glyphs can be predicted not to conform to the size-distance invariance hypothesis.

Since relative size is no longer a meaningful depth cue with CVA glyphs, it is interesting to investigate how other depth cues are affected. Necessarily, some depth cues should remain unaffected for users to be able to make any sense of the 3D structure of a CVA glyph visualization. In a VDM context, it is worth noting that size can still be used as a meaningful glyph attribute. This is true because CVA glyphs can be assigned different but constant visual angles depending on the data variable mapped to the size attribute. For example, some glyphs can be chosen to always cover 200 pixels and others 100 pixels in the same visualization.

3.2.1 The Linear Perspective Depth Cue. One important observation to make about CVA glyphs is that the CVA property only affects *individual glyphs*. As such, the position of each glyph in a cloud can still undergo the usual perspective transformation. This has two important consequences:

- (1) The linear perspective depth cue still works on the overall structure of a cloud of CVA glyphs. Vanishing points will still exist in the visualization, and the overall structures of the cloud will be transformed as humans are used to from everyday life. For instance, far-away glyph clouds will cover less screen space than nearby clouds.
- (2) From a VDM point of view, 3D position is still a meaningful visual attribute to assign data variables to.

When the familiar linear perspective depth cue is combined with the CVA property, two potentially very useful viewing properties are created: the clutter-dispersion property and the structure-solidification property.

Fig. 3. A CVA glyph cloud with the clutter-dispersion and structure-solidification properties illustrated. Yellow glyphs are clutter, and blue ones are structure. Note that all glyphs are the same size in pixels across all four images. (A) The cloud is viewed at such a large distance that none of the blue structure is revealed. Instead, the clutter appears solid. (B) Moving closer to the cloud reveals the blue structure, relatively unobscured by the clutter. This situation also represents a sweet spot, where clutter dispersion and structure solidification are both helping. (C) Standing inside a structure causes it to spread out and become imperceptible. (D) Moving away from the structure causes it to solidify into a recognizable shape.

The clutter-dispersion property. As a viewer moves closer to a CVA glyph cloud, the positions of the glyphs will spread out in the FoV due to linear perspective. However, the individual glyph's pixel coverage will *not* increase due to the CVA property. Thus, nearby CVA glyphs will *not* begin to fill out the entire FoV, thereby avoiding occlusion of the view of whatever is behind them. The problem of not having CVA glyphs was illustrated in Figure 1. This means that any structure that is currently obscured by clutter can be cleaned of the clutter by simply moving in closer. This feature seems attractive, as almost any viewing environment will offer some means of navigation/viewpoint manipulation. In many immersive virtual environments, this navigation becomes extremely natural because it can be achieved by moving the head or body naturally. The viewer will often be able to select some viewing distance, where the clutter glyphs between the viewer and the structure have spread out so much that the structure displayed behind them appears relatively clutter free. This property is illustrated in Figure 3(a) and (b). The main caveat is that the clutter and the structure have to cover more or less disjoint subsets of 3D space—that is, some boundary between clutter and structure should exist. If the clutter and structure overlap too much spatially, both will spread out simultaneously.

The structure-solidification property. If, on the other hand, the viewer is placed in the middle of what is believed to be a structure of interest, it is possible to move away from the structure until it appears to solidify in the FoV. This property exists because linear perspective will decrease the on-screen distance between glyphs as you move away from them. However, since CVA glyphs retain their pixel coverage, the structure glyphs will eventually touch and overlap inside the FoV, creating the impression of a large, connected structure. Thus, the structure-solidification property supports the perceptual grouping aspect of human structure recognition previously mentioned in Section 2.1. The solidification property is illustrated in Figure 3(c) and (d). Without CVA, the glyphs would simply become increasingly smaller and not solidify into a structure. As was the case with the clutter-dispersion property, it is a requirement that structure and clutter do not overlap too much in 3D space. The advantages of the clutter-dispersion and the structure-solidification properties can often be combined. This happens if the viewer finds a sweet spot, where the distance to the structure both causes the clutter to spread out and the structure to solidify at the same time. Such a sweet spot is illustrated in Figure 3(b).

3.2.2 Other Depth Cues. One powerful depth cue that is completely unaffected by the CVA property is occlusion. For instance, the CVA property does not prevent nearby glyphs from occluding distant ones. However, correct occlusion depends on the geometry of the CVA glyphs to be well defined relative to the other glyphs and geometry in the scene. That is, the CVA must be transformed into a meaningful, physical size in the space of the scene. One approach to achieving this is presented in Section 3.3.

Correct occlusion is likely to be quite important for precise interactions with the glyphs in the cloud, as the user can then rely on occlusion between glyphs and interaction objects such as 3D cursors or selection rays. Depth inferred from shading and cast shadows is also possible with CVA glyphs. As with occlusion, the surface geometry of the glyphs must be well defined to make correct shading and shadows.

Since correct perspective transformation of the positions of individual glyphs is possible, the depth cue gained from parallax caused by self-motion is present as well. According to Ware [2004], structure from motion is likely to be a strong cue when perceiving structures in 3D scatterplots. For example, this observation is utilized in spin techniques for desktop viewing of scatterplots, where the plot will continuously spin around some axis [Donoho et al. 1988]. Since both shading and correct perspective transformation of positions in space is possible, depth at short distances can also be inferred from stereoscopic viewing. In the experiment documented in this article, the use of a motion-tracked, stereoscopic HMD will ensure that the user can take advantage of all of the depth cues mentioned earlier, with the exception of cast shadows, which was not implemented in the visualization software used for testing.

3.2.3 Depth Cues and Viewing Distance. The strength of the relative size cue is known to increase with viewing distance [Cutting and Vishton 1995]. This fact is of interest because the relative size cue is the only cue affected by the CVA property. In Cutting and Vishton, a distance-dependent hierarchy of depth cues is given. These authors present a split of 3D space into three distinct sections, referred to as personal space, action space, and vista space. The experiments of this work take place mainly in personal space, which roughly corresponds to objects that are within reach of a stationary observer. In personal space, the relative size cue is dominated by the cues of occlusion, stereoscopic viewing, and motion parallax. This implies that for visualizations displayed in personal space, the lack of the relative size depth cue may very well be an acceptable trade-off for the added clutter-dispersion and structure-solidification properties of CVA. The larger the distance to the visualized structures, the less acceptable this trade-off will probably be—at least if correct inferences about depth and 3D structure are desired. The conclusions of this article should therefore be viewed in the context that the experiments took place in personal space.

3.3 Implementation Options

There are two options when implementing CVA glyphs in a visualization:

- (1) Use a polygonal representation as with most other geometry rendered by a GPU.
- (2) Use point sprites (i.e., images of the glyphs) in the rendering.

For several reasons, the latter option seems much more attractive than the former one. First of all, the CVA property comes for free if point sprites are just naïvely rendered without doing any sort of depth-dependent size correction, such as with the correction outlined in Gross and Pfister [2007]. This is true because a point sprite is simply a fixed-size texture image that is copied by the rendering pipeline's rasterizer to the screen space positions of the desired set of vertices. That is, each visible vertex becomes a copy of the sprite image on-screen. Since the pixel size of the sprite is fixed, glyphs rendered in this way are automatically CVA.

If the same result is to be achieved with polygon-based objects, then each object would have to have a unique depth-dependent scaling transformation applied for each frame to make them always appear with the same screen space size. Moreover, depending on the shape, size, and viewing angle of the glyphs, the parameters of this scaling may only be possible to compute by doing a separate render pass for each glyph to find its current screen space extents. This would be unattractive from a computational

point of view. An alternative would be to render the glyphs using an orthographic projection. However, this would ruin the clutter-dispersion and structure-solidification features, as these features depend on perspective projection. Using point sprites also has the advantage of only requiring one vertex per glyph to be processed in the vertex and geometry stages of the rendering pipeline. Using even a simple shape like a tetrahedron would require four vertices per glyph to be processed.

The disadvantage of using point sprites is that the variety of glyph shapes that can be used is quite limited, provided that the goal is to have the glyphs appear as a true 3D objects instead of flat 2D billboards. If the shape of the glyph does not feature full rotational symmetry, then each sprite will require a different view-dependent texture for each frame rendered. This is also quite unattractive from an efficiency point of view. For this reason, spheres are probably the only shapes that are currently practical to implement as CVA glyphs. Thus, glyph shape is not a viable visual attribute to use in VDM contexts using CVA glyphs. In fact, this seems to be the main immediate drawback of using CVA glyphs. The only other significant limitation seems to be that the missing relative size cue of CVA glyphs may be unacceptable to human viewers. Investigating the latter issue is one of the aims of the experiment of this work.

With spheres, the sprite image becomes identical across all glyphs and rendered frames. Thus, it becomes very simple to compute the surface geometry of the glyphs needed for correct shading and occlusion. The only outstanding issue is to compute the perceived radius of each sphere, depending on its current depth. This can be done simply in a vertex shader by reversing the formula usually used for approximating the on-screen size of spherical point sprites to appear as if they had undergone regular perspective transformation. The formula is presented here and on page 277 of Gross and Pfister [2007]:

$$s = 2r \cdot \frac{n}{p_z} \cdot \frac{h}{t-b} \Leftrightarrow r = \frac{s}{2} \cdot \frac{p_z}{n} \cdot \frac{t-b}{h}. \quad (1)$$

The version of Equation (1) on the right-hand side is the one of main interest here. r is the perceived radius that would currently result from a sphere of pixel diameter s , if the sphere were viewed at a distance of p_z with a vertical screen resolution of h pixels, using a perspective projection with top, bottom, and near clipping plane parameters t , b , and n .

Thus, the formula provides the information needed to precisely tell where the sphere boundary is in relation to the sphere center, and in relation to all other geometry. This information can then be passed to a fragment shader, which takes care of making correct per-pixel shading and depth testing for the glyphs. As mentioned in Section 3.2.2, getting this correct is necessary to get the depth-from-shading and depth-from-occlusion cues to work.

One final point to note about the technical implementation of CVA glyphs using points sprites is that it is computationally much more attractive than the alternative of using semitransparency to see through clutter—that is, splatting. This is because semitransparency implemented using alpha blending can only be correctly rendered if all glyphs are depth sorted on a frame-by-frame basis, or if multipass techniques or exotic frame buffers, such as A-buffers [Carpenter 1984], are employed. Currently, no graphics card with an A-buffer exists. No expensive depth sorting is needed for opaque CVA glyphs.

4. EXPERIMENT

4.1 Experimental Method and Hypotheses

As stated previously, the experiment is designed to take place in personal space, as this is the location where CVA glyphs have the greatest potential to work well. The goals of the experiment can be expressed through some research questions. One goal is to investigate the usefulness of CVA glyphs for

structure recognition purposes. In this context, it is of particular interest to investigate the following questions:

- (1) How well can people perceive structures using CVA glyphs compared to regular glyphs?
- (2) How does the presence of clutter affect user performance with and without CVA glyphs?
- (3) Under what conditions, if any, are CVA glyphs beneficial to use?

From these questions, a couple of testable hypotheses can be stated, which forms part of the basis of the design of the experiment:

H1: CVA glyphs will provide better structure perception and structure reproduction performance than regular glyphs in cluttered environments.

H2: CVA glyphs will not pose a problem to the perception and interactions of the users compared to regular glyphs.

Hypothesis 1 is justified by the fact that CVA glyphs have the potential of dispersing clutter, as described in Section 3.2. Furthermore, if the user can find a sweet spot from which to view the visualization, the structure of interest should solidify while remaining relatively unoccluded by clutter. None of this will happen with regular 3D glyphs.

The justification of hypothesis 2 is twofold. Although CVA glyphs are missing the natural relative size depth cue, the amount of other depth cues present is still large. Furthermore, the experiment will take place in personal space, where the relative size cue is of less importance. The structure-solidification property may also compensate for the depth perception lost due to the missing relative size cue in some cases. This is true because structure solidification may help stimulate the perceptual grouping aspect of human perception, which is something that will not happen with regular glyphs.

A secondary objective is to investigate various density levels of the presented glyphs. This is necessary because the statements made in the first two hypotheses can only be reasonably expected to be true, if there are sufficiently many glyphs present to form any structure at all. The work of Ulinski et al. [Ulinski et al. 2007; Ulinski 2008] supports this as a necessary parameter, as they already identified that there is a significant difference in selection performance in sparse and dense clouds of splats. However, it seems important to reinvestigate this result in terms of quantitative density levels. Therefore, the necessary research questions about density must be as follows:

- (1) Where is the lower density threshold, if any, that makes perception of structures difficult or even impossible?
- (2) Where is the upper density threshold, if any, that makes structure perception trivial?
- (3) How does the use of CVA glyphs interact with these density thresholds—that is, do the thresholds change if CVA glyphs are used?

These questions can also be rephrased as testable hypotheses:

H3: There exists a lower threshold below which structure perception becomes difficult and structure selection performance decreases.

H4: There exists an upper threshold above which structure perception becomes easy and structure selection performance increases.

H5: There is a statistically significant interaction between density and the use of CVA glyphs.

Hypotheses 3 and 4 can both be justified through the results of the experiments presented in Ulinski et al. [2007] and by observing that the importance of perceptual grouping must mean that to see a structure, there must be enough glyphs present to visually form a structure. The solidification property

Table I A Table of All Experimental Conditions

	Without Clutter					With Clutter				
	Density [glyphs/m ³]									
CVA	1,000	5,000	10,000	50,000	100,000	1,000	5,000	10,000	50,000	100,000
Yes	0	4	8	12	16	2	6	10	14	18
No	1	5	9	13	17	3	7	11	15	19

Note: The numerical entries in the table are the unique IDs assigned to each condition. For example, condition 4 has a density of 5,000 glyphs/m³ with CVA glyphs enabled and no clutter.

justifies why hypothesis 5 should be correct. The possibility of CVA solidification should mean that structures can be perceived as continuous surfaces at lower density levels than without CVA.

4.1.1 Experimental Factors. From these hypotheses, the factors of the experiment can be identified as such:

- Density* of the structure glyphs, measured in glyphs/m³. This factor has five levels: 1,000 to 5,000 to 10,000 to 50,000 to 100,000. These specific levels were chosen because they clearly show differences when inspected on a regular screen. The upper limit of 100,000 was chosen, such that the structure is fairly straightforward to see while still allowing reasonable frame rates (approximately 20 to 25 fps) on the available hardware platform. The lower limit of 1,000 is chosen since it is judged to be difficult to see any structure to the glyphs at this level. Thus, any upper and lower structure shape detection thresholds should be located somewhere in the chosen range.
- Clutter* is a two-level factor with the levels “yes” and “no” depending on the presence (yes) or absence (no) of clutter. The amount of clutter scales proportionally to the density of the structure, such that there is always four times more clutter glyphs than structure glyphs. However, the clutter glyphs are always spread out over a larger volume of space around the structures, effectively producing a clutter density of around 20% of the structure density.
- CVA* is another two-level factor with the levels “yes” and “no” depending on the use of CVA glyphs (yes) or regular glyphs (no). In cases where CVA glyphs are not used, the on-screen radius of the glyphs will change in the natural way, as described in the left-hand version of Equation (1).

This implies that the experiment is going to be a three-way factorial design with $5 \times 2 \times 2 = 20$ unique experimental conditions. To collect more data from the participants, the experiment is carried out using a within-subjects design with replication. A small pilot experiment was performed, revealing that two replications of each condition were achievable in 1 to 1.5 hours per participant, which was judged to be reasonable. Having two replications means that each participant will go through a total of 40 trials. An overview of all test conditions is summarized in Table I.

4.1.2 Task and Response Variables. The task to be carried out by the user is closely related to the desired response variables that we wish to measure. The task in all trials is as follows. The participant is presented with a structure rendered as blue glyphs. If the scenario also involves clutter, then the clutter glyphs will be yellow. The choice of colors is based on getting good contrast between structure and clutter. The task of the user is then to first explore the glyph cloud to locate the structure and then to reproduce the shape of the structure in situ by indicating three of its extreme points. This implies that the task given is not just a perceptual one but rather a full loop of perception followed by action. This scheme allows for several response variables—objective and subjective—to be measured. Without the indication of extreme points, only subjective measurements of the structure perception performance could be given. The positions of the participant’s hands were marked by spherical 3D

Fig. 4. A screenshot from one of the trials with 10,000 glyphs/ m^3 and CVA enabled. The green (structure) and red (clutter) glyphs are those currently selected by the semitransparent, user-created box. The user-controlled cursors are represented by small cyan and magenta spheres. The cyan cursor is seen in the lower right foreground of the image.

cursors rendered using regular perspective projection. The colors of the cursors were distinct (magenta and cyan) from the colors used on the glyphs (blue, yellow). One of the cursors can be seen as a cyan sphere in the foreground of Figure 4.

To be able to get reliable, quantitative results about the reproduction of a structure, the structure shape needs to be one that is both simple to recognize and reproduce. For this reason, we have chosen that the structure in all cases is a random 3D box featuring all possible nine degrees of freedom (DoF). That is, the structures are randomly translated, rotated, and scaled boxes made up from discrete 3D glyphs. The possible box volumes are constrained to manageable sizes, their dimensions being forced to lie between 25 and 75 cm. This choice is also well in line with the notion that the task should take place in personal space. Thus, the task of the participants is to create the perceived, minimal oriented bounding box (OBB) [Ericson 2005] around the blue structure glyphs. The reasons for this choice are as follows:

- (1) Since 3D boxes are geons, they are well known and recognizable to most people.
- (2) 3D boxes have enough DoF to be nontrivial to reproduce—in other words, you have to know the locations of the faces and corners of the box to have perceived it correctly and subsequently be able to reproduce it accurately.
- (3) Effective and precise 3D box construction interaction techniques exist, such as hand on corner (HOC) introduced in Ulinski et al. [2007] or three corners (3C) introduced in Stenholt and Madsen [2011]. The latter is chosen since it has been tested to be more precise than HOC. Specifically, the technique presented as $3 + 3 + 3$ DoF 3C is employed. Briefly put, this means that the user creates a box by indicating three corners on the desired box, pointing them out one at a time. See Stenholt and Madsen [2011] for more details on 3C techniques.
- (4) It is possible to quantitatively measure the difference between two boxes [Stenholt and Madsen 2011].

An alternative approach to the experiment would be to include different classes of 3D structures (geons), such as spheres, tetrahedrons, and cylinders. However, we have chosen not to do so for several reasons. First of all, this would add an additional multilevel factor to the experiment, which would increase the number of unique test conditions to an unmanageable amount for this experiment. Second, immersive virtual environment versions of construction techniques for most other shapes do not

seem to have been extensively studied in literature, making efficient, standardized production of them a study of its own. Finally, one may conjecture that the findings of an experiment with box structures may generalize to other structures—at least to those of comparable complexity, such as with a comparable amount of DoF and symmetry.

Considering that we wish to understand the difficulty of perceiving structures, both from an objective and a subjective viewpoint, the following response variables were chosen:

- Box error*: The box error is an objective, metric measurement of the difference between the correct and the perceived, user-created box. It is measured in meters. The metric is computed as described in Stenholt and Madsen [2011].
- Perception difficulty*: This is a subjective measurement. The user is asked how difficult it was to perceive the minimal bounding box of the structure on a scale from 1 (trivial) to 10 (impossible).
- Selection difficulty*: Reproducing the box structure in situ effectively amounts to a MOS task, where the user selects the glyphs in the structure using a tightly fitting box. The user is asked how difficult it was to reproduce the box structure on a subjective 1 to 10 scale like the one mentioned earlier.
- Selection sensitivity*: In selection experiments, the sensitivity is often used as a measure of quality. In this context, the sensitivity is the ratio of the number of selected structure glyphs to the total number of structure glyphs. This is an objective, quantitative measurement. It does not necessarily correlate with the box error, as the whole structure can be selected by many different boxes. Thus, any deviation from the minimal bounding box will adversely affect the box error, although not necessarily the sensitivity.
- Completion time*: The completion time of each trial is also measured as an objective response variable. Although the hypotheses do not deal explicitly with completion time, it is nonetheless interesting to see if there are any differences between the test conditions with respect to speed.

The clutter and structure glyphs are spatially placed such that the structure is completely embedded inside a spherical cloud of clutter. The spherical shape ensures that no inferences about the shape of the structure box can be made from the shape of the clutter. However, the structure volume must at the same time be completely free of clutter, allowing the user, in principle, to make a perfect selection of structure glyphs. To help the participants' perception of the volume covered by their currently created box, the color of structure glyphs changes from blue to green when selected. Similarly, false positives—that is, selected clutter glyphs—change from yellow to red once selected. This produces an alternative way of judging the current correctness of the constructed box by maximizing the amount of green glyphs and minimizing the amount of red ones. An example illustration from the experiment is shown in Figure 4. The virtual room was designed to fulfill several purposes. First of all, the dimensions of the virtual room corresponded approximately to the dimensions of the physical room in which the experiment was conducted to discourage participants from walking into physical walls. Second, the room featured textured tiles to give additional depth cues, particularly linear perspective, to the participants about the surroundings of the structure. The textures used were kept in shades of grey to avoid any confusion with the colors of the glyphs. A portion of the virtual room is seen in Figure 3.

The final consideration to make is the sampling strategy for the structure box. The glyph structure must necessarily have the sampled target box as its minimal bounding box. Otherwise, the ground truth to measure the box error against disappears. This means that either four glyphs must always be placed at four noncoplanar corners of the box, or that some glyphs must be placed on each face of the box. We have chosen the latter approach because pilot tests revealed that always having four box corners explicitly represented as glyphs was too conspicuous. It made the box shape identification a

Fig. 5. The HMD and mice used in the experiment. The HMD and both mice are fitted with infrared reflective markers to be identifiable by the OptiTrack system.

trivial task of indicating three already visualized locations in space, thus defeating the goal of finding density-level thresholds. The approach is therefore to always sample three glyphs on each of the six faces of the structure box. Apart from the 18 face glyphs, all other glyphs are randomly sampled from the entire structure's volume. Thus, there are no guarantees that the structure glyphs will be evenly spread out inside the volume of the box, such as could be achieved using sampling strategies such as n-rooks or Hammersley sampling [Suffern 2007]. The random sampling approach was chosen to best approximate the most likely conditions if the glyph positions had been mapped from an arbitrary real-world database. In any case, random sampling also represents the worst-case scenario for the purposes of human structure perception. Any kind of regularity in the positioning of the glyphs would most likely help in perceiving the correct structure.

4.2 Experimental Equipment

The experiment was performed using an nVisor SX111 HMD and a 24-camera OptiTrack setup for motion tracking. The HMD features a 66% partial overlap FoV of 102 degrees (horizontal) by 64 degrees (vertical). The HMD's resolution is 1280×1024 pixels per eye. The setup of the tracking system allows the participants to walk freely within a 2.25m radius from the center of the tracked space. Two off-the-shelf wireless presenter mice were fitted with tracking markers and used as interaction devices. All necessary actions—that is, indications of perceived box corners—were performed by clicking the left mouse buttons. The HMD and mice are shown in Figure 5. All of the scene, including cursors, was rendered stereoscopically in the HMD.

The visualization software used was a custom-made OpenGL renderer running on a 32-bit Windows XP computer with an Intel Q8200 2.33GHz quad-core CPU and 4GB of memory. The motion tracking software was implemented using the NaturalPoint Tracking Tools API. The graphics card used was an nVidia GTX 285 GPU.

4.3 Experimental Procedure

Before doing any trials of the experiment, each participant answered a few demographical questions about age, gender, and prior 3DUI experience. Then, a brief explanation of the equipment followed to allow the participants to familiarize themselves with the adjustment options of the HMD and the left button position on the mice. All participants were made aware that they could request breaks or termination of the experiment at any time. After donning the HMD, the participants were presented with a few easily perceptible practice structures of high density to explain the procedure of each trial.

In connection with this, a short training session in the use of the 3+3+3 DoF 3C box shaping technique was given, where participants could practice replicating some easy box structures. Participants were also explained about the difference between blue structure glyphs and yellow clutter glyphs, as well as the color changes induced by selecting glyphs. During the experiment, the participants were standing up and allowed to walk freely within the confines of the tracking volume.

A total of 19 unpaid participants took part in the test—17 male and 2 female. However, two of the male participants had to quit the experiment during the training phase due to cybersickness. Thus, all further analysis is only based on the 17 participants who completed the entire experiment. In the present case, the 17 participants tried all conditions twice, so the sample size is 34 for all conditions. The collected dataset represents 680 completed trials in total. All participants were recruited at the local university campus. The mean age was 28.5 years. The median 3DUI experience level reported on a subjective 1 (novice) to 5 (expert) scale was 2, meaning that the typical participant had experience with 3D computer games but not with advanced 3D viewing or interaction equipment such as the HMD used.

The instruction given to the participants was to approximate the minimal bounding box as well as possible in all cases. The participants were not made actively aware of the introduction of CVA glyphs prior to or during the trials. Thus, they remained ignorant of the peculiarities of CVA glyphs unless they noticed anything by themselves. Participants were allowed to restart individual trials if they felt that they had made a mistake when pointing out a box corner. In cases where the structure's bounding box was not perceptible, participants were instructed to make their best guess, nonetheless.

After the practice session, the actual experiment with logging of data began. The sequence of test conditions was randomized to avoid any possible bias induced by a fixed sequence. Both replications of the same test condition were in immediate sequence. All target structures were randomized on a trial-by-trial basis. After the two replications of each condition, participants were asked about their subjective opinion of (1) the difficulty of perceiving the bounding box around the blue glyphs and (2) the difficulty of reproducing the perceived bounding box, both on the previously mentioned 1 to 10 scale. As explained earlier, each participant tried all 20 test conditions twice. The test conditions are presented in Table I. On completion of all trials, an informal debriefing took place.

5. RESULTS

5.1 Preanalysis

All statistical analyses were performed using R [R Development Core Team 2011], and a significance level of $\alpha = 0.05$ was used in all tests. A preanalysis was performed to see if the data conformed to the requirements of ANOVA analysis—that is, independence, equal variance, and normality. Independence is secured by the fact that the trials were independently performed by different participants, as well as the fact that the sequence of the test conditions was randomized for each participant. However, Bartlett's and Levene's tests of equal variance reveal that the equal variance assumption is likely violated for all measured response variables. Although some controversy exists on the issue of robustness to violations of this assumption, recent literature on the issue, such as Norman [2010], indicates that even quite large violations of ANOVA's assumptions do not lead to wrong conclusions in general.

Based on the conclusions of Norman [2010], the relatively large total sample size of 680 trials, and a comparison of the outputs of a Friedman test (a nonparametric version of ANOVA) and a regular ANOVA, which revealed no noteworthy differences, we chose to go ahead with ANOVA regardless of the violations. All post hoc tests are made using Tukey's honest significant difference (HSD). The means and standard errors of the means for the box errors and subjective perceptual difficulty ratings for each tested condition are reported in Tables II through V.

Table II Mean ($\bar{x}$) and Standard Error of the Mean (SEM) for Box Errors [m] for All Experimental Conditions Encoded as $\bar{x}$, SEM

	Without Clutter					With Clutter				
	Density [glyphs/m ³]									
CVA	1,000	5,000	10,000	50,000	100,000	1,000	5,000	10,000	50,000	100,000
Yes	1.05, 0.074	0.62, 0.066	0.61, 0.073	0.34, 0.042	0.35, 0.061	1.02, 0.067	1.02, 0.077	0.76, 0.085	0.66, 0.077	0.69, 0.081
No	1.06, 0.072	0.66, 0.064	0.61, 0.070	0.35, 0.051	0.32, 0.053	1.19, 0.051	0.71, 0.066	0.67, 0.077	0.99, 0.12	1.23, 0.12

Table III Mean ($\bar{x}$) and SEM for Perceptual Difficulty Rankings [1 to 10] for All Experimental Conditions Encoded as $\bar{x}$, SEM

	Without Clutter					With Clutter				
	Density [glyphs/m ³]									
CVA	1,000	5,000	10,000	50,000	100,000	1,000	5,000	10,000	50,000	100,000
Yes	6.94, 0.307	4.88, 0.342	3.88, 0.292	2.18, 0.161	1.82, 0.200	8.00, 0.260	7.65, 0.254	6.12, 0.279	5.18, 0.364	6.00, 0.422
No	7.65, 0.310	4.41, 0.216	4.59, 0.293	1.94, 0.140	1.88, 0.178	8.59, 0.216	6.94, 0.340	6.06, 0.375	6.65, 0.454	8.94, 0.345

Table IV Mean ($\bar{x}$) and SEM for Selection Difficulty Rankings [1 to 10] for All Experimental Conditions Encoded as $\bar{x}$, SEM

	Without Clutter					With Clutter				
	Density [glyphs/m ³]									
CVA	1,000	5,000	10,000	50,000	100,000	1,000	5,000	10,000	50,000	100,000
Yes	5.94, 0.313	5.12, 0.337	4.47, 0.268	2.76, 0.250	3.21, 0.337	7.65, 0.391	7.35, 0.286	6.82, 0.306	5.94, 0.313	7.00, 0.435
No	7.47, 0.305	5.29, 0.303	4.59, 0.363	2.94, 0.301	2.82, 0.269	7.12, 0.321	6.65, 0.280	6.18, 0.410	7.76, 0.380	9.24, 0.341

Table V Mean ($\bar{x}$) and SEM for Selection Sensitivity [0 to 1] for All Experimental Conditions Encoded as $\bar{x}$, SEM

	Without Clutter					With Clutter				
	Density [glyphs/m ³]									
CVA	1,000	5,000	10,000	50,000	100,000	1,000	5,000	10,000	50,000	100,000
Yes	0.65, 0.034	0.73, 0.024	0.77, 0.020	0.85, 0.015	0.85, 0.014	0.55, 0.039	0.65, 0.031	0.65, 0.027	0.70, 0.026	0.72, 0.022
No	0.60, 0.036	0.73, 0.029	0.76, 0.023	0.84, 0.031	0.85, 0.018	0.60, 0.034	0.63, 0.032	0.70, 0.025	0.63, 0.035	0.46, 0.044

We decided to apply variance stabilizing transformations on each of the response variables as outlined in Montgomery [2008] to improve the situation. However, we could not entirely fix the issue of equal variance for most of the response variables. The selection difficulty did not indicate any need for variance stabilization. The data analysis is based on the transformed responses. The following transformations were estimated from the data and applied to the box errors, δ , the perceptual difficulties, p , the completion times, t , and the sensitivities, s :

$$\delta' = \log(\delta), p' = \sqrt{p}, t' = \frac{1}{\sqrt{t}}, s' = s^{2.5}. \quad (2)$$

As a final preliminary analysis, a boxplot of the transformed variables reveals that, at least visually, there appears to be interesting differences among the tested conditions. As an example, in Figure 6, the transformed box errors and perceptual difficulties are displayed grouped by condition.

5.2 Analysis

5.2.1 The Big Picture. A three-way within-subjects ANOVA analysis of the transformed box error reveals that all factors and interactions between two and three factors, except for the main effect of

Fig. 6. Box plots grouped by condition of box error (top), an objective measure of task difficulty, and of reported perceptual difficulty (bottom), a subjective measure of task difficulty. Note that the vertical axis in the box error plot is logarithmic. The left-most 10 conditions are all without CVA glyphs, whereas the right-most 10 conditions are all with CVA glyphs. The conditions within each colored group show the result of changing the density while keeping the other two factors constant. Note that all four low density conditions (1,000 glyphs/m³) appear both objectively and subjectively difficult. Also note the apparent beneficial difference that CVA glyphs make in the densest, cluttered conditions (i.e., compare the green boxes to the purple ones). Furthermore, it is interesting to see that the participants' subjective perceptions of difficulty are similar to the objectively measured errors.

using CVA, are significant. The significant p -values here range from $F_{1,16} = 5.48$, $p = 0.032$ (clutter-CVA interaction) to $F_{4,64} = 50.54$, $p = 3.95 \cdot 10^{-19}$ (density main effect). For instance, there are significant differences among the specific test conditions (three-factor interactions), as well as significant overall differences caused by the main effects and various combinations of the main effects. The results with respect to perception difficulty, selection difficulty, and sensitivity show significant effects and interactions in mostly the same places as the box error.

The completion time shows few significant effects compared to the other response variables. The main significant effect is that trials with clutter are significantly slower to complete than those without clutter, $F_{1,16} = 27.51$, $p = 8.02 \cdot 10^{-5}$. One main trend of the results is that participants found it difficult to perceive the box structures in the lowest density cases with 1,000 glyphs/m³. These cases were significantly more difficult than the other tested conditions, except for the cluttered, 100,000 glyphs/m³ case without CVA glyphs, which turned out to be the most difficult tested condition, no matter what response variable is considered. Another main trend is that the scenarios with the highest densities and no clutter were significantly easier than most of the other tested conditions.

The main goal of the experiment was to investigate the five proposed hypotheses. This means that instead of trying to understand every significant difference in the dataset, we have chosen to only investigate those in detail that matter the most to the hypotheses.

5.2.2 Hypotheses 1 and 2. H1 stated that the use of CVA glyphs would help users to perceive and reproduce the displayed structures more precisely in cluttered environments.

Investigating this with respect to the structure perception first, the main variable to look at is the reported perceptual difficulty. With respect to perceptual difficulty, CVA in itself is a significant effect, $F_{1,16} = 5.98$, $p = 0.026$. Interestingly, this means that across all tested conditions, the participants actually felt that it was easier to perceive structures with CVA glyphs than without. If we check how CVA glyphs fare in the most cluttered environments—that is, those with clutter and the highest structure density, we conclude that CVA glyphs do indeed make a significant difference. A paired t -test reports $t_{33} = 7.39$, $p = 1.75 \cdot 10^{-8}$. The visual difference caused by CVA in this situation was the one illustrated in Figure 1.

If H1 is investigated using the objective box error, we find that CVA makes a significant beneficial difference ($t_{33} = 5.46$, $p = 4.75 \cdot 10^{-6}$) in the most dense, cluttered condition, although not as a general main effect. Investigating H1 with respect to the reported selection difficulty shows the same result ($t_{33} = 5.27$, $p = 8.4 \cdot 10^{-6}$). In terms of sensitivity, we get $t_{33} = -4.29$, $p = 1.48 \cdot 10^{-4}$. This means that once the amount of clutter becomes very high, effectively covering the view of an otherwise well-defined structure, CVA glyphs make a significant difference no matter how we have chosen to measure the difficulty of the task. However, contrary to the statement of H1, CVA does not make a significant difference in general in cluttered environments. The correct statement of a revised H1 would therefore be that CVA glyphs help the perception of structures in cluttered environments where occlusion of the structure otherwise becomes dominant.

H2 stated that using CVA glyphs would not cause perceptual problems compared to rendering the glyphs in the natural way. From the preceding analysis of H1, H2 is already supported by the data. Perhaps the most surprising result here is that generally when using the participants' subjective opinions about the difficulty of perceiving and selecting structures, CVA glyphs turn out to be significantly better to use than regular glyphs. When measured using the objective box error, there is no significant difference across all scenarios. This means that CVA glyphs do not reduce people's objective ability to create the correct box around a structure, whereas they subjectively perceive it as being significantly easier with CVA glyphs. Thus, H2 is supported based on the available data.

5.2.3 Hypotheses 3 and 4. H3 stated that there exists a lower density threshold below which structure perception becomes more difficult. To test this, the main factor of interest is density. The performed ANOVA shows that this factor is significant with respect to all measured response variables except for completion time, all $p \ll 0.001$. Digging into the pairwise comparisons made by a Tukey's HSD comparison shows that the participants find the density level of 1,000 glyphs/m³ more difficult than all other density levels no matter what response variable is used (except for completion time). The p -values here range from $p = 0.005$ to $p < 0.001$. The only interpretation of this must be that a lower threshold level (or, more likely, a threshold range) exists somewhere between 1,000 and 5,000 glyphs/m³, which causes the perception of a shape from a discrete set of glyphs to become much more difficult, if not impossible. H3 is therefore unambiguously supported by the data.

H4 postulates the existence of an upper density threshold, above which structure perception becomes easier. Investigating the data shows significant differences between 10,000 and 50,000 glyphs/m³ for the response variables box error, perception difficulty, and sensitivity— p -values in the range $p = 0.044$ to $p < 0.001$. This means that the desired threshold probably lies between 10,000 and 50,000 glyphs/m³. The only response variables not indicating the threshold to be at the same location is

the self-reported selection difficulty that produces a significant difference between 5,000 and 50,000 glyphs/m³. Since this range mostly overlaps the range indicated by the other response variables, it seems most reasonable to say that more than 10,000 glyphs/m³ are needed to significantly improve structure perception. Thus, H4 is found to be in agreement with the data.

5.2.4 Hypothesis 5. The final hypothesis, H5, stated that the use of CVA glyphs would change the density thresholds postulated by H3 and H4. To investigate this hypothesis, the interaction term between CVA and density is of interest. This interaction was significant with respect to all response variables except for completion time. The box error shows that the threshold between 1,000 and 5,000 glyphs/m³ vanishes when CVA glyphs are used. Conversely, the significant difference between 10,000 and 50,000 glyphs/m³ vanishes without CVA glyphs. The broad interpretation of this must be that at low densities, the familiar relative size depth cue is more useful than the CVA property, but at higher densities CVA provides better structure perception. For instance, the thresholds have been moved or smoothed out by the use of CVA. The reported perception difficulty also shows that with CVA glyphs enabled, the significant difference between 1,000 and 5,000 glyphs/m³ disappears. However, the threshold between 10,000 and 50,000 glyphs/m³ remains significant. Thus, perceptually, the use of CVA glyphs has made the low-density cases more similar. With the subjective selection difficulty, the results of the analysis gives the same conclusion as the box error—in other words, that CVA glyphs make it easier to select glyphs at high densities but make no difference at low densities. The sensitivity thresholds are unaffected by the use of CVA glyphs. Thus, the conclusion about H5 must be that the use of CVA glyphs generally has an effect on structure perception density thresholds.

6. DISCUSSION AND CONCLUSION

In summary, this article has made the following contributions:

- (1) The introduction of CVA glyphs—a novel approach to 3D glyph rendering—where glyphs are rendered such that the individual glyphs always appear with the same pixel coverage on-screen. This property is very useful when combined with perspective projection.
- (2) An analysis of the perceptual consequences of CVA glyphs based on current theories, particularly why CVA glyphs have the potential to be useful as a clutter management technique in VDM of 3D scatterplots.
- (3) A discussion of various issues related to the technical implementation of CVA glyphs was given, arriving at the conclusion that point sprites are probably the best, currently feasible implementation option.
- (4) A demonstration through a rigorous experiment carried out in an immersive virtual environment that CVA glyphs do indeed have the expected advantages over regular glyphs, especially in conditions featuring a lot of occlusion from clutter.
- (5) Identification of possible threshold regions of glyph densities that significantly change participants' ability to perceive structures inside a set of discrete 3D glyphs.

The fact that CVA glyphs not only turned out to be usable, but in many cases superior to glyphs rendered in a more natural way, is an interesting result in itself. This means that the loss of the real-world relative size depth cue can be an advantage in applications, where data is visualized in personal space—that is, within a couple of meters from the viewer. Further testing is needed to see how CVA glyphs perform in the action and vista spaces. This result is made even more valuable by the fact that the proposed CVA glyphs not only outperformed traditional glyphs in very cluttered conditions but also generally did not perform any worse than realistic glyphs in other conditions. One reservation to keep in mind about these results is that the recruited participants were predominantly male university

students. It is therefore an interesting future prospect to see if these results also generalize to a broader population.

One interesting but informal observation made during the course of all trials of the experiment is that only 3 out of 17 participants noticed anything unnatural about the CVA glyphs. Furthermore, nobody made any statement that the CVA glyphs were annoying or unpleasant. The participants had no problems in using the clutter-dispersion and structure-solidification properties of CVA glyphs, which seemed quite natural to them. CVA glyphs also turned out to be advantageous in that the interactions needed for reproducing the structure boxes were more manageable in the very cluttered cases. This observation, made by several of the participants, is caused by the fact that the clutter-dispersion property of CVA glyphs also causes the user-controlled cursors to be more easily visible through the clutter.

Finally, the density thresholds found through the experiment are not only useful for CVA glyphs; as well, the thresholds provide general knowledge that can be used in planning the sampling density of data for any 3D scatterplot rendering if the goal is to have the viewers perceive shapes and structures in the plot.

REFERENCES

- Irving Biederman and Peter C. Gerhardstein. 1993. Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance* 19, 6, 1162.
- William W. Blessing, Ali A. Landauer, and Max Coltheart. 1967. The effect of false perspective cues on distance and size judgments: An examination of the invariance hypothesis. *American Journal of Psychology* 80, 2, 250–256. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=pubmed&cmd=Retrieve&dopt=AbstractPlus&list_uids=1967-11235-001.
- Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. 2005. *3D User Interfaces: Theory and Practice*. Addison-Wesley.
- Andreas Buja, John Alan McDonald, John Michalak, and Werner Stuetzle. 1991. Interactive data visualization using focusing and linking. In *Proceedings of the 2nd Conference on Visualization (VIS'91)*. IEEE, Los Alamitos, CA, 156–163. <http://dl.acm.org/citation.cfm?id=949607.949633>.
- M. Sheelagh T. Carpendale, David J. Cowperthwaite, and Frank D. Fracchia. 1996. Distortion viewing techniques for 3-dimensional data. In *Proceedings of the IEEE Symposium on Information Visualization (INFOVIS'96)*. IEEE, Los Alamitos, CA, 46. <http://dl.acm.org/citation.cfm?id=857187.857608>.
- Loren Carpenter. 1984. The A-buffer, an antialiased hidden surface method. *ACM SIGGRAPH Computer Graphics* 18, 3, 103–108. DOI: <http://dx.doi.org/10.1145/964965.808585>
- James E. Cutting and Peter M. Vishton. 1995. Perceiving layout and knowing distances: The integration, relative potency and contextual use of different information about depth. In *Handbook of Perception and Cognition*, W. Epstein and S. Rogers (Eds.). Academic Press, 69–117.
- Andrew W. Donoho, David L. Donoho, and Mitiam Gasko. 1988. MacSpin: Dynamic graphics on a desktop computer. *IEEE Computer Graphic and Applications* 8, 4, 51–58. DOI: <http://dx.doi.org/10.1109/38.7749>
- Niklas Elmqvist and Philippas Tsigas. 2007. View-projection animation for 3D occlusion management. *Computers and Graphics* 31, 6, 864–876. DOI: <http://dx.doi.org/10.1016/j.cag.2007.09.006>
- Emil Emmert. 1881. Größenverhältnisse der Nachbilder. *Klinische Monatsblätter für Augenheilkunde und für augenärztliche Fortbildung* 19, 443–450.
- Christer Ericson. 2005. *Real-Time Collision Detection*. Morgan Kaufmann.
- David Feng, Yueh Lee, Lester Kwock, and Russell M. Taylor II. 2009. Evaluation of glyph-based multivariate scalar volume visualization techniques. In *Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization (APGV'09)*. ACM, New York, NY, 61–68. DOI: <http://dx.doi.org/10.1145/1620993.1621006>
- Camilla Forsell, Stefan Seipel, and Mats Lind. 2006. Surface glyphs for efficient visualization of spatial multivariate data. *Information Visualization* 5, 2, 112–124. DOI: <http://dx.doi.org/10.1057/palgrave.ivs.9500119>
- Erik Granum and Peter Musaeus. 2002. *Constructing Virtual Environments for Visual Explorers*. Springer-Verlag, London, UK, 112–138. <http://dl.acm.org/citation.cfm?id=778632.778642>.
- Markus Gross and Hanspeter Pfister. 2007. *Point-Based Graphics*. Morgan Kaufmann. <http://books.google.com/books?id=9LCQ86u9Sc4C>.

- Christopher G. Healey, Laura Tateosian, James T. Enns, and Mark Remple. 2004. Perceptually based brush strokes for nonphotorealistic visualization. *ACM Transactions on Graphics* 23, 1, 64–96. DOI: <http://dx.doi.org/10.1145/966131.966135>
- Pourang Irani and Colin Ware. 2003. Diagramming information structures using 3D perceptual primitives. *ACM Transactions on Computer-Human Interaction* 10, 1, 1–19. DOI: <http://dx.doi.org/10.1145/606658.606659>
- Petra Isenberg and Danyel Fisher. 2009. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum* 28, 3, 1031–1038. DOI: <http://dx.doi.org/10.1111/j.1467-8659.2009.01444.x>
- Kurt Koffka. 1935. *Principles of Gestalt Psychology*. Vol. 7. Harcourt, Brace, and World. <http://books.google.com/books?hl=en&lr=&id=cLnqI3dvi4kC&oi=fnd&pg=PA3&dq=Principles+of+Gestalt+Psychology&ots=vVxmaQ-ehm&sig=pT63mGsLA3uz5hqC9EEYf0IS0uY>
- Andrea Kratz, Markus Hadwiger, Rainer Splechtna, Anton Fuhrmann, and Katja Buhler. 2006. GPU-based high-quality volume rendering for virtual environments. In *Proceedings of the Workshop on Augmented Environments for Medical Imaging and Computer-Aided Surgery (AMI-ARCS'06)*. <http://www.vrvis.at/publications/PB-VRVis-2006-034>.
- Jens Krüger, Peter Kipfer, Polina Kondratieva, and Rüdiger Westermann. 2005. A particle system for interactive visualization of 3D flows. *IEEE Transactions on Visualization and Computer Graphics* 11, 6, 744–756.
- Arturas Mazeika, Michael H. Böhlen, and Peer Mylov. 2008. Using nested surfaces for visual detection of structures in databases. In *Visual Data Mining: Theory, Techniques, and Tools for Visual Analytics*, S. Simoff, M. H. Böhlen, and A. Mazeika (Eds.). Springer, 91–102.
- Douglas C. Montgomery. 2008. *Design and Analysis of Experiments* (7th ed.). Wiley.
- Henrik R. Nagel. 2005. *Exploratory Visual Data Mining in Spatio-Temporal Virtual Reality*. Ph.D. Dissertation. Faculty of Engineering and Science, Aalborg University.
- Henrik R. Nagel, Erik Granum, Søren Bovbjerg, and Michael Vittrup. 2008. Immersive visual data mining: The 3DVDM approach. In *Visual Data Mining: Theory, Techniques, and Tools for Visual Analytics*, S. Simoff, M. H. Böhlen, and A. Mazeika (Eds.). 281–311.
- Geoff Norman. 2010. Likert scales, levels of measurement and the “laws” of statistics. *Advances in Health Sciences Education Theory and Practice* 15, 5, 625–632. <http://www.ncbi.nlm.nih.gov/pubmed/20146096>.
- Nicholas F. Polys, Doug A. Bowman, and Chris North. 2011. The role of depth and Gestalt cues in information-rich virtual environments. *International Journal of Human-Computer Studies* 69, 1–2, 30–51. DOI: <http://dx.doi.org/10.1016/j.ijhcs.2010.05.007>
- Wen Qi, Russell M. Taylor II, Christopher G. Healey, and Jean-Bernard Martens. 2006. A comparison of immersive HMD, fish tank VR and fish tank with haptics displays for volume visualization. In *Proceedings of the 3rd Symposium on Applied Perception in Graphics and Visualization (APGV'06)*. ACM, New York, NY, 51–58. DOI: <http://dx.doi.org/10.1145/1140491.1140502>
- R Development Core Team. 2011. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/> ISBN 3-900051-07-0.
- Timo Ropinski, Steffen Oeltze, and Bernhard Preim. 2011. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers and Graphics* 35, 2, 392–401. DOI: <http://dx.doi.org/10.1016/j.cag.2011.01.011>
- Priti Shah and Patricia A. Carpenter. 1995. Conceptual limitations in comprehending line graphs. *Journal of Experimental Psychology: General* 124, 1, 43–61. DOI: <http://dx.doi.org/10.1037/0096-3445.124.1.43>
- Simeon J. Simoff, Michael H. Böhlen, and Arturas Mazeika (Eds.). 2008. In *Visual Data Mining: Theory, Techniques, and Tools for Visual Analytics*. Lecture Notes in Computer Science, Vol. 4404. Springer.
- Rasmus Stenholt and Claus B. Madsen. 2011. Shaping 3-D boxes: A full 9-degree-of-freedom docking experiment. In *Proceedings of the IEEE Virtual Reality Conference (VR'11)*. 103–110. DOI: <http://dx.doi.org/10.1109/VR.2011.5759445>
- Kevin Suffern. 2007. *Ray Tracing from the Ground Up*. A. K. Peters, Ltd., Natick, MA.
- Anne Treisman. 1982. Perceptual grouping and attention in visual search for features and objects. *Journal of Experimental Psychology: Human Perception and Performance* 8, 194–214.
- Edward R. Tufte. 2001. *The Visual Display of Quantitative Information* (2nd ed.). Graphics Press. <http://www.worldcat.org/isbn/0961392142>.
- Amy Ulinski, Catherine Zambaka, Zachary Wartell, Paula Goolkasian, and Larry F. Hodges. 2007. Two handed selection techniques for volumetric data. In *Proceedings of the IEEE Symposium on User Interfaces (3DUI'07)*. DOI: <http://dx.doi.org/10.1109/3DUI.2007.340782>
- Amy Catherine Ulinski. 2008. *Taxonomy and Experimental Evaluation of Two-Handed Selection Techniques for Volumetric Data*. Ph.D. Dissertation. University of North Carolina at Charlotte, Charlotte, NC.
- John Viega, Matthew J. Conway, George Williams, and Randy Pausch. 1996. 3D magic lenses. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology (UIST'96)*. ACM, New York, NY, 51–58. DOI: <http://dx.doi.org/10.1145/237091.237098>
- ACM Transactions on Applied Perception, Vol. 11, No. 4, Article 19, Publication date: December 2014.

- Nan Wang, Alexis Paljic, and Philippe Fuchs. 2010. A study of Perception of volumetric rendering for immersive scientific visualization. In *Proceedings of the 20th International Conference on Artificial Reality and Telexistence (ICAT'10)*. 145–152.
- Colin Ware. 2004. *Information Visualization: Perception for Design* (2nd ed.). Morgan Kaufmann.
- Lee A. Westover. 1991. *SPLATTING: A Parallel, Feed-Forward Volume Rendering Algorithm*. Technical Report. Chapel Hill, NC.
- Shumin Zhai. 1995. *Human Performance in Six Degree of Freedom Input Control*. Ph.D. Dissertation. University of Toronto, Toronto, Canada.
- Haitao Zhang and Airee Kaufman. 2006. Interactive point-based isosurface exploration and high-quality rendering. *IEEE Transactions on Visualization and Computer Graphics* 12, 5, 1267–1274. DOI: <http://dx.doi.org/10.1109/TVCG.2006.153>
- Elena V. Zudilova and Peter M. A. Slood. 2003. Virtual reality and desktop as a combined interaction-visualisation medium for a problem-solving environment. In *Proceedings of the 2003 International Conference on Computational Science: Part III (ICCS'03)*. Springer-Verlag, Berlin, Heidelberg, 1025–1034. <http://dl.acm.org/citation.cfm?id=1762418.1762531>.
- Matthias Zwicker, Hanspeter Pfister, Jeroen van Baar, and Markus Gross. 2001. Surface splatting. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'01)*. ACM, New York, NY, 371–378. DOI: <http://dx.doi.org/10.1145/383259.383300>

Received August 2013; revised June 2014; accepted August 2014