Structure Perception in 3D Point Clouds

Preprint

Kenny Gruchalla,[1] Sunand Raghupathi,[2] and
Nicholas Brunhart-Lupo[1]

1 National Renewable Energy Laboratory
2 Columbia University

Presented at the ACM Symposium on Applied Perception (ACM SAP)
September 16-17, 2021

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-80547
September 2021

Structure Perception in 3D Point Clouds

Preprint

Kenny Gruchalla,[1] Sunand Raghupathi,[2] and
Nicholas Brunhart-Lupo[1]

1 National Renewable Energy Laboratory
2 Columbia University

Suggested Citation
Gruchalla, Kenny, Sunand Raghupathi, and Nicholas Brunhart-Lupo. 2021. *Structure Perception in 3D Point Clouds: Preprint*. Golden, CO: National Renewable Energy Laboratory. NREL/CP-2C00-80547. https://www.nrel.gov/docs/fy21osti/80547.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Conference Paper
NREL/CP-2C00-80547
September 2021

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program. This work was supported in part by the Laboratory Directed Research and Development (LDRD) Program at NREL. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Structure Perception in 3D Point Clouds

Kenny Gruchalla
kenny.gruchalla@nrel.gov
National Renewable Energy Laboratory
USA

Sunand Raghupathi
sunand.r@columbia.edu
Columbia University
USA

Nicholas Brunhart-Lupo
nicholas.brunhart-lupo@nrel.gov
National Renewable Energy Laboratory
USA

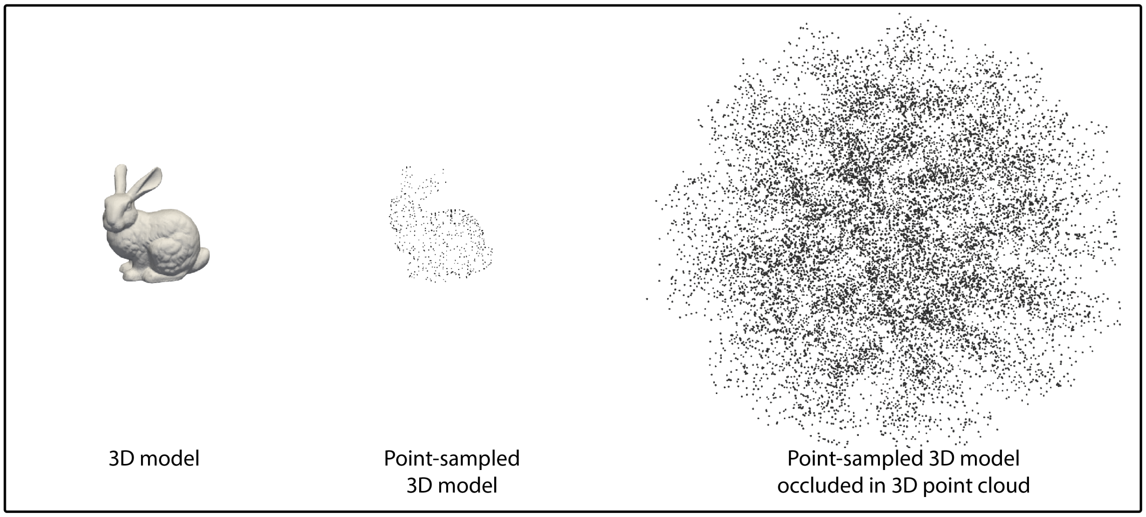

ABSTRACT

Understanding human perception is critical to the design of effective visualizations. The relative benefits of using 2D versus 3D techniques for data visualization is a complex decision space, with varying levels of uncertainty and disagreement in both the literature and in practice. This study aims to add easily reproducible, empirical evidence on the role of depth cues in perceiving structures or patterns in 3D point clouds. We describe a method to synthesize a 3D point cloud that contains a 3D structure, where 2D projections of the data strongly resemble a Gaussian distribution. We performed a within-subjects structure identification study with 128 participants that compared scatterplot matrices (canonical 2D projections) and 3D scatterplots under three types of motion: rotation, xy-translation, and z-translation. We found that users could consistently identify three separate hidden structures under rotation, while those structures remained hidden in the scatterplot matrices and under translation. This work contributes a set of 3D point clouds that provide definitive examples of 3D patterns perceptible in 3D scatterplots under rotation but imperceptible in 2D scatterplots.

CCS CONCEPTS

• **Human-centered computing → Empirical studies in HCI**; **Empirical studies in visualization**.

KEYWORDS

Visual perception, Scatterplots, Human factors, Data visualization, Data analysis, Encoding, Image generation

ACM Reference Format:
Kenny Gruchalla, Sunand Raghupathi, and Nicholas Brunhart-Lupo. 2021. Structure Perception in 3D Point Clouds. In *ACM Symposium on Applied Perception 2021 (SAP '21), September 16–17, 2021, Virtual Event, France.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3474451.3476237

1 INTRODUCTION

In practice, data analysts use both 2D and 3D visualization techniques to examine point clouds of abstract data. While 2D visualization is pervasive, 3D visualization is becoming more commonplace with the increasing adoption of tools like ParaView [Ahrens et al. 2005] and immersive virtual environments. When examining 3D point cloud data, the literature offers little clarity on the approach to use: multiple 2D projections versus a 3D projection. We examine the question, can structures exist that are only perceptible in 3D, or are multiple 2D projections always sufficient?

Scatterplots are one of the most common visualization techniques for data analysis [Friendly and Denis 2005; Tufte 1986]. A 2D scatterplot encodes two quantitative variables as points in a two-dimensional graph, mapping one variable to the x-axis and a second variable to the y-axis. Three-dimensional data is often mapped onto three separate 2D scatterplots: an x-y scatterplot, an x-z scatterplot, and a y-z scatterplot. Alternatively, 3D data can be mapped directly into 3D space and visualized with depth cues. There is very little empirical evidence to suggest whether 2D or 3D scatterplots — and in the case of 3D, which depth cues — better support data analysis. As such, there is generally no consensus in the visualization literature as to when or if 3D scatterplots should be used; however, it is generally accepted that the use of 3D visualization for abstract data can be problematic and requires careful justification [Munzner 2014].

The human vision system evolved to view a three-dimensional world, using various depth cues to interpret 3D structure, including structure-from-motion, stereopsis, vergence, perspective, occlusion, shading, and texture gradient [Howard and Rogers 2012a,b]. Visualization environments that support motion and stereo cues are becoming more accessible with the commoditization of virtual reality. Additionally, the field of immersive analytics [Skarbez et al. 2019] is growing. We have both empirical [Gruchalla 2004] and strong anecdotal evidence [Gruchalla and Brunhart-Lupo 2019] of improved data analysis in real-world settings in immersive environments. However, it is not clear if using 3D scatterplots is justified, particularly with some studies finding a benefit to 3D scatterplots [Arns et al. 1999; Kraus et al. 2019] and others studies showing 2D scatterplots and 2D scatterplot matrices should be preferred [Filho et al. 2017; Sedlmair et al. 2013]. We sought a simple, definitive, and reproducible example that would demonstrate a difference in the perception of a feature in 2D versus 3D scatterplots.

Hypothesis: Structures or patterns can exist in 3D point clouds that are imperceptible in 2D projections but are readily perceptible in 3D scatterplots using some combination of depth cues.

There is a long history of synthesizing data sets to demonstrate the importance of data visualization. In 1973 F.J. Anscombe developed *Anscombe's Quartet* [Anscombe 1973] to demonstrate the value of visualizing data compared to only using summary statistics. The Quartet is a set of four datasets with identical summary statistics (i.e., mean, standard deviation, and correlation), suggesting the datasets are markedly similar. However, visualizing the dataset with 2D scatterplots reveals they are markedly different. Matejka and Fitzmaurice have expanded on this notion, describing an optimization method to develop the *Datasaurus Dozen* [Matejka and Fitzmaurice 2017]: twelve visually distinct data sets with equivalent (to two decimal places) summary statistics to a data set with a 2D scatterplot that reveals the outline of a dinosaur.

Inspired by Anscombe's Quartet and the Datasaurus Dozen, we describe a method to develop 3D point clouds with structures that are visible in a 3D scatterplot, but are occluded in any 2D projection. We have developed three data sets by point sampling popular 3D models (the Stanford Bunny [Turk and Levoy 1994], the Utah Teapot [Blinn and Newell 1976], and the Viewpoint Animation Engineering Cow [Schroeder et al. 2006] packaged with Open Scene Graph). We then occlude those models with amorphous clusters of points (see Fig. 1). We developed a method to scale the density of these clusters to resemble some arbitrary probability distribution. For our three models, we chose the Gaussian distribution. Since searching for a rabbit is one of the objectives of these data, we have coined these the *Caerbannog Point Clouds* downloadable from https://data.nrel.gov/submissions/153.

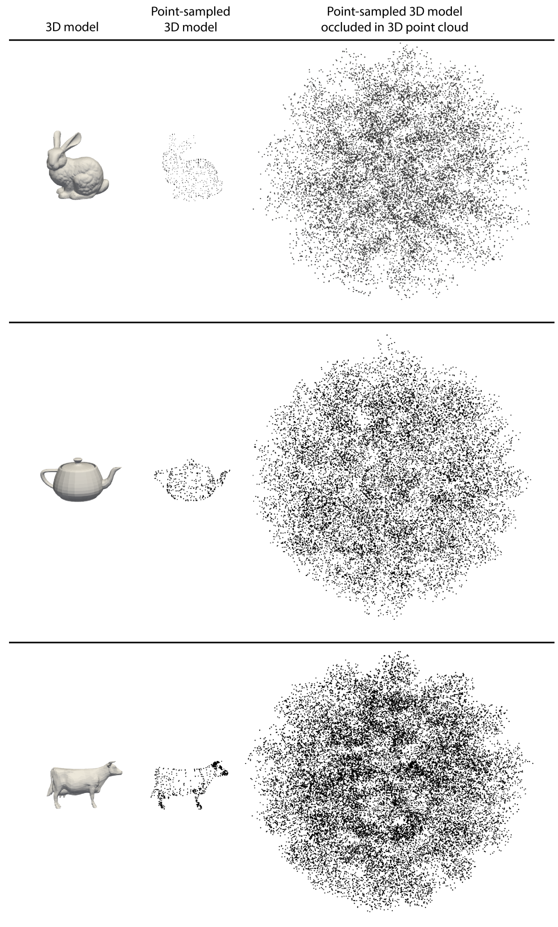

Figure 1: Caerbannog Point Clouds provide point-sampled 3D models occluded in amorphous clouds of points. In a user study, users were significantly better at identifying the three models when visualized as 3D scatterplots under rotation than in any other condition. Stanford bunny: 81% of users identified the bunny in the 3D point cloud under rotation compared to 6% users identifying the bunny in 2D projections of the point cloud, $\chi^2(1, 128)$=143.5, p<2.2e-16. Utah Teapot: 72% of users identified the teapot in the 3D point cloud under rotation compared to 8% users identifying the teapot in 2D projections of the point cloud, $\chi^2(1, 128)$=106.93, p < 2.2e-16. OSG Cow: 46% of users identified the cow in the 3D point cloud under rotation compared to 2% users identifying the cow in 2D projections of the point cloud, $\chi^2(1, 128$=64.383, p=1.024e-15.

We performed an Amazon Mechanical Turk user study with 128 participants, asking participants to identify occluded structures in point clouds. Users were presented with one of four models (the Stanford Bunny, the Utah Teapot, the OSG Cow, and a noise control with no hidden structure) in one of four conditions 3D scatterplot

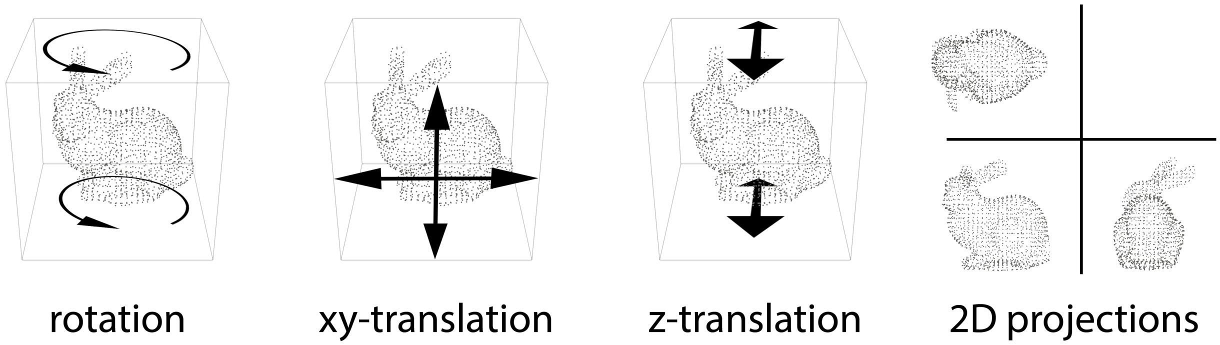

Figure 2: Four conditions studied shown with an unoccluded Stanford bunny: rotation, xy-translation, z-translation, and static 2D projections.

under motion (rotation, xy-translation, and z-translation) and 2D scatterplots of the canonical projections in a scatterplot matrix (see Figure 2). We arranged the models and conditions in a Graeco-Latin square, intermixed with attention tests to filter out careless and inattentive users. We found that users were significantly better at identifying the three models when visualized with a 3D scatterplot under rotation than in any other condition. Users generally could not identify the models from the 2D scatterplots or when visualized as a 3D scatterplot under translation.

2 RELATED WORK

If there is a benefit of visualizing 3D data with a 3D scatterplot, it would be a function of one or more depth cues that differentiates the 3D data from its 2D projection. There are many possible depth cues. Monocular depth cues include structure-from-motion, relative size, oculomotor accommodation, curvilinear perspective, texture gradient, defocus blur, lighting, and shading [Howard and Rogers 2012b]. Biocular depth cues include stereopsis, convergence, and shadow stereopsis [Howard and Rogers 2012a]. In this work, we only consider structure-from-motion to impart depth to our 3D scatterplots.

2.1 Structure-from-motion

Perceiving depth from motion is one of the strongest depth cues, and scientists and philosophers have been considering the importance of this depth cue for millennia [Todd 2004]. One form of structure-from-motion is the motion parallax, where an observer can judge the depth of stationary objects by the movement of the observer [Gibson et al. 1959]. Objects closer to the observer move faster across the visual field than objects farther away. Motion parallax is one of the primary depth cues provided by virtual reality displays [Cruz-Neira et al. 1993].

In this work, we consider a stationary observer and moving objects. Wallach and O'Connell [1953] performed one of the earliest published experiments on the visual perception of structure-from-motion of this form. Observers were able to identify rotating objects from the shadows cast from those objects. They described this as the *kinetic depth effect*. The original kinetic depth experiments were limited to solid and wire-frame objects. Later work demonstrated that users could perceive 3D structure from unconnected rotating points [Braunstein 1962; Green Jr. 1961; Ullman 1979]. In addition to rotation, early work also showed that users could perceive depth from translation [Braunstein 1966, 1976].

Since the original kinetic depth experiments, the empirical findings suggest users are reliably able to judge the topological, ordinal, and affine properties of objects under rigid motion. Furthermore, empirical results have shown that users can infer some depth information from an arbitrary configuration of points with as few as two motion-sequence frames [Todd 1995].

Despite the many decades of cognitive science research, there is limited empirical evidence on the relative benefits of depth cues in understanding 3D data visualizations [Ware 2012]. The vision and cognitive science literature tells us we can perceive structure from a cloud of 3D points under motion, but it doesn't tell us if we *should*, particularly if we have good 2D alternatives.

2.2 2D vs 3D

After decades of empirical studies comparing 2D and 3D visualizations, the results are mixed. St. John et al. [2001] reviewed 16 studies that compared 2D and 3D visualizations and suggested that 3D displays can improve spatial understanding but may inhibit judging relative positions and distances. And they confirmed these findings with an experiment using simple block shapes. A more recent review of 162 publications describing 184 experiments by McIntire et al. [2014] focused on 2D visualization versus 3D visualization with stereoscopic displays. Here too, the results were mixed. In 60% of the studies, 3D showed a definitive benefit, while the remaining studies found a benefit to 2D, mixed results, or inconclusive results. The benefit of 3D varied by tasks. In the judgments of positions or distances, 57% of studies found a clear benefit for 3D. In navigation tasks, 42%. In tasks related to finding, identifying, and classifying objects, 65%. And 52% of studies with spatial understanding tasks found a benefit for stereoscopic 3D visualization.

The empirical evidence suggests that a 3D visualization is not always better, and the applicability of using 3D is highly dependent on the nature of the data, the tasks, and the combination of depth cues employed. Therefore, we consider what is specifically known about viewing 3D point cloud data.

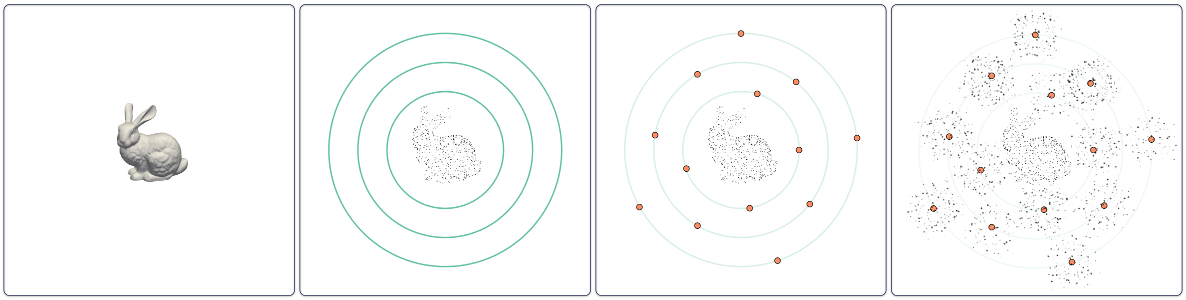

(a) Initial object to obscure **(b) Concentric spheres added** **(c) Sampling on spheres** **(d) Point clouds on samples**

Figure 3: Detail of point cloud generation procedure. In *(a)*, we begin with the object to be obscured. Moving to *(b)*, a sampling is made of the points on the surface of the object. Three concentric spheres are positioned around the object (we use circles in this diagram to aid explanation). In *(c)*, a sampling is made on the surfaces of each sphere. These are seed points. These points are used as shown in *(d)*, where a point cloud is centered and oriented at each of these points. As described in Section 3, the models' point density is modulated so that any arbitrary 2D projection of the complete point cloud will resemble a 2D Gaussian distribution, and the 3D density of the cloud resembles a 3D Gaussian distribution.

2.3 Scatterplots

3D scatterplots are widely used [Brunhart-Lupo et al. 2020; Bugbee et al. 2019; Donoho et al. 1988; Kosara et al. 2004; Piringer et al. 2004; Sanftmann and Weiskopf 2012; Zeckzer et al. 2016], but there is comparatively little research that empirically investigates the benefit of 3D visualization for 3D point clouds. And once again, the 2D versus 3D scatterplot research provides mixed results.

Two notable studies have shown 2D scatterplots outperform 3D scatterplots. Sedlmair et al. [2013] performed a data study to evaluate the analysis of data produced from dimensionality reduction with 2D scatterplots, 2D scatterplot matrices, and interactive 3D scatterplots. Two trained coders evaluated 816 scatterplots and concluded that 2D scatterplots are often sufficient, while 3D scatterplots rarely helped and occasionally hurt. Wagner Filho et al. [2017] compared 2D scatterplots with screen-based and VR-based 3D approaches. Their tasks included finding nearest neighbors and classes, identifying classes and outliers, and comparing classes. In this study, users were faster using the 2D scatterplot and reported 2D to be more intuitive for the given tasks.

Conversely, there are studies with empirical evidence of 3D scatterplots outperforming 2D scatterplots. Kraus et al. [2019] performed a user study with 18 participants in a cluster identification task, comparing a scatterplot matrix with 3D scatterplots in three different visualization environments. The 3D scatterplots outperformed the 2D scatterplot matrix in task time and correctness. Arns et al. [1999] compared statistical data analysis between XGobi on a desktop to 3D scatterplot in an immersive environment. They evaluated the identification and brushing of data clusters. Users were able to identify clusters in the immersive environment twice as well as on the desktop. Raja et al. [2004] compared various 3D scatterplot tasks (e.g., trend determination, cluster identification, outlier identification) between immersive and non-immersive environments. The additional depth cues afforded by the immersive visualization seem to improve the task; however, the authors reached no definitive statistical conclusions due to a small subject population.

We contribute additional evidence in this area of investigation through an empirical user study, evaluating different structure-from-motion depth cues of 3D scatter plots against 2D projections with a clear result. We also contribute the point clouds used for the study as a publicly available dataset.

3 DATA SYNTHESIS

We have developed a data synthesis method[1] to generate 3D point clouds from polygonal models, such that the shape of the model will be occluded in any 2D projection. We transform the polygon model into a point model by randomly sampling the polygon vertices. We then choose random points around the model location as seed locations for amorphous point clouds whose density we modulate such that 2D projections of the full set of points strongly resemble a Gaussian distribution.

We chose to obscure our model of interest by surrounding it with a second model of an amorphous cloud. We wanted to surround the points of interest with a point cloud of similar density and texture, which ruled out random 3D noise. Therefore, we used another 3D polygonal mesh, an amorphous cloud, as an occluding model sampling its vertices in the same method we sampled the model of interest. To seed our occluding model locations, we placed the point-sampled model-of-interest (e.g., Stanford Bunny points) in the center of three concentric spheres. Next, we sampled points randomly across the surfaces of those spheres, using those points as the seed locations of the occluding model. With the model-of-interest and the occluding models arranged, we assign the point densities of those objects (see Figure 3).

We modulate the models' point densities so that any arbitrary 2D projection of the dataset will resemble a 2D Gaussian distribution, and the 3D density of the dataset resembles a 3D Gaussian distribution. First, we choose a basis density, which we use to scale the subsequent densities. In our case, we use the density of the model-of-interest as the basis density. Then, we choose some probability

[1] Source code is available at https://www.github.com/kgruchal/bring-me-a-shrubbery

Table 1: Table of study questions. The study was prefaced with four pretest questions, followed by questions, counter-balancing models and conditions with periodic attention tests.

Question	Grouping	Model	Condition	Occlusion	Image
1	pretest	Dragon	rotation	None	https://i.imgur.com/8A5Nzb7.gif
2	pretest	Bunny	rotation	Minimal	https://i.imgur.com/HFseNN2.gif
3	pretest	Teapot	2D projections	Minimal	https://i.imgur.com/w2wKQWj.png
4	pretest	Cow	xy-translation	Minimal	https://i.imgur.com/kCmg9jR.gif
5	row 1	Noise	rotation	Medium	https://i.imgur.com/vnuK7ck.gif
6	row 1	Teapot	xy-translation	Medium	https://i.imgur.com/pi0eqR9.gif
7	row 1	Cow	z-translation	Heavy	https://i.imgur.com/lR3NTgY.gif
8	row 1	Bunny	2D projections	Light	https://i.imgur.com/BqXYTZ3.png
9	attention	Cow	rotation	Minimal	https://i.imgur.com/2jQJ4dy.gif
10	row 2	Teapot	z-translation	Medium	https://i.imgur.com/xgH87sw.gif
11	row 2	Noise	2D projections	Medium	https://i.imgur.com/Cb5CUFH.png
12	row 2	Bunny	rotation	Light	https://i.imgur.com/zLjOq6y.gif
13	row 2	Cow	xy-translation	Heavy	https://i.imgur.com/ndkm89G.gif
14	attention	Teapot	rotation	Medium	https://i.imgur.com/nEiGHjH.gif
15	row 3	Cow	2D projections	Heavy	https://i.imgur.com/OYfFRB2.png
16	row 3	Bunny	z-translation	Light	https://i.imgur.com/pcJone0.gif
17	row 3	Noise	xy-translation	Medium	https://i.imgur.com/24CdVOH.gif
18	row 3	Teapot	rotation	Medium	https://i.imgur.com/8z2c22K.gif
19	attention	Bunny	z-translation	Minimal	https://i.imgur.com/J6lv06O.gif
20	row 4	Bunny	xy-translation	Light	https://i.imgur.com/NgRp7IS.gif
21	row 4	Cow	rotation	Heavy	https://i.imgur.com/rxErFcw.gif
22	row 4	Teapot	2D projections	Medium	https://i.imgur.com/SO9qmcm.png
23	row 4	Noise	z-translation	Medium	https://i.imgur.com/ukkdMcb.gif
24	attention	Teapot	2D projections	Minimal	https://i.imgur.com/w2wKQWj.png

distribution function (in our case, a 3D Gaussian). For each of the sampled points, we calculate the value of the function at that point and scale the seed density by this value. This scaling defines the density of the object placed at that point. As a result, the 3D density plot of the dataset resembles a 3D Gaussian, and any arbitrary 2D projections resemble a 2D Gaussian.

4 STUDY

To empirically evaluate the perceptibility of the models in our point clouds, we conducted a user study[2] on Amazon Mechanical Turk. The study used a within-subjects design, with 24 questions that presented the users with static images or an animation of a point cloud and asked, *"What shape, if any, do you see hidden in the points?"* with multiple choice answer options of ***Dragon, Bunny, Cow, Teapot,*** *or* ***Nothing***. The 24 questions included pretest, study, and attention questions (see Table 1).

4.1 Pretest Questions

We used the first four questions as training and pretest with obvious answers to identify and filter out any users that might not understand the instructions. The first training-pretest question presented a rotating point-sampled version of the Stanford Dragon [Curless and Levoy 1996] unoccluded, asking, *"In the following set of questions we will ask you to identify shapes that are made of points. What shape do you see?"*. Followed by a rotating cloud of points partially occluding the bunny model with the following text, *"The shape, if it exists, will be hidden in a cloud of points. What shape, if any, do you see hidden in the points?"*. Pretest question 3: three on-axis 2D projections of the teapot model, *"Here we show the front, top, and side views of a cloud of points. What shape, if any, do you see hidden in the points?"*

[2]The study was reviewed and approved by our Institutional Review Board, *IRB00000067*

4.2 Study Questions

In the primary questions, users were presented with one of four models (the Stanford Bunny, the Utah Teapot, the OSG Cow, or a noise control with no identifiable structure) in one of four conditions: 3D scatterplot under motion (xy-translation, z-translation, rotation) or 2D scatterplots of the canonical projections in a scatterplot matrix (see Figure 2). The density of the point clouds occluding the three models varied with the three models. The bunny was the least occluded, followed by the teapot, and finally, the cow in the highest density. We counter-balanced the models and conditions in a Graeco-Latin square. We presented the animated conditions as GIFs, which we captured using MayaVi [Ramachandran and Varoquaux 2011] using the default perspective projection. We presented the static 2D projections in a group of three principal projections: front, top, side.

4.3 Attention Questions

We also intermixed attention tests with the study questions as an attempted mechanism to filter out careless and inattentive users.

The attention questions mimicked the study questions but at a much lower level of occlusion. A level that we believed participants would be able to detect the model structure easily.

5 RESULTS

We recruited participants through TurkPrime [Litman et al. 2017]. All participants had completed over 100 Human Intelligence Tasks (HITs) with a HIT approval of at least 90%. We recruited a total of 143 participants; 128 users completed all 24 questions in an average time of 25 minutes. We discarded the partial results of the 15 participants that did not complete all 24 questions from the analysis.

We began with an analysis of the quality of the responses from the remaining 128 participants, based on the *attention* questions. Of the *attention* questions, question 14 (see Table 1) was an outlier with less than 69% of the participants correctly identifying the teapot. On analysis, we realized question 14 was incorrectly coded, which we had intended to be an attention test with a minimal amount of occlusion; however, it was miscoded with an occlusion equivalent to the study questions. Therefore, we disregarded question 14 as an attention test. Of the remaining seven attention and pretests, 91 participants (71%) answered all seven of these tests correctly, while 118 participants (92%) answered at least six of the questions correctly (see Figure 4). The number of correct responses for the individual attention and pretest is inconsistent (see Figure 5), ranging from 99% correct to 78% correct. Measuring the attentiveness of workers on Mechanical Turk is a challenging problem with no accepted standard to what classifies a worker as *attentive* or *inattentive* [Hauser et al. 2018]. Furthermore, the inconsistency across our attention tests may suggest that the difficulty of these tests were unsuccessfully chosen and may be measuring something other than attentiveness. Therefore, in the following analysis, we consider both the results for the 91 *attentive* participants and the results for the full complement of 128 participants. While the attentive participants were generally more successful, we see the differences were slight and always within the confidence interval (see Figures 6 and 8).

Table 2: Table of study results, comparing the rotation condition to the other conditions with a 2-sample test for equality of proportions.

Model	Condition	No. correct	Significance $\chi^2(1, 128)$
Bunny	Rotation	104 (81.25%)	
Bunny	xy-translation	26 (20.31%)	92.663, p<2.2e-16
Bunny	z-translation	6 (4.69%)	149.98, p<2.2e-16
Bunny	2D Projections	8 (6.25%)	143.25, p<2.2e-16
Teapot	Rotation	92 (71.88%)	
Teapot	xy-translation	2 (1.56%)	133.16, p<2.2e-16
Teapot	z-translation	1 (1.56%)	133.16, p<2.2e-16
Teapot	2D Projections	10 (7.81%)	106.93, p<2.2e-16
Cow	Rotation	59 (46.09%)	
Cow	xy-translation	7 (5.47%)	53.099, p=3.17e-13
Cow	z-translation	2 (1.56%)	67.492, p<2.2e-16
Cow	2D Projections	3 (2.34%)	64.383, p=1.024e-15

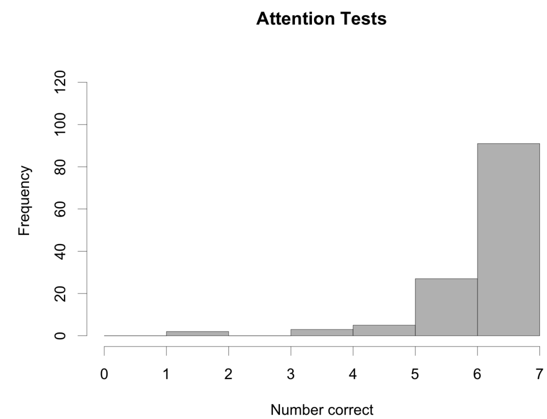

Figure 4: Histogram of the number of attention and pretest questions the participants answered correctly. Ninety-one participants (71%) answered all seven correctly. 118 (92%) answered six or more correctly.

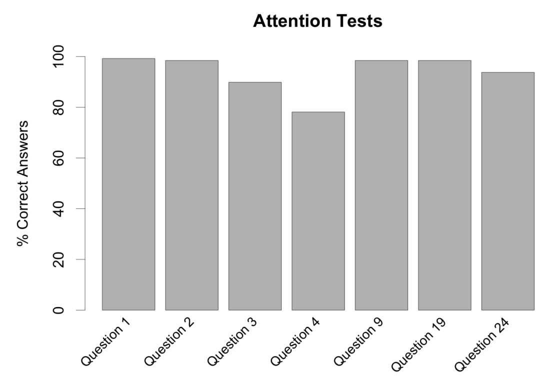

Figure 5: Bar graph of the percent of correct answers for the attention and pretest questions.

Across the three models, users were significantly better at identifying the hidden object when visualized as a 3D scatterplot under rotation than in any other condition (see Figure 6). We compare rotation against the other conditions with a 2-sample test for equality of proportions (see Table 2). The Standford bunny was the least occluded of the three models, and had the highest number of correct detections. Considering the population of 128 participants, 104 participants identified the bunny under rotation compared to 26 correct identifications under xy-translation, $\chi^2(1, 128)$=92.663, $p < 2.2\text{e-}16$, only 6 participants correctly identified the bunny under z-translation, $\chi^2(1, 128)$=149.98, p<2.2e-16, and 8 participants were correct when using the 2D projections, $\chi^2(1, 128)$=143.25, p<2.2e-16.

The teapot had more occlusion and was detected by 92 participants, which is significant when compared 2 correct participants in both xy-translation and z-translation, and 10 correct detections using 2D projections. Finally, the cow was the most heavily occluded and was correctly detected by 59 participants under rotation. With

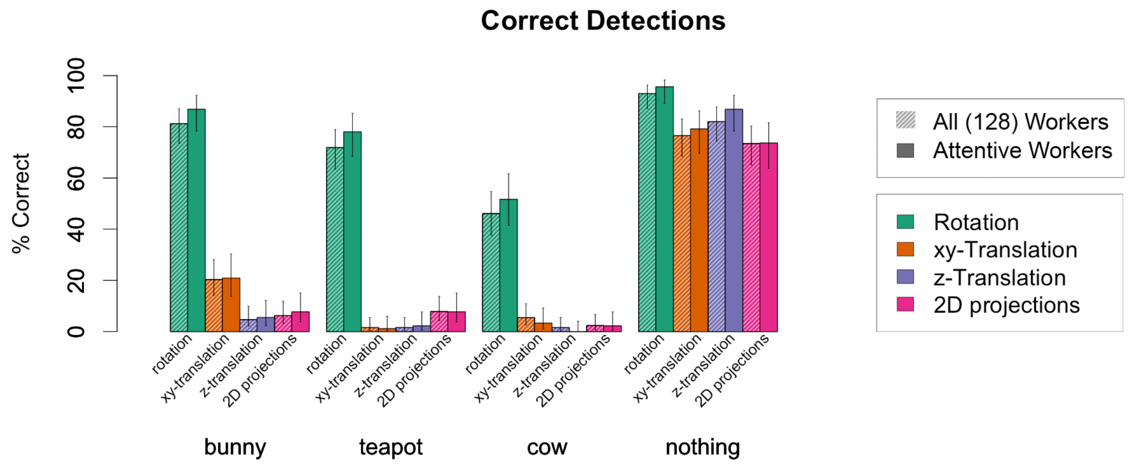

Figure 6: The graph shows the percentage of correct responses for the sixteen study questions. The error bars show the 95% binomial proportion confidence intervals. Hashed bars show the percentage of correct responses for all 128 participants, and solid bars show correct responses for the "attentive" participants.

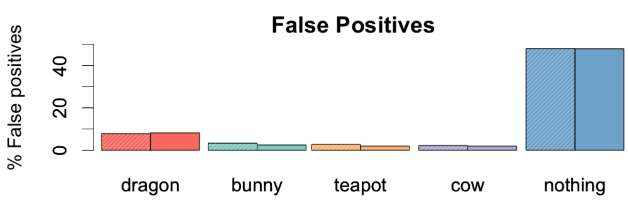

Figure 7: Percentage of false positive answers. Participants incorrectly reported noise when there was hidden model for 48% of the questions.

only 7 correct under xy-translation, 2 correct under z-translation, and 3 correct using 2D projections. See Table 2 for details.

While 3D rotation was unambiguously better than the other three conditions, there is no clear difference between the other three conditions, with exception to the xy-translation for the bunny. Translation in the xy-plane did significantly better for the bunny than z-translation, $\chi^2(1, 128)$=12.893, p=0.00033, and 2D projections, $\chi^2(1, 128)$=9.8018, p=0.001743. For the teapot and cow models, there are no clear differences between the translation and 2D projection conditions.

We might construe the number of correct detections for noise control as participants identifying a lack of a model correctly; however, *nothing* was generally the default choice, as can be seen by the number of noise false positives in Figure 7 and in the distribution of answers in Figure 8.

6 DISCUSSION

This work contributes definitive evidence that there can be patterns that are identifiable in 3D scatterplots under rotation but not identifiable in 2D scatterplots; however, more work is needed to understand how that result generalizes. We derived the patterns of interest from popular 3D models, which provided shapes recognizable for our participant population—Mechanical Turk workers we assumed would not be data analysis experts. These shapes may not be representative of features or clusters in real-world point clouds; obviously, it would be unlikely to find a bunny-shaped cluster in a dimension-reduction data set. However, the use of these objects is not wholly unreasonable, as complex shapes can be observed in real-world point-cloud data. For example, Bugbee et al. [2019] describe a zoomorphic shape having "horns" in a 3D point cloud generated from t-SNE—a dimension-reduction technique.

In our experiment, we occluded each structure by varying degrees of point density, which corresponds to the number of correct detections under rotation 81%, 72%, and 46%, respectively. These point clouds provide three separate examples that support our hypothesis. However, without varying the density condition, we cannot discern if the differences between conditions are a function of the model (i.e., bunny shapes are easier to see than cow shapes) or a function of the point density. We suspect there are elements of both at play. Furthermore, based on the bunny results, the models might become perceptible under translation at lower occlusions. By systematically lowering the density, would we eventually see an inflection point when translation becomes a viable depth cue? At that point, would the models also be perceptible in the 2D projections? Understanding the differences between shapes and deepening the investigation of 2D projection versus 3D translation as a function of occlusion is future work.

Our axis of rotation was roughly parallel to the image, which is consistent with rotation in VR-based motion parallax (i.e., moving around an object). However, we know from prior work that the user's perceptions can be significantly affected by the axis of orientation [Todd 2004]. Additionally, we located all of our structures of interest central to the cloud of points near that rotational axis. What influence does moving that axis of rotation away from the structure of interest have? These questions will be explored in future work.

This work was in part motivated by our desire to understand when there might be value in visualizing 3D scatterplots immersively. We used our immersive environment during the development of our datasets. We were readily able to perceive the hidden

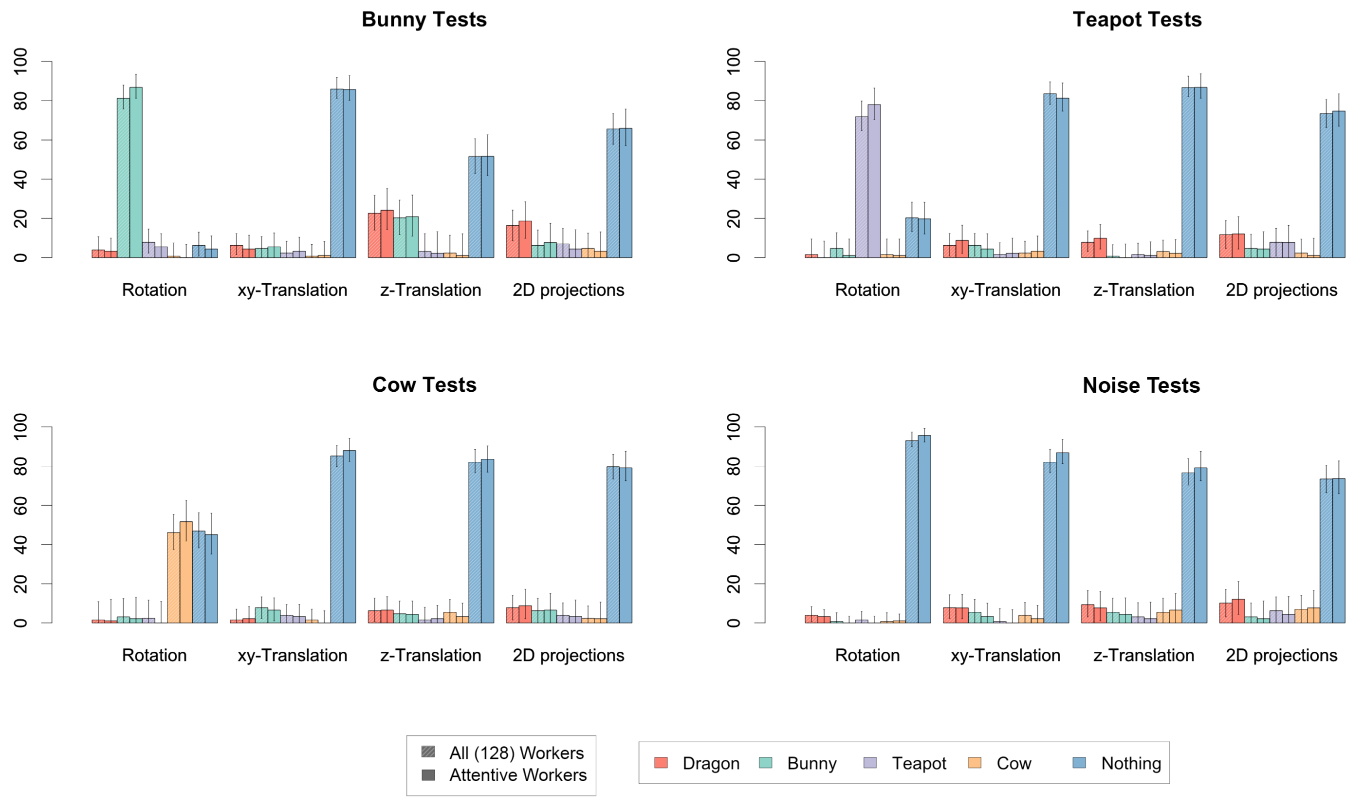

Figure 8: The graphs show the percentage of responses for each model. The error bars show the 95% multinomial proportions confidence intervals calculated by the Sison and Glaz [Sison and Glaz 1995] method. Hashed bars show all 128 participants' responses, and solid bars show the percentage of responses for the "attentive" participants.

structures when viewing these data immersively with or without stereopsis. In retrospect of the study results, this immersive perception is somewhat surprising. There is very little rotation in our head positions – tracker logs confirm that the head movement is dominated by translation. However, there was some small amount of rotation in the orientation, and we know from prior work that very small head movements can aid depth perception [Aytekin and Rucci 2012; de la Malla et al. 2016]. The perception could be the cognitive aid of being embodied, or maybe a minimal amount of rotation is sufficient to aid perception. A user study of immersive visualization of 3D scatterplots, controlling the amount and types of movement is future work.

7 CONCLUSION

We have synthesized point cloud data that definitively demonstrate that 3D visualization *can* reveal some structures in 3D point clouds under rotation that are not perceptible in 2D projections, supporting our hypothesis. We have shown three separate examples where 2D projections were insufficient to identify structures in 3D point clouds. Synthesizing our results and the mixed results from the literature, we conclude that it is critical to always examine your data in multiple ways. Others have advocated taking multiple views (both 2D and 3D) [Tory et al. 2006]. Particularly for larger complex data, no one perspective is likely sufficient, whether that perspective a purely quantitative statistical view or a view afforded by some visualization technique.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy under Contract No. DE-AC36-08GO28308 with Alliance for Sustainable Energy, LLC, the Manager and Operator of the National Renewable Energy Laboratory. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program. This work was supported in part by the Laboratory Directed Research and Development (LDRD) Program at NREL. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes. NREL is a national laboratory of the U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, operated by the Alliance for Sustainable Energy, LLC.

REFERENCES

James Ahrens, Berk Geveci, and Charles Law. 2005. ParaView: An End-User Tool for Large-Data Visualization. In *The Visualization Handbook*.

F. J. Anscombe. 1973. Graphs in Statistical Analysis. *The American Statistician* 27, 1 (Feb. 1973), 17–21. https://doi.org/10.1080/00031305.1973.10478966

Laura Arns, Dianne Cook, and Carolina Cruz-Neira. 1999. The benefits of statistical visualization in an immersive environment. In *Proceedings IEEE Virtual Reality*. 88–95. https://doi.org/10.1109/VR.1999.756938

Murat Aytekin and Michele Rucci. 2012. Motion parallax from microscopic head movements during visual fixation. *Vision Research* 70 (Oct. 2012), 7–17. https://doi.org/10.1016/j.visres.2012.07.017

James F. Blinn and Martin E. Newell. 1976. Texture and Reflection in Computer Generated Images. *Commun. ACM* 19, 10 (Oct. 1976), 542–547. https://doi.org/10.1145/360349.360353

Myron L. Braunstein. 1962. Depth perception in rotating dot patterns: Effects of numerosity and perspective. *Journal of Experimental Psychology* 64, 4 (1962), 415–420. https://doi.org/10.1037/h0048140

Myron L. Braunstein. 1966. Sensitivity of the observer to transformations of the visual field. *Journal of Experimental Psychology* 72, 5 (1966), 683–689. https://doi.org/10.1037/h0023735

Myron L. Braunstein. 1976. *Depth Perception Through Motion*. Academic Press.

Nicholas Brunhart-Lupo, Brian Bush, Kenny Gruchalla, Kristi Potter, and Steve Smith. 2020. Collaborative Exploration of Scientific Datasets Using Immersive And Statistical Visualization. In *Proceedings of the 2020 Improving Scientific Software Conference*, W. Hu, D. Del Vento, and S. Su (Eds.). 15–29.

Bruce Bugbee, Brian W. Bush, Kenny Gruchalla, Kristin Potter, Nicholas Brunhart-Lupo, and Venkat Krishnan. 2019. Enabling immersive engagement in energy system models with deep learning. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 12, 4 (2019), 325–337. https://doi.org/10.1002/sam.11419 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/sam.11419

Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. 1993. Surround-screen Projection-based Virtual Reality: The Design and Implementation of the CAVE. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM, New York, NY, USA, 135–142. https://doi.org/10.1145/166117.166134

Brian Curless and Marc Levoy. 1996. A Volumetric Method for Building Complex Models from Range Images. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. ACM, New York, NY, USA, 303–312. https://doi.org/10.1145/237170.237269

Cristina de la Malla, Stijn Buiteman, Wilmer Otters, Jeroen B. J. Smeets, and Eli Brenner. 2016. How various aspects of motion parallax influence distance judgments, even when we think we are standing still. *Journal of Vision* 16, 9 (July 2016), 8–8. https://doi.org/10.1167/16.9.8

Andrew W. Donoho, David L. Donoho, and Miriam Gasko. 1988. MacSpin: dynamic graphics on a desktop computer. *IEEE Computer Graphics and Applications* 8, 4 (July 1988), 51–58. https://doi.org/10.1109/38.7749

Jorge A. Wagner Filho, Marina Fortes Rey, Carla Maria Dal Sasso Freitas, and Luciana Porcher Nedel. 2017. Immersive Analytics of Dimensionally-Reduced Data Scatterplots. In *2nd Workshop on Immersive Analytics*.

Michael Friendly and Daniel Denis. 2005. The early origins and development of the scatterplot. *Journal of the History of the Behavioral Sciences* 41, 2 (2005), 103–130. https://doi.org/10.1002/jhbs.20078

Eleanor J. Gibson, James J. Gibson, Olin W. Smith, and Howard Flock. 1959. Motion parallax as a determinant of perceived depth. *Journal of Experimental Psychology* 58, 1 (1959), 40–51. https://doi.org/10.1037/h0043883

Bert F. Green Jr. 1961. Figure coherence in the kinetic depth effect. *Journal of Experimental Psychology* 62, 3 (1961), 272–282. https://doi.org/10.1037/h0045622

Kenny Gruchalla. 2004. Immersive well-path editing: investigating the added value of immersion. In *IEEE Virtual Reality 2004*. 157–164. https://doi.org/10.1109/VR.2004.1310069

Kenny Gruchalla and Nicholas Brunhart-Lupo. 2019. The Utility of Virtual Reality for Science and Engineering. In *VR Developer Gems*, William R. Sherman (Ed.). Taylor Francis, Chapter 21, 383–402. https://doi.org/10.1201/b21598-21

David Hauser, Gabriele Paolacci, and Jesse J. Chandler. 2018. *Common Concerns with MTurk as a Participant Pool: Evidence and Solutions*. preprint. PsyArXiv. https://doi.org/10.31234/osf.io/uq45c

Ian P. Howard and Brian J. Rogers. 2012a. *Perceiving in Depth, Volume 2: Stereoscopic Vision*. Oxford University Press, USA.

Ian P. Howard and Brian J. Rogers. 2012b. *Perceiving in Depth, Volume 3: Other Mechanisms of Depth Perception*. Oxford University Press, USA.

Mark St. John, Michael B. Cowen, Harvey S. Smallman, and Heather M. Oonk. 2001. The use of 2D and 3D displays for shape-understanding versus relative-position tasks. *Human Factors* 43, 1 (2001), 79–98. https://doi.org/10.1518/001872001775992534

Robert Kosara, Gerald N. Sahling, and Helwig Hauser. 2004. Linking Scientific and Information Visualization with Interactive 3D Scatterplots. In *WSCG*.

M. Kraus, N. Weiler, D. Oelke, J. Kehrer, D. A. Keim, and J. Fuchs. 2019. The Impact of Immersion on Cluster Identification Tasks. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1. https://doi.org/10.1109/TVCG.2019.2934395

Leib Litman, Jonathan Robinson, and Tzvi Abberbock. 2017. TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods* 49, 2 (01 Apr 2017), 433–442. https://doi.org/10.3758/s13428-016-0727-z

Justin Matejka and George Fitzmaurice. 2017. Same Stats, Different Graphs: Generating Datasets with Varied Appearance and Identical Statistics Through Simulated Annealing. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 1290–1294. https://doi.org/10.1145/3025453.3025912

John P. McIntire, Paul R. Havig, and Eric E. Geiselman. 2014. Stereoscopic 3D displays and human performance: A comprehensive review. *Displays* 35, 1 (Jan. 2014), 18–26. https://doi.org/10.1016/j.displa.2013.10.004

Tamara Munzner. 2014. *Visualization Analysis and Design* (1 edition ed.). A K Peters/CRC Press, Boca Raton.

Harald Piringer, Robert Kosara, and Helwig Hauser. 2004. Interactive focus+context visualization with linked 2D/3D scatterplots. In *Proceedings. Second International Conference on Coordinated and Multiple Views in Exploratory Visualization, 2004.* 49–60. https://doi.org/10.1109/CMV.2004.1319526

Dheva Raja, Doug A. Bowman, John Lucas, and Chris North. 2004. Exploring the Benefits of Immersion in Abstract Information Visualization. In *In proceedings of Immersive Projection Technology Workshop*.

Prabhu Ramachandran and Gaël Varoquaux. 2011. Mayavi: 3D Visualization of Scientific Data. *Computing in Science Engineering* 13, 2 (March 2011), 40–51. https://doi.org/10.1109/MCSE.2011.35

Harald Sanftmann and Daniel Weiskopf. 2012. 3D Scatterplot Navigation. *IEEE Transactions on Visualization and Computer Graphics* 18, 11 (Nov. 2012), 1969–1978. https://doi.org/10.1109/TVCG.2012.35

Will Schroeder, Ken Martin, and Bill Lorensen. 2006. *The Visualization Toolkit: An Object-oriented Approach to 3D Graphics*. Kitware.

Michael Sedlmair, Tamara Munzner, and Melanie Tory. 2013. Empirical Guidance on Scatterplot and Dimension Reduction Technique Choices. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (Dec. 2013), 2634–2643. https://doi.org/10.1109/TVCG.2013.153

Cristina P. Sison and Joseph Glaz. 1995. Simultaneous Confidence Intervals and Sample Size Determination for Multinomial Proportions. https://doi.org/10.1080/01621459.1995.10476521

Richard Skarbez, Nicholas F. Polys, J. Todd Ogle, Chris North, and Doug A. Bowman. 2019. Immersive Analytics: Theory and Research Agenda. *Frontiers in Robotics and AI* 6 (2019). https://doi.org/10.3389/frobt.2019.00082

James T. Todd. 1995. The visual perception of three-dimensional structure from motion. In *Perception of space and motion*. Academic Press, San Diego, CA, US, 201–226. https://doi.org/10.1016/B978-012240530-3/50008-0

James T. Todd. 2004. The visual perception of 3D shape. *Trends in Cognitive Sciences* 8, 3 (2004), 115–121. https://doi.org/10.1016/j.tics.2004.01.006

Melanie Tory, Arthur E. Kirkpatrick, M. Stella Atkins, and Torsten Moller. 2006. Visualization Task Performance with 2D, 3D, and Combination Displays. *IEEE Transactions on Visualization and Computer Graphics* 12, 1 (Jan. 2006), 2–13. https://doi.org/10.1109/TVCG.2006.17

Edward R. Tufte. 1986. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, USA.

Greg Turk and Marc Levoy. 1994. Zippered Polygon Meshes from Range Images. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. ACM, New York, NY, USA, 311–318. https://doi.org/10.1145/192161.192241

S. Ullman. 1979. The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 203, 1153 (Jan. 1979), 405–426. https://doi.org/10.1098/rspb.1979.0006

H. Wallach and D. N. O'Connell. 1953. The kinetic depth effect. *Journal of experimental psychology* 45, 4 (1953), 205–217. https://doi.org/10.1037/h0056880

Colin Ware. 2012. *Information Visualization: Perception for Design* (3 edition ed.). Morgan Kaufmann, Waltham, MA.

Dirk Zeckzer, Daniel Gerighausen, and Lydia Muller. 2016. Analyzing Histone Modifications in iPS Cells Using Tiled Binned 3D Scatter Plots. In *2016 Big Data Visual Analytics (BDVA)*. 1–8. https://doi.org/10.1109/BDVA.2016.7787042