

AUTOMATING USER-PREFERRED CAMERA PLACEMENT FOR
VOLUME RENDERED SCIENTIFIC VISUALIZATION

by

GINNI GALLAGHER

A THESIS

Presented to the Department of Computer Science
and the Robert D. Clark Honors College
in partial fulfillment of the requirements for the degree of
Bachelor of Science

May 2023

An Abstract of the Thesis of

Ginni Gallagher for the degree of Bachelor of Science
in the Department of Computer Science to be taken June 2023

Title: Automating User-Preferred Camera Placement for Volume Rendered Scientific
Visualization

Approved: Hank Childs, Ph.D.
Primary Thesis Advisor

Camera placement is essential in the world of scientific visualization. Different camera placements expose different information about the data. Viewpoint quality (VQ) metrics are one method of evaluating the quality of camera placement in visualizations. VQ metrics also have the potential to direct the automation of camera selection in scientific visualization, an important issue as the computational capacity to produce data is far outpacing the capacity to save data to storage. Previous research has used VQ metrics to successfully predict camera placements that match user preferences in images generated by a popular visualization technique, called isosurfacing. With this study, we extend the previous research and investigate the efficacy of VQ metrics in predicting user preferred camera placements for images created using another popular visualization method, known as volume rendering.

This study involves two main components: (1) gathering user preferences of camera placements and (2) a performance analysis of how accurately VQ metrics were able to predict user preferences. We found that the top performing VQ metric was able to correctly predict user preferences of volume rendered images up to 66% of the time. This result supports previous findings about the efficacy of VQ metrics in predicting user preferred images generated using

isosurfaces. Together, these findings provide further evidence that VQ metrics are a promising approach for guiding the automation of selecting camera placements for scientific visualization methods.

Acknowledgements

First of all, I would like to thank Professor Hank Childs for enthusiastically supporting me throughout this process. I also want to thank Nicole Marsaglia for providing me with all of the resources and materials I needed to carry on her work. It has been an honor to work with you both and to grow and learn under your guidance. I want to thank my best friend, Liz, for being there with me every step of the way these past four years in the Clark Honors College. It wouldn't have been the same without you. And last but not least, I want to thank my family for always being my biggest supporters. I will forever be grateful for the extra time I got to spend so close to home.

Table of Contents

Introduction	8
Camera Placement	8
Automation	8
Consideration of User Preference	10
Background	11
What Makes a Good Image?	11
How to Find Good Images?	12
Viewpoint Quality Metrics	13
Isosurfacing and Volume Rendering	15
My Study	18
Data Corpus for Comparing Viewpoints	18
Determining User Preferences	18
Viewpoint Analysis	24
Evaluating Results	25
Results	27
Viewpoint Quality Metric Findings	27
Efficacy of Top Performing Metric	28
Potential Influences on Results	32
Conclusion and Future Work	34
References	36

List of Figures

- Figure 1: Example of image rendered using different camera viewpoints [5]. This figure demonstrates how some images are better representations of a tea kettle than others. Certain camera viewpoints omit important features of the image, such as the handle or the spout. 11
- Figure 2: An example of isosurfacing. This image depicts the hurricane data set rendered using isovalues of 10 and 90 [18]. These parameters result in an image that contains a blue layer representing field data of value 10 on top of a red layer that represents field data of value 90. 16
- Figure 3: An example of the Mantle data set with a fully opaque transfer function (left) and the transfer function used in this study (right). 17
- Figure 4: The image collection for our corpus. A 3D visualization of each data set was generated using volume rendering and uniquely designed transfer functions. Then, using a Fibonacci Lattice pattern, images were taken of each data set from ten evenly spaced camera locations. 21
- Figure 5: A depiction of what a Fibonacci Lattice looks like [2]. 22
- Figure 6: An example question from the user survey. Participants selected the image they preferred to represent the simulation. If participants did not have a strong preference for either image, they may select “Neutral.” Participants answered questions for as long as they wanted, then selected “Done” to exit the survey. 23
- Figure 7: An example of a volume rendered image translated into isosurfaces to be used in the Ascent analysis. The image on the left was produced using volume rendering operations in VisIt. The image on the right was produced by Ascent, using the isovalues given as input. 25
- Figure 8: A depiction of the logic used in a single-metric oracle [2]. Two different cameras, C1 and C2, and the selected metric, M1, are given as input to the oracle. The oracle then compares the metric scores for each camera, makes a decision about which camera has a better viewpoint based on the metric provided, and outputs its decision. 26
- Figure 9: Mantle rendered from Camera #6 [19]. This camera was assigned as a VERY GOOD camera. 30
- Figure 10: Mantle rendered from Camera #9 [19]. This camera was assigned as a POOR camera. 30

List of Tables

Table 1: Correct predictions for each metric using Student data corpus.	27
Table 2: Correct predictions for each metric using Professional data corpus.	27
Table 3: This table details the performance analysis of the top performing metric, DDS Entropy, on the Professional survey data corpus. A cross section entry of this table describes the winning percentage, i.e., the number of times the DDS oracle correctly predicted user preferences, based on the categorization of each type of camera match up. For instance, the (POOR, GOOD) rating of 80% means that in pairings of one “poor” camera and one “good” camera, the oracle made correct predictions 80% of the time. The table cells are colored based on performance; camera pairings that were correctly predicted more than 70% of the time are colored green, pairings correctly predicted 60%-70% of the time are colored yellow, and pairings that were correctly predicted less than 60% of the time are colored orange.	31
Table 4: This table details the performance analysis of the top performing metric, DDS Entropy, on the Student survey data corpus. This table follows the same coloring schema as Table 3 above.	31
Table 5: DDS prediction rate by data set using Student data corpus.	33
Table 6: DDS prediction rate by data set using Professional data corpus.	33

Introduction

Camera Placement

Scientific visualization provides necessary tools to a wide range of scientific domains. Camera placement is a critical element of scientific visualization. Deciding where to place the camera in a 3D visualization can dramatically impact the visible information within an image. Different camera positions can expose or hide different elements of the visible data. Thus, it is important that camera positions are selected in a way that ensures necessary information is clear and comprehensible to the viewer. There are two main approaches for selecting a camera position in a visualization. Visualization software typically display scenes from a default camera position. Then, users modify the data by inputting their own camera positioning or using their mouse to drag data around the X, Y, and Z axes. Alternatively, and much less frequently, the camera position can be set using automation. One way to perform an automatic approach relies on the utilization of metrics or preset measurements that guide the selection of a camera position without the need for human interaction.

Automation

Many domain scientists rely on computational simulations to avoid the high costs associated with running real experiments. In many cases, large scientific simulations require a level of computational power that can only be obtained by running on supercomputers. Supercomputers have the power to quickly generate enormous amounts of data. The rate at which computers can produce data has started to outpace the rate at which data can be saved to memory. This problem is referred to as the I/O (input/output) gap [1].

Traditionally, simulations are run in a post-hoc setting. Post-hoc processing is when data is first saved to storage, then afterwards read in by visualization software. However, with the increasing I/O gap, this process is becoming less practical. In situ processing is an alternative method that allows data to be processed as it is generated, thus reducing the I/O costs of a visualization. In situ processing typically occurs with no human in the loop, meaning completely independent from human intervention [1]. As a result, automating the processes in scientific visualization is increasingly relevant. Specific to this work, the goal of automating camera selection is to find camera placements that produce useful images while allowing visualizations to run in an in situ setting.

Research on automated approaches of selecting camera positions aim to discover methods of effectively setting acceptable camera positions that generate good images. Some notable research studies have examined viewpoint quality (VQ) metrics as a method of setting camera positions. VQ metrics are designed to evaluate the quality of a visualization's camera placement based on metric-specific formulas. Viewpoints with higher VQ metric values are in theory "better" than camera views with lower metric values [2]. One research study in particular, Marsaglia et al. [2], investigated user preferences in correlation with new and existing VQ metrics. This novel method of evaluating user-preferred viewpoints guided research on automating camera placement selection towards a solution that analyzes what domain users consider a "good" viewpoint of the visualization. Marsaglia et al. [2] found that existing VQ metrics are not significantly effective in predicting what camera positions users prefer. However, Marsaglia et al. [2] contributed their own entropy-based metrics that were significantly better predictors of user preferences, making a correct prediction 68% of the time.

Consideration of User Preference

Research that considers how user preferences can be used to guide automated selection of camera placement has not been widely investigated. For instance, Marsaglia et al. [2] only considered images generated through isosurfacing. While isosurface representations of data are commonly used in the scientific visualization community, it is not the only technique for displaying data. Volume rendering is another popular method of scientific visualization. While few published research studies have investigated automating camera placements for volume rendered imagery, even fewer have considered how user preferences relate to camera positions in volume rendered images. This research aims to fill that gap by extending the work of Marsaglia et al. [2] to explore whether existing VQ metrics can effectively predict user preferred camera placements in volume rendered images.

Background

What Makes a Good Image?

In scientific visualization, “good” images depend on the type, context, and intended use of a data set. Typically, a user will input their desired camera placement that renders a preferred image. Human direction allows for special considerations, such as what the data captures and what situation the data represents. While it is agreed upon by visualization researchers that no single camera placement can consistently and unanimously produce the most optimal viewpoint, humans generally agree that certain types of viewpoints are better than others [3]. Healey et al. [4] describes the goal of a desirable image as being able to support visual analysis, exploration, and discovery of novel insights. They further state that images should be generated in a way such that they are memorable, salient, and draw attention to areas of interest in the data [4].




Figure 1: Example of image rendered using different camera viewpoints [5]. This figure demonstrates how some images are better representations of a tea kettle than others. Certain camera viewpoints omit important features of the image, such as the handle or the spout.

How to Find Good Images?

Polonsky et al. [5] aimed to answer this question by establishing a methodology which consists of defining view descriptors. Later research refers to view descriptors as viewpoint quality (VQ) metrics, which is the term used for the remainder of this paper. Essentially, each VQ metric assigns a score to a particular view, which can then be used as a comparator against the scores of other, different views of that same data set. Polonsky et al. [5] attempted to quantify good views by basing their VQ metrics on the following principles: (1) measure of geometric complexity, (2) inherent view-dependent features, and (3) assigning value to non-primitive elements of the model. By evaluating images from different vantage points under formulas derived from these three principles, images can be compared with one another based on how they scored on different VQ metrics. In terms of findings, Polonsky et al. [5] found that no one VQ metric they considered in their study can perfectly evaluate image “goodness.”

Several other researchers have also investigated the efficacy of VQ metrics in relation to automatic camera selection, specifically with a focus on entropy based VQ metrics [5, 6, 7]. Takahashi’s [7] investigations found that viewpoint entropy can serve as a useful tool for naïve users to automate good viewpoints in the case of volume rendered images. Marsaglia et al. [2] also tackles entropy-based VQ metrics and additionally looks at the performance of certain VQ metrics in combination with other existing metrics. Their work supports the notion that entropy-based VQ metrics are good predictors of what camera placements generate good images.

Marsaglia et al. [2] explored how user preferences can be used to guide the selection of good images. In combining existing and newly discovered VQ metrics with data on user preferred images, VQ metrics that most closely align with images that users deem optimal can be determined. This is a recent area of automated camera placement research that has motivated and

directed the focus of my work. However, Marsaglia et al. [2] solely focused on the use of VQ metrics in respect to images produced by isosurfacing. As isosurfaces are only one of several different methods for rendering images, we believe their study is missing some perspective on other applications of visualization. Our work contrasts with Marsaglia et al. [2] as we aim to extend their work and consider how VQ metrics can be used to predict good images for volume rendered images.

Viewpoint Quality Metrics

VQ metrics are one of the leading methods for guiding automatic selection of camera placements in scientific visualization. This section describes the specific VQ metrics relevant to this work. Other existing metrics have been excluded because they are not suitable for the scope of this study. In particular, Marsaglia et al. [2] constructed their analysis of these metrics such that their results could be applied in situ. This work applies similar constraints for comparison purposes.

Number of Visible Triangles

This metric originates from Plemenos [8] and was further developed by Plemenos and Benayada [9]. In computer graphics, triangles are used to draw images by approximating areas of information, like color. This metric scores images containing a higher number of visible triangles, thus a higher quantity of visible information, with higher scores.

Projected Area

This metric also originates from Plemenos and Benayada [9]. Projected area measures the total area of visible polygons that make up an image. Similar to the Number of Visible Triangles metric, Projected Area favors images with a greater quantity of visible information.

Plemenos and Benayada (PB)

This metric is a combination of Number of Visible Triangles and Projected Area metrics.

Visibility Ratio

This metric, also from Plemenos and Benayada [9], measures the visibility ratio of an image. The visibility ratio is defined as the visible surface area divided by the total surface area of a visualization.

Viewpoint Entropy

This metric was first utilized for selecting camera viewpoints by Vasquez et al. [6]. This metric considers projected area of data and the number of visible faces to evaluate the entropy at a given viewpoint.

Viewpoint Kullback-Leibler Distance (VKL)

This metric was developed by Sbert et al. [10]. VKL measures the Kullback-Leibler distance between the normalized distribution of projected area and the real area visible from a given viewpoint [2].

Maximum Depth

This metric was created by Stoev and Strasser [11], but first applied to viewpoint exploration by Secord et al. [12]. This metric measures the maximum depth of the model at a specific viewpoint. With a greater depth, it is possible that more information will be visible to the user.

The following metrics were created by Marsaglia et al. [1, 2]. These metrics utilize Shannon Entropy, which calculates the average amount of information in some entity [13].

Higher levels of entropy reflect a higher quantity and variety of information. Marsaglia only considers visible data in these metrics, that is the area of an image that is not occupied by empty or background space.

Data Entropy

This metric calculates the entropy of visible field data from a specific viewpoint.

Depth Entropy

This metric calculates “the entropy of the distances from the camera to the visible field of data” at a specific viewpoint [2].

Shading Entropy

This metric calculates the entropy of the visible shading coefficients. This metric uniquely considers how the light and shading of a visualization might affect the quality of an image from certain camera placements.

DDS

This metric combines the previous three metrics: data entropy, depth entropy, and shading entropy. Marsaglia et al. [1, 2] found this combined metric to be the most accurate in predicting user preferences, with a 68% success rate. Marsaglia’s three entropy-based metrics were combined by adding the raw scores of each individual metric to get one score. In this study, DDS will be used as a single metric.

Isosurfacing and Volume Rendering

The distinction between isosurfacing and volume rendering is very relevant to this work. Isosurfacing is a visualization operation that generates images of three-dimensional data using

“surfaces.” An isosurface represents a single value within the data set. Everywhere the data set contains data equivalent to the isovalue, the data is included in the rendered image. Setting multiple isovalues creates multiple isosurfaces, which generates a visualization made up of layers at varying field values. An example of an image rendered using isosurfaces can be seen below in Figure 2.




Figure 2: An example of isosurfacing. This image depicts the hurricane data set rendered using isovalues of 10 and 90 [18]. These parameters result in an image that contains a blue layer representing field data of value 10 on top of a red layer that represents field data of value 90.

Volume rendering is another operation widely used in scientific visualization. Instead of using isovalues to define what information should be included in an image, volume rendering uses a specification called a transfer function to determine what values of field data should be included. Volume rendered images also consider an additional component of opacity, which is contained in a visualization’s transfer function. This is helpful for images that have useful or salient data that is obstructed or contained within exterior data.

Transfer functions describe the level of opacity and color a volume rendering should have at different field data values. Figure 3 demonstrates a visualization of Mantle where a fully opaque transfer function is applied and another where a specific transfer function is applied. The figure shows that when the transfer function is applied, most of the field data in the middle of the block becomes nearly transparent, making the blue, red, and yellow protrusions visible.




Figure 3: An example of the Mantle data set with a fully opaque transfer function (left) and the transfer function used in this study (right).

My Study

The design of this study matches Nicole Marsaglia's [2] work on investigating an entropy-based approach for identifying user-preferred camera positions in scientific visualizations rendered using isosurfacing. With the goal of extending Marsaglia's study to explore how VQ metrics can be used to predict user preferences in volume rendered images, the data collection, methods, and evaluation procedures discussed below closely follow those outlined in their study.

Data Corpus for Comparing Viewpoints

This section details the data corpus used for evaluating viewpoint metrics. The data corpus is comprised of three elements: (1) a collection of volume rendered images to be used as comparators in the user survey, (2) the results of a user survey detailing user preferences, and (3) VQ metric scores for ten different images of each data set used in the study.

Determining User Preferences

This section describes the process of building the collection of volume rendered images used as the basis for the user survey. Subsequently, the user survey is used to generate data on preferences of various camera viewpoints. The images discussed in this section are also used later to inform the selection of parameters for image analysis.

Data Sets

The data sets used in this study were obtained from the collection of data sets used in Marsaglia's initial work [2]. Marsaglia selected these data sets because they represent a diverse

assortment of images and domains, so the results of their study could possibly be applied to a greater realm of scientific data sets. Marsaglia utilized 10 large-scale scientific data sets gathered from the IEEE Visualization Conference's Scientific Visualization Contest and the Exascale Computing Project from the United States' Department of Energy. In this study, only nine of these data sets are evaluated. The one data set excluded from this research was determined to be unsuitable for volume rendering because the images captured from different camera positions were similar and lacked discernable qualities. The data sets used in the study are described below:

- Asteroid: A data set of the impact of an asteroid hitting water [14].
- Constit: A data set of the deformation response in polycrystalline materials [15].
- ExaAm Truchas: A data set that examines micro-structures of Additive Manufacturing [16].
- Fluid Dynamics: A data set that models the flow of water through a cylindrical shape [17].
- Hurricane: A data set of Hurricane Isabel [18].
- Mantle: A data set that models the Earth's mantle [19].
- Miranda: A data set of hydrodynamic turbulence [20].
- S3D-N2: A data set of combustion's N2 field data [21].
- S3D-UVEL: A data set of combustion's U Velocity field data [21].

A specific time slice of each data set described above was selected. These time slices were chosen because they were determined to sufficiently represent the simulation as a whole.

Rendering Images

Images of each data set were created using volume rendering. For the purpose of this study, transfer functions were selected through a combination of trial-and-error testing and referring to the isovalues used to generate the images in Marsaglia et al. If the selection of transfer functions was implemented poorly, problems would have emerged during this process. The volume rendered images in this study are designed to resemble the isosurface images used in Marsaglia et al. and highlight comparable features of the data. Therefore, if the selected transfer functions did not adequately mirror Marsaglia's isovalues, comparisons between the Marsaglia study and this work would not be applicable.

The main goal of generating an image library was to create images that were aesthetic, comprehensible, and highlighted important elements of the data set. Different color schemes, lighting arrangements, and transfer functions were tested on each individual data set before deciding on what schema was selected. The final collection of volume rendered images used in our corpus is displayed in Figure 4.




Figure 4: The image collection for our corpus. A 3D visualization of each data set was generated using volume rendering and uniquely designed transfer functions. Then, using a Fibonacci Lattice pattern, images were taken of each data set from ten evenly spaced camera locations.

Camera Placement

For every data set, images were rendered from each of ten different camera positions. Camera positions were selected using a Fibonacci Lattice, which ensured the camera is placed at ten evenly spaced locations circling around the three-dimensional visualization, as shown in Figure 5. This method of camera placement was selected as it was the same procedure used in

Marsaglia et al. [2]. Additionally, using a Fibonacci Lattice ensured that images are rendered from unique and dissimilar vantage points, and that every data set is captured from comparable positions.

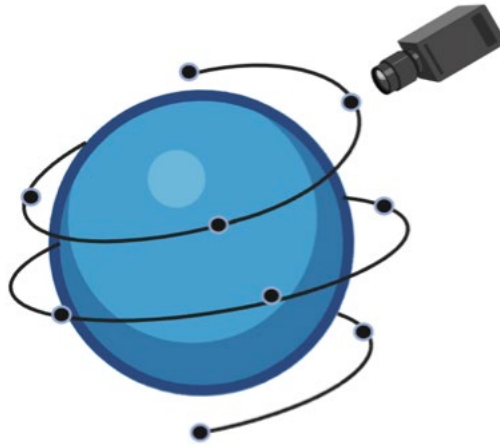


Figure 5: A depiction of what a Fibonacci Lattice looks like [2].

User Survey

The next element included in the corpus of data is user preferences of camera viewpoints. Two user surveys were conducted to gather these user preferences. The first survey consisted of Scientific Visualization Students at the University of Oregon—a class comprised mainly of juniors, seniors, and graduate students in the University’s Computer Science program. The second survey consisted of professionals in the field of scientific visualization—the same group that was invited to participate in the user survey conducted by Marsaglia et al. [2].

To take part in the survey, participants were given access to a website where they could indicate their preferences for different image pairings. Each survey question was comprised of two images randomly selected from the image collection. For each survey question, the two images displayed belonged to the same data set but were captured from different camera

viewpoints. Participants of the survey were asked to select which image they felt best represents the simulation. If participants could not determine if one image is superior to another, they were given a third option that indicated the participant had no preference between the images. After making a preference selection, a new random pairing of images displayed on the page. This sequence of events repeated until the participant exited the survey. Participants did not have any prior knowledge of what real-world situation the visualization is illustrating, so they were not privy to what elements of a visualization may be more relevant or important than others. See Figure 6 for an example of the survey design.




Figure 6: An example question from the user survey. Participants selected the image they preferred to represent the simulation. If participants did not have a strong preference for either image, they may select “Neutral.” Participants answered questions for as long as they wanted, then selected “Done” to exit the survey.

The first survey comprised of UO scientific visualization students received 610 responses. The second survey comprised of scientific visualization professionals received another 1548 responses. This data constitutes the user preference data corpus for this study.

Viewpoint Analysis

This section describes the process of preparing and running the analysis for each data set. The analysis procedure utilized a high-performance computing, in situ visualization infrastructure called Ascent. The Ascent library contains implementations to calculate the eleven VQ metrics used in this study. When running an example for a particular VQ metric, Ascent reads in information about the data set time slice used, camera placements, and calculation parameters, and outputs a score for each of the 10 images rendered from different viewpoints.

The VQ metric implementations in Ascent are designed to work on solid surfaces. This means that images generated using isosurfaces are preferable to volume rendered images, which typically contain masses of transparent data. In order to reuse the pipeline utilized in Marsaglia et al. [2], while also ensuring the scope of this research remains relevant to images generated through volume rendering, images were constructed by selecting approximate isovalues that mimic the attributes of the volume rendered images used in the user survey.

Selecting Isovalues to Represent Volume Rendered Images

The transfer functions created for each of the volume rendered images included in the user survey are used as a guideline for the selection of isovalues. Isovalues were selected at the approximate narrow range in the transfer function where the rendered image is nearly or fully opaque. While not explicitly evaluating volume rendered imagery, this method allows Ascent to read in isosurface image data that visually represents the volume rendered version of that image.

Figure 7 below demonstrates how a volume rendered image can be translated to an isosurface image, which is interpretable by Ascent.

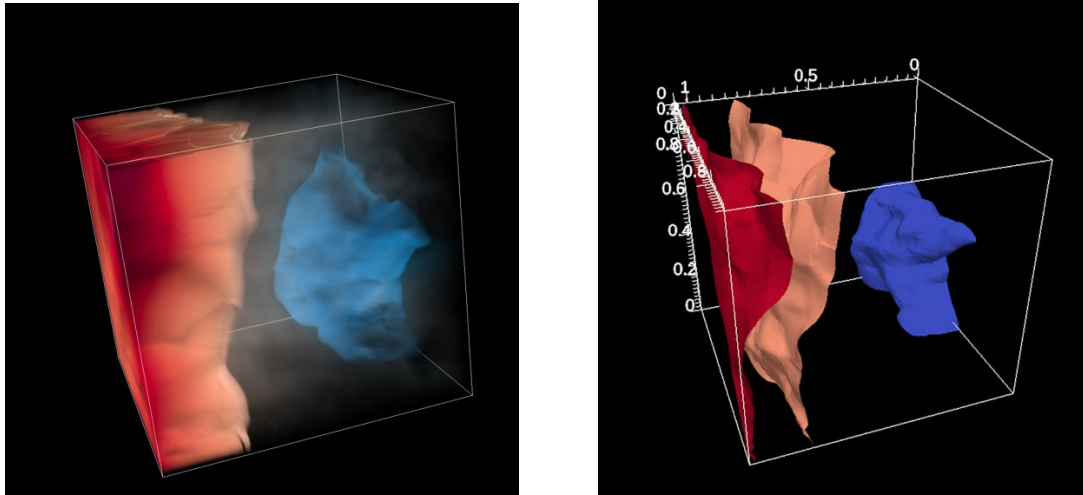


Figure 7: An example of a volume rendered image translated into isosurfaces to be used in the Ascent analysis. The image on the left was produced using volume rendering operations in VisIt. The image on the right was produced by Ascent, using the isovalues given as input.

The isovalues were selected using a combination of transfer function values and trial and error of what images looked most representative of the volume rendered data set. This mapping process could have introduced problems if the isovalues selected were not sufficiently representative of the accompanied volume rendered image. However, through careful assessment and repetitive testing, satisfactory isovalues were chosen for analysis and issues related to this process were avoided.

Evaluating Results

The efficacy of each VQ metric used in this study was evaluated using comparator oracles. The oracles use VQ metric scores to predict what camera placement users prefer. This study utilizes single-metric oracles that only consider one VQ metric at a time when making

predictions. This contrasts with the Marsaglia study, which additionally considered multi-metric oracles. That said, the Marsaglia study found that a combination of three oracles (data entropy, depth entropy, shading entropy) was the most effective, and that combination is used here as a single metric with a corresponding oracle. Thus, eleven single-metric oracles are utilized in this study, one for each of the eleven VQ metrics described previously in the Background section.

A single-metric oracle takes two scores for a given VQ metric as input, one score from camera position C_1 and one score from camera position C_2 . If the score for C_1 is greater than the score for C_2 , the oracle will select C_1 as the winning viewpoint. For all metrics used in the study, a higher score represents a “better” image, so the oracle will always select the image producing the highest score.

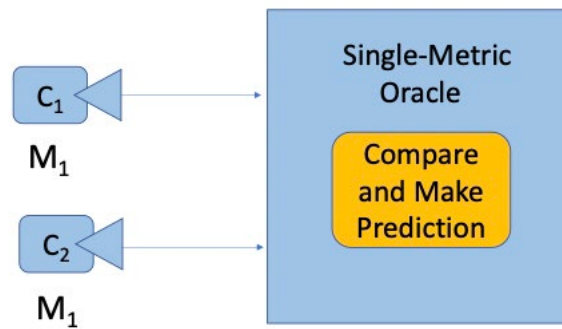


Figure 8: A depiction of the logic used in a single-metric oracle [2]. Two different cameras, C_1 and C_2 , and the selected metric, M_1 , are given as input to the oracle. The oracle then compares the metric scores for each camera, makes a decision about which camera has a better viewpoint based on the metric provided, and outputs its decision.

Oracles that perform the highest are the oracles that make the most accurate predictions. Our evaluation measures the number of times each oracle can correctly predict the user’s preference. If an oracle makes a correct prediction, it receives one point. If the oracle does not make a correct prediction, it receives zero points. At the end of our analysis, the oracle with the highest sum of points is determined to be the best oracle at predicting user preferences.

Results

Viewpoint Quality Metric Findings

This section discusses the efficacy of metric-based oracles in predicting user preferences of camera placements for various scientific visualization data sets. In this discussion, user preferences are partitioned into two groups: Student preferences and Professional preferences. The evaluation of each metric for both groups is detailed below in Tables 1 and 2.

Student Data		
Metric	Correct Predictions	% (/498)
Data Entropy	278	55.8%
DDS Entropy	313	62.9%
Viewpoint Entropy	175	35.1%
Shading Entropy	267	53.6%
Depth Entropy	296	59.4%
Visibility Ratio	172	34.5%
Max Depth	310	62.2%
# of Visible Triangles	182	36.5%
Plemenos & Benayada	200	40.2%
VKL Distance	178	35.7%
Projected Area	203	40.8%

Table 1: Correct predictions for each metric using Student data corpus.

Professional Data		
Metric	Correct Predictions	% (/973)
Data Entropy	572	58.8%
DDS Entropy	643	66.1%
Viewpoint Entropy	375	38.5%
Shading Entropy	600	61.7%
Depth Entropy	605	62.1%
Visibility Ratio	377	38.7%
Max Depth	549	56.4%
# of Visible Triangles	398	40.1%
Plemenos & Benayada	389	40.0%
VKL Distance	368	37.8%
Projected Area	384	39.5%

Table 2: Correct predictions for each metric using Professional data corpus.

The data collected from the Student Group and the Professional Group resulted in 610 and 1548 responses, respectively. However, some of the responses provided by participants indicated that they felt neutral about the two images presented to them. Responses that signified the participant did not have a strong preference of one image over the other were discarded in the

evaluation of results. As such, the evaluation only considers the 498 Student responses and 973 Professional responses where the participant indicated they did have a preference.

The rate of correct oracle predictions was generally similar between the Student and Professional responses. The top scoring metric for both the Student survey and the Professional survey data corpus was DDS Entropy, with an accurate prediction rate of 62.9% and 66.1%, respectively. With the highest prediction rate of approximately 66%, DDS Entropy was the most accurate of the metrics considered in this study at predicting human preference for both surveyed groups. This suggests that DDS Entropy is the best VQ metric to use when predicting what users may believe to be good viewpoints in a data visualization.

Several of the other VQ metrics evaluated in this study performed significantly worse than a random-selection oracle, which would make correct predictions 50% of the time. Viewpoint entropy made correct predictions only 35.1% and 38.5% of the time for Student and Professional preferences. Visibility ratio, VKL distance, and projected area also performed poorly in predicting both Student and Professional preferences. Such low accuracy rates indicate that these metrics would likely be inadequate in informing the selection of camera placements based on user preferences.

Efficacy of Top Performing Metric

This section explores the performance of DDS Entropy, the most accurate metric at predicting user preferences. Tables 3 and 4 detail the efficacy of DDS selections for different types of cameras. Each of the ten camera viewpoints used in every data set were classified as one of four types: POOR, FAIR, GOOD, and VERY GOOD. These categorizations were determined by the rate at which each individual camera was selected as the preferred camera in the user survey. Remember, the user survey asked participants to select their preferred image between

two images of the same data set captured from different camera viewpoints. For example, Camera #6 of the Mantle data set was selected by participants of the user survey 29 times out of the 30 times it was included in a question. This camera had a winning percentage of 96.7%. Alternatively, Camera #9 of the Mantle data set was selected only one out of 30 times, having a winning percentage of just over 3%. Each camera position classified as POOR had a winning rate of 0%-25%, FAIR had a winning rate of 25%-50%, GOOD had a winning rate of 50%-75%, and VERY GOOD had a winning rate of 75%-100%.




Figure 9: Mantle rendered from Camera #6 [19]. This camera was assigned as a VERY GOOD camera.




Figure 10: Mantle rendered from Camera #9 [19]. This camera was assigned as a POOR camera.

The DDS oracle made its own selections regarding what camera viewpoint it predicted would best fit users' preferences, independent of the winning rate classification assigned to each camera. Then, these predictions were compared to users' true preferences, and evaluated based on the type of camera. These findings are detailed in Table 3 and 4 below.

	POOR	FAIR	GOOD	VERY GOOD	SUM
POOR	47% (11/23)	73% (53/72)	80% (194/242)	93% (90/100)	79.6% (348/437)
FAIR	73% (53/72)	62% (25/40)	52% (96/182)	68% (45/66)	60.8% (219/360)
GOOD	80% (194/242)	52% (96/182)	44% (56/127)	47% (53/111)	60.3% (399/662)
VERY GOOD	93% (90/100)	68% (45/66)	47% (53/111)	50% (5/10)	67.2% (193/287)

Table 3: This table details the performance analysis of the top performing metric, DDS Entropy, on the Professional survey data corpus. A cross section entry of this table describes the winning percentage, i.e., the number of times the DDS oracle correctly predicted user preferences, based on the categorization of each type of camera match up. For instance, the (POOR, GOOD) rating of 80% means that in pairings of one “poor” camera and one “good” camera, the oracle made correct predictions 80% of the time. The table cells are colored based on performance; camera pairings that were correctly predicted more than 70% of the time are colored green, pairings correctly predicted 60%-70% of the time are colored yellow, and pairings that were correctly predicted less than 60% of the time are colored orange.

	POOR	FAIR	GOOD	VERY GOOD	SUM
POOR	77% (7/9)	79% (38/48)	70% (54/77)	87% (61/70)	78.4% (160/204)
FAIR	79% (38/48)	54% (32/59)	43% (39/89)	69% (36/52)	58.5% (145/248)
GOOD	70% (54/77)	43% (39/89)	43% (16/37)	36% (18/49)	50.4% (127/252)
VERY GOOD	87% (61/70)	69% (36/52)	36% (18/49)	62% (5/8)	67.0% (120/179)

Table 4: This table details the performance analysis of the top performing metric, DDS Entropy, on the Student survey data corpus. This table follows the same coloring schema as Table 3 above.

These findings demonstrate that DDS Entropy works best when predicting selections that consist of one POOR camera. The DDS oracle was most accurate when deciding between a POOR camera viewpoint and a VERY GOOD camera viewpoint, correctly predicting user preferences 87% of the time for Students and 93% of the time for Professionals. A higher

success rate between POOR and VERY GOOD camera views is expected, as it should be easier to differentiate between an image with high levels of visible information and detail versus images with very little. It is also relevant to highlight the high performance of the DDS oracle in making predictions between POOR and FAIR cameras, as well as POOR and GOOD cameras.

Despite high performances in pairings where POOR cameras are involved, the DDS oracle is far less accurate when making selections between images that are a combination of FAIR, GOOD, and VERY GOOD. This may be a result of images being equally sufficient in representing the simulation that it becomes difficult to select a “best” image. These results do not suggest that DDS Entropy works reliably in all cases; however, there is an indication that the DDS metric is effective in uncovering images that are *not* good.

Potential Influences on Results

This section provides some explanation regarding aspects of this study that may have influenced the outcome of this evaluation.

Comparisons of the Student and Professional survey response data corpus demonstrate how user preferences can oftentimes vary. Ultimately, users are human, and humans do not regularly agree on things. Participants in the Professionals’ survey have more experience and expertise viewing and analyzing large, scientific data sets than the Students, so they may have unique insights that lead them to make different decisions about camera placements. Different users may also have differing preferences when it comes to image aesthetics, orientation, and level of detail.

Additionally, the efficacy of a given VQ metric may be dependent on the particulars of the data sets used in this study. For instance, a data set that has more obvious and distinguishable features, like a clear orientation or focal point, may yield different results than data sets that are

more complex and disorienting. The tables below detail the correct prediction rate of DDS Entropy, the top performing metric, by data set for both the Student and Professional survey groups.

Student Data	
Data Set	Prediction Rate
Asteroid	58.8%
Constit	70.5%
ExaAm Truchas	75.0%
Fluid Dynamics	54.0%
Hurricane	46.8%
Mantle	56.9%
Miranda	67.1%
S3D-N2	88.6%
S3D-UVel	51.9%

Table 5: DDS prediction rate by data set using Student data corpus.

Professional Data	
Data Set	Prediction Rate
Asteroid	72.6%
Constit	69.3%
ExaAm Truchas	66.1%
Fluid Dynamics	77.7%
Hurricane	34.8%
Mantle	65.5%
Miranda	67.5%
S3D-N2	72.6%
S3D-UVel	68.7%

Table 6: DDS prediction rate by data set using Professional data corpus.

The tables above demonstrate that the Student and Professional responses reflect conflicting preferences of camera placement for some of the data sets. DDS Entropy was able to predict Student preferences a shocking 88.6% of the time in respect to the S3D-N2 data set. Meanwhile, DDS Entropy predicted Professionals' preferences a measly 34.8% of the time in respect to the Hurricane data set. Many factors may contribute to the incongruent prediction rates contained in Table 5 and Table 6, but DDS Entropy still proves to be at or above its overall prediction rate of 66% in Constit, ExaAm Truchas, Miranda, and S3D-N2 data sets for both survey groups.

Conclusion and Future Work

This research aimed to further explore the automation of selecting camera viewpoints for scientific visualizations based on user preferences. One of the primary goals of this study was to expand and further the research conducted by Marsaglia et al., which focused on the automation of camera placements in images rendered using isosurfacing. Marsaglia found that some VQ metrics did work in effectively predicting what images users prefer. In this study, we found that Marsaglia's top performing metric, DDS Entropy, is also effective when used on images generated through volume rendering.

This research found that in respect to volume rendered images, DDS Entropy is most effective when used to detect bad images. DDS Entropy was particularly useful when determining selections between image pairings of one good image and one distinctly bad image. Conducting exhaustive searches and analyses on a set of images to find what image best represents the simulation is costly and ineffective. However, using a tool like DDS Entropy to reject images that are discernably not good may be a quick and effective way to support the automation of camera viewpoint selection during in situ scientific visualizations.

In this work, we observed that DDS Entropy was able to predict user preferences up to 66% of the time. In consideration with Marsaglia's findings of a successful prediction rate of 68%, this discovery provides additional evidence that DDS Entropy is an effective metric in predicting camera placements that match user preferences. Furthermore, these findings provide encouraging support that DDS Entropy could possibly be applied to more scientific visualization methods in the future.

This research advances our understanding of whether VQ metrics can be applied to various methods of rendering images in the field of scientific visualization. The findings of this

study suggest that certain VQ metrics do work as good indicators of user preference for both isosurface and volume rendered images. However, isosurfacing and volume rendering are just two approaches to image rendering. Future work on this topic may consider other image rendering methods. For instance, future work may include exploration into whether VQ metrics, such as DDS Entropy, are also effective in predicting the user preferred viewpoints in visualizations rendered using other popular methods like streamlines or domain clipping.

References

- [1] N. Marsaglia, "Automating Camera Placement For In Situ Visualization," Ph.D. dissertation, Univ. Oregon, 2022.
- [2] N. Marsaglia, Y. Kawakami, S. D. Schwartz, S. Fields, and H. Childs, "An Entropy-Based Approach for Identifying User-Preferred Camera Positions," 2021 IEEE 11th Symposium on Large Data Analysis and Visualization (LDAV), New Orleans, LA, USA, 2021, pp. 73-83, doi: 10.1109/LDAV53230.2021.00015.
- [3] V. Blanz, M. J. Tarr and H. H. Bühlhoff, "What object attributes determine canonical views?", *Perception*, vol. 28, no. 5, pp. 575-99, 1999.
- [4] C. G. Healey, J. T. Enns. Attention and visual memory in visualization and computer graphics. *IEEE Trans Vis Comput Graph.* 2012 Jul;18(7):1170-88. doi: 10.1109/TVCG.2011.127. PMID: 21788672.
- [5] O. Polonsky, G. Patané, S. Biasotti, C. Gotsman, and M. Spagnuolo. What's in an image? *The Visual Computer*, 21(8):840-847, Sep 2005.
- [6] P. P. Vázquez, M. Feixas, M. Sbert and W. Heidrich, "Viewpoint selection using viewpoint entropy", *Proceedings of the Vision Modeling and Visualization Conference 2001 VMV '01*, pp. 273-280, 2001.
- [7] S. Takahashi, I. Fujishiro, Y. Takeshima and T. Nishita, "A feature-driven approach to locating optimal viewpoints for volume visualization," in *Visualization Conference, IEEE*, Minneapolis, Minnesota, 2005 pp. 495,496,497,498,499,500,501,502. doi: 10.1109/VISUAL.2005.1532834
- [8] D. Plemenos. *Contribution a l'étude et au développement des techniques de modélisation, generation et visualization de scenes : le projet multiformes*. PhD Thesis, 1991. 1991NANT2031.
- [9] D. Plemenos and M. Benayada. Intelligent display in scene modelling. new techniques to automatically compute good view. In *GraphiCon '96*. Saint Petersburg (Russia), July 1996.
- [10] M. Sbert, D. Plemenos, M. Feixas, and F. González. Viewpoint quality: Measures and applications. In *Proceedings of the First Eurographics Conference on Computational Aesthetics in Graphics, Visualization and Imaging*, Computational Aesthetics'05, pp. 185-192. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2005.
- [11] S. L. Stoev and W. Strasser. A case study on automatic camera placement and motion for visualizing historical data. In *IEEE Visualization, 2002. VIS 2002.*, pp. 545-548, Oct 2002.
- [12] A. Secord, J. Lu, A. Finkelstein, M. Singh, A. Nealen. Perceptual models of viewpoint preference. *ACM Trans. Graph.*, 30(5):109:1-109:12, Oct. 2011.

- [13] S. Vajapeyam. “Understanding Shannon's Entropy Metric for Information.” *Arxiv.org*, Cornell University, <https://arxiv.org/pdf/1405.2061.pdf>.
- [14] J. Patchett and G. Gisler. Deep water impact ensemble data set. 4, 2017.
- [15] R. A. Carson, S. R. Wopschall, and J. A. Bramwell. ExaConstit, Aug 2019.
- [16] J. Belak, J. Turner, and E.T. Team. Exaam: Additive manufacturing process modeling at the fidelity of the microstructure. In *APS March Meeting Abstracts*, vol. 2019, pp. C22-010, 2019.
- [17] Hurricane Isabel simulation data. p. online, 2004.
- [18] J. Kuhnert. Meshfree numerical schemes for time dependent problems in fluid and continuum mechanics. *Advances in PDE modeling and computation*, pp. 119–136, 2014.
- [19] M. H. Shahnas, W. R. Peltier, Z. Wu, and R. Wentzcovitch. The high- pressure electronic spin transition in iron: Potential impacts upon mantle mixing. *Journal of Geophysical Research: Solid Earth*, 116(B8), 2011.
- [20] A. W. Cook, W. H. Cabot, P. L. Williams, B. J. Miller, B. R. de Supinski, R. K. Yates, and M. L. Welcome. Tera-scalable algorithms for variable-density elliptic hydrodynamics with spectral accuracy. In *SC'05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*, pp. 60–60. IEEE, 2005.
- [21] S. Treichler, M. Bauer, A. Bhagatwala, G. Borghesi, R. Sankaran, H. Kolla, P. McCormick, E. Slaughter, W. Lee, A. Aiken, and J. H. Chen. S3d-legion: An exascale software for direct numerical simulation of turbulent combustion with complex multicomponent chemistry. 11 2017.