

# On Visual Quality of Optimal 3D Sampling and Reconstruction

Tai Meng, Benjamin Smith, Alireza Entezari, Arthur E. Kirkpatrick, Daniel Weiskopf, Leila Kalantari, Torsten Möller\*

School of Computing Science, Faculty of Applied Sciences  
Simon Fraser University

## ABSTRACT

This paper presents a user study of the visual quality of an imaging pipeline employing the optimal body-centered cubic (BCC) sampling lattice. We provide perceptual evidence supporting the theoretical expectation that sampling and reconstruction on the BCC lattice offer superior imaging quality over the traditionally popular Cartesian cubic (CC) sampling lattice. We asked 12 participants to choose the better of two images: one image rendered from data sampled on the CC lattice and one image that is rendered from data sampled on the BCC lattice. We used both synthetic and CT volumetric data, and confirm that the theoretical advantages of BCC sampling carry over to the perceived quality of rendered images. Using 25% to 35% fewer samples, BCC sampled data result in images that exhibit comparable visual quality to their CC counterparts.

**CR Categories:** G.1.2 [Numerical Analysis]: Approximation—Spline and piecewise polynomial approximation; H.5.m [Information Interfaces and Presentation]: Miscellaneous; I.4.5 [Image Processing and Computer Vision]: Reconstruction; I.4.10 [Image Processing and Computer Vision]: Image Representation—Volumetric.

**Keywords:** optimal sampling and reconstruction, perceptual image quality, body-centered cubic (BCC) lattice, Cartesian lattice

## 1 INTRODUCTION

For sampling volumetric data, the body-centered cubic (BCC) lattice possesses attractive theoretical advantages over the commonly used Cartesian cubic (CC) lattice. Although both lattices provide a regular sampling of 3D space, for a given number of samples, the BCC lattice preserves more of the signal. Further, in an important sense, BCC sampling is theoretically optimal: a mathematically equivalent signal quality can be achieved with fewer BCC samples than CC samples. In this paper, we provide empirical evidence to complement these mathematical arguments, showing that BCC sampling can provide visually comparable results using only about 70% of the samples required for a CC lattice. A consequence of practical importance is that, if the same number of samples are taken on the BCC lattice, the resulting visualization provides substantial improvements in accuracy.

A detailed discussion of the mathematical similarities and differences between BCC and CC lattices is provided in the papers by Entezari et al. [2, 3]. To summarize, the arguments focus on signals with an isotropic low-pass spectrum. The low-pass isotropic assumption implies that the underlying data has important information in all directions. This is a reasonable assumption when there is no additional knowledge about the phenomenon to be sampled, and is typically the case for visualization and rendering applications.

\*e-mail: {tmeng,brsmith,aentezar,ted,weiskopf,lkalanta,torsten}  
@cs.sfu.ca

For a signal with an isotropic low-pass spectrum, the densest arrangement of the replicas of the spectrum in the Fourier domain will allow for the optimal sampling rate in the spatial domain. The optimal arrangement is to pack these replicas in a face-centered cubic (FCC) lattice [5]. The FCC periodic arrangement of replicas in the Fourier domain demands periodic sampling on the dual to the FCC lattice, which is the BCC lattice. Thus, the BCC lattice is optimal in the spatial domain for signals meeting the low-pass isotropic assumption.

Recently, this theoretical advantage has led to practical algorithms that include high-quality reconstruction filters adapted to the BCC lattice [2] as well as efficient implementations of those filters that are twice as fast computationally as their CC counterparts [4]. These mathematical and algorithmic results set the stage for advocating BCC-based sampling, reconstruction, and processing of volumetric data. It is important to point out that the arguments for accuracy comparisons have been made on the basis of asymptotic error behavior—the standard numerical analysis reasoning. However, our goal is to produce images for perceptual consumption. Numerical reasoning is insufficient for this task, and a perceptual evaluation of the effectiveness of BCC and CC lattices is required.

### 1.1 Research Question

Basic perception research typically restricts user studies to simply structured stimuli in 2D images. In contrast, an image of a 3D volumetric dataset can only be produced by applying a pipeline of processing steps that includes components such as reconstruction filter, transfer function, illumination, camera location, projection from 3D space to a 2D image, and other effects depending on the particular volume rendering process chosen [10]. All these components may influence the image quality. Ideally, their effects should be modeled mathematically and analyzed in a rigorous manner. However, no one has yet been able to provide such a comprehensive model of the volume rendering process.

Furthermore, it is extremely difficult to investigate the perceptual effects of these aforementioned parameters of the volume rendering process. Varying these parameters would lead to a high-dimensional parameter study that is infeasible. Therefore, we restrict our study to a sub-region of this vast parameter space. Accordingly, the question we are trying to answer in this paper is: “At what relative sampling density will BCC-based and CC-based volume data become indistinguishable for a human observer, provided we use numerically comparable reconstruction filters as well as the same lighting conditions, transfer functions, and rendering pipeline?”

### 1.2 Related Work on Visual Perception

The idea of validating reconstruction and rendering algorithms via a user study is not new. Most related to this paper is the work by Mitchell and Netravali [11], who showed that the numerical analysis of a class of filters they studied is independent of the perceptual variety of effects that this same class of filters achieves. Hence, they engage in a perceptual study of 9 users to evaluate their images according to image artifacts such as blurring, ringing, and anisotropic

behavior. By extension of their result, we cannot simply predict the degree of visual similarity between images created by reconstruction techniques with the same asymptotic error bounds. Therefore, even though the BCC and CC reconstruction filters we employ have the same asymptotic error behaviour (Section 3.4.1), a user study is still needed to investigate questions of visual similarity in the images produced.

There are also several user studies that evaluate various perceptual aspects in visualization. Examples include the evaluation of textures [6], surface layers [1], flow visualization methods [7] and shape perception [12]. However, these previous works do not directly investigate the accuracy of the underlying rendering pipeline and the image quality related to data reconstruction.

We also considered adopting an existing 2D image metric to answer our research question. Metrics for 2D images typically gauge the blurriness, sharpness, ringiness, or related characteristics of images. In particular, the blurriness metric recently developed by Marziliano et al. [9] seemed to have the potential to answer our research question. Marziliano et al.’s blur metric connects average edge width to perceptual blurriness. In our preliminary analysis, the problem of visual comparability is related to gauging the amount of details in two images. We further thought that “the amount of detail” could be measured by a blurriness metric. However, this turned out not to be the case—visual quality for 3D renderings cannot be measured by the amount of details alone. For a 3D signal, as the sampling resolution decreases, sharp features begin to disintegrate into tiny pieces upon reconstruction. This results in an increase in details according to Marziliano et al.’s metric, when in fact visual quality has degraded. Another problem with gauging the visual quality of a set of 3D samplings of the same signal is that sampling might cause some features to slightly shift. Trying to apply a rigid-body image registration algorithm to correct the shifts seems infeasible because features could move in opposite directions. This essentially prevents the use of any full-reference metric that assumes feature alignment between the truth image and the test image. To the knowledge of the authors, there is no existing visual quality metric—in 2D or 3D—that could measure visual quality of images while accommodating these problems.

### 1.3 Contribution

The main contribution of this paper is to confirm that to achieve images of comparable visual quality, we need approximately 30% fewer samples on the BCC lattice than the CC lattice. This is the best we could hope for according to theoretical expectation. Our study shows that for both synthetic and CT data, comparable visual quality is achieved between BCC and CC sampled data at a relative BCC resolution (relative to CC) of 65% to 75%.

## 2 HYPOTHESIS

We coin the term *sampling comparability* to describe two different sampling patterns that represent the same amount of information of an underlying function. As discussed earlier, signal-processing theory says that a BCC lattice of approximately 70% of the samples of a CC lattice is sampling-comparable to the CC lattice. Given sampling comparability between BCC and CC sampled data, we reason that visual comparability of these data should occur at a relative BCC resolution of 70% (relative to CC).

We define visual comparability as perceived similarity in appearance. Corresponding to our research question, the following hypothesis is investigated in this paper: “Given a fixed set of rendering parameters, BCC and CC sampled data appear to be visually comparable to a viewer when a relative sampling density of 70% is used for BCC.”

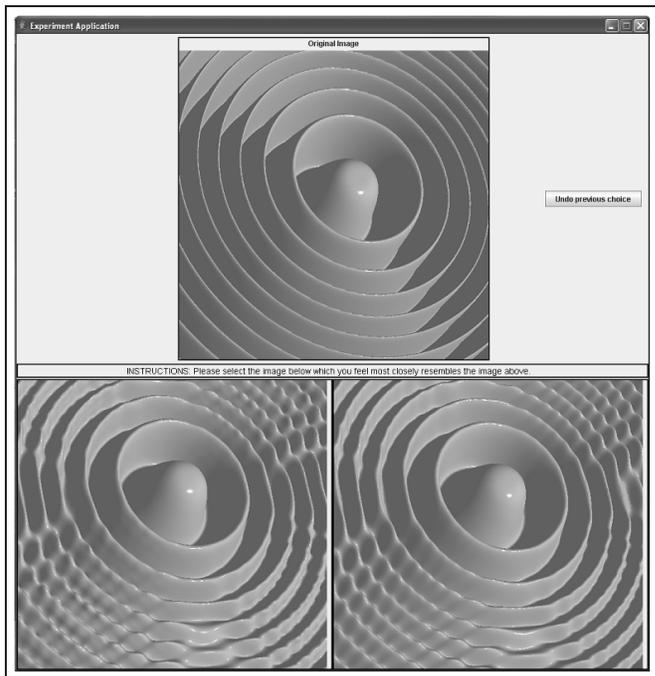


Figure 1: Screenshot of the application used to perform the image discrimination task. The image at the top is the ground truth (or original) image. Images at the bottom are CC and BCC downsampled versions of the ground truth. The participant selects the better of the BCC/CC images by mouse-clicking within the chosen image.

We choose to test visual comparability for a small set of rendering parameters. As discussed later in Section 3, a few different rendering parameters—signal selection, sampling resolution, and camera viewpoint—are varied one at a time with the purpose of covering a range of typical viewing conditions. In order to yield quantitative results, we restrict all other rendering parameters to comparable values. In particular, lighting conditions, transfer functions, and rendering algorithm are not varied. While these conditions are easy to enforce, the comparability of the reconstruction techniques on CC and BCC sampled data requires closer scrutiny (Section 3.4).

## 3 VISUAL COMPARABILITY EXPERIMENT

We conducted an image matching study with 12 participants to determine the range of relative sampling resolutions at which human observers find BCC images comparable to CC images.

### 3.1 Task Description

Participants performed an image matching task. They were shown three images (Figure 1), comprising one “original” or “ground truth” image at the top, and a pair of BCC and CC images at the bottom. Each image is  $500 \times 500$  pixels in dimension. In this visual discrimination experiment, participants were asked to select the lower image that “most closely resembles the image above”. If the participant could not determine which image more closely resembled the original, they were instructed to choose arbitrarily. A choice was made by clicking on one of the two bottom images.

The CC-sampled image was randomly assigned to be the left or right lower image, and the BCC image was placed in the other slot. In case the participant made a choice inadvertently, the application provided an undo button that would return to the previous trial.

However, no participant used this undo feature.

Participants were first informed of the nature of the task and signed a consent form. The experiment began with a training phase, designed to give the participant practice with the application and the image matching task. Image pairs for the training phase presented a sequence of 13 gradually more difficult choices. As with actual trials, each pair contained a CC and a BCC image, and participants were not given feedback about the correctness of their choices. Training images were taken from the set used in the main experiment, and consisted of 9 images of the synthetic signal and 4 images of the CT scan, using different rendering parameters.

The main phase of the experiment consisted of 8 blocks of 24 trials each, for a total of 192 trials per participant. Trials alternated between blocks of the synthetic signal and blocks of the CT scan. Within each block, images were presented for all resolutions and views of that signal. Participants were shown each CC/BCC image pair 4 times. Trials within blocks were randomized. Participants were encouraged to rest between blocks if they wished, to help alleviate boredom and maintain focus on the task. Each participant took between 30 minutes and an hour to complete the experiment.

### 3.2 Image Selection

The goal of image selection is to prepare a representative sample of images. In the selection process, we consider such secondary properties as signal selection, sampling resolution, and camera view.

#### 3.2.1 Signal Selection

We chose two signals. The first is the Marschner-Lobb function [8], a synthetic dataset commonly used in the volume graphics literature as a benchmark for testing reconstruction algorithms. Hereafter, we refer to this signal as “ML” (Figure 2a and Figure 2b). This signal is attractive since the local frequencies increase as we move out from the center. Hence, within one image, the impact of the rendering algorithm on widely varying frequencies can be observed.

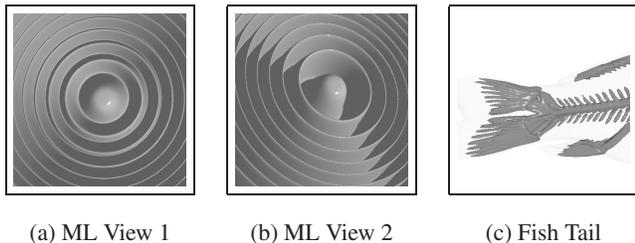


Figure 2: To account for the effect of different signal properties, both the synthetic ML signal as in (a, b) and a real-life Fish Tail signal as in (c) were presented. To account for possible bias due to camera view, the 3D ML was presented from both a “straight on” view, or View 1 as in (a), and a “tilted” view, or View 2 as in (b).

The second signal is a subset of the CT scan of a carp. We focused on the back half of the fish, including the tail fin. We refer to this signal as “Fish Tail” (Figure 2c). Compared to ML, Fish Tail is a real-world signal obtained through a CT scan reconstructed at a resolution of  $256 \times 256 \times 256$ . Fish Tail was chosen to represent non-synthetic data typical for many applications. Fish Tail also represents non-abstract signals with some “meaning” to participants.

Both of these signals exhibit high frequency features along different orientations and hence constitute viable candidates for testing our hypothesis.

The ground truth image of the ML signal was rendered directly from the analytic function definition [8]. Since ML is analytic, it can be thought of as an infinite-resolution ground truth. The ground

ML, CC80		Fish Tail, CC140		Fish Tail, CC180	
BCC47	41%	BCC87	48%	BCC105	40%
BCC52	55%	BCC91	55%	BCC113	49%
BCC54	62%	BCC95	62%	BCC118	56%
BCC55	65%	BCC97	67%	BCC122	62%
BCC56	69%	BCC98	69%	BCC125	67%
BCC57	72%	BCC99	71%	BCC128	72%
BCC58	76%	BCC100	73%	BCC130	75%
BCC59	80%	BCC102	77%	BCC132	79%
BCC60	84%	BCC103	80%	BCC136	86%
BCC62	93%	BCC105	84%	BCC140	94%
BCC64	102%	BCC111	100%	BCC143	100%
BCC66	112%	BCC114	108%	BCC147	109%

Table 1: For each CC sampling as indicated by heading, selected BCC resolutions are shown with their relative resolutions as a percentage of the CC sampling.

truth image for the Fish Tail signal was the  $256^3$  CC reconstruction from the CT scan.

#### 3.2.2 Sampling Resolution

Another secondary property of interest is the sampling resolution of the signals used in the discrimination task. As we are interested in the visual comparability of BCC samplings relative to a single CC sampling, one possible confounding variable is the CC sampling resolution used in the experiment.

The ML signal is a ripple-like pattern with an infinite series of concentric rings. The further a ring is from the center of the ML signal, the higher its frequency content. Traditionally, only a few of the rings of the ML signal that are confined to the domain of  $[-1, 1]^3$  were sampled at the rate of  $40 \times 40 \times 40$  for testing purposes. This was inadequate, since we wanted to have a broader mixture of high and low frequency content in our sampled data. Therefore, we chose to sample more rings of the ML up to the interval of  $[-2, 2]^3$ . To maintain the same sampling rate on the CC lattice, we therefore cast  $80 \times 80 \times 80$  samples. This will be denoted as CC80 in the remainder of this paper. Using a single CC resolution allowed us to feature two views of the ML samplings. To keep the number of trials manageable within the one hour limit of our user experiment, we did not include a second CC sampling of the ML. For the ML signal, both the CC and BCC datasets were sampled directly from the analytic function.

For the Fish Tail, we used two CC resolutions:  $140 \times 140 \times 140$  and  $180 \times 180 \times 180$ . We denote these samplings as CC140 and CC180. When downsampling the original dataset to non-dyadic CC resolutions, we applied a proper rational resampling by zero-padding in the Fourier domain followed by a downsampling in the spatial domain. The BCC-sampled versions were created in a similar way. After a Fourier zero-padding step, downsampling from CC to BCC in the spatial domain is achieved by keeping all datapoints of the CC dataset that have the same parity in their coordinates.

We denote BCC samplings using the notation  $BCC_n$  for a BCC resolution of  $n \times n \times (2n)$  samples. For the ML and the Fish Tail, the chosen BCC resolutions along with their relative percentages to the CC resolutions are enumerated in Table 1.

Figure 10 illustrates ML and Fish Tail signals rendered at various CC and BCC resolutions. It offers intuition that BCC 70% (relative to CC) is where visual comparability occurs. Note that the diagonal artifacts apparent in the ML images are due to the lighting direction; if the lighting direction were straight-on, the two diagonals would appear symmetric.

### 3.2.3 Camera View

Due to the nature of the sampling patterns used by CC and BCC, the position of the camera used to create an image can be a significant secondary property. For example, a CC sampling may be visually comparable to BCC when it is displayed using an axis-aligned view, but appear different from a second, non-aligned camera view. Therefore, we used two views of ML: one straight-on view and one from a tilted viewpoint (Figure 2a and Figure 2b). To keep the number of trials manageable within the one hour limit of our user experiment, a second view of the Fish Tail was not used.

### 3.3 Rendering Pipeline

We used a raycaster to render the images for this study. The raycaster uses an opaque transfer function in order to extract isosurface information from the volumetric data. The isosurfaces provide an appropriate level of detail from the volumetric data for the users to make comparisons. The isovalue was chosen so that meaningful isosurfaces with sufficient visual information were shown. Transparent volumetric renditions would clutter the images with too much detailed information and would make the judgement process unduly difficult.

### 3.4 Numerical Comparability of Reconstruction Techniques

A precondition in our hypothesis is that the reconstruction techniques we employ for CC and BCC are numerically comparable. Failing to meet this precondition may confound the independent variables. Therefore, we ensure that the precondition is met from two fronts.

#### 3.4.1 Theoretical Inference of Numerical Comparability

Our volume-rendered images are produced using asymptotically comparable BCC and CC reconstruction filters in the rendering pipeline. For CC reconstruction, we chose the widely used tri-cubic B-spline filter. The tri-cubic B-spline produces approximations to an original function  $f$  which are twice differentiable (i.e.  $C^2$ ). For BCC reconstruction, we chose the box spline proposed in [2], which also allows for a  $C^2$  reconstruction.  $C^2$  reconstruction is a typical requirement for volume rendering, since it results in a smooth reconstruction of the gradient field which is in turn used to compute the lighting in the rendering process.

The approximation order for both of the above-mentioned reconstruction methods is four [4]. This is an asymptotic measure of approximation. When we scale the sampling lattices (BCC or CC) by the scalar  $h < 1$  so that the sampling points become closer to each other, we take a higher number of samples from the original function. Therefore, our reconstructions  $\tilde{f}_{hBCC}$  and  $\tilde{f}_{hCC}$  provide a better approximation to the original function  $f$  as the  $L_2$  norm of the error in the approximation decays to zero. The approximation order determines the rate at which the error in approximation decays to zero. Given the original, unknown function  $f$ :

$$\|f - \tilde{f}_{hBCC}\| \in O(h^4), \quad (1)$$

$$\|f - \tilde{f}_{hCC}\| \in O(h^4). \quad (2)$$

Based on the asymptotic comparability of the chosen CC and BCC filters, we infer that the two filters are roughly numerically comparable.

#### 3.4.2 Heuristic Evaluation of Numerical Comparability

To validate our inference of numerical comparability of the filters, we assume that numerical comparability is primarily characterized by comparable  $L_2$  errors, and designed a heuristic accordingly. In

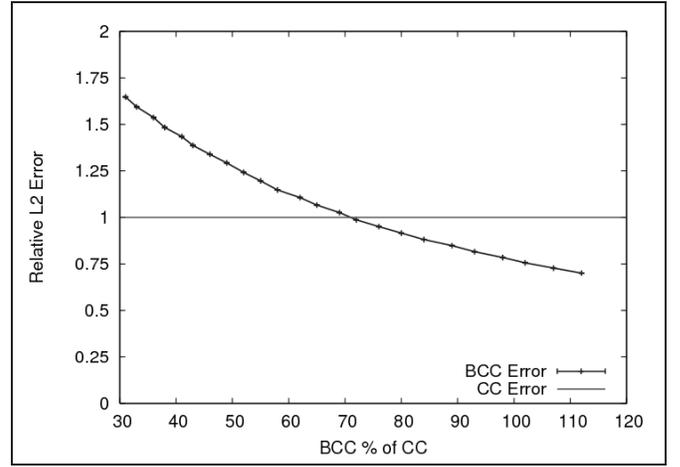


Figure 3: Relative  $L_2$  error for BCC and CC reconstruction methods when applied to ML samplings where the CC resolution is set at CC80. BCC resolution relative to CC is plotted along the x axis, while  $L_2$  error relative to CC is plotted along the y axis. Error bars indicate a 95% confidence interval for BCC. The 95% confidence interval for CC is too narrow to be plotted.

brief, we measure  $L_2$  errors of BCC samplings with respect to the ground truth signal, and the  $L_2$  errors of CC samplings with respect to the ground truth signal. We then determine at what relative BCC resolutions the  $L_2$  error becomes comparable to that of CC.

Computing the  $L_2$  error involves the integration of 3D functions. Analytic integration is complicated in this case, and typical deterministic quadrature rules are computationally expensive. Therefore, we use a stochastic approach by Monte-Carlo integration. We base each  $L_2$  computation on 10,000 points randomly positioned within the domain according to a uniform probability density function. Twenty of these  $L_2$  computations are evaluated for each dataset.

The mean  $L_2$  errors across the trials are plotted with respect to each signal under study (Figure 3 and 4). A 95% confidence interval (CI) is overlaid on each data point to show the variability of the Monte-Carlo results. The mean  $L_2$  errors, along with the 95% CI, give an accurate approximation to the range within which the true  $L_2$  errors are found.

The light grey (straight) lines in the plots denote the  $L_2$  errors recorded for CC samplings, whereas the dark grey lines denote the  $L_2$  errors recorded for the BCC samplings. The vertical axis has been normalized with respect to the CC  $L_2$  error and therefore demonstrate relative  $L_2$  errors. The point at which the light and dark grey curves intersect in each plot tells us the relative BCC sampling resolution at which the BCC  $L_2$  error is comparable to that of CC. However, due to the approximate ranges afforded by the 95% CI's, we need to consider the intersections of the CI's instead.

In the plot for ML with a CC resolution of CC80 (Figure 3), the CI's are so narrow that they are invisible given the limited resolution of the plot. The afore-described intersection occurs between BCC 69% and BCC 72%. This means that we are 95% confident that the point of numerical comparability for ML resides somewhere between BCC 69% and BCC 72%.

In the plot for Fish Tail with a CC resolution of CC140 (Figure 4), the CI's are wider than in the ML plot. Repeating the CI intersection analysis, we determine with a 95% certainty that the point of numerical comparability resides between BCC 69% and BCC 75%.

Since the sampling comparability result summarized in Section 2 suggests that CC and BCC are sampling-comparable at a relative

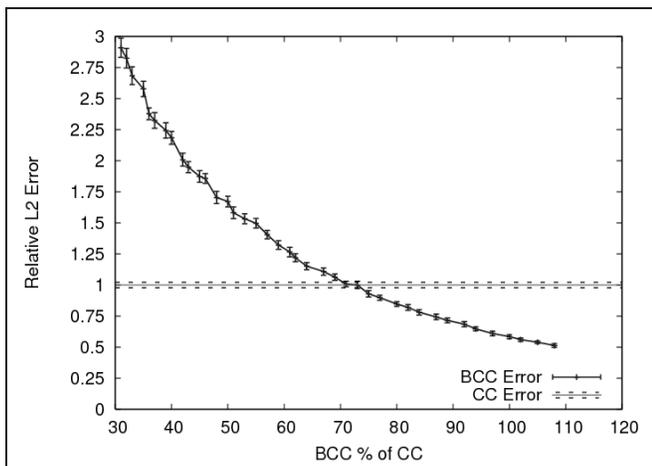


Figure 4: Relative  $L_2$  error for BCC and CC reconstruction methods when applied to Fish Tail samplings where the CC resolution is set at CC140. BCC resolution relative to CC is plotted along the x axis, while  $L_2$  error relative to CC is plotted along the y axis. Error bars indicate a 95% confidence interval for BCC. The gray dotted lines mark the 95% confidence interval for CC.

resolution of BCC 70%, we deduce that applying numerically comparable reconstruction filters would result in numerically comparable reconstructions at a relative BCC resolution of 70%. We adopt this result to the signals in our user study; since the ranges of comparability in both plots include the theoretical prediction of BCC 70%, we have a strong indication that the reconstruction filters we employ are indeed numerically comparable.

### 3.5 Experimental Setup and Administration

The experiment was conducted in a small room that was generally insulated from outside distractions. Participants sat at a desk with images displayed on an LCD screen. The experimental software ran on a desktop computer connected to the screen. During pilot runs, an illusion of motion was sometimes observed when images changed. To eliminate this effect, blank rectangular regions were displayed between trials for 0.25 seconds, and flashed over the spaces where the new images were to appear.

To minimize the possibility of influencing participant choices, the experiment administrators had minimal exposure to and minimal knowledge of the CC and BCC image generation process. Administrators assumed a passive role, and did not interact with participants, except to record participant comments regarding the image discrimination task.

### 3.6 Participant Selection

A total of 12 participants were selected. All were graduate students from the Computing Science department or the Engineering department of Simon Fraser University. Age and gender were not considered to affect a person's ability to detect relative differences in images, so no attempt was made to balance for these variables across experimental conditions.

We required that participants be unfamiliar with the process of BCC sampling and reconstruction, and not involved in the generation of CC or BCC images. Individuals with expertise in color science, medical imaging, and computer graphics in general were also excluded due to their expertise in visual quality and imaging artifacts. Participants were also required to have sufficiently good eyesight to allow them to perceive the signal reconstructions clearly.

This was evaluated by asking participants to read a short sentence displayed on the LCD screen prior to the experiment.

### 3.7 Expected Results

The primary dependent measure was the participant's choice of the CC or BCC sampling as being closer in appearance to the original image. During the study, participants were also asked to discuss the image features that motivated their choice.

We anticipated that if a given image pair exhibited visual comparability, a participant would choose BCC as the better image roughly 50% of the time. The region of visual comparability between BCC and CC should be characterized by regions of uncertainty and variability in participant preference. We expected these regions to occur near a relative resolution of BCC 70%. In addition, we expected that if the relative resolutions of BCC to CC were lower than 70%, CC would be the preferred choice, and for relative resolutions higher than the point of visual comparability, BCC would be the preferred choice.

In terms of secondary properties, we expected that different signals with different secondary properties would exhibit different ranges of visual comparability, but that these ranges would contain the ideal of BCC 70% described in our hypothesis. In particular, we anticipated that the ranges of visual comparability should be similar across the different settings for signal selection, sampling resolution, and camera view.

## 4 RESULTS AND DISCUSSION

### 4.1 Quantitative

We collected the image preferences for our 12 participants and separated them according to signal, CC resolution, and camera view. BCC resolution as a relative percentage of CC resolution was calculated as the ratio of the total number of samples on the BCC lattice over the total number of samples on the CC lattice, rounded to the nearest percentage.

Aggregate plots for all participants were created for data analysis (Figures 5 - 9). Mean user preferences for BCC resolutions are plotted with 95% confidence intervals. Individual plots of participant preferences were also examined to check for statistical outliers and ensure consistency with the aggregate analysis. No significant problems were identified in the results of any individual participant, and the results from all 12 participants are fully incorporated into the aggregate plots.

When the user preference confidence interval for a particular BCC resolution crosses the 50% user preference line, we say that visual comparability occurs at that particular BCC resolution with respect to the fixed CC resolution. We further define the range of visual comparability as all relative BCC resolutions falling between the two visually incomparable resolutions immediately before (after) the first (last) visually comparable resolutions.

Figure 5 and Figure 6 describe the results for ML View 1 and View 2, respectively. These plots show a general trend of increasing participant preference for BCC with respect to increasing relative resolution. The point at which participant preference passes the range of visual comparability is between 65% and 84% for View 1, and between 55% and 72% for View 2. Aggregating over both views for ML, Figure 7 indicates a range of comparability between 65% and 72%.

Visual comparability for the BCC-sampled Fish Tail signal falls between 62% and 71% for CC140 (Figure 8), and between 56% and 75% for CC180 (Figure 9). Variability in participant preference, as measured by the calculated confidence intervals, increases as relative resolution approaches the range of visual comparability. These results are in excellent agreement with our hypothesis that visual comparability occurs at a relative resolution of BCC 70%. Also in

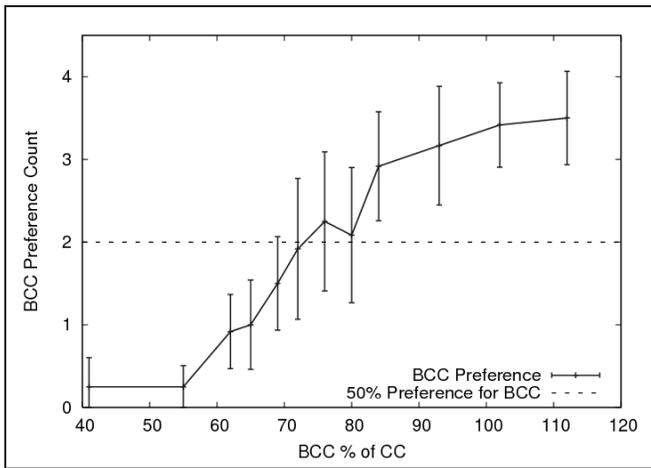


Figure 5: Mean BCC preference calculated from the 12 participant results for ML View 1, at a resolution of CC80. Error bars represent a 95% confidence interval about the mean.

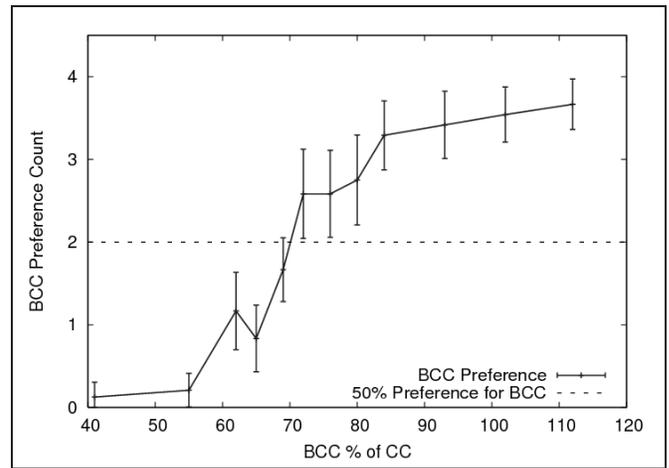


Figure 7: Mean BCC preference calculated from the 12 participant results for ML aggregated over views, at a resolution of CC80.

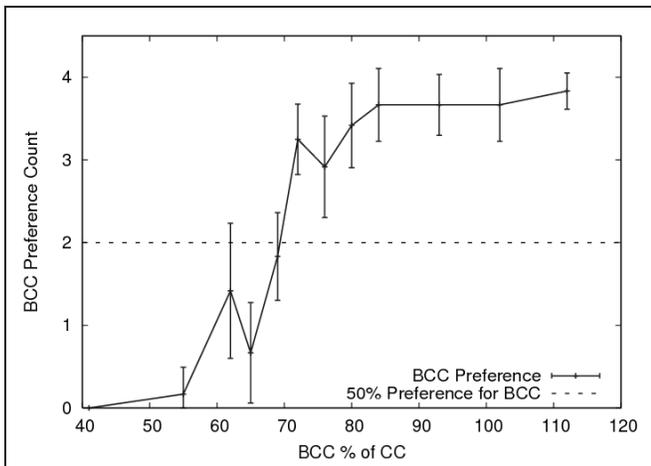


Figure 6: Mean BCC preference calculated from the 12 participant results for ML View 2, at a resolution of CC80.

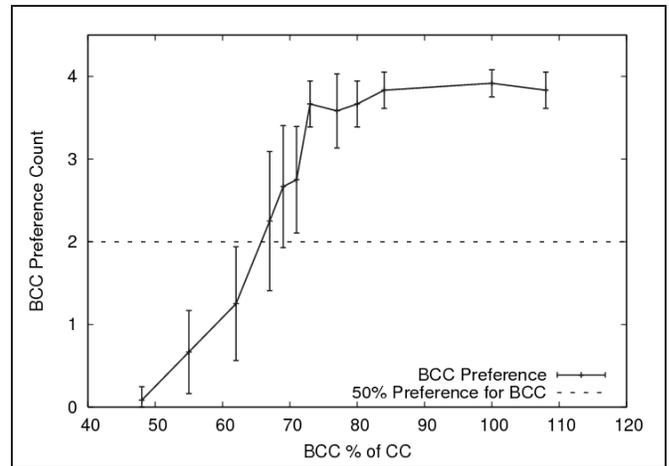


Figure 8: Mean BCC preference calculated from the 12 participant results for Fish Tail, at a resolution of CC140.

agreement with our expectations, variability in participant preferences decreases as relative resolution moves away from the range of visual comparability.

As discussed above, the plots for both ML views (Figure 5 and Figure 6) exhibit a general upward trend. Figure 6 exhibits one unexpected feature: a “dip” in preference around 65%. The dip is not statistically significant because the large confidence intervals in that region still allow us to fit a monotonically increasing user preference curve. However, a close investigation of the corresponding BCC images of ML revealed a wave interference pattern that is particularly pronounced at the resolution at which the dip occurred. A more advanced evaluation of this interference pattern as part of a more detailed user study could be future work.

Another discrepancy was observed with regards to camera view. For the ML signal, the ranges of visual comparability were determined to be between 65% to 84% for View 1 and between 55% and 72% for View 2. The two ranges do not appear to coincide. This indicates that camera views may have an effect on perceived visual quality. Further work is required to fully explore the significance of this disparity.

## 4.2 Qualitative

During the experiment, participants were encouraged to discuss the criteria they used to discriminate between images. Participants commented on a number of aspects of the discrimination task.

For both the ML and the Fish Tail signals, a few participants noted that the shading effects were different across different image pairings, but that they did not use this as the primary criterion for their discrimination. The majority of participants also remarked that the discrimination task was easier when Fish Tail images were presented.

While examining the ML images, participants remarked that they found curvature, symmetry, and the degree of distortion along edges to be important characteristics for comparison. One participant noted that the task could be difficult in some cases: sometimes the two images were clearly different, but neither image was “better” than the other.

Specific to the Fish Tail images, participants mentioned focusing on the ribs of the fish and evaluating their connectivity compared to the ground truth. In addition, participants commented that the shape of the fish’s larger solid fin sections provided a useful criterion.

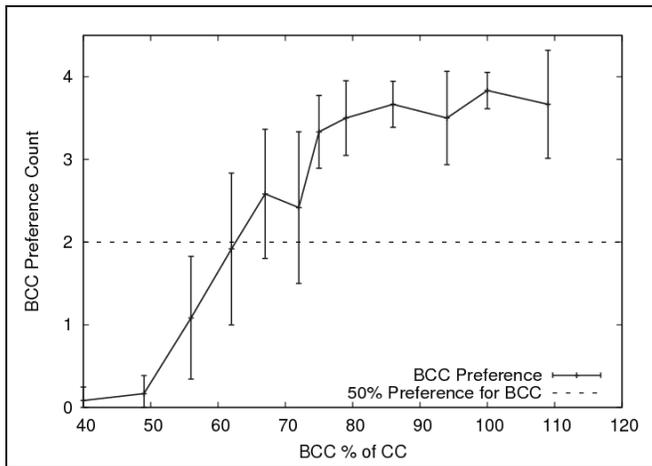


Figure 9: Mean BCC preference calculated from the 12 participant results for Fish Tail, at a resolution of CC180.

## 5 CONCLUSIONS AND FUTURE WORK

The results of this experiment provide strong evidence in support of our hypothesis that it takes the BCC lattice 30% fewer samples than the CC lattice to achieve comparable visual quality. For the conditions examined, BCC sampled data exhibit visual comparability to CC sampled data at a relative sampling density around the hypothesized 70%.

It is important to note that although the results here are encouraging, generalization from these results must be done carefully. Visual comparability, even for simple images, may be different for different applications. Domain specific applications must be examined on a case by case basis. The important features of volume data, and by extension the criteria for visual comparability, are determined by the application. We hope to have provided a guideline for those who wish to pursue domain specific extensions of our work.

It should also be noted that the parameter space of signals, sampling resolutions, and camera positions examined in this study was relatively small. It is our belief that in the general case, the results of this study are applicable. That is, we believe that the visual comparability results we obtained in this paper will hold so long as we employ BCC and CC reconstruction filters that are numerically comparable. Other parameters of the rendering pipeline such as lighting or transfer function are not believed to have a significant perceptual effect on the images rendered.

Due to the high number of parameters involved in a volume rendering process, our general belief encompasses a vast parameter space. The experiment described in this paper only investigated a small sub-region of this vast space. One direction for future work then, is to continue exploring this vast space of parameters. For example, the effect of camera placement could be more thoroughly investigated. As another example, the effect of varying iso-values could be evaluated. As the research community develops a better understanding of this parameter space, a metric may be developed that predicts the region of BCC/CC visual comparability for a given signal. Until then, the  $L_2$  error analysis presented in Section 3.4.2 may act as a reliable heuristic.

Based on this user study and previous theoretical and algorithmic results [2, 4], we now have evidence that BCC sampling is substantially more accurate, more computationally efficient, and perceptually better than the traditionally popular Cartesian sampling technique. This study is a milestone, completing the argument that BCC-sampled data is preferable in every aspect relevant to volume rendering. If these perceptual results are confirmed for a variety of

other signals and rendering pipelines, it will make a strong case for BCC as the sampling method of choice for volume rendering.

## 6 ACKNOWLEDGMENT

The research in this paper has been supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada and the Canadian Foundation for Innovation (CFI).

## REFERENCES

- [1] A. S. Bair, D. House, and C. Ware. Texturing of layered surfaces for optimal viewing. *IEEE Transactions on Visualization and Computer Graphics (Proceedings of IEEE Visualization 2006)*, 12(5):1125–1132, 2006.
- [2] A. Entezari, R. Dyer, and T. Möller. Linear and cubic box splines for the body centered cubic lattice. In *Proceedings of IEEE Visualization 2004*, pages 11–18, 2004.
- [3] A. Entezari, T. Meng, S. Bergner, and T. Möller. A granular three dimensional multiresolution transform. In *Proceedings of Eurographics/IEEE-VGTC Symposium on Visualization 2006*, pages 267–274, 2006.
- [4] A. Entezari, D. V. D. Ville, and T. Möller. Practical box splines for volume rendering on the body centered cubic lattice. *Submitted to: IEEE Transactions on Visualization and Computer Graphics*, 2006.
- [5] T. Hales. Cannonballs and honeycombs. *Notices of the AMS*, 47(4):440–449, Apr. 2000.
- [6] V. Interrante, H. Fuchs, and S. M. Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):98–117, 1997.
- [7] D. H. Laidlaw, R. M. Kirby, C. Jackson, J. S. Davidson, T. Miller, M. DaSilva, W. Warren, and M. Tarr. Comparing 2D vector field visualization methods: A user study. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):59–70, 2005.
- [8] S. R. Marschner and R. J. Lobb. An evaluation of reconstruction filters for volume rendering. In *Proceedings of IEEE Visualization 1994*, pages 100–107, 1994.
- [9] P. Marziliano, F. Dufaux, S. Winkler, and T. Ebrahimi. Perceptual blur and ringing metrics: Application to JPEG2000. *Signal Processing: Image Communication*, 19:163–172, 2004.
- [10] M. Meißner, J. Huang, D. Bartz, K. Mueller, and R. Crawfis. A practical evaluation of popular volume rendering algorithms. In *Proceedings of the 2000 IEEE Symposium on Volume Visualization*, pages 81–90, 2000.
- [11] D. P. Mitchell and A. N. Netravali. Reconstruction filters in computer graphics. *Computer Graphics (ACM SIGGRAPH 1988 Proceedings)*, 22(4):221–228, 1988.
- [12] G. Sweet and C. Ware. View direction, surface orientation and texture orientation for perception of surface shape. In *Proceedings of Graphics Interface 2004*, pages 97–106, 2004.

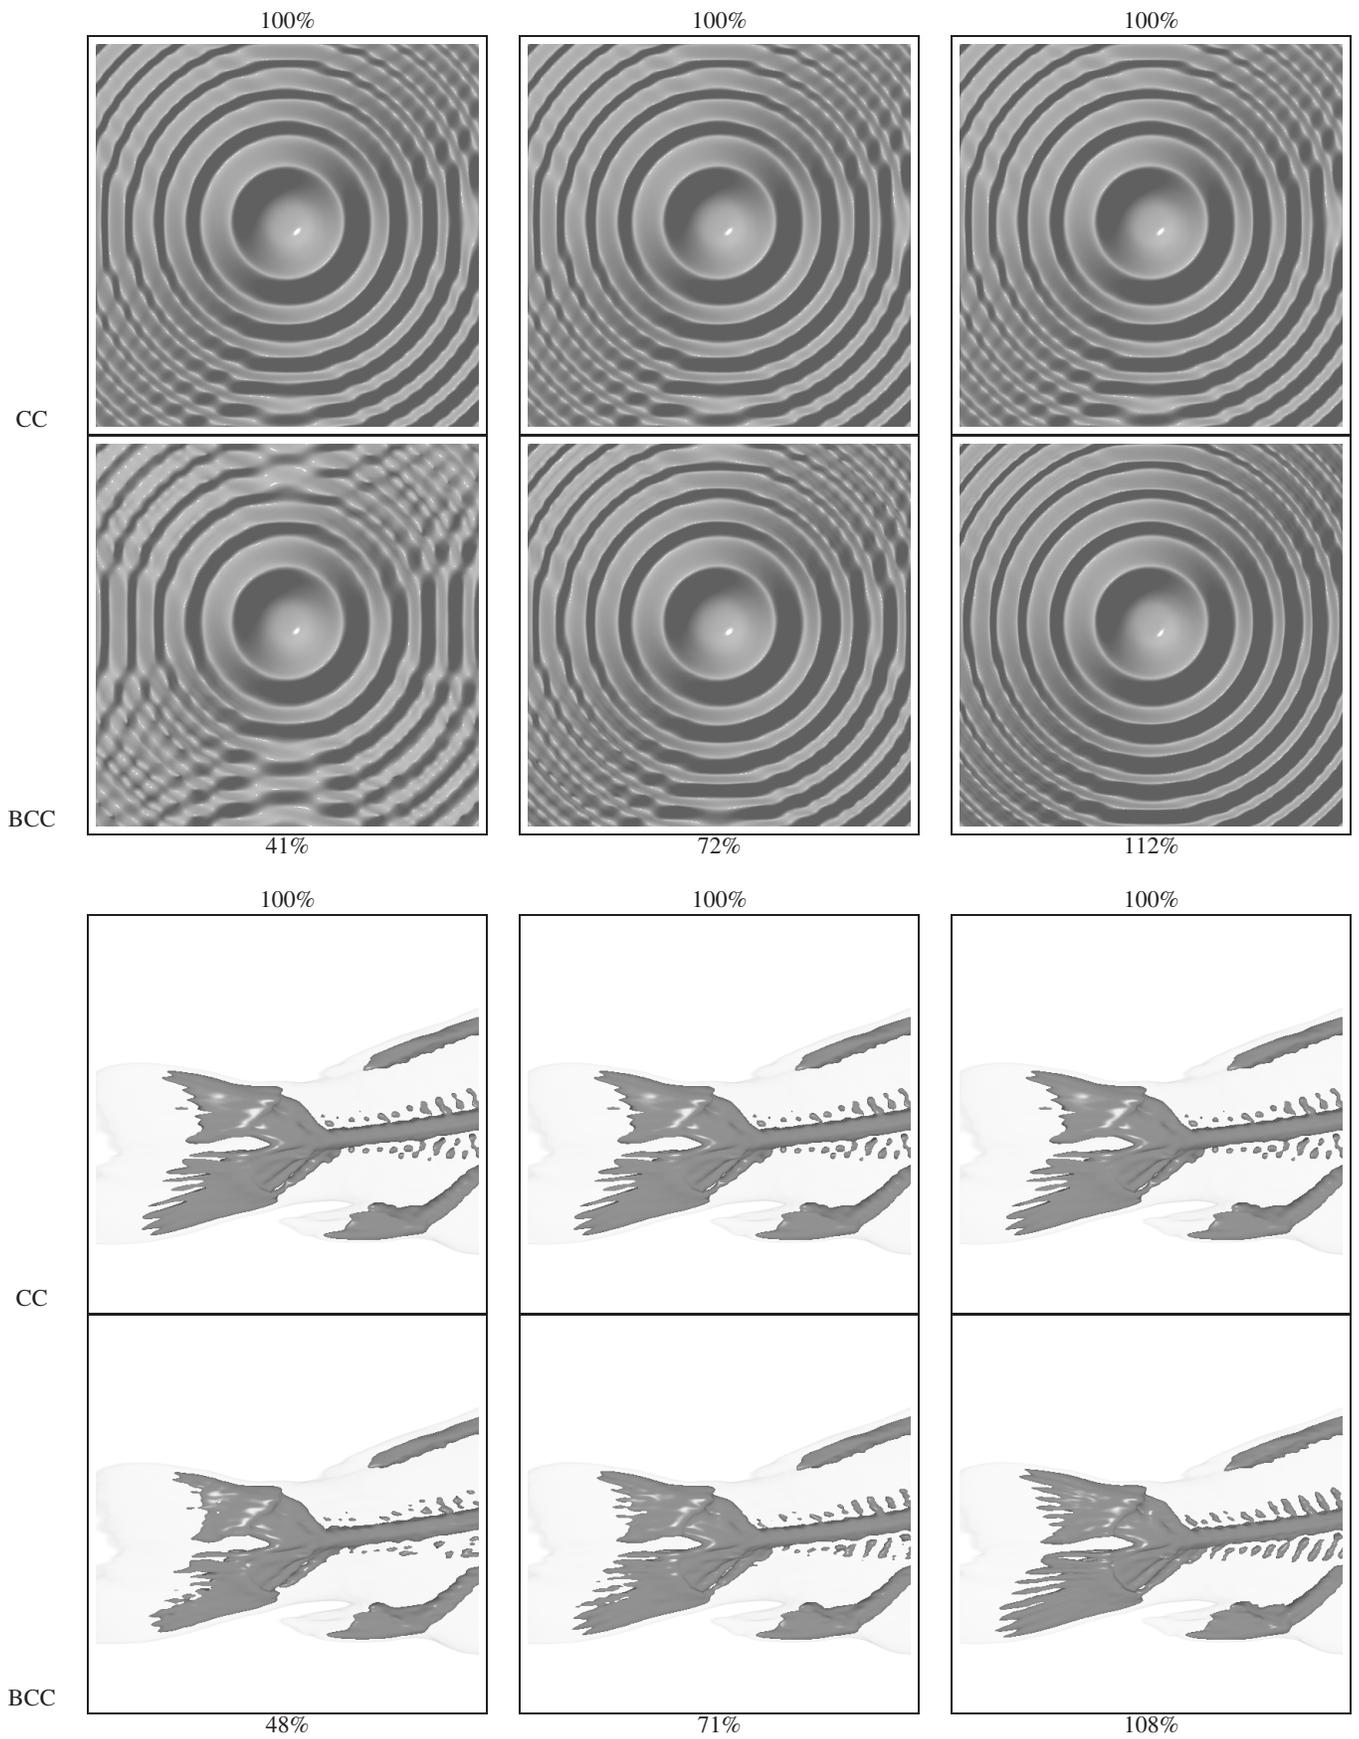


Figure 10: Illustration of ML and Fish Tail signals rendered at various resolutions. In each table, resolution percentages are relative to a fixed CC resolution. The CC image is duplicated three times in the top rows of each table. The tables offer intuition that BCC 70% (relative to CC) is where visual comparability occurs for both the ML and the Fish Tail (compare the middle columns to the other columns).