Volume Denoising for Visualizing Refraction

David Rodgman Min Chen
University of Wales Swansea, UK *

Abstract

Rendering refraction in volume visualization requires smoothly distributed normals to synthesize good quality visual representations. Such refractive visualizations are more susceptible to noise in the data than visualizations that do not involve refraction. In this paper, we present the first study to investigate convolution and diffusion methods with respect to this problem, highlighting the need for improving the smoothness of dataset normals while minimizing unnecessary geometrical distortion. We outline the mathematical specifications of four volume denoising filters, namely (a) *Gaussian convolution*, (b) *inhomogeneous linear diffusion*, (c) *anisotropic nonlinear diffusion* and (d) *regularized anisotropic nonlinear diffusion*. We apply these filters to volumetric scenes involving refractive objects, and introduce two metrics for measuring distortion and smoothing effects. Our results indicate that it is possible to make significant improvements in image quality in refractive visualization without introducing excessive distortion.

Keywords: Volume visualization, discrete ray tracing, refraction, denoising, smoothing, nonlinear diffusion.

1 Introduction

Refraction is all around us. Without it we would not be able to see. The failure of the lenses in our eyes to refract incoming light would render the world an indecipherable blur. Externally, we also see refractive objects everywhere. Our vision system is accustomed to the visual effects caused by refraction, and is able to benefit from the visual effects present in refractive scenes by acquiring additional visual cues from these effects.

Figure 1 shows a volumetric refractive scene, which is rendered using ray casting (a) without refraction, and (b) with refraction assuming uniform refractive index for the ”water” block. Although (a) may be able to fulfil a basic illustrative function in many circumstances, it offers viewers an incorrect and unrealistic visualization. (b) gives a typical refractive visualization which can be achieved by many surface rendering systems, a small number of which are also able to simulate heterogeneous refractive materials using attribute mapping techniques. It can be seen that (b) gives greatly improved perception of depth (of the rod) and shape (of the ”water”). In order to keep up with the rapid advances in volumetric data capture and volumetric scientific simulation, Rodgman and Chen recently proposed and studied a collection of methods for rendering refraction in discrete ray tracing [Rodgman and Chen 2001]. By associating a volume object with a refractive index field, one can use discrete ray tracing to directly render complex refractive conditions. This development is particularly important to volume visualization where the intermixing of translucent and opaque objects is commonplace.

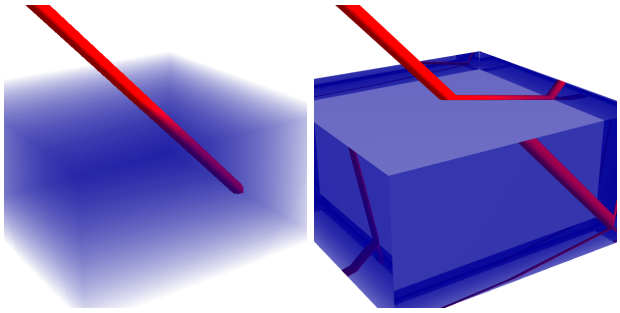

(a) (b)

Figure 1: A water container and a rod, rendered (a) without refraction, (b) with refraction.

In order to produce high quality images featuring volume objects which are refractive, we must concern ourselves with the quality of the datasets from which the image is synthesized. For these purposes, mathematically defined scalar fields provide unsurpassed quality: they are free from both noise and discretization errors, and therefore allow us to analytically derive precise normals. Consequently, the images produced are of excellent quality.

The situation is somewhat different, however, with discrete datasets. Even if the dataset has been sampled above the Nyquist limit, theoretically permitting a perfect reconstruction of the original signal, in practice aliasing errors are hard to avoid. This is particularly the case where a low rendering time is desirable, since it imposes limits on the complexity and degree of the interpolation function. Further, if the dataset has been acquired from a real world source, such as an MRI scanner or similar, the introduction of noise in the scanning process poses considerable problems specific to refractive volume rendering, as discussed in [Rodgman and Chen 2001].

As Möller *et al* point out in [Möller et al. 1997a; Möller et al. 1997b], the quality of the normals has a greater effect on image quality than the quality of the data itself. They were referring only to image quality in the context of opaque, non-refractive objects, but this observation is even more applicable to transparent, refractive objects. This is because when a ray refracts, its direction after refraction is directly governed by the normal at the point of refraction. Consequently, an error in the normal estimation process affects not only the shading at that point (as is the case where refraction is not involved), but also the location of every subsequent sample point. A means of obtaining high quality normals is therefore critical to good image quality with refractive volume rendering.

Errors in the normals may be attributed to three main sources. Reconstruction errors occur when a dataset is under-sampled (pre-aliasing), or when the reconstruction method used is inadequate (post-aliasing). Additionally, the scanning process typically introduces noise into the data. Post-aliasing may be reduced to acceptable levels by use of higher order filters; this problem has been

* e-mail: {csdave, m.chen}@swansea.ac.uk

thoroughly addressed in the literature [Theußl 2000; Marschner and Lobb 1994] and will not be discussed further here. Pre-aliasing and noise in the dataset, however, have been dealt with in general less satisfactorily, and not at all in the specific context of refraction.

In addition to issues caused by pre-aliasing errors and noise in the data, an excess of fine detail in the dataset can also reduce the effectiveness of refraction in delivering depth cues and other information about the shape of the volume object to the viewer, with similar results to noisy data. We demonstrate these problems with Figure 2, where a CT head dataset is visualized in conjunction with a brain dataset. In Figure 2(a) which was was rendered using volume ray casting without refraction, the CT head has its top part removed and is shown as a translucent object. Whilst it clearly displays the internal structure of the brain, it offers limited depth cues, and poor information about the spatial relationship of the brain and the skull. Figure 2(b), which results from direct surface rendering, provides slightly better depth cues, but incorrect visual representation of the spatial relationship.

Figure 2(c), also rendered using volume ray casting, introduces refraction into the visualization, which permits much better depth perception and appreciation of geometrical structure. However, it introduces at the same time several artefacts, including scattered, noisy specular highlights, large regions of noise attributable to incoherent refraction, and a cracking effect directly above the brain stem, all attributable to "noisy" refraction. In Figure 2(d), a denoising filter is applied to the CT dataset prior to the rendering; as a result, the smoothed head dataset facilitates a much clearer visualisation of internal structure (i.e., the brain). Moreover, it gives a considerably less noisy impression of the neck area, it conveys a sense of curvature via the refractive patterns in the "glass" which is not present in Figure 2(c), and it correctly shows the brain stem as an unbroken section of tissue. Artefacts such as scattered, noisy specular highlights present in the unprocessed data are also eliminated in Figure 2(d).

In order to address these problems, we begin by demonstrating their effects in Section 3, after a brief review of related work. In Section 4 we propose a number of different methods for "improving" normals; specifically convolution and non-linear diffusion, and in Section 5 we develop a means of measuring the effects of any dataset-processing technique on both image quality and accuracy. The denoising methods are tested and evaluated according to our metrics in Section 6; and finally, we make some concluding remarks in Section 7.

2 Related Work

There are a diversity of modalities, such as CT, MRI and ultrasound, for the acquisition of volume datasets. The aliasing and noise introduced during digitization is of concern to many algorithms for processing and rendering such datasets. Much effort has been placed on the correct reconstruction of normals; a number of normal estimation methods have been studied and compared by Yagel et al [Yagel et al. 1992] and Möller *et al* [Möller et al. 1997a]. Recent work by Möller *et al* [Möller et al. 1998], Neumann *et al* [Neumann et al. 2000], Persoon *et al* [Persoon et al. 2003] and Rössl *et al* [Rössl et al. 2003] represents attempts to reconstruct accurately the scalar function associated with a volume dataset and its gradient.

In image processing and computer vision, we commonly see a different approach to the problem of noise, and a variety of filters were designed for smoothing or denoising images [Lee 1980; Kuan et al. 1985; Perona and Malik 1990; Weickert 1998]. In recent years, some of these filters have been successfully generalized for smoothing surfaces. For example, Peng *et al* employed Wiener filtering for smoothing noisy triangular meshes [Peng et al. 2001]. Desbrun *et al* employed anisotropic diffusion for denoising height

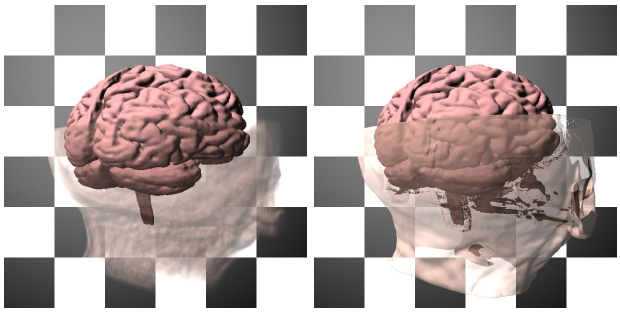

(a) volume rendering (b) surface rendering

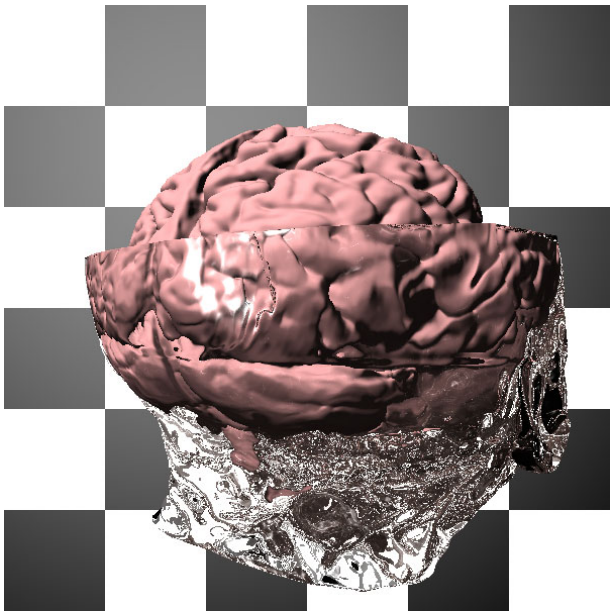

(c) the original CT head with refraction

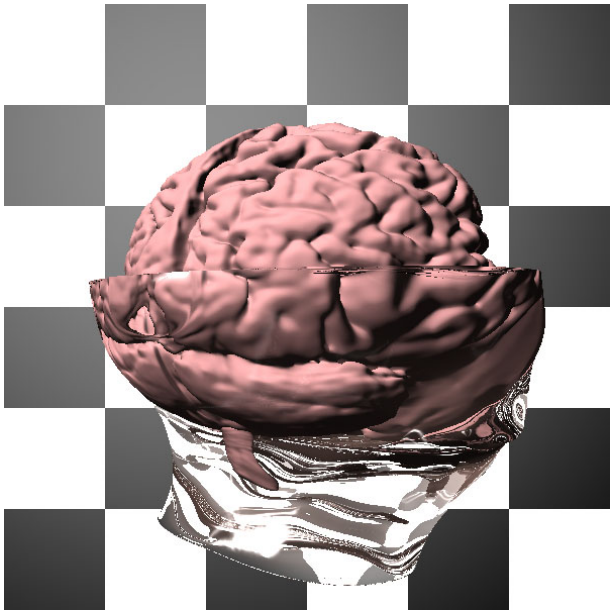

(d) after denoising using a diffusion method

Figure 2: Combinational visualization of a CT head dataset and a brain dataset. (a) and (b): the original CT head is rendered without refraction. (c): the original CT head is rendered with refraction. (d): the CT head is rendered after denoising.

functions and bivariate data [Desbrun et al. 2000]. Tasdizen *et al* used an anisotropic diffusion method for smoothing surfaces via normal maps [Tasdizen et al. 2002]. In addition, there was much other effort in surface fairing, smoothing and denoising [Moreton and Sequin 1992; Welch and Witkin 1994; Whitaker 1994; Desbrun et al. 1999; Clarenz et al. 2000; Ohtake et al. 2000; Sapiro 2001; Guskov and Wood 2001].

In terms of volume visualization, surface denoising can be applied to a particular iso-surface contained in a volume data set, but is not suitable for direct volume rendering. There have been some attempts in direct volume smoothing. Hilton *et al* proposed two wavelet-based noise removal algorithms for denoising MRI data [Hilton et al. 1996], and Angelin *et al* gave a brief description of a 3D implementation of brushlet basis functions for Fourier domain denoising [Angelini et al. 1999]. The former was a study based on statistical analysis and the latter was an investigation in the context of segmentation. Recently Bertram proposed an iterative fairing method for smoothing contours inside volume datasets using bicubic parametric functions [Bertram 2003].

It is necessary to investigate the effectiveness of denoising methods in the context of direct volume rendering. The successful deployment of nonlinear diffusion methods in surface denoising offers a good starting point for volume denoising, and the demand for smooth gradient in visualizing refraction [Rodgman and Chen 2001] makes it a suitable case study for evaluating the effectiveness of volume denoising methods.

3 Motivation

We demonstrate the importance of having good quality normals for refraction by means of the following test case. Figure 3(a, b) shows an opaque sphere rendered with analytically derived normals, and with randomly perturbed normals[1]. The same experiment is repeated with a translucent sphere in Figure 3(c, d).

It can be seen by visual inspection that image quality is better preserved in (b), in the absence of refraction, than in (d); this observation is endorsed by a comparison of the root mean square (RMS) difference between the opaque image pair, and the transparent image pair, which have RMS difference equal to 39.1% and 99.0% respectively.

Improving dataset normals does not necessarily connote greater accuracy; rather, our goal is to manipulate the dataset normals in such a way that we may produce refractive visualizations of the dataset which use refraction to convey three dimensional cues such as depth perception in a meaningful way, without overwhelming the viewer with noisy, over-complicated images. In some senses, we wish to simplify the dataset, and smooth the appearance of the object in question somewhat, in order to improve the subjective image quality.

4 Denoising Methods

In this section we discuss a number of methods for denoising datasets. We consider the dataset to have resolution $R = (R_1, R_2, R_3) \in \mathbb{N}^3$. We thus denote the unprocessed dataset by $F(\boldsymbol{x})$ where $\boldsymbol{x} = (x_1, x_2, x_3) \in \mathbb{N}^3$, such that $0 \leq x_i < R_i$ for $i = 1, 2, 3$, and the processed dataset derived from F by $F'(\boldsymbol{x})$. We also define X to be the set of all such $\boldsymbol{x}$.

[1] Each normal is perturbed by a random angle θ, selected such that θ has an average value of 20° and a normal (Gaussian) distribution.

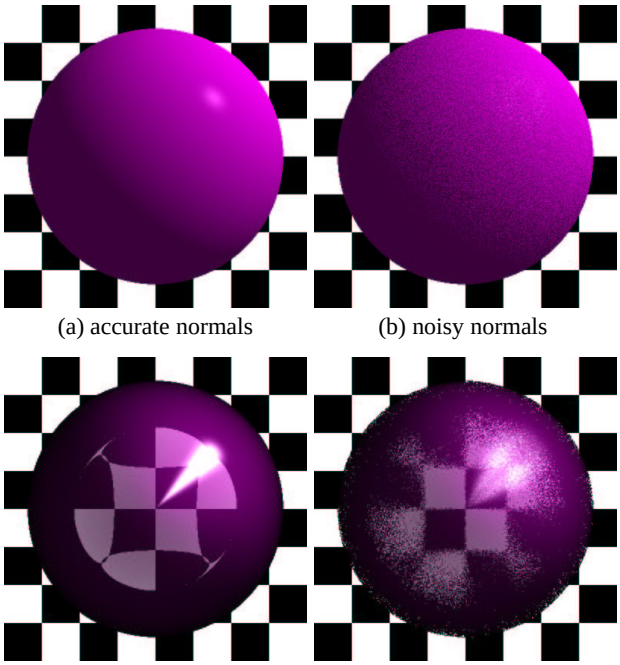

(a) accurate normals (b) noisy normals

(c) accurate normals (d) noisy normals

Figure 3: An opaque sphere and a translucent sphere rendered, respectively, with (a, c) accurate normals, and (b, d) randomly perturbed normals.

4.1 Convolution with a Gaussian filter

A simple method for removing high frequencies in the dataset is to convolve with an ideal low pass filter (the sinc function); that is, to remove all frequencies above a certain threshold. This method is not realizable in practice, since the sinc function is of infinite width; it follows that convolving the sinc function with a dataset of finite width is equivalent to convolving with a truncated sinc function, which gives poor performance in the stop band, and leads to ringing effects (Gibb's phenomenon) [Gibbs 1899].

For this reason, we do not discuss use of an ideal low pass filter further, but instead examine a low pass filter which has better performance in the stop band.

By using a filter which has better stop-band performance, we may reduce the artefacts noted above. The natural choice here is the Gaussian filter, $K_\sigma(\boldsymbol{x})$, defined as

$$k_\sigma(y) = \frac{1}{2\pi\sigma^2} \cdot e^{\left(-\frac{y^2}{2\sigma^2}\right)} \quad (1)$$

$$K_\sigma(\boldsymbol{x}) = k_\sigma(x_1) \cdot k_\sigma(x_2) \cdot k_\sigma(x_3) \quad (2)$$

where σ denotes the standard deviation of the Gaussian filter; higher values cause increased reduction of high frequencies. Whilst in theory this function has infinite width, thus rendering it susceptible to the artefacts previously observed with the ideal filter, in practice it closely approaches zero sufficiently quickly for this issue not to arise. We thus define $F'(\boldsymbol{x})$ by convolution with K_σ:

$$F'(\boldsymbol{x}) = (K_\sigma * F)(\boldsymbol{x}). \quad (3)$$

By selecting different values for σ, we may control the degree of smoothing, as demonstrated in the first row of Figures 4 and 5. These images, and in particular the opaque visualizations in Figure 4(a)-(d), show that the *Gaussian convolution* approach is unable to preserve fine detail. This is not surprising, since this operation

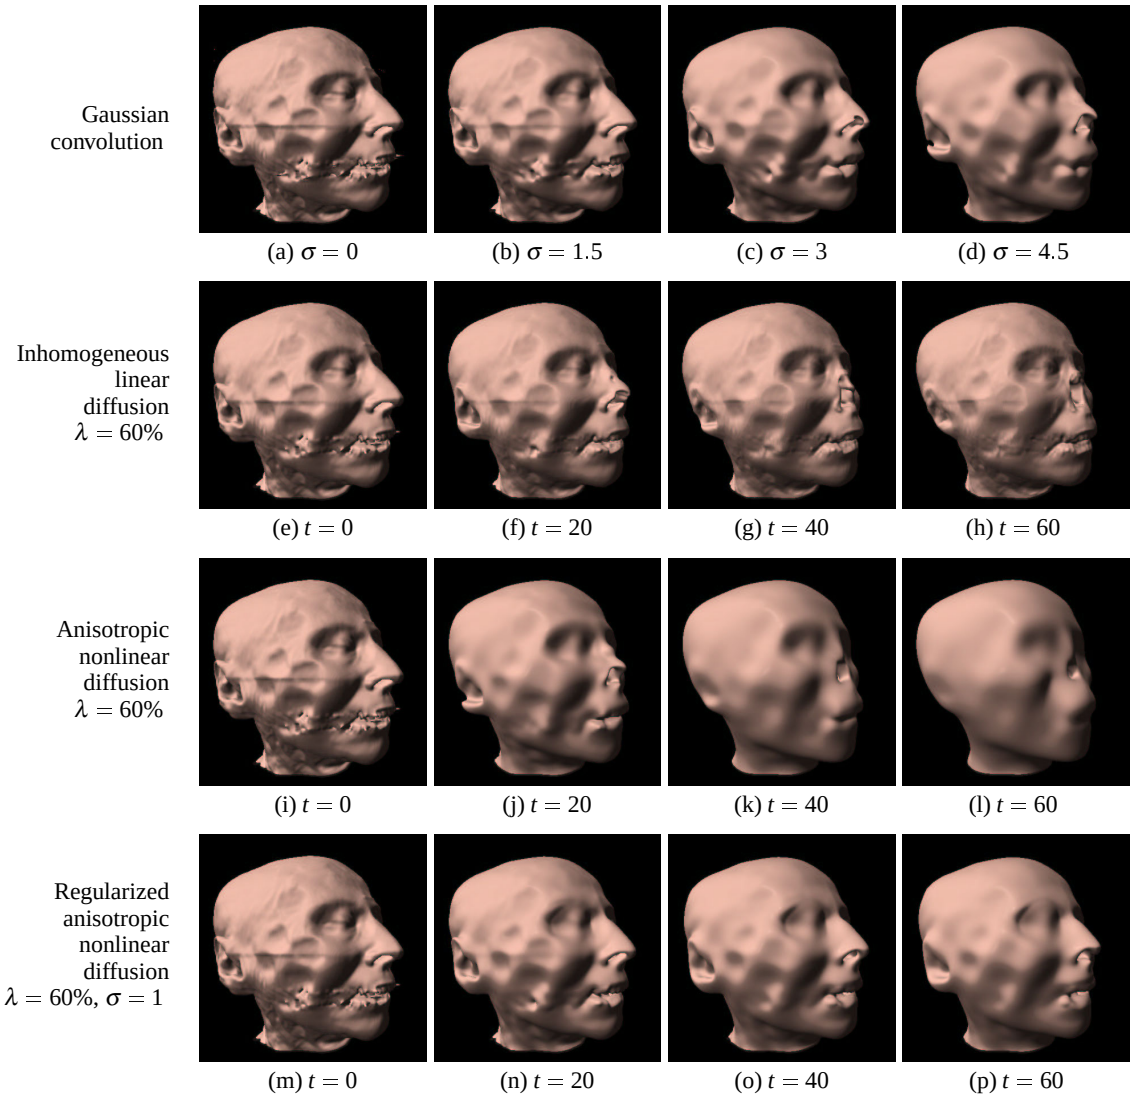

Figure 4: Results of applying different denoising methods to an opaque object.

has no ability to locally adapt to the dataset. Consequently, there is some degradation of the overall shape of the head, and fine structures such as the nose are especially adversely affected.

Comparing this degradation to the improved, "cleaner" appearance of the refractive visualizations in Figure 5(a)-(d) clearly shows that this method, whilst certainly effective at removing noise from the dataset, offers a tradeoff between noise reduction and shape preservation.

These results appear to confirm that the source of the noise in the refractive visualization is indeed mostly high frequency noise in the original dataset, since application of a low-pass Gaussian filter is able to reduce these visualization artefacts. However, the shape degradation present when high frequencies are filtered out indicates that this approach has significant drawbacks.

4.2 Inhomogeneous Linear Diffusion

A significant drawback of convolution with the Gaussian filter is that this method does not permit any local adaption to the dataset; it is not possible to inhibit smoothing in non-boundary regions, or to control the orientation or degree of smoothing at boundaries. *Inhomogeneous linear diffusion* [Weickert 1997] offers more control in this respect.

Firstly, we define the *scale-space* as $U(\boldsymbol{x},t)$, where t is the time. We consider the scale-space to be evolving such that at time $t = 0$, $U(\boldsymbol{x},t) = F(\boldsymbol{x})$, and so that as t increases, U represents a more heavily processed version of F.

We begin by defining[2] the general nonlinear diffusion equation in terms of the rate of change of $U(\boldsymbol{x},t)$ with respect to t:

$$\frac{dU}{dt} = \nabla \cdot (D \cdot \nabla U). \tag{4}$$

Here, $D \cdot \nabla U$ is the *flux*, which is equal to the *diffusion tensor* D, a positive definite symmetric matrix, multiplied by the normal of U. This flux can be thought of as representing the flow of "particles" from a region of high density to a region of low density.

When D is a constant, this process is said to be linear and homogeneous, and is in fact equivalent to convolution with a Gaussian

[2] In a small abuse of notation, ∇U refers to the gradient of the scalar field U, whereas $\nabla \cdot V$ refers to the divergence of the vector V.

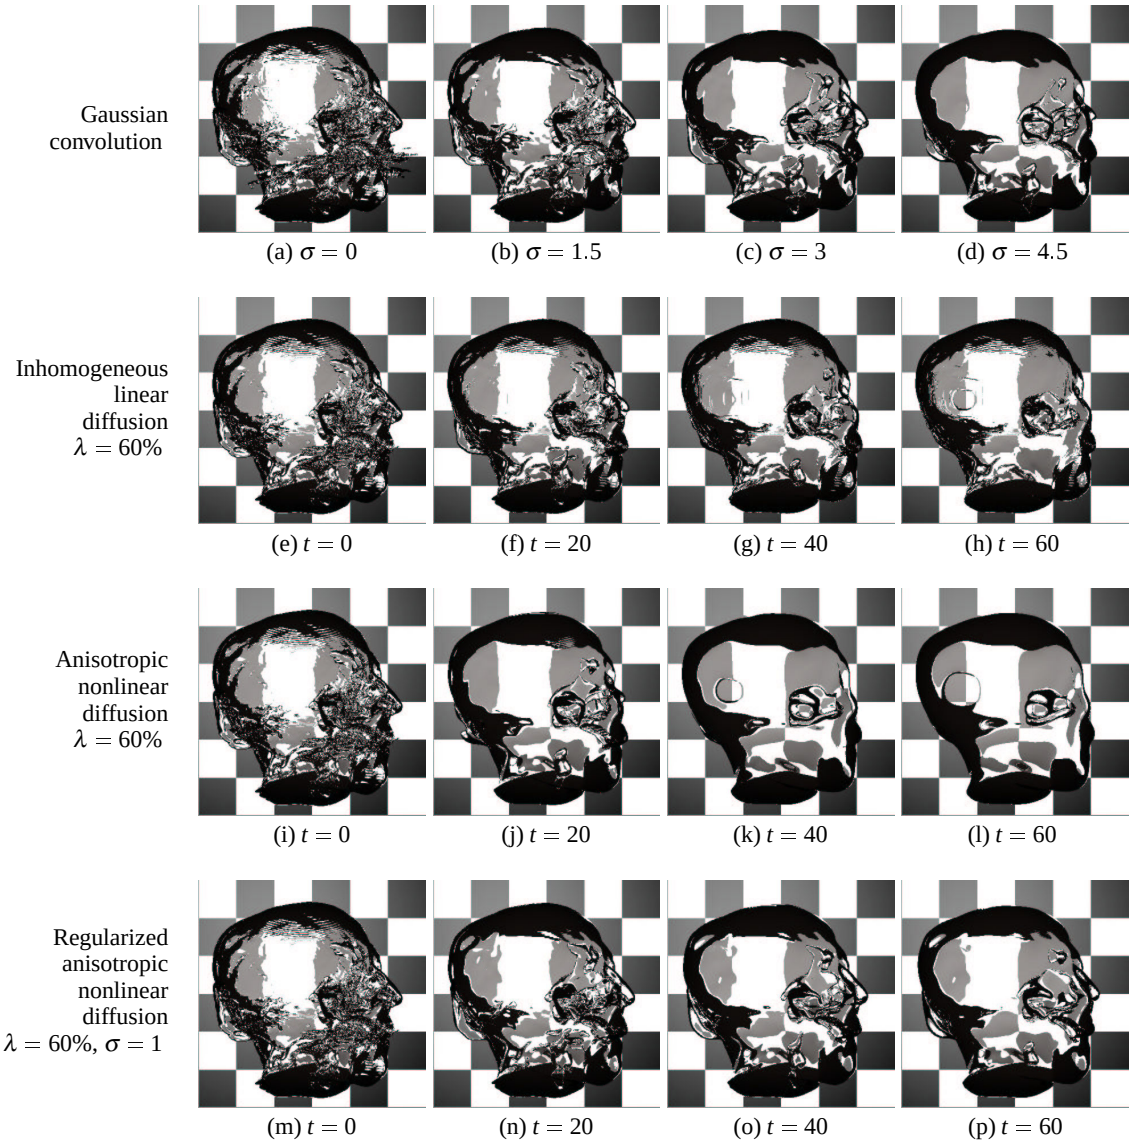

Figure 5: Results of applying different denoising methods to a translucent object.

filter, as discussed in Section 4.1. In the inhomogeneous case, however, D is a function of $\boldsymbol{x}$; typically, of $F(\boldsymbol{x})$.

Here we discuss isotropic *inhomogeneous linear diffusion*; that is, where the flux is always parallel to ∇F, and we therefore replace the diffusion tensor D with a scalar-valued *diffusivity* function, g. There are many possible choices for g; here we use the Perona-Malik diffusivity function,

$$g(|\nabla F|^2) = \frac{1}{\sqrt{1+|\nabla F|^2/\lambda^2}} \qquad (\lambda \neq 0). \tag{5}$$

This results in the following diffusion equation:

$$\frac{dU}{dt} = \nabla \cdot \left(g(|\nabla F|^2) \cdot \nabla U \right). \tag{6}$$

That is, at any given point $\boldsymbol{x}$ in the dataset, there is a flow of density along the direction of the original dataset normal, ∇F, the magnitude of which is proportional to the diffusivity, $g(|\nabla F|^2)$.

This diffusivity function inhibits diffusion as $|\nabla F|^2$ increases. The effect is to better maintain well-defined surfaces (by inhibiting smoothing for large values of $|\nabla F|^2$), whilst permitting smoothing in inner regions.

Some example results of applying *inhomogeneous linear diffusion* to the CT head are given in the second row of Figures 4 and 5. The effectiveness of such a process is of course strongly dependant on the value chosen for λ. Selecting suitable values for λ is not straightforward, and is discussed in detail in Section 4.5.

4.3 Anisotropic Nonlinear Diffusion

In order to preserve boundaries whilst permitting smoothing in these regions, it is necessary to use an anisotropic method [Weickert 1998]. This method derives its non-linearity from the fact that g is a function of $U(\boldsymbol{x},t)$. We therefore have the diffusion equation

$$\frac{dU}{dt} = \nabla \cdot \left(g(|\nabla U|^2) \cdot \nabla U \right). \tag{7}$$

This has the effect of allowing the flow to increase for large values of t at the location of small structures that have been smoothed at

lower values of t. We might therefore expect more smoothing at the location of these small structures for sufficiently large values of t.

Some results of applying *anisotropic nonlinear diffusion* to the CT head are given in the third row of Figures 4 and 5. In a similar manner to *inhomogeneous linear diffusion*, this method preserves surfaces by inhibiting smoothing at boundaries. However, for the purposes of refraction, it is the normals at the boundaries that are most critical, and where smoothing is required most; whilst we might expect these two methods to preserve fine detail and overall shape better than *Gaussian convolution*, the theoretical basis for improving the smoothness of normals in boundary regions is weak.

4.4 Regularized Anisotropic Nonlinear Diffusion

A significant problem with *anisotropic nonlinear diffusion* is that it gives rise to so-called *staircasing artefacts*, as shown in Figure 6. These artefacts manifest themselves as small, regularly spaced discontinuities in the processed dataset. Whilst they do not significantly affect the overall shape of the processed dataset, they have a substantial effect on the normals of the processed dataset, which are critical to obtaining high quality refractive images.

Figure 6: Staircasing effect with *anisotropic nonlinear diffusion*.

In order to improve the performance of *anisotropic nonlinear diffusion*, we make use of a *regularization* step [Alvarez et al. 1992] to control the diffusion process, giving rise to the following diffusion equation:

$$\frac{dU}{dt} = \nabla \cdot \left(g(|\nabla U_\sigma|^2) \cdot \nabla U \right) \tag{8}$$

where we replace U in Equation 7 with U_σ, which is the convolution of U with a Gaussian filter with fixed $\sigma > 0$ as discussed in Section 4.1. In other words, we have $U_\sigma(\boldsymbol{x},t) = (U * K_\sigma)(\boldsymbol{x},t)$

The effect of this *regularization step* is to improve the stability of the diffusion process. In particular it improves the performance of the diffusion process in the presence of noise, and dramatically reduces staircasing artefacts. It is therefore well suited to applications such as refractive visualization, where the smoothness of the normals is critical. Some results of applying *regularized anisotropic nonlinear diffusion* to the CT head are given in the fourth row of Figures 4 and 5.

4.5 Determining a Value for Lambda

The choice of a value for λ strongly determines the effects of the diffusivity function. However, fine-tuning values for λ by hand is a time consuming trial and error process. Specifically, very low values for λ result in smoothing being suppressed everywhere, and the resulting images do not convey the benefits of having been smoothed. In the case where λ is set too high, smoothing is not sufficiently inhibited at surfaces, and the result is that the surfaces, as well as the internals of the dataset, become blurred; distortion is therefore increased.

A useful approach for structuring the choice of λ somewhat is to compute all the normals in the dataset, sort them (in ascending order) by magnitude, discarding all zero-length normals, and then set λ equal to the magnitude of the normal at the nth percentile, where n is controlled by the user[3]. Selecting a suitable percentile for λ is discussed further in Section 6.

5 Metrics for Measuring Distortion and Smoothing Effects

In order to properly compare these different methods for improving refractive image quality, it is necessary to devise an objective scheme for quantifying the ways in which both the data and the images are affected. Since in each method, the smoothing process represents a tradeoff between smoothing and volume shape degradation, we separate this task into two separate problems: evaluating the degree to which the overall shape of a volume is affected, and quantifying the improvement in image quality. We address these two problems in the following two sections, and proceed to discuss how these schemes may be used to interpret the efficacy of the various methods in Section 4.

5.1 Distortion Metric

In order to evaluate the level of distortion present, we must first introduce a *transfer function* τ, in the form of a lookup table. This corresponds to the transfer function responsible for extracting the data of interest during visualization.

Our motivation in using the transfer function used to visualize the data is to enable us to ignore small changes in the data which do not affect our interpretation of it; for example, if $F(p) \neq F'(p)$, but both $F(p)$ and $F'(p)$ lie in the range taken to represent air, so that $\tau(F(p)) = \tau(F'(p))$, no distortion is considered to have occurred at p. This allows the distortion metric to take advantage of the *a priori* knowledge of the dataset encapsulated in τ.

We then define P, the set of points at which $F(\boldsymbol{x})$ and $F'(\boldsymbol{x})$ are considered to represent different features, by

$$P = \left\{ \forall p \in X \;\middle|\; \tau\big(F(p)\big) \neq \tau\big(F'(p)\big) \right\}. \tag{9}$$

We may then define m_d as the ratio of the number of points in P to the total number of points occupied by $F(\boldsymbol{x})$:

$$m_d = \frac{|\mathrm{P}|}{|\mathrm{X}|} = \frac{|P|}{R_1 \cdot R_2 \cdot R_3}. \tag{10}$$

It follows that low values of m_d indicate lower distortion, and that $0 \leq m_d \leq 1$ (since $P \subseteq X$). The ideal smoothing method will have m_d close to zero, indicating no distortion.

An alternative measure of distortion might be to weight the metric, imposing a proportionally higher penalty in regions where a point on an iso-surface of interest in F lies further from the corresponding point in F'; that is, in areas where the processed dataset

[3] We make use of this method in other parts of this paper where we refer to $\lambda = n\%$ indicating lambda equal to the magnitude of the normal at the n'th percentile.

differs strongly from the original. A suitable metric might be based on the following equation for m_d:

$$m_d = \sqrt{\sum \psi_{F,F'}(p)^2} \tag{11}$$

$\forall p \in X$ such that p lies on the iso-surface of interest in F, and where $\psi_{F,F'}(p)$ represents the distance from p to the corresponding point in F'. However, there are some practical issues with this metric, not least the problem of establishing the location of the corresponding point for p. Whilst nonlinear diffusion theory guarantees continuity in the succession of datasets $U(\boldsymbol{x},0)$ to $U(\boldsymbol{x},t)$ (that is, features smoothly and continuously evolve into their final state at time t), and that in principle features may be tracked, in practice implementing this would be considerably more complex than the previous approach. Moreover, it is not clear that such a weighted metric is desirable: there is no obvious order of preference between a processed dataset with a moderate level of shape degradation everywhere and one with low degradation generally and a few areas of high degradation, given an equal score on the first proposed metric.

For these reasons we use the metric proposed in Equation 10 to evaluate distortion.

5.2 Coherence Improving Metric

It is harder to design a coherence metric, since the goal of this work is to some extent subjective. However, proceeding from the stated goal that we wish to produce refractive images which are free from noise, it follows that an object-space metric might be somewhat contrived, on the grounds that this metric aims to measure the improvements made in the image domain. (The previous metric for distortion measures shape degradation, which occurs in the volume domain, and is therefore more appropriately measured in this domain). We will therefore consider image-space methods of evaluating improvement.

Since object shape degradation is considered in the preceding metric, we ignore this here and focus only on coherence. Comparing, for example, the images in Figure 3, we see that the images rendered with unperturbed normals contains less energy at high frequencies than their counterparts with perturbed normals. This observation has been observed to hold for a range of images. We may measure improvements, therefore, by considering the frequency distribution of the image; images with a high proportion of coherent rays (that is, rays which travel along similar paths to adjacent rays) will tend to contain less high frequency energy.

We consider the frequency distribution by applying the discrete Fourier transform to the images. The result is a two dimensional complex array; in order to reduce this to a single value which can be easily compared for different images, we multiply the magnitude of each element in the array, which corresponds to the amount of energy present at a certain frequency, by the frequency represented by that element, and take the average of the resulting array. The result is a number which gives an indication of frequency distribution for the entire image, with values close to zero indicating few high frequencies, and large values indicating a large quantity of high frequency energy. We express this coherence metric m_c mathematically in the following way, given an image $i : \mathbb{N}^2 \to \mathbb{N}$ of resolution M, N:

$$I(u,v) = \frac{1}{MN} \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} i(x,y) e^{-2\pi j(ux/M + vy/N)} \tag{12}$$

$$m_c = \frac{4}{MN} \sum_{u=0}^{\frac{M-1}{2}} \sum_{v=0}^{\frac{N-1}{2}} |I(u,v)| \cdot \sqrt{\frac{u^2}{M^2} + \frac{v^2}{N^2}}. \tag{13}$$

According to the periodicity property of the discrete Fourier transform (given by Equation 12), the discrete Fourier transform of $i(x,y)$, denoted here by $I(u,v)$, is symmetric about the origin, and repeats horizontally and vertically respectively, at distances $\frac{M}{2}$ and $\frac{N}{2}$ from the origin. For this reason, in order to analyze the frequencies present in the discrete Fourier transform we need only consider the values in the range $(0,0)$ to $(\frac{M}{2}, \frac{N}{2})$; therefore we sum only these frequencies, and multiply the result by $\frac{4}{MN}$, which is the reciprocal of the number of values considered. Note that in computing the frequency represented by some point in $I(u,v)$ – that is, in the last term of Equation 13 – we scale u, v by M, N respectively in order to ensure that the value produced by this metric is (approximately) invariant with respect to image size and resolution, modulo changes in precision attributable to change in resolution.

As will be seen in Section 6, this metric results in values which tend to decrease as λ or t increase. This decrease is not exactly monotonic, since a few data points do not fit this pattern; however, the trend is otherwise significant enough to suggest that this metric has some correspondence with increased ray coherence in the rendered images. Whilst it is not meaningful to use this metric to compare unrelated images, it is nonetheless useful in comparing similar images from a series of datasets processed with different parameters but rendered using identical parameters.

We demonstrate the validity of this metric in Figure 7. Here, the CT Head dataset is processed using *anisotropic nonlinear diffusion* (described in Section 4.3) with varying parameters. The parameters are selected such that m_c has a similar value in each case; it can be seen that the resulting images exhibit comparable levels of coherence, especially when contrasted with the images in Figure 5, which exhibit substantial diversity.

(a) $\lambda = 50\%$ $t = 150$ $m_c = 2.53$ (b) $\lambda = 60\%$ $t = 50$ $m_c = 2.56$ (c) $\lambda = 90\%$ $t = 30$ $m_c = 2.56$

Figure 7: Images with similar values for m_c.

6 Results and Remarks

Figures 4 and 5 show the results of applying the methods discussed in Section 4 to the CT Head dataset, and for each method, examples of both opaque and refractive rendering are given. The methods vary widely in both the degree to which image quality is improved, and in the level of distortion introduced. In the following sections, we discuss these images subjectively, and with respect to the metrics introduced in Section 5.

The CT Head dataset is a demanding test for refractive visualization, in that it features significant levels of noise in all regions of the dataset, as a consequence of the scanning process, and particularly high levels of noise around the teeth, attributed to metal fillings in the teeth interfering with the original scanning process. Additionally, the pockets of air in the nasal cavities cause frequent, incoherent refraction in this region. There are some pre-aliasing artefacts present at the top of the head, visible as curved bands in the images in the first row of Figure 5. The presence of fine structures, such as the nose and the ear pose additional problems for smoothing algorithms which aim to preserve the overall shape.

6.1 Qualitative Analysis of Results

The opaque images in Figure 4 demonstrate that both *inhomogeneous linear diffusion* and *anisotropic nonlinear diffusion* are poor at preventing distortion. *Inhomogeneous linear diffusion* performs particularly badly with respect to fine structures, as can be seen in the erosion of the nose. Moreover, it fails to smooth the noise around the teeth. *Anisotropic nonlinear diffusion* distorts heavily everywhere; large features such as the dimples in the side of the face are smoothed away. *Gaussian convolution* and *regularized anisotropic nonlinear diffusion* both perform well, although the images in the fourth column indicate that *Gaussian convolution* performs less well at preserving thin structures in the data (such as the nose or the ears). For both of these methods, the overall shape of the head is well preserved; both deal well with the noise around the teeth.

Improvements in coherence can be observed in Figure 5. *Gaussian convolution* performs well; as a consequence of filtering out all high frequencies, noise is reduced, and adjacent rays tend to travel along more closely matched paths, leading to significant overall improvement in image quality. It is also effective at reducing pre-aliasing artefacts. *Inhomogeneous linear diffusion* shows some improvements; noise in all parts of the dataset has been significantly reduced. Some new artefacts have been introduced, however; these may be observed in the region above the right ear. This method also fails to smooth the pre-aliasing artefacts. *Anisotropic nonlinear diffusion* is extremely effective at smoothing both types of artefact, even for low values of t; however, this method also introduces the same artefacts as *inhomogeneous linear diffusion*. *Regularized anisotropic nonlinear diffusion* performs well, smoothing both pre-aliasing artefacts and noise, without introducing new artefacts.

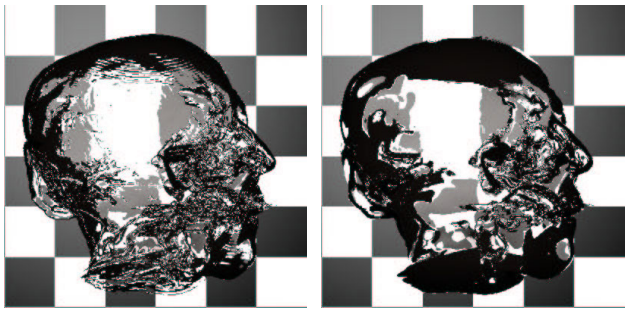

(a) original dataset

(b) hybrid rendering $\lambda = 60\%, t = 50$

(c) hybrid rendering $\lambda = 90\%, t = 100$

(d) Smoothed rendering as (b)

Figure 8: Hybrid rendering method.

6.2 Hybrid Rendering Method

A possible approach for improving image quality without affecting distortion is to use the smoothed dataset only for computation of normals. This is demonstrated in Figure 8, using *regularized anisotropic nonlinear diffusion* ($\lambda = 60\%, t = 50$). Figure 8(a) shows the original (un-smoothed) dataset; Figure 8(b) shows the hybrid approach, and Figure 8(c) both dataset values and normals computed from the smoothed dataset.

The benefits of this approach are clear, in that it offers a means of improving image quality without distorting the original data at all. However, the smoothed dataset may not always have appropriate normals in the location of boundaries in the original data; this problem prevents it from being as effective in increasing coherence as the conventional method of rendering; and in fact as F' diverges further from F, artefacts are introduced, as may be seen in Figure 8(c).

6.3 Quantitative Analysis of Results

The data in Table 1, which is obtained from the metrics discussed in Section 5, suggests that *regularized anisotropic nonlinear diffusion* performs better than all other methods. Direct comparisons are difficult, since values for both m_d and m_c do not exactly correspond.

Gaussian convolution

	σ	0.0	1.5	3.0	4.5
Gaussian	m_d	0	0.0027	0.0108	0.0287
	m_c	23.19	19.40	16.34	13.88

Diffusion methods with $\lambda = 60\%$

	t	0	20	40	60
IHLD	m_d	0	0.0040	0.0195	0.0365
	m_c	23.19	18.36	16.55	15.96
ANLD	m_d	0	0.0122	0.0398	0.0617
	m_c	23.19	14.12	10.80	10.44
R-ANLD	m_d	0	0.0024	0.0066	0.0102
	m_c	23.19	17.22	14.70	13.33

Table 1: Results of applying distortion and coherence metrics to datasets and refractive images from Figure 5. Methods considered include *Gaussian convolution*, *inhomogeneous linear diffusion* (IHLD), *anisotropic nonlinear diffusion* (ANLD), and *regularized anisotropic nonlinear diffusion* (R-ANLD).

Method	m_d	m_c
IHLD (t = 35, $\lambda = 60\%$)	0.0147	16.88
Gaussian (σ = 3.18)	0.0151	16.09
ANLD (t = 15, $\lambda = 70\%$)	0.0153	13.64
R-ANLD (t = 95, $\lambda = 60\%$, $\sigma = 1$)	0.0152	12.11

Table 2: Comparison of results for different methods in a condition where parameters selected to give similar values for m_d. Methods are listed in the descending order of the noise level (i.e., incoherence) indicated by m_c.

We therefore select values for t, λ and σ such that m_d remains approximately constant in Table 2, so that we may directly compare the effectiveness of different methods in terms of m_c. When m_d is fixed at around 0.015 giving similar level of distortion for all methods, the *regularized anisotropic nonlinear diffusion* method (and, to a lesser extent, the *anisotropic nonlinear diffusion* method) offers better improvement of smoothness as measured by the coherence metric.

Method	m_d	m_c
IHLD (t = 30, $\lambda = 70\%$)	0.0261	15.12
Gaussian ($\sigma = 3.75$)	0.0253	15.08
ANLD (t = 75, $\lambda = 40\%$)	0.0204	15.13
R-ANLD (t = 135, $\lambda = 40\%$, $\sigma = 1$)	0.0067	15.06

Table 3: Comparison of results for different methods in a condition where parameters are selected to give similar values for m_c. Methods are listed in the descending order of the amount of distortion indicated by m_d.

Similarly, in Table 3, we select parameter values such that m_c remains approximately constant to enable us examine results in terms of m_d. It is clear that when m_d is fixed at around (for instance) 15.1, *regularized anisotropic nonlinear diffusion* results in much less distortion than any of the other three methods.

This confirms the outcome of qualitative analysis in Section 6.1 based on the visual results in Figures 4 and 5. *Regularized anisotropic nonlinear diffusion* is clearly most effective at improving image quality whilst preserving the shape of the original data, whilst preserving the shape of the original data. *Anisotropic nonlinear diffusion* performs less well, and *Gaussian convolution* and *inhomogeneous linear diffusion* both perform relatively poorly.

Considering the overall ability to improve ray coherence with minimal cost of distortion, Figure 9 gives an intriguing illustration of the relative performance of these methods, each depicted by a line graph with m_d as its x-axis. *Gaussian convolution* was applied to the CT datasets with a series of σ values, $\sigma = 0, 1, 2, 3, \ldots$ marked by the corresponding data points. All three diffusion methods are computed with $\lambda = 40\%$ and $t = 0, 20, 40, \ldots, 140$. $\sigma = 4$ was chosen for computing the line graph of *regularized anisotropic nonlinear diffusion*. The y-axis indicates the relative improvement of coherence m_c in the context of a rendered refractive image with a view similar to those in Figure 5. For example, we compute $m_c(\sigma = 0)/m_c(\sigma)$ for the *Gaussian convolution*, and $m_c(t = 0)/m_c(t)$ for the diffusion methods. Figure 9 supports the conclusions drawn above; in particular, the limitations of Gaussian convolution are revealed, and the strong performance of *regularized anisotropic nonlinear diffusion* is made apparent.

Some more detailed observations about the performance of *regularized anisotropic nonlinear diffusion* may be derived from the graphs in Figure 10. For values of λ below roughly 50%, the rate of growth of m_d is low. The rate of reduction of m_c, however, with respect to λ, is fairly uniform in the range $30\% - 70\%$. For these reasons, it seems reasonable to suggest that suitable choices for λ typically lie around the 50th percentile. This is borne out informally by experiences using these methods, although this cannot be proven without further experimentation on a range of datasets.

7 Conclusions and Future Work

We have shown that it is possible to achieve good results when using refraction to render datasets which contain noise, pre-aliasing and small features. Figure 2, which depicts the CT head dataset both (a) unprocessed and (b) after application of *regularized anisotropic nonlinear diffusion*, gives a compelling example of deploying such techniques in medical visualization.

By introducing metrics which allow us to quantify distortion and improvements in image quality, we are able to make objective comparisons between a number of different methods for smoothing datasets. *Regularized anisotropic nonlinear diffusion* appears to be the most effective of the methods tested, both from a theoretical standpoint, and in practice. We believe by using these methods for refractive visualizations that depth perception can be improved, and

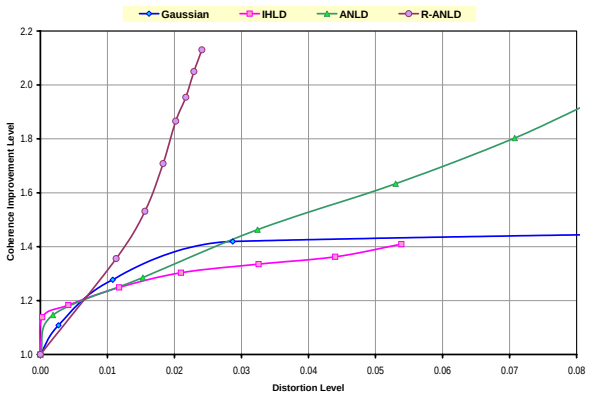

Figure 9: Relative performance of four different denoising methods. The measurement of coherence improvement is the result of dividing the m_c value obtained from the unprocessed dataset by the m_c value obtained from the processed dataset.

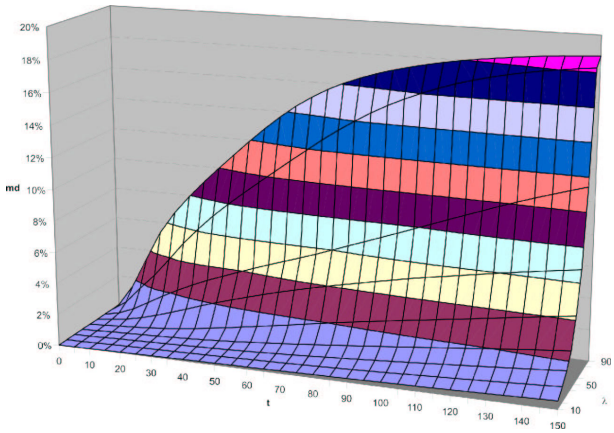

(a) Plotting m_d as a function of t (x-axis) and λ (y-axis)

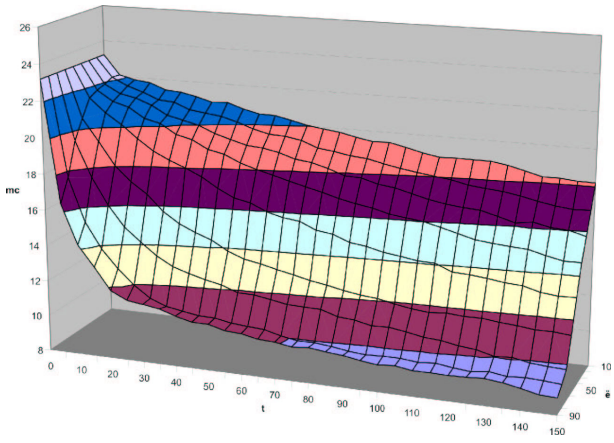

(b) Plotting m_c as a function of t (x-axis) and λ (y-axis)

Figure 10: Metrics for *regularized anisotropic nonlinear diffusion*.

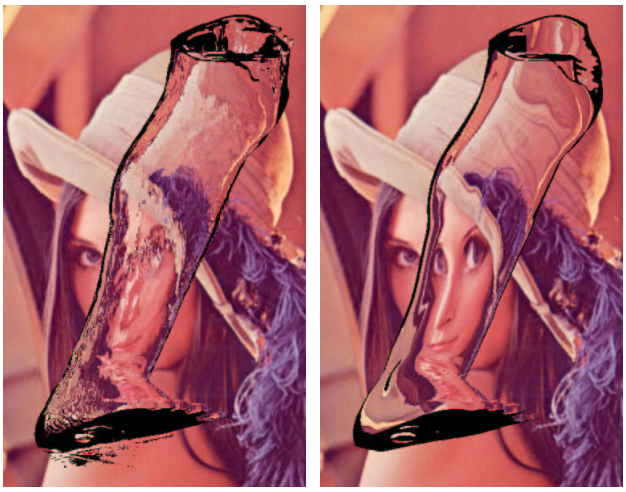

(a) original (b) with denoising

Figure 11: A statue leg rendered (a) without and (b) with application of *regularized anisotropic nonlinear diffusion*.

that significant improvements in image quality may be attained, as Figure 11 shows.

It is possible to eliminate distortion by means of the hybrid rendering method; however, this method cannot achieve the same subjective levels of quality as directly rendering smoothed datasets.

There exists a clear tradeoff between image quality and distortion. However, given the large parameter space and the complex relationship between these parameters and the resulting image quality (with respect to both coherence and distortion), it is difficult to present a "best" approach that is suitable for all applications.

We aim to continue this work by focusing more strongly on *regularized anisotropic nonlinear diffusion*; in particular, by examining different diffusion functions, and considering the effects of varying σ in more detail. In addition, other methods for denoising may be able to improve on these results.

8 Acknowledgments

The first author is grateful for his PhD studentship from University of Wales Swansea for the period 1999-2003. The CT Head dataset is the courtesy of University of North Carolina, Chapel Hill, the statue leg dataset courtesy of German Federal Institution for Material Research and Testing, Berlin, Germany. All images in this paper are rendered using *vlib* [Winter and Chen 2001].

References

ALVAREZ, L., LIONS, P., AND MOREL, J. 1992. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal. 29*, 845–866.

ANGELINI, E., LAINE, A., TAKUMA, S., AND HOMMA, S. 1999. Directional representations of 4D echocardiography for temporal quantification of LV volume. In *Proceedings of MICCAI'99*, 430–440.

BERTRAM, M. 2003. Fairing scalar fields by variational modeling of contours. In *Proceedings of IEEE Visualiszation 2003*, 387–392.

CLARENZ, U., DIEWALD, U., AND RUMPF, M. 2000. Anisotropic geometric diffusion in surface processing. 397–405.

DESBRUN, M., MEYER, M., SCHRDER, P., AND BARR, A. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of SIGGRAPH 1999*, 317–324.

DESBRUN, M., MEYER, M., SCHRDER, P., AND BARR, A. 2000. Anisotropic feature-preserving denoising of height fields and bivariate data. In *Proceedings of Graphics Interface 2000*, 145–152.

GIBBS, J. W. 1899. Fourier series. *Nature 59*.

GUSKOV, I., AND WOOD, Z. 2001. Topological noise removal. In *Proceedings of Graphics Interface 2001*, 19–26.

HILTON, M., OGDEN, T., HATTERY, D., EDEN, G. F., AND JAWERTH, B. 1996. Wavelet denoising of functional mri data. In *Wavelets in Biology and Medicine*, A. Aldroubi and M. Unser, Eds. CRC Press, 93–112.

KUAN, D. T., SAWCHUK, A. A., STRAND, T. C., AND CHAVEL, P. 1985. Adaptive noise smoothing filter for images with signal-dependent noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-7*, 165–177.

LEE, J. S. 1980. Digital image enhancement and noise filtering by use of local statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-2*, 165–168.

MARSCHNER, S., AND LOBB, R. 1994. An evaluation of reconstruction filters for volume rendering. In *IEEE Visualization '94*, IEEE Computer Society Press, R. D. Bergeron and A. Kaufman, Eds., 100–107.

MÖLLER, T., MACHIRAJU, R., MUELLER, K., AND YAGEL, R. 1997. A comparison of normal estimation schemes. In *IEEE Visualization '97*, IEEE, 19–26. ISBN 0-58113-011-2.

MÖLLER, T., MACHIRAJU, R., MUELLER, K., AND YAGEL, R. 1997. Evaluation and design of filters using a Taylor series expansion. *IEEE Transactions on Visualization and Computer Graphics 3*, 2 (April - June), 184–199. ISSN 1077-2626.

MÖLLER, T., MÜLLER, K., KURZION, Y., MACHIRAJU, R., AND YAGEL, R. 1998. Design of accurate and smooth filters for function and derivative reconstruction. In *Proceedings of IEEE Symposium on Volume Visualization'98*, 143–151.

MORETON, H. P., AND SEQUIN, C. H. 1992. Functional optimization for fair surface design. In *Proceedings of SIGGRAPH '92*, 167–176.

NEUMANN, L., CSÉBFALVI, B., KÖNIG, A., AND GRÖLLER, E. 2000. Gradient estimation in volume data using 4D linear regression. *Computer Graphics Forum 19*, 3, C351–C357.

OHTAKE, Y., BELYAEV, A. G., AND BOGAEVSKI, I. A. 2000. Polyhedral surface smoothing with simultaneous mesh regularization. In *Proceedings of Geometric Modeling and Processing 2000*, 229–237.

PENG, J., STRELA, V., AND ZORIN, D. 2001. A simple algorithm for surface denoising. In *Proceedings of IEEE Visualiszation 2001*, 107–112.

PERONA, P., AND MALIK, J. 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence 12*, 7.

PERSOON, M. P., SERLIE, I. W. O., POST, F. H., AND VOS, F. M. 2003. Visualization of noisy and biased volume data using first and second order derivative techniques. In *Proceedings of IEEE Visualiszation 2003*, 379–385.

RODGMAN, D., AND CHEN, M. 2001. Refraction in discrete raytracing. In *Volume Graphics 2001*, Springer, New York, K. Mueller and A. Kaufman, Eds. ISBN 3-211-83737-X.

RÖSSL, C., ZEILFELDER, F., NÜRNBERGER, G., AND SEIDEL, H.-P. 2003. Visualization of volume data with quadratic super splines. In *Proceedings of IEEE Visualiszation 2003*, 393–400.

SAPIRO, G. 2001. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press.

TASDIZEN, T., WHITAKER, R., BURCHARD, P., AND OSHER, S. 2002. Geometric surface smoothing via anisotropic diffusion of normals. In *Proceedings of IEEE Visualiszation 2002*, 125–132.

THEUSSL, T. 2000. *Sampling and Reconstruction in Volume Visualization*. PhD thesis, Vienna University of Technology.

WEICKERT, J. 1997. A review of nonlinear diffusion filtering. In *Scale-Space Theories in Computer Vision*, 3–28.

WEICKERT, J. 1998. *Anisotropic Diffusion in Image Processing*. ECMI Series, Teubner, Stuttgart.

WELCH, W., AND WITKIN, A. 1994. Free-form shape design using triangulated surfaces. In *Proceedings of SIGGRAPH '94*, 247–256.

WHITAKER, R. T. 1994. Volumetric deformable models: active blobs. In *Visualization in Biomedical Computing, SPIE*, R. A. Robb, Ed.

WINTER, A. S., AND CHEN, M. 2001. vlib: A volume graphics API. In *Volume Graphics 2001*, Springer, New York, 133–147. ISBN 3-211-83737-X.

YAGEL, R., COHEN, D., AND KAUFMAN, A. 1992. Normal estimation in 3D discrete space. *The Visual Computer*, 278–291.