

Regularised Anisotropic Nonlinear Diffusion for Rendering Refraction in Volume Graphics

David Rodgman* and Min Chen[†]

* ARM Limited, Cambridge, United Kingdom
[†] University of Wales Swansea, United Kingdom

Abstract

Rendering refraction in volume graphics requires smoothly distributed normals to synthesise good quality visual representations. Such refractive visualisation is more susceptible to noise in the data than visualisations that do not involve refraction. In this paper, we address the need for improving the continuity of voxel gradients in discretely sampled volume datasets using nonlinear diffusion methods, which was originally developed for image denoising. We consider the necessity for minimising unnecessary geometrical distortion, detail the functional specification of a volumetric filter for regularised anisotropic nonlinear diffusion (R-ANLD), discuss the further improvements of the filter, and compare the efficacy of the filter with an anisotropic nonlinear diffusion (ANLD) filter as well as a Gaussian filter and a linear diffusion filter. Our results indicate that it is possible to make significant improvements in image quality in refractive rendering without excessive distortion.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation, I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism, I.4.3 [Image Processing and Computer Vision]: Enhancement, J.3 [Computer Applications]: Life and Medical Sciences.

1. Introduction

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 images are synthesised. For these purposes, mathematically defined scalar fields provide unsurpassed quality. They are free from both noise and discretisation 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 ren-

dering 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 [RC01].

As Möller *et al* point out in [MMMY97], 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.

We demonstrate the importance of having good quality

[†] Contact Author: Min Chen, Department of Computer Science, University of Wales Swansea, Singleton Park, Swansea SA2 8PP, United Kingdom. Tel: +44 1792 295663. Fax: +44 1792 295708. Email: m.chen@swansea.ac.uk.




Figure 1: Importance of accuracy in refractive images. Normals of the sphere are perturbed by an average of 15° . In (a) an opaque sphere where the normal perturbation statistically follows a normal distribution; in (b) a translucent sphere with the same normal perturbation as (a), but with no refraction; in (c) a translucent sphere, as (b), with refraction.

normals for refraction with a simple test case. Figure 1 shows three spheres; in each case, the normals are perturbed by an average of 15° , and the level of perturbation at every point has a normal distribution. It can be seen by visual inspection that image quality is better preserved in the absence of refraction, than in its presence. This observation is also endorsed by a numerical image comparison, confirming that an image of an object featuring refraction is more severely affected by the presence of noise than an image of a similar but non-refractive object.

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 thoroughly addressed in the literature (e.g., [ML94]) 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. Hence, the modelling and rendering of discrete, noisy volume datasets with refractive effects is a challenging problem in volume graphics. Figure 2(a) shows a typical example of this challenge, where a noisy dataset is rendered with and without refraction. The background image helps highlight the impact of noise in the rendered image. Our objective is to synthesise a smooth refractive image with a minimal level of distortion as in Figure 2(b).

Denosing is a classic problem in image processing and computer vision. Nonlinear diffusion filtering is a powerful and effective tool built on a solid theoretical foundation [Wei97]. In this paper, we address the need for im-



Figure 2: The solid objects illustrate the level of noise in the original digitised dataset and the ability of a volume denoising method in preserving the geometrical features of the object. The transparent objects, aided by a background image, highlight the significant impact of noise upon the quality of the rendered refraction and necessity of reducing noise.

proving the continuity of voxel gradients in discretely sampled volume datasets using anisotropic nonlinear diffusion. After a brief review of previous work in related areas in Section 2, we outline the algorithmic process for rendering refraction using discrete ray tracing in volume graphics in Section 3. We then detail the functional specification of a volumetric filter for regularised anisotropic nonlinear diffusion in Section 4. We discuss the further improvement of the filter through effective controlling the regularisation step in Section 5 and selecting the diffusivity function in Section 6. We present a set of experimental results in Section 7, and highlight the efficacy of regularised anisotropic nonlinear diffusion in comparison with a Gaussian filter and a linear diffusion filter. Finally, we offer our observations and some concluding remarks in Section 8.

2. Related Work

The concept of refraction with which we are familiar nowadays was developed in 1621 by the Dutch astronomer and

mathematician Willebrord van Roijen Snell (1580-1626). Snell's law [BW75], that is, $\eta_i \sin(\theta_i) = \eta_r \sin(\theta_r)$, remains a basic principle in modelling the behaviour of light.

Yagel *et al.* [YCK92] suggested a limited approach to implementing refraction in volume graphics in 1992, requiring the analytical intersection locations, and the analytical derived normals. Neither of these data are usually available in volume datasets. This approach may be characterised as using discrete ray tracing for colour accumulation, but using surface ray tracing for rendering refraction.

Refraction was first introduced in discrete ray tracing in a general manner by Rodgman and Chen in 2001 [RC01]. A comprehensive treatment of refraction and some related problems was given by Rodgman in 2003 [Rod03].

There are a diversity of modalities, such as CT and MRI, for the acquisition of volume datasets. The aliasing and noise introduced during digitisation is of concern to many algorithms for processing and rendering such datasets. It is necessary to investigate the use of denoising methods in the context of direct volume rendering with refraction. The successful deployment of nonlinear diffusion methods in surface denoising offers a good starting point for volume denoising, and the demand for a smooth gradient in rendering refraction, as shown in Figure 2, makes it a suitable case study for evaluating the effectiveness of volume denoising methods.

Much effort has been placed on the correct reconstruction of normals. A number of normal estimation methods have been studied and compared [YCK92, MMY97]. There has been recent attempts [MMK*98, NCKG00] to reconstruct accurately the continuous 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 (e.g., [PM90, WtV98]). In recent years, some of these filters have been successfully generalised for smoothing surfaces. For example, Peng *et al.* employed Wiener filtering for smoothing noisy triangular meshes [PSZ01]. Desbrun *et al.* employed anisotropic diffusion for denoising height functions and bivariate data [DMSB00]. Tasdizen *et al.* used an anisotropic diffusion method for smoothing surfaces via normal maps [TWBO02]. In addition, there was much other effort in surface fairing and denoising (e.g., [MS92, WW94, DMSB99]).

In terms of volume visualisation, surface denoising can be applied to a particular iso-surface contained in a volume dataset, 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 [HOH*96] and Angelini *et al.* gave a brief description of a three-dimensional implementation of brushlet basis functions for Fourier domain denoising [ALTH99]. The former was a study based on statisti-

cal analysis and the latter was an investigation in the context of segmentation.

3. Refraction in Volume Graphics

3.1. Modelling Refraction

A *scalar field* is a function $F : \mathbb{E}^3 \rightarrow \mathbb{R}$. A *spatial object* is defined as a tuple of scalar fields, (A_0, A_1, \dots, A_k) , $k \geq 0$, where each field specifies a particular attribute of the object, such as colour and reflection coefficients. In volume graphics, an important attribute is *opacity*, which defines the visible geometry of the object. We write a spatial object as (O, A_1, \dots, A_k) , where O is a scalar field specifying the opacity of every point $p \in \mathbb{E}^3$.

From this general definition of spatial object, we can treat a *volume object* as a spatial object whose opacity field is bounded [CT00]. This also facilitates the construction of a volume object from one or more volume datasets coupled with some interpolation functions and transfer functions. By using scalar fields as the basis of all spatial objects, we can treat digitised volume datasets in a manner consistent with mathematically or procedurally defined scalar fields.


In surface graphics, we typically model refraction by associating each object with two parameters, *transmission coefficient* k_t and *refractive index* η [Whi80]. In the above definition of spatial object, the opacity field O facilitates a more accurate and flexible specification of transmission coefficient. For each object, we model its optical density in the form of a scalar field $H : \mathbb{E}^3 \rightarrow \mathbb{R}$, which gives a refractive index $\eta = H(p)$ for every point $p \in \mathbb{E}^3$.

Of particular importance in modelling refraction is the change in refractive index. When modelling the path of a ray of light from one continuous medium to another, we use the *relative refractive index*, η_r to refer to the relative change in refractive index. In classical ray tracing, as shown in Figure 3(a), given the respective refractive indexes of the medium that the ray is leaving, and the medium that the ray enters, η_i and η_r , we have $\eta_r = \frac{\eta_i}{\eta_r}$.

As we do not assume the existence of a surface of infinitesimal thickness in volume graphics, the definition of a signal η_r at every point on a surface is not directly applicable to spatial objects. As illustrated in Figure 3(b), a more tangible specification of η_r would be $\eta_r = \frac{H(u)}{H(v)}$, where u and v are two arbitrary points in \mathbb{E}^3 . However, from both modelling and rendering perspectives, there is no longer a unique η_r defined at an individual point in \mathbb{E}^3 . Obtaining a sensible value of η_r at individual samples hence becomes a critical problem for rendering refraction in discrete ray tracing.

3.2. Rendering Refraction

When directly rendering a volume graphics scene, for each ray, as it passes through a spatial object, we sample at discrete points, $s_0, s_1, \dots, s_i, \dots$, along the ray, accumulating



(a) η_r on a surface (b) η_r in a volume (c) SSFL sampling

Figure 3: In volume graphics, while defining a refractive index field $H(p)$, $p \in \mathbb{E}^3$ is physically meaningful and appropriate (b), the specification η_r is not as straightforward as in surface graphics (a). In the process of sampling a refractive index field, we assume some uniformity in a small region around a sampling point. One basic but effective sampling approach is the SSFL method (c).

colour and opacity at each point. With the specification of refractive indices, the ray may change its direction at each sampling position if the sampled relative refractive index, η_r , is not equal to one.

When the sampling step is sufficiently small, we may assume some uniformity in a small domain around a sampling position s_i . If the refractive index field is not constant in this domain, then we consider the domain to be divided into two sub-domains, as shown in Figure 3(c), each with a uniform optical density. It is obviously sensible to separate the two sub-domains by the tangent plane perpendicular to the normal N at s_i . We may compute or estimate the normal based on an appropriate scalar field of the spatial object, typically a geometry field, an opacity field or the refractive index field itself. For example, given a mathematically defined geometry field (e.g., a spherical function), we can calculate normals analytically. Given an opacity field defined upon a raw dataset (e.g., through an opacity mapping), we can estimate normals using central differences.

Within the small sampling domain around s_i , we compute the relative refractive index at s_i by obtaining two representative refractive indices, one for each sub-domain. For a known incident ray $I = \text{normalise}(s_i - s_{i-1})$, we obtain the direction of refraction vector T in the normal way (as discussed by Foley *et al.* [FvFH96]). This is consistent with the Snell's law.

$$T = \left(\eta_r (N \bullet I) - \sqrt{1 - \eta_r^2 (1 - (N \bullet I)^2)} \right) N - \eta_r I.$$

In the following discussions, we assume that vectors N , I , and T are normalised.

Rodgman and Chen proposed several algorithmic approaches to the sampling process in discrete ray tracing [RC01]. One simple but effective approach is the SSFL (Single-sided Sampling, Fixed interval Length) method. In order to focus on the issue of denoising in a consistent man-

ner, all experimental results presented in this paper were rendered using this method.

The SSFL approach takes samples at intervals of a pre-defined length λ as shown in Figure 3(c). The relative refractive index η_r at s_i is approximated with the two refractive indices sampled at s_{i-1} and s_i , that is:

$$\eta_r = \frac{\text{a representative } \eta_i}{\text{a representative } \eta_t} = \frac{H(s_{i-1})}{H(s_i)}.$$

From η_r , we compute the refraction vector T (and consequently s_{i+1}) as $s_{i+1} = s_i + \lambda T$. The critical angle in relation to total internal reflection can be determined in the same way as in traditional ray tracing.

This method is consistent with the standard sampling mechanism for accumulating colour and opacity in volume ray casting, and thus it intrinsically facilitates an efficient implementation. It does not suffer from any anomalies as some other approaches [RC01], and its accuracy can easily be achieved by depth-supersampling, that is, reducing the pre-defined interval length λ .

4. Regularised Anisotropic Nonlinear Diffusion

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 images 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.

In the following discussion, we adopt the terminology and notation of Weickert [Wei97] with a specific focus on 3D scalar fields and volume datasets. Consider a volume dataset of the resolution $R = (R_1, R_2, R_3) \in \mathbb{N}^3$. We thus denote the unprocessed dataset by $F(\mathbf{x})$ where $\mathbf{x} = (x_1, x_2, x_3) \in \mathbb{N}^3$, such that $0 \leq x_i < R_i$ for $i = 1, 2, 3$, and denote the processed dataset derived from F by $F'(\mathbf{x})$. We also define X to be the set of all such \mathbf{x} .

Firstly, we define a *scale-space* of scalar fields as $U(\mathbf{x}, t)$, where t is the scale. We consider the scale-space to be evolving such that at scale $t = 0$, $U(\mathbf{x}, t) = F(\mathbf{x})$, and so that as t increases, $U(\mathbf{x}, t)$ represents a more heavily processed version of $F'(\mathbf{x})$.

Given a specific scalar field in the scale-space, $U(\mathbf{x}, t)$, ∇U refers to the gradient of U , which is a vector field. Given an arbitrary vector field V , $\nabla \cdot V$ refers to the divergence of V . Hence $\nabla \cdot \nabla U$ is the divergence of the gradient of U .

We begin by defining the *general diffusion equation* in terms of the rate of change of $U(\mathbf{x}, t)$ with respect to t as:

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

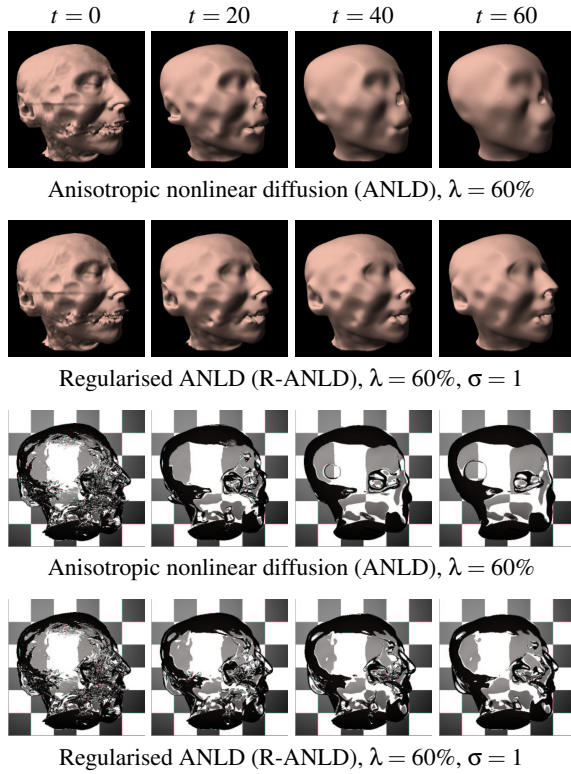


Figure 4: Results of applying the ANLD and R-ANLD methods to the CT head dataset, and rendering the processed datasets as opaque and transparent objects respectively.

where, $D \cdot \nabla U$ is the *flux*, which is equal to the *diffusion tensor* D , a positive definite symmetric matrix, multiplied by the gradient 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 over the entire X , this process is said to be *linear* and *homogeneous*, and is in fact equivalent to convolution with a Gaussian filter. However, if D is a function of \mathbf{x} , or more typically, a function of $F(\mathbf{x})$, D is said to be *inhomogeneous* or *space-dependent*. When D is a function of the differential structure of the evolving X , that is, a function of $U(\mathbf{x}, t)$, D is said to be *nonlinear*. When $D \cdot \nabla U$ is *parallel* to ∇U , this condition is referred to as *isotropic*, and D can be replaced by a positive scalar-valued *diffusivity* g , otherwise it is said to be *anisotropic*.

Weickert proposed the *Anisotropic Nonlinear Diffusion* (ANLD) method for image enhancement [Wei98]. This method derives its non-linearity from the fact that g is a function of $U(\mathbf{x}, t)$. We therefore have the diffusion equation as:

$$\frac{dU}{dt} = \nabla \cdot \left(g(|\nabla U|^2) \cdot \nabla U \right) \quad (1)$$

where g is a scalar-valued *diffusivity* function. There are many possible choices for g ; here we use the diffusivity function proposed by Charbonnier *et al.* [CBAB94], that is:

$$g(|\nabla U|^2) = \frac{1}{\sqrt{1 + |\nabla U|^2 / \lambda^2}} \quad (\lambda \neq 0).$$

With such a diffusivity function, at a given time t , there is a flow of density along the direction of the current voxel gradient, ∇U , the magnitude of which is proportional to the diffusivity, $g(|\nabla U|^2)$. This diffusivity function inhibits diffusion as $|\nabla U|^2$ increases. The effect is to better maintain well-defined surfaces (by inhibiting smoothing for large values of $|\nabla U|^2$), whilst permitting smoothing in inner regions.

As U is also a function of t , 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 ANLD to the CT head are given in the first and third rows of Figure 4. 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 will be discussed in detail in the last paragraph of this section. ANLD preserves surfaces by inhibiting smoothing at boundaries.



Figure 5: Staircasing effect with ANLD.

A significant problem with ANLD is that it gives rise to so-called *staircasing artifacts*, as shown in Figure 5. These artifacts manifest themselves as small, regularly spaced discontinuities in the processed dataset. Whilst they do not significantly affect the shape of the processed dataset on a large scale, they introduce energy near the Nyquist frequency, which has a substantial effect on the normals of the processed dataset and is critical to obtaining high quality refractive images.

In order to improve the performance of ANLD, we make use of a *regularisation* step [ALM92] to control the diffusion process, giving rise to the following equation for *Regularised Anisotropic Nonlinear Diffusion* (R-ANLD):

$$\frac{dU_\sigma}{dt} = \nabla \cdot \left(g(|\nabla U_\sigma|^2) \cdot \nabla U_\sigma \right)$$

where we replace U in Equation 1 with U_σ , which is the convolution of U with a Gaussian filter with a fixed $\sigma > 0$.

In other words, we have

$$U_{\sigma}(\mathbf{x}, t) = (K_{\sigma} * U)(\mathbf{x}, t)$$

where

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

$$K_{\sigma}(\mathbf{x}) = k_{\sigma}(x_1) \cdot k_{\sigma}(x_2) \cdot k_{\sigma}(x_3)$$

The effect of this *regularisation 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 visualisation, where the smoothness of the normals is critical. Some results of applying R-ANLD to the CT head are given in the second and fourth rows of Figure 4.

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 is to compute the magnitudes of the normals in the dataset, sort them (in ascending order) by magnitude, discard zero-length normals, and then set λ equal to the magnitude of the normal at the n th percentile, where n is controlled by the user. We make use of this method in this paper where we refer to $\lambda = n\%$ indicating lambda equal to the magnitude of the normal at the n th percentile. This is similar to the method used by Canny in [Can86].

5. Controlling the Regularisation Step

The presence of a Gaussian filtering stage in R-ANLD aims to prevent noise in the data from interfering with the direction of smoothing. However, by Gaussian filtering the data, an undesirable averaging effect is also had on the magnitude of the normal; thus for sufficiently large values of σ , smoothing is not sufficiently inhibited in boundary regions. Conversely, for values of σ which are too low, sharp boundaries are well preserved, but the resulting quality of the normals is diminished. This results in staircasing artefacts. The effect that these staircasing artefacts, which are most clearly visible below the eye, can have on refraction is obvious.

Figure 6 shows a selection of visualisations of the CT head dataset for different values of σ . The images show clearly the trade-off between excessive smoothing across boundary regions for higher values of σ (which is best displayed by the 2D slices in the first column), and the diminished quality of normals for low values of σ (in the opaque

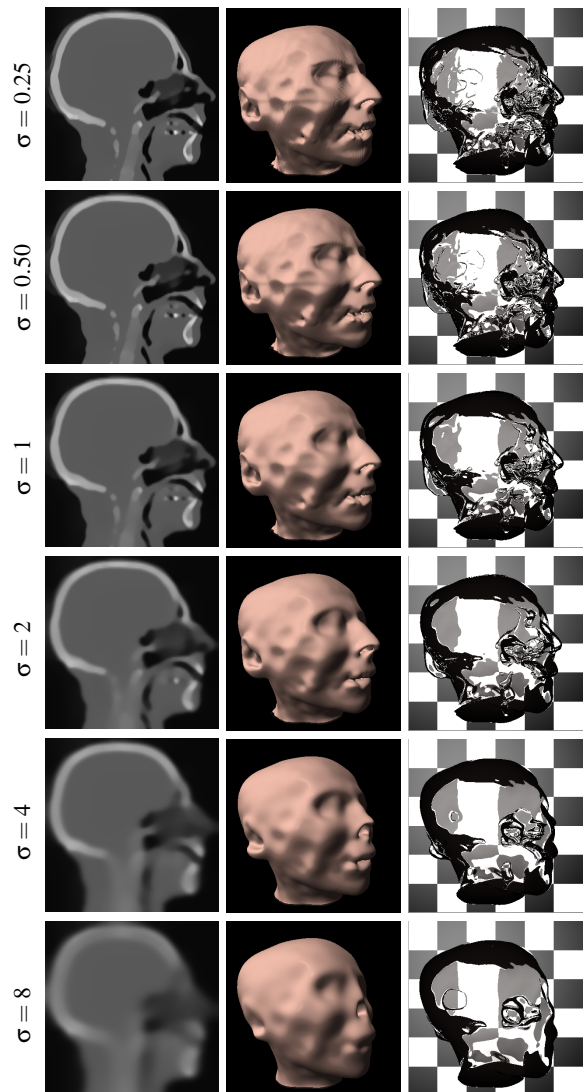


Figure 6: Effects of varying σ in R-ANLD ($t = 50$, $\lambda = 50\%$).

images in the second column). The images in the third column reveal the effects of σ on transparent images.

Results suggest that an acceptable balance is obtained for values of σ in the range one to two. For values below one, some banding is visible in the opaque images; the staircasing in the data responsible for this effect is also responsible for the disturbance in the upper left part of the head in the transparent set of images. This effect is typical of anisotropic diffusion methods where the regularisation step is insufficient [CM96]; smooth edges evolve into discontinuous piecewise linear segments.

It is clear that, for the CT head dataset, higher values of σ can cause significant distortion to the head, particularly to

thin structures such as the nose and the ear. However, determining a suitable value for σ is not straightforward; the acceptable range for σ is data-dependent. Datasets containing few small features (relative to the sampling interval; that is, datasets sampled at a rate much higher than the Nyquist frequency) will clearly be less affected by the Gaussian smoothing step, since this step corresponds to the application of a low-pass filter. Clearly, the proportion of high frequency energy present in the dataset must inform the selection of a value for σ ; however, the trade-off between staircasing artifacts and distortion still remains.

6. Selecting a Diffusivity Function

Several diffusivity functions are mentioned in the literature. Perona and Malik proposed two possible functions for g in [PM90], which we refer to here as g_{PM1} and g_{PM2} . Others include g_C by Charbonnier *et al.* [CBAB94], g_B by Black *et al.* [BSMH98], and g_W by Weickert [WtV98]. Equations for these functions are given below:

$$\begin{aligned} g_{PM1}(x) &= \frac{1}{1 + (x/\lambda)^{1+\alpha}} \quad \text{where } \alpha > 0 \\ g_{PM2}(x) &= e^{-(x/\lambda)^2} \\ g_C(x) &= \frac{1}{\sqrt{1 + (x/\lambda)^2}} \\ g_B(x) &= \begin{cases} \frac{1}{2} [1 - (\frac{x}{\lambda})^2]^2 & |x| \leq \lambda \\ 0 & |x| > \lambda \end{cases} \\ g_W(x) &= 1 - e^{-\frac{3.315}{(x/\lambda)^4}} \end{aligned}$$

One can study the behaviours of these diffusivity functions by plotting each $g(x)$ and its corresponding flow function $\Phi(x) = g(x) \cdot x$. In fact $\Phi(x)$ is an abstraction of $g(|\nabla F(\mathbf{x})|) \cdot |\nabla F(\mathbf{x})|$, which provides a useful means for visualising the degree of flow between adjacent voxels permitted by the diffusion function, as a function of the normal magnitude.

Due to space constraints, it is not feasible to show all these graphs, which can be reproduced trivially. Instead we outline what one can observe from the graphs.

The flow function for g_{PM1} has a long ‘tail’ section after the peak, indicating that flow is permitted, even for large normal magnitudes, thus implying poor edge preservation compared to the other functions. The location of the peak in the flow function for g_{PM2} corresponds more closely to the value chosen for λ ; flow quickly ceases for larger values of x . For this reason, we would expect this function to better preserve edges than g_{PM1} .

The diffusivity function proposed by Charbonnier *et al.*, g_C , has a unique flow function, which rapidly, asymptotically, increases to a maximum value and does not decrease. For small values of λ , this function would behave similarly to linear filtering. In general, this flow function suggests that

it should smooth most heavily in the region of normals of high magnitude, whilst potentially preserving subtler features in a dataset. Despite this odd behaviour, good results have been obtained with this function [Wei97]. This is one of the reasons why g_C was used to synthesise most of the images in this paper.

Function g_W has, uniquely, a significant ‘flat’ section, indicating an equal response to all normals below the threshold (here, the end of this section corresponds to the value chosen for λ); that is, uniform smoothing in regions with low normal magnitudes. Peak flow is achieved for x equal to λ ; the long ‘tail’ of the flow function indicates that there will be some smoothing parallel to the normal even for normals of high magnitude. This is expected to improve surface smoothness, albeit with some cost in terms of distortion.

Function g_B has, by far, the least area of all these functions, indicating a considerably lower degree of smoothing for a given value of λ . It is similar in shape to the g_{PM2} ; however, the corresponding flow function reveals a ‘hard’ cut-off point for values of $x > \lambda$. Smoothing parallel to the normal is therefore entirely inhibited for such values of x ; this would imply less distortion.

7. Results and Analysis

In order to evaluate the denoising methods and parameter variations, it is necessary to devise an objective scheme for quantifying the ways in which both the data and the images are affected. Since the smoothing process represents a trade-off between smoothing and volume shape degradation, we separate this task into two separate sub-tasks, namely evaluating the degree to which the overall shape of a volume is affected, and measuring the improvement in image quality. We utilise a distortion metric, m_d , for quantifying the former, a coherence metric, m_c , for the latter.

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 visualisation. Our motivation in using the same transfer function also used for visualisation is to enable us to ignore small changes in the data which do not affect our visual interpretation. For example, if $F(\mathbf{x}) \neq F'(\mathbf{x})$, but both $F(\mathbf{x})$ and $F'(\mathbf{x})$ lie in the range taken to represent air, so that $\tau(F(\mathbf{x})) = \tau(F'(\mathbf{x}))$, no distortion is considered to have occurred at \mathbf{x} . 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 $F(\mathbf{x})$ and $F'(\mathbf{x})$ are considered to represent different features, by

$$P = \{\forall \mathbf{x} \in X \mid \tau(F(\mathbf{x})) \neq \tau(F'(\mathbf{x}))\}.$$

The distortion metric m_d is then defined as the ratio of the

number of points in P to the total number of voxels in X :

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

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.

Since the goal of this work is to produce refractive images which are free from noise, a coherence metric should thereby measure the improvements made in the image domain. We therefore consider the frequency distribution of an image by applying the discrete Fourier transform to 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. Given an image $i: \mathbb{N}^2 \rightarrow \mathbb{N}$ of resolution M, N , we have:

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

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

According to the periodicity property of the discrete Fourier transform, 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 analyse 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 the equation of m_c – 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.

Applying the two metrics to Figure 4, we obtain a set of quantitative results in Table 1, indicating that R-ANLD performs better than ANLD.

Direct comparisons are difficult, since values for m_d and m_c do not exactly correspond. 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 . In this experiment, we also included two other filters, namely Gaussian and *inhomogeneous linear diffusion* [Wei97]. When m_d is fixed at around 0.015 giving similar level of distortion for all methods, the R-ANLD method (and, to a lesser extent, the ANLD method) offers better improvement of smoothness than Gaussian and linear diffusion filters.

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

	t	0	20	40	60
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 in Figure 4. R-ANLD is shown to produce more coherent images with less distortion than ANLD.

Method	m_d	m_c
IHLD ($t = 35, \lambda = 60\%$)	0.0147	16.88
Gaussian ($\sigma = 3.18$)	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 .

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 .

around (for instance) 15.1, R-ANLD results in much less distortion than any of the other three methods.

Considering the combined effects of further improvements discussed in Sections 5 and 6, we conducted a large experiment with a wide range of parameters (for 2250 different combinations in total). Firstly, for each diffusivity function, the diffusion process was applied for $10\% \leq \lambda \leq 90\%$ (with lambda increasing in steps of 10%), for $0 \leq t \leq 150$, with $\sigma = 1$. Secondly, the effects of increasing σ in the range $2 \leq \sigma \leq 4$ were explored by generating datasets as above, but with λ restricted to the two values giving the most promising results from the data generated by the first set of parameters.

The set of parameters giving the lowest level of distortion for a particular maximum level of coherence are shown in Table 4, for a range of maxima for m_c . The results clearly show that g_C typically produces least distortion for some desired level of coherence. Whilst there is some variation for extreme maximum values of m_c , the middle range of the table is unequivocal. It should be noted that the extrema of the data represent levels of smoothing that are of little practical use; for values of m_c above 22, there is little observable ef-

Max. m_c	λ	t	Func.	σ	m_c	m_d
24.5	80	085	g_B	1	24.10	0.0032
24.0	80	045	g_B	1	23.66	0.0028
23.5	80	020	g_B	1	23.10	0.0024
23.0	50	035	g_B	1	22.52	0.0001
22.5	20	005	g_{PM2}	1	22.41	0.0000
22.0	10	010	g_C	1	21.91	0.0015
... another 20 records of g_C omitted here ...						
11.5	40	130	g_C	4	11.33	0.0236
11.0	60	130	g_{PM1}	4	10.98	0.0275
10.5	80	055	g_C	1	10.46	0.0356
10.0	80	075	g_C	1	10.00	0.0420
9.5	70	120	g_W	4	9.45	0.0477

Table 4: Diffusion parameters giving rise to least distortion for some maximum level of coherence.

fect on the levels of noise present in the resultant images, and for levels of m_c below 11.5, the distortion is considerable, although of course the acceptability of such results is application dependent.

When excluding g_C from the results, we find that g_{PM1} produces the next best results. In this way, we observed a strong ordering on the diffusivity functions. With the help of m_d and m_c , we may place the diffusivity functions in the following order of preference (in the context of rendering refractive images): $g_C \prec g_{PM1} \prec g_W \prec g_{PM2} \prec g_B$.

As the complexity of the algorithm depends primarily on the scale t and the dataset resolution R , the selection of the diffusivity function has relatively minor effects on the run-time.

8. Conclusions

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, shown at the beginning of this paper, depicts a volume dataset both (a) unprocessed and (b) after application of *regularised anisotropic nonlinear diffusion* (R-ANLD), and gives a compelling example of deploying such techniques in volume graphics.

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 and parameter variations for smoothing datasets. R-ANLD appears to be the most effective of the methods tested, both from a theoretical standpoint, and in practice. The diffusivity function proposed by Charbonnier *et al.* was also shown to offer more efficacious results, in comparison with other diffusivity functions in the literature.

The effective deployment of an appropriate denoising method in volume graphics allows the direct rendering of volume datasets with refractive effects, which can improve

depth perception significantly. As shown Figure 7, a noisy CT head dataset is visualised in conjunction with a brain dataset, and the CT head has its top part removed and is shown as a translucent object. Figure 7(a) was rendered using volume ray casting without refraction. Whilst it clearly displays the internal structure of the brain, it offers limited depth cues, and poor spatial perception of the brain and the skull. Figure 7(b), which results from direct surface rendering, provides slightly better depth cues, but incorrect visual representation of the spatial relationship.




Figure 7: Combinational visualisation of a CT head dataset and a brain dataset. The original CT head rendered with refraction in (c) offers more depth cue than those without refraction in (a) and (b). The visualisation produced after denoising in (d) conveys more information about the internal structure than (c).

Figure 7(c), also rendered using volume ray casting, introduces refraction into the visualisation, 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 7(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 7(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 7(d).

Refraction is a powerful tool in visualisation. We aim to develop denoising based on R-ANLD into a process requiring minimal user involvement, and integrate such a process into a general purpose discrete ray tracing pipeline for enabling high quality refraction rendering in volume graphics.

Acknowledgements

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. All images in this paper are rendered using *vlib* [WC01].

References

- [ALM92] ALVAREZ L., LIONS P., MOREL J.: Image selective smoothing and edge detection by nonlinear diffusion. *SIAM Journal Numerical Analysis* 29 (1992), 845–866.
- [ALTH99] ANGELINI E., LAINE A., TAKUMA S., HOMMA S.: Directional representations of 4D echocardiography for temporal quantification of LV volume. In *Proc. MICCAI'99* (Cambridge, 1999), pp. 430–440.
- [BSMH98] BLACK M., SAPIRO G., MARIMONT D., HEEGER D.: Robust anisotropic diffusion. *IEEE Transactions on Image Processing* 7, 3 (1998), 421–432.
- [BW75] BORN M., WOLF E.: *Principles of Optics*, fifth ed. Pergamon Press, New York, 1975.
- [Can86] CANNY J.: A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-8* (1986), 679–698.
- [CBAB94] CHARBONNIER P., BLANC-FÉRAUD L., AUBERT G., BARLAUD M.: Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proc. IEEE International Conference on Image Processing* (1994), vol. 2, pp. 168–172.
- [CM96] CONG G., MA S.: Nonlinear diffusion for early vision. In *Proc. 13th International Conference on Pattern Recognition* (August 1996), vol. A, pp. 403–406.
- [CT00] CHEN M., TUCKER J.: Constructive volume geometry. *Computer Graphics Forum* 19, 4 (2000), 281–293.
- [DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. SIGGRAPH 1999* (1999), pp. 317–324.
- [DMSB00] DESBRUN M., MEYER M., SCHRÖDER P., BARR A.: Anisotropic feature-preserving denoising of height fields and bivariate data. In *Proc. Graphics Interface 2000* (2000), pp. 145–152.
- [FvFH96] FOLEY J., VAN DAM A., FEINER S., HUGHES J.: *Computer Graphics: Principles and Practice*, second ed. Addison Wesley, September 1996.
- [HOH*96] HILTON M., OGDEN T., HATTERY D., EDEN G. F., JAWERTH B.: Wavelet denoising of functional MRI data. In *Wavelets in Biology and Medicine*, Aldroubi A., Unser M., (Eds.). CRC Press, 1996, pp. 93–112.
- [ML94] MARSCHNER S., LOBB R.: An evaluation of reconstruction filters for volume rendering. In *Proc. IEEE Visualization '94* (October 1994), pp. 100–107.
- [MMK*98] MÖLLER T., MUELLER K., KURZION Y., MACHIRAJU R., YAGEL R.: Design of accurate and smooth filters for function and derivative reconstruction. In *Proc. IEEE Symposium on Volume Visualization* (October 1998), pp. 143–151.
- [MMMY97] MÖLLER T., MACHIRAJU R., MUELLER K., YAGEL R.: A comparison of normal estimation schemes. In *Proc. IEEE Visualization '97* (Phoenix, AZ, November 1997), pp. 19–26.
- [MS92] MORETON H., SEQUIN C.: Functional optimization for fair surface design. In *Proc. SIGGRAPH '92* (1992), pp. 167–176.
- [NCKG00] NEUMANN L., CSÉBFAI B., KÖNIG A., GRÖLLER E.: Gradient estimation in volume data using 4D linear regression. *Computer Graphics Forum* 19, 3 (2000), C351–C357.
- [PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 7 (July 1990), 629–639.
- [PSZ01] PENG J., STRELA V., ZORIN D.: A simple algorithm for surface denoising. In *Proc. IEEE Visualization 2001* (San Diego, California, 2001), pp. 107–112.
- [RC01] RODGMAN D., CHEN M.: Refraction in discrete raytracing. In *Proc. Volume Graphics 2001* (New York, 2001), Mueller K., Kaufman A., (Eds.), Springer.
- [Rod03] RODGMAN D.: *Refraction in Volume Graphics*. PhD thesis, University of Wales, Swansea, Singleton Park, Swansea, SA2 8PP, Wales, UK, September 2003.
- [TWBO02] TASHDIZEN T., WHITAKER R., BURCHARD P., OSHER S.: Geometric surface smoothing via anisotropic diffusion of normals. In *Proc. IEEE Visualization 2002* (Boston, Massachusetts, 2002), pp. 125–132.
- [WC01] WINTER A., CHEN M.: *vlib*: A volume graphics API. In *Proc. Volume Graphics 2001* (New York, 2001), Springer, pp. 133–147.
- [Wei97] WEICKERT J.: A review of nonlinear diffusion filtering. In *Scale-Space Theories in Computer Vision* (1997), pp. 3–28.
- [Wei98] WEICKERT J.: *Anisotropic Diffusion in Image Processing*. ECMI Series, Teubner-Verlag, Stuttgart, Germany, 1998.
- [Whi80] WHITTED T.: An improved illumination model for shaded display. *Communications of the ACM* (1980), 343–349.
- [WTV98] WEICKERT J., TER HAAR ROMENY B., VIERGEVER M.: Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Transactions on Image Processing* 7 (1998), 398–410.
- [WW94] WELCH W., WITKIN A.: Free-form shape design using triangulated surfaces. In *Proc. SIGGRAPH '94* (1994), pp. 247–256.
- [YCK92] YAGEL R., COHEN D., KAUFMAN A.: Normal estimation in 3D discrete space. *The Visual Computer* (1992), 278–291.