

The Occlusion Spectrum for Volume Classification and Visualization

Carlos D. Correa, *Member, IEEE*, and Kwan-Liu Ma, *Senior Member, IEEE*




Fig. 1. Left: MRI of a meningeoma. Transfer functions (TF) based on boundaries (insets on the left: top, 1D TF with gradient modulation, bottom, 2D TF of intensity vs. gradient magnitude) cannot separate the tumor from the vessels (where gradients are also strong). A transfer function based on occlusion separates the tumor from the vessels and the ventricular structures from skull and skin. Right: CT dataset with contrast agent. Classification is difficult due to overlap between bone and vessel structures (see red color in ribcage in the insets). An occlusion-based TF properly classifies bone and also highlights internal structure (blue).

Abstract— Despite the ever-growing improvements on graphics processing units and computational power, classifying 3D volume data remains a challenge. In this paper, we present a new method for classifying volume data based on the ambient occlusion of voxels. This information stems from the observation that most volumes of a certain type, e.g., CT, MRI or flow simulation, contain occlusion patterns that reveal the spatial structure of their materials or features. Furthermore, these patterns appear to emerge consistently for different data sets of the same type. We call this collection of patterns the *occlusion spectrum* of a dataset. We show that using this occlusion spectrum leads to better two-dimensional transfer functions that can help classify complex data sets in terms of the spatial relationships among features. In general, the ambient occlusion of a voxel can be interpreted as a weighted average of the intensities in a spherical neighborhood around the voxel. Different weighting schemes determine the ability to separate structures of interest in the occlusion spectrum. We present a general methodology for finding such a weighting. We show results of our approach in 3D imaging for different applications, including brain and breast tumor detection and the visualization of turbulent flow.

Index Terms—Transfer functions, Ambient Occlusion, Volume Rendering, Interactive Classification

1 INTRODUCTION

One of the issues in classifying and visualizing 3D volumetric data sets is the lack of explicit geometric information and limited semantics. However, the reliance on 3D images to describe complex 3D objects and processes is what permits scientists to quickly visualize the results of MRI scans or flow simulations without time-consuming pre-processing or segmentation. Due to the speed and parallelization of texture processing, most visualization systems rely on classification schemes that extract *local* information to better represent the materials in a volume data set. In the simplest case, the value at each voxel suffices. In most cases, however, additional information, such as gradient and higher order derivatives, are necessary for classification [16, 11]. Nonetheless, the information required to compute these

derivatives still can be categorized as local. In this paper, we advocate for the use of more global metrics to characterize the structures in a 3D volumetric object. This global metric should encode enough information to distinguish between features that appear coherently within a certain spatial neighborhood.

In this work, we study the issue of using occlusion information to classify volumetric objects. Specifically, we use the ambient occlusion of a voxel as a metric for classification. Ambient occlusion has the advantage of being viewpoint independent and encodes the average contribution of the surrounding neighborhood to the visibility of every voxel in the volume. We noticed that the distribution of occlusion for certain features varies coherently depending on the relationship between these features and their surroundings. For example, in medical images, bones have an occlusion distribution clearly differentiated from skin tissue or contrast-enhanced vessels, which may have the same intensity in the image modality. Therefore, we can characterize some of these components more clearly when considering their ambient occlusion contribution. Traditionally, ambient occlusion has been used for improving the rendering of volumetric models, particularly isosurfaces. In our paper, we use a more general notion of ambient occlusion, which includes the contribution of all voxels around

Carlos D. Correa and Kwan-Liu Ma are with the University of California, Davis, E-mail: {correac,ma}@cs.ucdavis.edu.

Manuscript received 31 March 2009; accepted 27 July 2009; posted online 11 October 2009; mailed on 5 October 2009.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

a given point to its visibility. Rather than a rendering quantity, we use the result as an independent variable that can be combined with intensity value to provide more meaningful transfer functions.

We refer to the distribution of ambient occlusion in a data set as the *occlusion spectrum* of the data set. When combined with the intensity values, the 2D distribution provides a classification space that separates features that are highly occluded, e.g., those at the interior of objects, from those that are not occluded, such as the outer layers of an object. For example, MR images often exhibit the same intensity values for certain features that are clearly internal, e.g., bones, with others that are at the boundaries, such as skin. An example is shown in Figure 2, where we highlight some of the structures that appear when we select different regions in the spectrum, such as ventricular anatomy, skull, brain and skin. Analogously, flow simulations often exhibit regions where the internal and external characteristics of flow differ greatly. The occlusion spectrum of these data sets enables scientists to separate regions of interest depending on their overall spatial characteristics and formulate hypotheses about the spatial nature of the quantity they are visualizing.

Using our method for classification is advantageous because ambient occlusion: (1) encodes the contribution of the voxels in the neighborhood of a given point with a single scalar value, rather than an n -dimensional vector or histogram; (2) is easy to compute for sampled volumetric data and can be implemented rather effortlessly in current programmable hardware, and (3) exhibits spatial coherence, important for identifying features and their spatial relationships.

At first, the occlusion spectrum leads easily to 2D transfer functions, where one dimension is intensity value and the other is occlusion. This type of classification proves very useful for a large number of datasets from a variety of domains. We show that the occlusion spectrum helps isolate regions that are spatially concentrated in regions of varying occlusion, and represents a different classification space. Instead of highlighting boundaries, as many current classification methods do, they highlight structure. Nonetheless, 2D transfer function editors may not be easy to understand for non-expert users. For this reason, we also present a simpler editor that preserves the simplicity of 1D transfer function editors with simple manipulations to control the desired effect of occlusion. The issue of occlusion also poses a question of what exactly makes a given intensity value to be more occluding than another. In this paper, we present a method that automatically finds a visibility mapping that results in the occlusion spectrum that maximizes the variance along certain intensity intervals of interest. This ensures that the features along those intensity values are more likely to be separated than others. Through a number of examples, we show that the occlusion spectrum, when used for classification, is a powerful technique aimed at extracting features that share certain spatial characteristics. Currently, this requires either segmentation or cutaway views. We show that we can obtain similar or better results without expensive data pre-processing.

2 RELATED WORK

The classification of volumetric models has been the focus of research since the inception of visualization systems and volume graphics. The ever-ubiquitous 1D transfer function seems still the most common approach, despite the advances made towards higher dimensional transfer functions, such as those based on first (i.e., gradient) [16, 11] and second derivatives (i.e., curvature) [8, 12]. Yet these methods are popular in visualization systems because they use mostly local information, which makes them easy to implement and readily fit for parallel computation in modern GPUs. Recently, a number of techniques, which attempt to move towards gathering more global information, have been proposed, leading to a continuous spectrum of techniques. In one end, next to convolution-based approaches (e.g., gradient and curvature) are lighting transfer functions [17] and shape detection filters [26]. These methods are still based on derivatives and highlight material boundaries. The latter, however, uses the matrix of second derivatives to better identify the local shape of tissues (i.e., line, thin plate and blob). Several researchers have recognized the limitations of these gradient-based approaches. For example, Pfister et al. [20]




Fig. 2. The Occlusion Spectrum for an MRI data set. Anatomical structures (e.g., skull, brain, ventricles) appear depending on how internal they are.

and Lundstrom et al. [18] noted that noise usually distorts material boundaries. Rottger et al. also recognized that spatial information is inherently lost in traditional 1D or 2D histograms, and expanded these with spatial information [23]. As a solution to the blurring of gradient-based 2D histograms, Sereda et al. propose a different way of selecting boundaries, searching for high and low values in paths that follow the gradient near the voxels in a boundary [28].

Lundstrom et al. proposed the use of local histograms to better represent the distribution of intensity values in a given neighborhood [18]. This departs from convolution-based approaches in that it requires a larger neighborhood. Their results show that the use of local histograms greatly improve tissue separation for the case of overlapping intensity ranges. The search for better tissue separation has motivated a series of more global approaches, which require additional information about the entire data set to aid classification. These approaches rely more on finding *structure* than highlighting boundaries. Takahashi uses topology to guide the transfer function design [4, 29]. In particular, Takahashi et al. has used topological relationships between structures to measure the inclusion of isosurfaces [30]. They propose a similar dimension, called *inclusion level*, that focuses on structure rather than boundaries. Their method requires finding critical points, which may not be robust to noise, so they must rely on a thresholding mechanism. Zhang and Bajaj follow a similar approach for the visualization of protein pockets, using signed distance transforms to quantify the inclusion of isosurfaces [34]. Due to the issue with noise, they restrict their application to smooth free-form surfaces. In this paper, our notion of occlusion is more general and not restricted to nested structures. Correa and Ma use size-based transfer functions to classify features based on their size [3]. Huang and Ma use region growing to guide the definition of 2D transfer functions [9].

One issue with these approaches is the reliance on complex high dimensional spaces for classification. A number of user interface (UI) mechanisms have been proposed, including the contour spectrum [1], which reduces classification to handling a series of curves, transfer function widgets [14] and user painting [32]. Rezk-Salama et al. use high-level semantics to define and adapt widgets from one data set to another [21]. In this paper, we explore a novel dimension for classification, namely the ambient occlusion of individual voxels. Ambient occlusion is a single dimension that summarizes the contribution of voxels in a large neighborhood of a given point and is spatially coherent.

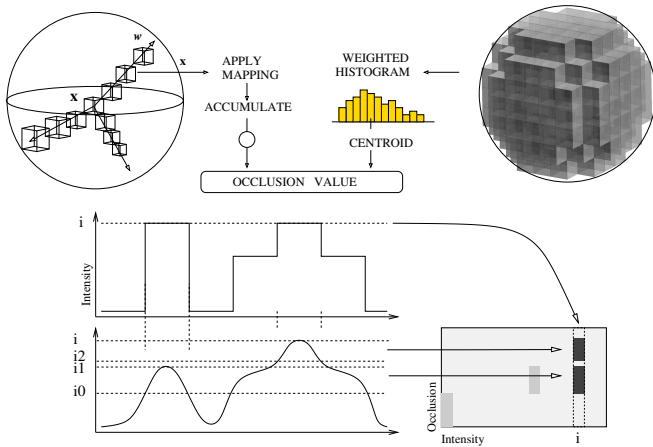


Fig. 3. Top: Ambient occlusion is obtained as the average of occlusion contribution from samples around a voxel. It is equivalent to finding the centroid of the weighted histogram of intensities around that voxel. Bottom: An example of an occlusion spectrum for a simple 1D profile. Two structures of intensity i are depicted, one surrounded by a low intensity, the other with a medium intensity. The ambient occlusion is essentially an averaging operation. As a result, the intensities for the first structure are “smoothed” down to a low intensity in the range $(i0, i1)$, while the other structure is smoothed down to the range $(i2, i)$. When we plot these intervals in a 2D histogram, we see a clear separation of structures.

Ambient occlusion was introduced by Zhukov [35] as the notion of *obscurance* to model the ambient illumination of an object without costly global illumination operations. Since then, ambient occlusion has become a fast technique for obtaining high quality renderings of illuminated objects and has been successfully implemented in the GPU [27, 5, 2]. Since ambient occlusion requires evaluating the visibility of a point with respect to a number of occluders, some have proposed acceleration techniques such as occlusion fields [15] and the use of pre-computed information such as local histograms [24] or mutual probabilities [6]. For a complete survey on ambient occlusion techniques, refer to Knecht [13], and Mendez and Sbert [19].

Recently, there has been a growing interest in improving volume rendering with ambient occlusion and other approximation of global illumination. Ritschel uses the GPU to compute the visibility of volume data sets and provide natural illumination effects such as soft shadows and attenuation [22]. Wyman et al. [33] use precomputed illumination volumes to incorporate global illumination of isosurfaces. Ropinski et al. [24] present a system for dynamic ambient occlusion that adapts to changes in visibility such as transfer function manipulation. The use of ambient occlusion was exploited by Tarini et al. to enhance the visualization of large molecules. They show that ambient occlusion results in an enhanced perception of the 3D shape of large proteins [31]. Ruiz et al. also use obscurances to generate high quality renderings of volumetric objects at a low cost. Coupled with color bleeding, their framework results in realistic volume rendered images effective for visualization [25]. For more details on ambient occlusion for volume rendering, see the tutorial notes by Hadwiger et al. [7]. In our work, although the process of constructing the occlusion spectrum is borrowed from ambient occlusion computation, our focus is quite different. Rather than using occlusion exclusively for rendering, we use occlusion for classification. This implies a fundamental difference with our predecessors, in that we do not rely on pre-classification of tissue to determine visibility.

3 THE OCCLUSION SPECTRUM

The occlusion spectrum refers to the distribution of ambient occlusion of a given intensity value in a 3D volumetric object. We can represent this as a 2D histogram, where one axis is the intensity value and the other axis is the occlusion. Unlike 2D histograms based on

| | | |
|-----------------------|---|--|
| Linear Ramp | $M(\mathbf{x}) = S(\mathbf{x})$ | |
| Truncated Linear Ramp | $M(\mathbf{x}) = \begin{cases} S(\mathbf{x}) & \tau_0 < S(\mathbf{x}) < \tau_1 \\ 0 & \text{otherwise} \end{cases}$ | |
| Distance weighted | $M(\mathbf{x}) = S(\mathbf{x})e^{-\ \mathbf{x}-\mathbf{x}_c\ ^2}$ | |
| Opacity weighted | $M(\mathbf{x}) = \alpha(\mathbf{x})$ | |

Table 1. Visibility mappings for the computation of occlusion, where $S(\mathbf{x})$ is the intensity of a point \mathbf{x} .

gradient magnitude, this spectrum does not highlight boundaries, but *structures*. A concentration in the 2D spectrum towards the higher occlusion values indicates structures that are more “internal” since they are more likely to be occluded, whereas concentrations in the lower occlusion values indicate more “external” structures. An example of the occlusion spectrum of an MRI data set is shown in Figure 2. Here, we see an intensity value interval where there are structures such as skin, skull and lateral ventricles. Once they are plotted in terms of their occlusion, they can be separated clearly. Similarly, brain intensities can be separated from occluding skin tissue. As a result, we can obtain a visualization that isolates the brain without the need for segmentation.

To compute the occlusion spectrum, we turn to the *ambient occlusion* of a point, which represents the obscurance of the point due to the neighboring voxels in a volume. Unlike traditional ambient occlusion, which only computes this quantity for visible points, we define an adaptive visibility mapping that considers every voxel in a neighborhood around the voxel. We show that this quantity is equivalent to computing the centroid of the weighted histogram of intensities around the voxel. We then construct the spectrum as the 2D histogram of intensity vs. occlusion. One of the key properties of this histogram is the ability to separate structures of interest. Depending on the visibility mapping function, this separation may be impaired. For this reason, we present a general methodology that finds the best parameters for the visibility function that maximizes the likelihood of separation. These steps are detailed in the following sections.

3.1 Ambient Occlusion

To quantify and measure the occlusion of a voxel, we turn to ambient occlusion, used widely to approximate the ambient attenuation of a point given the surrounding scene. This can be expressed as:

$$AO(\mathbf{x}) = \frac{1}{\pi} \int_{\Omega} (1 - V(\mathbf{x}, \omega)) (\omega \cdot \mathbf{n}) d\omega \quad (1)$$

where \mathbf{x} is the location of a point or voxel, \mathbf{n} represents the normal of the surface through this point and $V(\mathbf{x}, \omega)$ is the visibility of \mathbf{x} along a direction ω . The directions ω are taken to cover the hemisphere Ω defined by the normal of the point. When $AO(\mathbf{x}) = 0$, the point is unoccluded.

In this paper, we use a more general notion of occlusion, which also takes into account the intensities in the other hemisphere. Therefore, we define occlusion as the weighted average of the visibility of a point along directions in a sphere (Fig. 3a):

$$O(\mathbf{x}) \approx \frac{1}{N} \sum_{\phi=0}^{\pi} \sum_{\theta=0}^{2\pi} A(\mathbf{x}, \omega(\theta, \phi)) \quad (2)$$

where N is the number of neighbors and $A(\mathbf{x}, \omega)$ is the directional occlusion of a point along direction ω , defined as:

$$A(\mathbf{x}, \omega) = \sum_{t=0}^T M(\mathbf{x} + t\omega) \quad (3)$$

where T is the number of samples along direction ω and $M(\mathbf{x})$ is a visibility mapping function of a sample \mathbf{x} . In the case of traditional

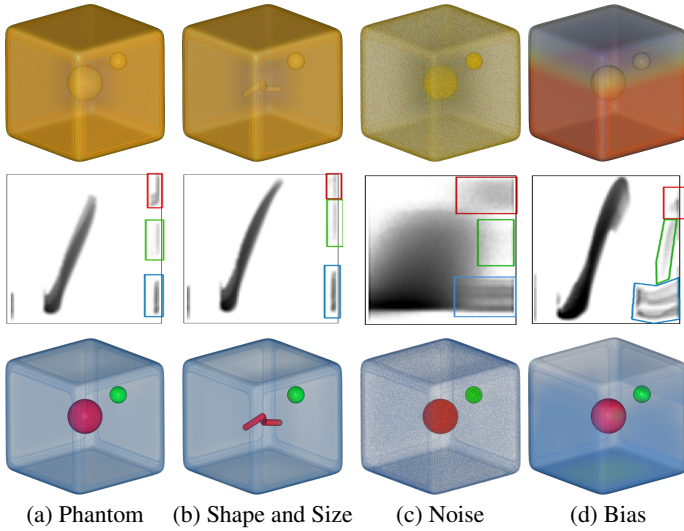


Fig. 4. Properties of the occlusion spectrum. (a) Phantom data set. Note the separation of structures even though they share the same intensity. (b) Structures of different shape and size. (c) Gaussian noise of $\sigma = 1.0$. (d) Multiplicative bias.

ambient occlusion on isosurfaces, the mapping function is a binary function that is 0 when the intensity is the isovalue of interest and the point is in the hemisphere defined by the gradient of the central voxel. Here, we define more general visibility mapping, which should retain as much information as possible about the distribution of intensities in the neighborhood of a voxel. Table 1 summarizes some of the mappings we have explored, including data-centric approaches such as linear ramps and Gaussian weighted neighborhoods, and rendering-centric approaches such as user-defined opacity mappings.

3.2 Rationale

This concept of occlusion can be considered as a weighted average of the neighborhood surrounding a voxel. The rationale behind this classification space is that occlusion can be considered as the convolution of the volume with a low pass filter. As long as the filter has a size larger than the structures we want to classify, the average is affected by the distribution of voxels surrounding this structure. Consider Figure 3, where we plot a 1D intensity profile, composed of two structures of intensity i , one of which is surrounded by a low intensity and the second surrounded by a medium intensity. When we convolve the profile with a low-pass filter, i.e., we compute the average response, the intensity of the first structure decreases to interval (i_0, i_1) and the second structure decreases to interval (i_2, i) , where $i_1 < i_2$. When we plot these in a 2D histogram, the two regions with the same intensity can be separated.

It can be seen that the occlusion is the centroid of the weighted histogram of the neighborhood of a voxel. Let f_i denote the frequency of occurrence of $M(\mathbf{x}) = i$. Therefore, $\sum_i f_i = N$, where N is the number of voxels in the spherical neighborhood $N_R(\mathbf{x})$ of voxel \mathbf{x} and radius R .

$$O(\mathbf{x}_0) = \frac{1}{N} \sum_{\mathbf{x} \in N_R(\mathbf{x}_0)} M(\mathbf{x}) = \frac{\sum_i i f_i}{\sum_i f_i} \quad (4)$$

3.3 Occlusion Properties

The occlusion spectrum has a series of properties essential for robust classification.

Coherence. The occlusion spectrum is coherent. That is, for a pair of neighboring points in the volume, it is expected that occlusion varies smoothly. This is achieved by its definition, since occlusion is essentially a convolution. Let G be a convolution filter describing the

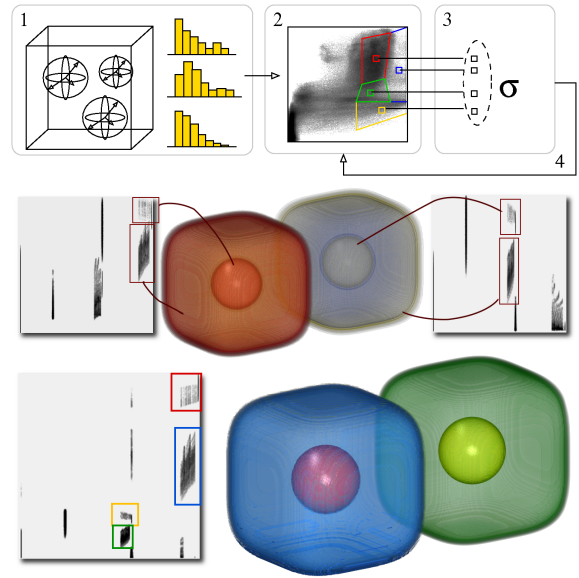


Fig. 5. Top: Adaptive Mapping Selection is computed in a series of steps: (1) Obtain local histograms (2) compute the occlusion spectrum for a set of mapping parameters and clustering (3) obtain variance of means (4) repeat for next set of parameters. The mapping is the one with maximum variance of means. Middle: A data set consisting of two sets of nested structures, the first of high intensity surrounded by medium intensity (left), the other medium surrounded by low intensity (right). Mapping selection that maximizes variance for the high intensity (left) fails to separate the other. Conversely, a mapping that maximizes for the medium intensity (right) fails to separate the high intensities. Bottom: An adaptive mapping separates the structures for both high and medium intensities, as seen in the occlusion spectrum.

occlusion operation. The occlusion can be described as the convolution: $O(\mathbf{x}) = G \circ S(\mathbf{x})$. The spatial derivative of the occlusion is then:

$$\frac{\partial O(\mathbf{x})}{\partial \mathbf{x}} = \frac{\partial G \circ S(\mathbf{x})}{\partial \mathbf{x}} = G \circ \frac{\partial S(\mathbf{x})}{\partial \mathbf{x}} \quad (5)$$

Thus, the occlusion retains the coherence properties of the volume, as long as the convolution filter, determined by the mapping function G is continuous.

Noise and bias. Occlusion is also inherently robust to additive noise and multiplicative bias, because of the averaging performed by the convolution. Let us consider a transformed scalar field $S'(\mathbf{x}) = \beta(\mathbf{x})S(\mathbf{x}) + E$, where β is a multiplicative bias and $E \sim N(0, \sigma^2)$ is an additive noise of mean 0 and variance σ^2 . Therefore, the occlusion on the transformed scalar field is

$$G \circ S'(\mathbf{x}) \leq \beta_{max} G \circ S(\mathbf{x}) \quad (6)$$

for a maximum bias β_{max} in the neighborhood of a point, which means that the occlusion distribution is not transformed by the additive noise but is scaled due to bias. As long as the variance of the bias does not exceed the size of the largest structure we can separate, this scaling does not affect the separability in the occlusion space. An example is shown in Figure 4.

Shape and Size. Occlusion is also robust to changes in shape and size, as shown in Figure 4, as long as the neighborhood is larger than the size of structures we wish to detect. That is, for larger structures, they voxels become self-occluded and the neighborhood intensities have little effect in the occlusion value.

To understand the effect of size, we ran a sensitivity plot of the occlusion computation for different neighborhood sizes. Figure 6 shows four sensitivity plots for two modalities. The horizontal axes represent the intensity value and radius of the occlusion sphere,




Fig. 6. Sensitivity plots of the occlusion spectrum for two data set types with respect to radius of ambient occlusion sphere. In the plane, we show intensity values vs. radius of the occlusion sphere. The vertical axis the variance of occlusion for each intensity. Notice the consistent peaks around radii of 50 – 100 voxels, for a volume of diameter 256.

and the vertical shows the (normalized) variance of the occlusion spectrum. The larger the area under a curve is, the larger the overall variance of the occlusion spectrum becomes, and therefore, so the likelihood of separability of features. As expected, small neighborhoods do not provide enough variance. Large neighborhoods, on the other hand, also drop variance when they are much larger than the structures we want to detect. Therefore, we consistently see a shape that indicates the best radius in the middle of a given interval. With the help of these profiles, one can estimate the extents of the occlusion neighborhood for detecting features of a certain size.

3.4 Adaptive Mapping Selection

One of the key advantages of the occlusion spectrum is its ability to separate structures with the same intensity based on the data distribution in its neighborhood. Because the occlusion depends on a visibility mapping function, the effectiveness of the spectrum depends on the parameters of this mapping function. For example, when we consider a truncated ramp function, the cutoff values τ_0 and τ_1 affect the average intensity of the local neighborhood. Furthermore, a given mapping can maximize the separability of certain structures in one intensity, but hinders it for another intensity. Let us consider a phantom data set consisting of two spheres surrounded by a hollowed structure. One of the spheres is a high intensity surrounded by a medium intensity, and the other is of medium intensity surrounded by low intensity, as shown in Figure 5. A mapping function is found that separates the sphere from the outer layer for the high intensity, but fails to separate the one for the medium intensity. Another mapping function works for the medium intensity but fails for the high one. Therefore, it becomes necessary to select a mapping adaptively, depending on the intensities of interest the user wishes to classify. In general, this can be accomplished by finding the best parameters that maximize the variance of means of the occlusion spectrum for the intensity interval of interest. This, however, requires to compute the occlusion for each combination of parameters, which is computationally expensive. Instead, we present a faster mechanism based on local histograms, depicted in Figure 5(Top). For each intensity interval of interest we compute the local histogram for each voxel. The occlusion can be found as the centroid of each of these histograms. Next, we compute the distribution of occlusion for the intensity intervals and cluster them. Finally, we find the variance of the means. We repeat this process for each parameter in the mapping function. The mapping is the one that maximizes the




Fig. 7. Occlusion-based Classification of a meningioma. Left: 2D transfer function based on intensity and gradient magnitude. Although transfer functions based on boundaries can highlight the tumor and vessels from noise and other unrelated structures, these intensities often overlap and the distribution of gradient is not associated with any particular structure. The result is a commingled visualization of tumor and vessels. Right: An occlusion-based transfer function gives higher visibility to the tumor, and it is clearly separated from the vessels. Note also a cleaner visualization that removes some of the noise in the forehead.

variance of means of the occlusion distribution. A result of such an adaptive mapping is seen in Figure 5, where we can clearly see the structures of interest clearly separated in the occlusion spectrum.

4 OCCLUSION TRANSFER FUNCTION (OTF)

The most immediate application of the occlusion spectrum is the classification based on occlusion. An occlusion transfer function is therefore a mapping from the space spanned by the scalar values S and the occlusion O into color and opacity, $S \times O \mapsto [0, 1]^4$. By tagging different regions of the resulting 2D histogram and assigning color and opacity, users can select regions with similar intensity values but in rather different locations within the data set.

An OTF differs from other gradient-based transfer functions in that material boundaries are not the main criteria for deciding what to classify. Although regions of certain homogeneity might exhibit different occlusion signatures than those on the boundary of materials, the main factor remains the degree of occlusion, independent (for most part) of whether they have large or small gradients. An example is shown in Figure 7, which depicts a meningioma or brain tumor. On the left, we show the result of classifying based on intensity and gradient magnitude using a 2D transfer function. A gradient-based transfer function does not help identification much, since the overlapping intensities also exhibit large gradients. With an OTF (right), we can isolate the tumor from the arteries and veins, and the occluding tissue due to noise and overlapping intensities. This works robustly for tumors and similar structures because these are features that are consistently surrounded by a certain intensity range (e.g., brain tissue), that is different from other structures of similar intensity (e.g., air in the case of the occluding tissue on the skull). Although tumors may appear in different locations (thereby changing the occlusion values), they are still separable from low occluded structures such as skin.

4.1 Occlusion Transfer Function Editor

One of the issues of high-dimensional transfer functions is the reliance on complex widgets and spaces to specify the color and opacity mappings. Even for 2D transfer functions, the detection of material boundaries as arcs in the resulting histogram may not be straightforward as affected by noise and bias. It is no surprise that most medical systems still rely on simple 1D editors to highlight features of interest, despite the obvious and well-studied limitations of such approach. The 2D occlusion transfer function is no exception. Effective classification requires the user to control widgets in 2D space, while keeping track of an elusive third dimension that tacitly represents opacity. Representing these in a 3D space would only complicate the matter. However,

we have found that the idea of occlusion fits well the idea of a 2D space. This is mainly due to the intrinsic orthogonality of these two dimensions. Any intensity value can potentially exhibit any degree of occlusion, so that the y-dimension can be grasped intuitively as the interiority of a material.

Nonetheless, we believe that our approach can better serve the general users of visualizations with a simpler interface. To achieve this, we make the following observations: (1) The 1D classification is ubiquitous and easy to understand. (2) Assigning different colors to the same intensity value at different occlusion levels may be misleading, where color is used to indicate the magnitude of a given quantity.

In our interface, we decouple the 2D classification space into two 1D spaces. This is motivated by the notion of improving the likelihood of a point being part of a given structure one dimension at a time. The opacity function of a sample point is defined as:

$$\alpha(\mathbf{x}) = \alpha_S(S(\mathbf{x}))\alpha_O(S(\mathbf{x}), O(\mathbf{x})) \quad (7)$$

where α_S is an opacity mapping based on intensity, and α_O is an opacity mapping based on occlusion. In our case, we define $\alpha_S(s) = \sum_i G_{\mu_i, \sigma_i}(s)$ as a sum of Gaussians. The first space (Figure 8) retains the characteristics of a typical 1D classification based on intensity values, including color selection. The X dimension denotes the intensity values, while the Y dimension denotes opacity. In the second space (bottom), we retain the X dimension as the intensity values and add the occlusion spectrum as a plot. The Y dimension corresponds to occlusion. To change the opacity based on occlusion, we use *occlusion curves*, which span the entire domain, but can be adjusted in the Y-dimension (occlusion). These curves represent the means of Gaussian bells and the size of the area around the curve represents the standard deviation of the Gaussian. Therefore, the opacity mapping is defined as: $\alpha_O(s, o) = G_{\mu_s, \sigma_s}(o)$ for an intensity value s and an occlusion o .

Interacting with these curves simplifies much of the complexity added by dimensional transfer functions, although it is limited to a single extra variable, in this case opacity. Figure 8 depicts the classification process. On top, the intensity intervals in red and orange are associated with high occlusion, isolating the meningioma. Next, the user highlights a different intensity interval corresponding to brain, but also to occluding tissue such as skin. Therefore, at the bottom, the user moves the curve towards the upper end, selecting the structures with high occlusion for that intensity interval, namely the brain tissue.

4.2 Implementation

The occlusion spectrum is the 2D histogram resulting from two volumes, the original intensity volume, and the ambient occlusion volume. In our implementation, since ambient occlusion is used for classification, we pre-compute it as a separate volume. This makes the rendering stage of our technique computationally comparable to that of 2D transfer functions based on gradient magnitude. The occlusion volume, since it is essentially an average, can be of lower resolution than the original data. To compute the ambient occlusion volume we use the programmability of current GPUs and the *render-to-3D-texture* capabilities. We render a number of quadrilaterals corresponding to the different slices of the ambient occlusion volume. Each of this quadrilaterals is processed in parallel in a fragment shader that computes the occlusion at every voxel by explicitly encoding Eq. 2. After all slices are processed, the result is a 3D volume containing the occlusion at every voxel. Finally, classification is incorporated into a GPU-based ray casting shader that uses a 2D texture lookup to obtain the color and opacity of each sample point according to their intensity and occlusion values.

5 CASE STUDIES

5.1 Cancer Diagnosis on Breast CT

According to the National Cancer Institute in the U.S [10], breast cancer incidence in women in the United States is 1 to 8, or about 13%. The need to diagnose early a breast cancer tumor is becoming increasingly pressing as the imaging modalities improve and their costs are reduced. The goal of screening exams is to find cancers before they




Fig. 8. Example of 1D editors for occlusion. Top: the user highlights the skin intensities (on the left, low opacity), vessels (red) and the meningioma (orange). The occlusion curve appears towards the top, highlighting only internal structures (the tumor). Middle: the user increases the opacity of brain intensities. Fatty tissue and other structures share the same intensity and occlude the brain. Bottom: The user moves the occlusion curve towards the top, so that only the internal structure (the brain) is shown while the outer layers are left transparent.

start to cause symptoms. Since cancers at this stage are usually small, the ability to extract the right information from the different imaging modalities becomes crucial. Current imaging methods include ultrasound, digital mammography, MRI, positron emission tomography (PET) and CT scans [10]. We used the occlusion spectrum to classify tumors at varying stages from a set of patients obtained using breast CT scans. Figure 9 shows a series of images from three different patients with different characteristics. In our experiments, we found that the occlusion spectra of these data sets are very similar, even with dramatic changes in the visibility. For example, Figure 9(middle) shows a patient with a breast implant. Due to the quantitative nature of CT, the implant shares the same intensity values of the tumor and glandular tissue. In Figure 9 we see zoomed-in views of regions of interest, presumably containing cancerous tissue, for traditional (using intensity and gradient modulation) and occlusion-based classification, on the left and right insets, respectively. In general, intensity and gradient-based classification of the 3D volume can only focus on the glandular tissue and requires low opacity to gain sufficient visibility. The result is a rather homogeneous collection of fuzzy blobs, often obscured by other non glandular tissue. If the radiologist requires to add other tissue for context, the tumor would not be seen unless a cutting plane is introduced. With occlusion-based classification, on the other hand, the results are more clearly segmented from the surrounding tissue and can even show some of the vascular structures that feed the growing tumor. In our preliminary evaluation of our technique, an expert in breast cancer imaging found the results of our classification most useful since we show, for the first time, a clear view of the tumor and the surrounding blood vessels. In traditional imaging, the radiologist goes back and forth between a set of 2D images, and, depending on the size and orientation of certain features, determines whether a particular blob is a vessel or a tumor.

We have also discussed our results with a surgeon specialized in cancer treatment. As suggested by the expert, one important implication of occlusion-based classification is the extraction of *layers*, essential for the planning of surgical procedures, especially for the treatment of cancer. Although radiologists are still used to 2D medical visualiza-




Fig. 9. Occlusion-based classification of CT scans for breast cancer detection. For each data set, we show at the bottom the result of classification based on intensity and gradient magnitude modulation vs. occlusion-based classification, for regions containing tumors. Boundary-based transfer functions get visibility of the tumor at the expense of reducing the opacity and clarity of the image. With an occlusion transfer function, we can render a more defined surface of the tumor, as well as the nearby vessels and skin for context.

tion metaphors, the use of 3D imaging often lets them see new structures that may imply updates in a planned surgical procedure. The ability to depict the different layers surrounding a tumor lets surgeons determine the best course of action and whether a particular surgical procedure is warranted. For example, certain tumors can be treated with thermal ablation as an alternative to open surgery, reducing the health risks of the patient.

5.2 Simulation of 3D Phenomena

Another interesting domain is the visualization of simulation data sets. Although the occlusion spectrum for anatomical data sets seems to correlate to specific organs and tissues, there is no evident analogy in simulation data sets. However, the occlusion spectrum tells us about the relative spatial distribution of quantities of interest. The occlusion transfer function therefore enables the scientist to formulate questions such as what is the distribution of a given range of values, e.g., velocity or vorticity, in the interior of the volume as compared to its exterior. For example, in combustion simulations, regions of weakly burning flames are often obscured by strongly burning regions. These obscured regions, however, are of most interest to scientists interested in understanding processes such as reignition.

A similar problem occurs in hydrodynamic simulations such as the simulation of core collapse supernovae. The scientists are interested in visualizing the interplay of different quantities such as pressure, density and energy that result in shock waves and the ensuing explosion. The direct visualization of entropy has been a common way of understanding the formation and behavior of these shockwaves. However, often times these shockwaves are formed in both the internal regions near the core and the outermost regions. Therefore, visualizing these quantities near the core becomes difficult. An example is shown in Figure 10, where the visualization of entropy at the interior is obscured by the outer shells of entropy. As an alternative, most visualization systems let the user define a cutaway, which provides visibility of the core and the surrounding entropy (Figure 10(middle)) However, this results in the formation of structures due to the cut that are not in the original data set and may be misleading. Here we can classify in terms of occlusion. In this transfer function, we maintain the same color mapping and intensity opacity, but modulate opacity with respect to occlusion, so that innermost regions are assigned a higher opacity. The result is shown in Figure 10(right). Note that the same internal structures are shown, but without the misleading issues introduced by a cutaway. An important consideration when applying occlusion to simulation datasets is their temporal nature. From time step to time

step, the occlusion spectrum varies. Therefore, an occlusion transfer function that highlights certain region may not be as effective for the next transfer function. This, however, is not exclusive of the occlusion spectrum but also to general 1D and 2D histogram, and general solutions to both must be sought after.

6 CONCLUSION

We have presented a novel technique for classifying volumetric objects based on occlusion. The issue of occlusion has been the focus of numerous efforts, mostly in an attempt to minimize it and provide visibility of otherwise obscured structures. In this paper, we have shown that we can understand occlusion in a rather different way, which is a scalar field that encodes, in a single dimension, the spatial structure of complex datasets. We presented the *occlusion spectrum* of a volumetric dataset, which encodes the 2D distribution of intensity values and occlusion. This space separates structures based on the distribution of intensities in their neighborhood. Therefore, the occlusion dimension directly maps, in most cases, to the internality of a structure. Features that appear isolated or at the exterior of larger structures can be clearly separated from those at the interior.

In our validation with medical imaging experts we found that *occlusion* (1) is an easily grasped concept that relates directly to the way they interact with anatomical data and that helps decide certain procedures such as surgical planning, and (2) leads to better quantification of cancer tumors due to the ability to isolate them without the need for segmentation. One of the issues of 2D transfer functions is the higher dimensionality of the classification space, which implies a number of user interface challenges. In our experiments, we found that the occlusion spectrum of data sets of certain type maintain some similarity. Therefore, it is possible to generate classification templates that can be re-targeted across data sets without much user intervention. Since ambient occlusion is spatially coherent and easy to implement in contemporary graphics hardware, we believe it can be deployed in most visualization systems without much effort. Combined with other capabilities, such as cutaway planes and advanced lighting, occlusion-based volume rendering can be used to obtain images with unprecedented clarity and quality.

ACKNOWLEDGEMENTS

The authors wish to thank the sources of our data sets: Prof. John M. Boone from the Biomedical Engineering department at UC Davis, the Terascale Supernova Initiative, the OsiriX Foundation and Jeff Orchard.




Fig. 10. Occlusion-based classification of a supernova simulation. Left: traditional classification via a 1D transfer function based on intensity value (entropy) enhanced with gradient magnitude modulation. Middle: To visualize internal dynamics, scientists often use cutaways, which introduce additional structure due to the cut. Right: an occlusion TF lets scientists see internal dynamics without losing information due to cuts.

REFERENCES

- [1] C. L. Bajaj, V. Pascucci, and D. R. Schikore. The contour spectrum. In *Proc. IEEE Visualization '97*, pages 167–173, Los Alamitos, CA, USA, 1997. IEEE Computer Society Press.
- [2] M. Bunnell. *GPU Gems 2*, chapter Dynamic Ambient Occlusion and Indirect Lighting, pages 223–233. Addison Wesley, 2005.
- [3] C. Correa and K.-L. Ma. Size-based transfer functions: A new volume exploration technique. *IEEE Trans. on Visualization and Computer Graphics*, 14(6):1380–1387, 2008.
- [4] I. Fujishiro, T. Azuma, and Y. Takeshima. Automating transfer function design for comprehensible volume rendering based on 3d field topology analysis. In *IEEE Visualization*, pages 467–470, 1999.
- [5] Y. Gong, Y. Zhang, W. Chen, and Q. Peng. Dynamic anisotropic occlusion. In *Eurographics 2006 Short Presentations*, 2006.
- [6] F. González, M. Sbert, and M. Feixas. An information-theoretic ambient occlusion. In *Computational Aesthetics*, pages 29–36, 2007.
- [7] M. Hadwiger, P. Ljung, C. R. Salama, and T. Ropinski. Advanced illumination techniques for GPU volume raycasting. In *SIGGRAPH Asia '08: ACM SIGGRAPH ASIA 2008 courses*, pages 1–166, 2008.
- [8] J. Hladůvka, A. König, and E. Gröller. Curvature-based transfer functions for direct volume rendering. In *Spring Conference on Computer Graphics 2000 (SCCG 2000)*, volume 16, pages 58–65, 2000.
- [9] R. Huang and K.-L. Ma. Rgvis: Region growing based techniques for volume visualization. In *Proc. Pacific Conference on Computer Graphics and Applications*, pages 355–363, 2003.
- [10] N. C. Institute. www.cancer.gov, 2009.
- [11] G. Kindlmann and J. W. Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *VVS '98: Proceedings of the 1998 IEEE symposium on Volume visualization*, pages 79–86, 1998.
- [12] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Moller. Curvature-based transfer functions for direct volume rendering: Methods and applications. In *Proc. IEEE Visualization 2003*, pages 513–520, 2003.
- [13] M. Knecht. State of the art report on ambient occlusion. techreport, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, 2007.
- [14] J. Kniss, G. Kindlmann, and C. Hansen. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proc. IEEE Visualization 2001*, pages 255–262, 2001.
- [15] J. Kontkanen and S. Laine. Ambient occlusion fields. In *Proc. Symposium on Interactive 3D graphics and games*, pages 41–48, 2005.
- [16] M. Levoy. Display of surfaces from volume data. *IEEE Comput. Graph. Appl.*, 8(3):29–37, 1988.
- [17] E. B. Lum and K.-L. Ma. Lighting transfer functions using gradient aligned sampling. In *Proc. IEEE Visualization '04*, pages 289–296, 2004.
- [18] C. Lundstrom, P. Ljung, and A. Ynnerman. Local histograms for design of transfer functions in direct volume rendering. *IEEE Trans. on Visualization and Computer Graphics*, 12(6):1570–1579, 2006.
- [19] A. Méndez-Feliu and M. Sbert. From obscurances to ambient occlusion: A survey. *Visual Computer*, 25(2):181–196, 2009.
- [20] H. Pfister, B. Lorenzen, C. Bajaj, G. Kindlmann, W. Schroeder, L. S. Avila, K. Martin, R. Machiraju, and J. Lee. The transfer function bake-off. *IEEE Comput. Graph. Appl.*, 21(3):16–22, 2001.
- [21] C. Rezk-Salama, M. Keller, and P. Kohlmann. High-level user interfaces for transfer function design with semantics. *IEEE Trans. on Visualization and Computer Graphics*, 12(5):1021–1028, 2006.
- [22] T. Ritschel. Fast GPU-based Visibility Computation for Natural Illumination of Volume Data Sets. In P. Cignoni and J. Sochor, editors, *Short Paper Proc. of Eurographics 2007*, pages 17–20, 2007.
- [23] S. Roettger, M. Bauer, and M. Stamminger. Spatialized transfer functions. In *EuroVis*, pages 271–278, 2005.
- [24] T. Ropinski, J. Meyer-Spradow, S. Diepenbrock, J. Mensmann, and K. Hinrichs. Interactive volume rendering with dynamic ambient occlusion and color bleeding. *Comp. Graphics Forum*, 27(2):567–576, 2008.
- [25] M. Ruiz, I. Boada, I. Viola, S. Bruckner, M. Feixas, and M. Sbert. Obscurance-based volume rendering framework. In *Proc. of Volume Graphics 2008*, pages 113–120, 2008.
- [26] Y. Sato, C.-F. Westin, A. Bhalerao, S. Nakajima, N. Shiraga, S. Tamura, and R. Kikinis. Tissue classification based on 3d local intensity structure for volume rendering. *IEEE Trans on Visualization and Computer Graphics*, 6(2):160–180, 2000.
- [27] M. Sattler, R. Sarlette, G. Zachmann, and R. Klein. Hardware-accelerated ambient occlusion computation. In B. Girod, M. Magnor, and H.-P. Seidel, editors, *Vision, Modeling, and Visualization 2004*, pages 331–338. Akademische Verlagsgesellschaft Aka GmbH, Berlin, 2004.
- [28] P. Sereda, A. Vilanova, I. W. O. Serlie, and F. A. Gerritsen. Visualization of boundaries in volumetric data sets using LH histograms. *IEEE Trans. on Visualization and Computer Graphics*, 12(2):208–218, 2006.
- [29] S. Takahashi, Y. Takeshima, and I. Fujishiro. Topological volume skeletonization and its application to transfer function design. *Graph. Models*, 66(1):24–49, 2004.
- [30] S. Takahashi, Y. Takeshima, I. Fujishiro, and G. M. Nielson. Emphasizing isosurface embeddings in direct volume rendering. *Scientific Visualization: The Visual Extraction of Knowledge from Data*, pages 185–206, 255.
- [31] M. Tarini, P. Cignoni, and C. Montani. Ambient occlusion and edge cueing for enhancing real time molecular visualization. *IEEE Trans. on Visualization and Computer Graphics*, 12(5):1237–1244, 2006.
- [32] F.-Y. Tzeng, E. B. Lum, and K.-L. Ma. A novel interface for higher-dimensional classification of volume data. In *Proc. IEEE Visualization 2003*, pages 505–512, 2003.
- [33] C. Wyman, S. Parker, P. Shirley, and C. Hansen. Interactive display of isosurfaces with global illumination. *IEEE Trans. on Visualization and Computer Graphics*, 12(2):186–196, 2006.
- [34] X. Zhang and C. Bajaj. Extraction, quantification and visualization of protein pockets. *Computational Systems Bioinformatics Conference*, 6:275–286, 2007.
- [35] S. Zhukov, A. Inoes, and G. Kronin. An ambient light illumination model. In G. Drettakis and N. Max, editors, *Rendering Techniques '98*, Eurographics, pages 45–56. Springer-Verlag Wien New York, 1998.