

Multi-Density Woodcock Tracking: Efficient & High-Quality Rendering for Multi-Channel Volumes

Alper Sahistan¹ , Stefan Zellmann² , Nate Morrical³ , Valerio Pascucci¹ , and Ingo Wald³ 

¹SCI Institute at the University of Utah, Salt Lake City, UT

²University of Cologne, Germany

³NVIDIA

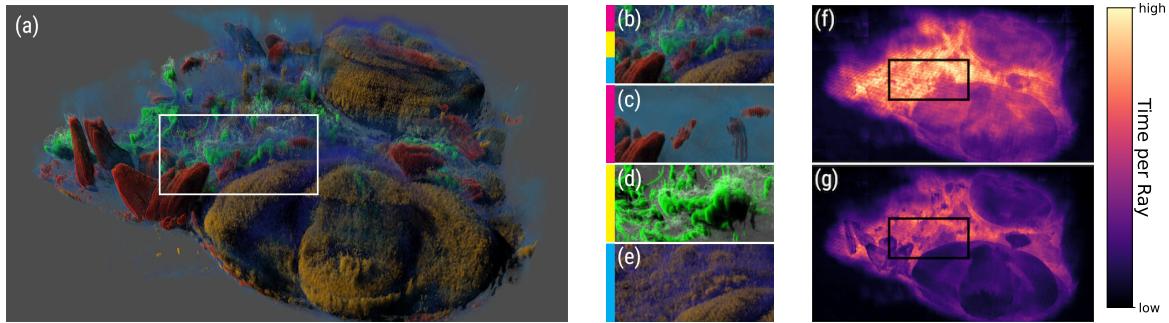


Figure 1: Multi-channel volume rendering of the Zebrafish dataset’s three channels ($640 \times 640 \times 121$): (a) Converged rendering at 2560×1440 resolution with 100 FPS using our multi-density Woodcock tracking. (b) A crop of the white rectangle in (a). For the same crop, (c), (d), and (e) display only the first, second, and third channels, respectively. (f) The added linear cost of sampling all three channels for every ray and (g) the amortized cost of our method over time per ray heatmap.

Abstract

Volume rendering techniques for scientific visualization have increasingly transitioned toward Monte Carlo (MC) methods in recent years due to their flexibility and robustness. However, their application in multi-channel visualization remains underexplored. Traditional compositing-based approaches often employ arbitrary color blending functions, which lack a physical basis and can obscure data interpretation. We introduce multi-density Woodcock tracking, a simple and flexible extension of Woodcock tracking for multi-channel volume rendering that leverages the strengths of Monte Carlo methods to generate high-fidelity visuals. Our method offers a physically grounded solution for inter-channel color blending and eliminates the need for arbitrary blending functions. We also propose a unified blending modality by generalizing Woodcock’s distance tracking method, facilitating seamless integration of alternative blending functions from prior works. Through evaluation across diverse datasets, we demonstrate that our approach maintains real-time interactivity while achieving high-quality visuals by accumulating frames over time.

CCS Concepts

• Computing methodologies → Ray tracing; Volumetric models; • Human-centered computing → Scientific visualization;

1. Introduction

From simulations to microscopy images, visualization of scientific data is essential for domain scientists to draw informed conclusions quickly and accurately. Modern scientific data are typically multifaceted, with each aspect represented as a distinct *channel* (field) corresponding to a specific attribute (velocity, temperature, pressure, etc.) of the phenomena within a volume. Understanding the causality and relationships between these attributes requires examining them not only in isolation but also in relation to each other. Multi-channel visualization offers a viable solution, providing a practical approach for observing the intricate connections among different fields in the data.

Volume rendering techniques have evolved from analytical compositing-

based methods—such as slicers [EKE01] and ray marching [TT84]—to Monte Carlo-based tracking solutions. Unlike compositing methods, which combine multiple partial samples via *opacity* [IKLH04b], tracking approaches interpret opacity as a physical *density*, allowing for a single collision per ray. A prominent technique in this domain is Woodcock tracking [WMPT65]. These Monte Carlo methods are generally easy to implement and provide greater interactivity and flexibility, but their application to multi-channel scientific volumes remains largely unexplored.

Multi-channel volume rendering presents challenges in color *blending* across multiple fields and colormaps. Unlike single-channel methods, rendering N fields involves N opacities that can sum to over 100%, complicating color contributions. Standard approaches normalize opacities

or blend using the maximum opacity, compositing over accumulated color. However, these methods can cause color discontinuities and out-of-colormap colors and increase the computational cost ($\Theta(N)$) as channels grow. Additionally, user interaction becomes more complex with more channels [RS04, Lia08], making it challenging to design blending operators that ensure proper depth ordering and intuitive usability [Kim11].

Despite these challenges and considering their popularity, blending functions from previous works continue to offer practical solutions in specific contexts. For example, disregarding occlusion can be helpful [CS99] when spatial relationships are nonessential. Likewise, observing mixed colors can hint at correlations when using primary-color transfer functions [RS04, Lia08]. To our knowledge, a formal way to combine these techniques to provide a unified blending modality does not exist. Formulating this modality becomes particularly counterintuitive when compositing-based rendering is involved.

In this work, following a similar idea to analog decomposition tracking [KHLN17], we extend Woodcock tracking to multi-channel scientific visualization (sci-vis), addressing the fundamental challenges of rendering multi-channel data. Our algorithm employs a Monte Carlo process to identify the closest sample by evaluating each channel's density along the ray, thereby avoiding the limitations of previous methods that rely on accumulating numerous partial samples. As illustrated in Figure 1, our approach enables accurate and efficient visualization of arbitrary combinations of N channels across a volume. The contributions of this paper are as follows:

- Two generalizations of the Woodcock tracking traversal to multi-channel volumes:
 - an easy-to-implement serial method that traverses and collects samples for N channels one at a time,
 - and an optimized method that prunes the redundant traversals by traversing N channels in a single traversal.
- A density-driven, physically based inter-channel blending function that avoids the pitfalls of prior color blending functions.
- A unified rendering modality within the Woodcock tracking that allows substituting inter-channel color blending functions from prior works—e.g., max opacity selection and weighted mixing.

2. Related Work

2.1. Volume Rendering

Earlier works in volume rendering utilize slicers [CCF94, CN93, EKE01, LKC05] where a 3D texture of view-aligned slices are re-sampled to approximate the volume rendering equation. Although slicers are actively used in multi-channel volume rendering, single-channel methods largely abandoned slicers with the advent of GPGPU frameworks like CUDA and OpenCL.

Nowadays, the standard way to render volumes is ray marching-based methods [PH89, KW03, RGW*03, HSS*05]. These methods produce high-quality, noise-free images by analytically approximating the volume rendering equation via accumulated contributions from multiple partial samples along a ray's path. However, when undersampling, marchers introduce bias and artifacts even with jittered sampling [RSK08]. Additionally, incorporating secondary effects such as volumetric shadows or gradient shading dramatically hinders the rendering performance as these mandate more sampling for each initial view-aligned sample [SDM*21]. To combat these costs, prior works have proposed space skipping and adaptive sampling strategies [LLY06, MUWP19, ME11, WZU*21]. Specifically for shadows, shadow maps offer a solution for achieving more

interactive rates, but they occupy extra memory and require recomputation whenever the transfer function or lighting changes [IKLH04a].

Tracking-based Monte Carlo estimators have gained traction in scientific and cinematic rendering [FWKH17]. Despite introducing variance (noise) and needing to converge over multiple samples, they offer advantages such as efficient handling of secondary effects like shadows due to their lower asymptotic complexity. They also enable artifact-free adaptive sampling [YIC*10, SKTM11]. Szirmay-Kalos et al. further improve efficiency by using a 3D digital differential analyzer (DDA) to adjust sampling rates based on local density estimates stored in a coarser grid.

Woodcock (delta) tracking has become popular in recent works [GKT16, MHK*19, HMES20, MZS*23], where volumes are homogenized using fictitious particles based on maximal density. Morrical et al. introduce adaptive Woodcock tracking for compressed clusters of unstructured meshes [MSG*23]. Zellmann et al. explore data structures and algorithms to efficiently path trace Adaptive Mesh Refinement volumes using Woodcock tracking [ZWS*24].

The use of tracking methods for rendering multiple overlapping volumes has been explored in prior work [NSJ14, KHLN17], predominantly within cinematic rendering. Novak et al.'s residual ratio tracking [NSJ14] reduces variance by splitting the volume into homogeneous control and heterogeneous residual volumes. Kutz et al.'s spectral and decomposition tracking [KHLN17] introduces the idea of selecting the minimum distance among free-flight samples from multiple volumes, a technique shown to fit within the integral framework of Galtier et al. [GBC*13] and later expanded further [GMH*19]. In contrast, we utilize these ideas for multiple overlapping heterogeneous volumes within the sci-vis applications.

2.2. Multi-Channel Visualization

The multi-channel visualization has two orthogonal problems: user interface [MJW*13] and rendering [CR08]. In this paper, we mainly tackle the rendering aspect while allowing the usage of traditional transfer functions. Nevertheless, the solutions to these orthogonal problems can sometimes be intertwined [KZX*23]. The thesis by Kim [Kim11] offers many insights into the multi-channel visualization domain.

Given the tedious nature of setting a precise transfer function for all the channels, Pan et al.'s work [PLL*24] concentrates on design galleries where users can select images resembling their target image to construct a transfer function implicitly. Likewise, Kim et al. [KSC*10a, KSC*10b] employ dimensionality reduction schemes to aid in the design of multidimensional transfer functions. Our work is compatible with these transfer function helper frameworks, given that their neural networks are trained with our renderer.

Khlebnikov et al. [KKSS13] propose a random-phase Gabor noise-driven cell-space redistribution pattern and filtering scheme for simultaneous multi-channel display. However, this technique can blur details or create counterintuitive renderings, particularly for untrained observers. Herzberger et al. [HHK*23] introduce the residency Octree for out-of-core mixed resolution in web-based multi-channel rendering, focusing on efficient data streaming with a ray-guided volume renderer from Crassin et al. [CNLE09].

FluoRender [WOCH09, WOCH12, WOH*17] is one of the most prominent tools for multi-channel visualization, using slicer-based rendering while enabling a variety of user interactions. Slicers render one slice at a time, loading only necessary transfer functions and channels, effectively serializing the rendering and reducing texture memory and VRAM usage. FluoRender users can select, highlight, or segment

structures using brushing tools with morphological diffusion, aiding in exploring correlations. To avoid visual clutter from overlapping fields, Fluorender offers non-physically based compositing modes, such as post-render compositing and weight-based color mixing with depth correction. Given Fluorender's broad appeal, we incorporate some of these modes into Woodcock tracking for our experiments.

Opacity-based rendering often relies on physically inaccurate models and struggles with opacity mixing. Previous efforts [CS99, RS04, Lia08] have addressed some issues with inter-channel color blending. We propose adopting density-based approaches, such as Woodcock tracking, to overcome to tackle these issues in a physically-motivated manner. Density-based approaches shift color blending to screen space, accumulating one sample at a time across frames. Our approach also supports integrating other blending techniques from prior works.

3. Woodcock Tracking Background

In Woodcock tracking [WMPT65], photon and particle behaviors are simulated through a Monte Carlo process. Generalizing the rendering equation [KWH84] for volumes and simplifying for the emission/absorption model gives us the simplified Volume Rendering Equation (VRE):

$$L(x, \omega) = \int_{t=0}^d T(t) \sigma_a(x) L_e(x_t, \omega) dt, \quad (1)$$

In Equation 1, radiance, $L(x, \omega)$, at point x looking in the direction of ω is calculated by taking the integral of transmittance $T(t)$ times the absorption coefficient $\sigma_a(x)$ and emitted radiance $L_e(x_t, \omega)$ between x and x_t . In the sci-vis context, L_e is determined by the colormap of the transfer function that maps a scalar number to an RGB value, and $\sigma_a(x)$ is influenced by the user-defined alpha component of the transfer function for the same scalar value.

The transmittance for homogeneous volumes follows the Beer-Lambert law where σ_t is constant:

$$T(t) = \exp(-\sigma_t t) \quad (2)$$

To simulate photon-particle collisions, we calculate a photon's *free-flight distance* until it hits a particle, denoted as t' . This calculation involves computing the probability density function (PDF) of $T(t)$, denoted as $p(t)$, and importance sampling $p(t)$ using ξ as a random number:

$$p(t) = \sigma_t \exp(-\sigma_t t), \quad t' = \frac{-\ln(1-\xi)}{\sigma_t} \quad (3)$$

To apply this formulation to heterogeneous volumes, we can imagine the heterogeneous volume homogenized by fictitious *null particles*. The ratio, and thus the probability, of encountering these null particles is dictated by the maximum density of the volume or subvolume. This coefficient, known as the *majorant*, is denoted as $\bar{\sigma}$ and can be substituted for σ_t in Equation 3.

Finally, the VRE needs to be adjusted to handle null collisions and normalize coefficients, essentially turning them into probabilities. Therefore we define the probability of hitting a real particle, $P_{real}(x)$ as the ratio of density at point x to maximum density (*majorant*). Therefore, the remainder of the probabilities is the null collision, $P_{null}(x)$, probability.

$$P_{real}(x) = \frac{\sigma_a(x)}{\bar{\sigma}}, \quad P_{null}(x) = \frac{\bar{\sigma} - \sigma_a(x)}{\bar{\sigma}}, \quad (4)$$

So we end up with the final form of Equation 1:

$$L(x, \omega) = \int_{t=0}^d p(t) [P_{real}(x) L_e(x_t, \omega) + P_{null}(x) L(x_t, \omega)] dt \quad (5)$$

This version of the VRE selects one of two paths when computing incoming color at x from direction ω : (i) With probability P_{real} the scalar field is sampled, assigning $L_e(x_t, \omega)$ a color via transfer function

f_{cm} . (ii) With probability P_{null} , the function $L(x, \omega)$ is invoked again (recursive) from a further point x_t . The integral relies on $p(t)$, a stepping function. Since this approach follows a Monte Carlo process, we use its importance-sampled version, the free-flight distance t' , from Equation 3 to determine the stepping distance.

While the majorant $\bar{\sigma}$ is presented as a global constant, subdividing the volume into regions with local majorants $\bar{\sigma}_i$ reduces null collisions by better constraining null particles, improving performance. Instead of a single global majorant, we use a grid of $\bar{\sigma}_i$ s.

We use a similar notation to [FWKH17] and point the readers to [PJH23] for further details of these derivations.

4. Method

As the VRE is defined for single channel volumes in Equation 5, we introduce our physically motivated formulation to render and blend N fields, $\theta_n | n \in \mathbb{N} \wedge n \leq N$, as a natural extension of Woodcock tracking in Subsection 4.1. Additionally, we propose another extension to Equation 5 to formally incorporate blending functions from prior works in Subsection 4.2.

We detail how to set up a Woodcock renderer that uses multiple densities in Subsection 4.3. Then, in Subsection 4.4 and Subsection 4.5, we explore the implementation space for the method proposed in Subsection 4.1.

4.1. Multi-density Woodcock Tracking

We can devise a physically based and convenient Woodcock tracking method by treating N fields as multiple densities. Doing so, we naturally extend the physically based Woodcock tracking to multiple channels.

We generalize the formulation Equation 5 to operate over multiple fields. We substitute the $P_{real} L_e(x_t, \omega)$ term with the sum of the multiplication between probabilities $P(\theta_n)$ and N radiances where θ_n is the n -th field of the volume to end up with a formulation for multi-density Woodcock tracking:

$$L(x, \omega) = \int_{t=0}^d p(t) \left[\sum_{n=1}^N (P(\theta_n, x) L_e(x_t, \omega, \theta_n)) + P_{null}(x) L(x_t, \omega) \right] dt \quad (6)$$

The modified probabilities are as follows:

$$P(\theta_n, x) = \frac{\sigma_a(x, \theta_n)}{\bar{\sigma}_n}, \quad P_{null}(x) = \sum_{n=1}^N \left(1 - \frac{\sigma_a(x, \theta_n)}{\bar{\sigma}_n} \right) \quad (7)$$

Where function σ_a yields the density for channel θ_n at point x , i.e., $\sigma_a(x, \theta_n) = \theta_n(x)$. We keep the property of probabilities adding up to 1 from Equation 4, so $\sum_{n=1}^N P(\theta_n) + P_{null} = 1$ for any point x .

One way to think about this formulation is that we are testing for a real collision (absorption) for N volumes. The remaining possibility is the null collision probability.

Mapping the set of N opacities to physical densities, we turn the process into a Monte Carlo process that blends the colors through accumulation in screen-space. Unlike prior blending/sampling strategies, our approach ensures physical correctness without arbitrary normalizations. Therefore, Equation 6 implicitly defines a distinct blending function. Figure 3 illustrates our density-based blending function applied to three overlapping volumes, each mapped to a flat primary-colored transfer function.

$$OURS \left(\begin{array}{c} \text{Red} \\ \text{Green} \\ \text{Blue} \end{array} \right) = \text{Rainbow Color}$$

Figure 3: Three single-channel renderings are shown from left to right: Red, green, and blue cube volumes with linearly decreasing densities on the X, -X, and -Y axes, respectively. Using our formulation from Equation 6, their multi-channel rendering is given on the right-most image.

	Post-render compositing	Mix w/ equal weights	Density-based (ours)	D.-based (ours) w/ shdw
Depth Ordering	✗	✓	✓	✓
Phys.-B. Colors	✗	✗	✓	✓

Figure 2: Visual comparison of various blending modes over a close-up of *Miranda* dataset with four channels (left-to-right): Compositing after rendering, mix function from Equation 10, our density-based blending function from Equation 6, and our blending mode with volumetric shadows. They are evaluated on their ability to achieve correct depth ordering and physically based color blending.

4.2. Generalization to Blending Functions in Woodcock

Previously defined methods for blending N channels are often exploratory rather than physically based. Researchers have introduced various blending functions, such as maximum opacity selection and user-weighted mixing [RS04, Lia08], which have been effective in practice. While these blending techniques are not physically motivated, they offer valuable heuristics and have an established user base accustomed to their results. This section demonstrates how these methods can be integrated within a Woodcock tracker, leading to alternative formulations.

To allow the substitution of other blending functions, we go back to Equation 5, and we exploit the fact that these blending functions serve the purpose of reducing N colors and N opacities to one color and opacity pair. By allowing redefinition of $\sigma_a(x, \theta_n)$, and $L_e(x, \omega, \theta_n)$, we can specialize to other blending functions for a Woodcock tracker. Note that, since N fields co-exist in the same volume, the majorant, $\hat{\sigma}$, is now the sum of N maximum densities.

The max opacity sample selection (or *max*) always chooses the sample with the highest opacity at the given sample point. Therefore, the maximum opacity selection can be defined as:

$$\sigma_a(x, \theta_n) = \begin{cases} \theta_n(x), & \text{if } \theta_n(x) > \theta_i(x) \mid \forall i \in \mathbb{N} \wedge i < N \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

$$L_e(x, \omega, \theta_n) = \begin{cases} f_{cm}(\theta_n(x)), & \text{if } \theta_n(x) > \theta_i(x) \mid \forall i \in \mathbb{N} \wedge i < N \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

As *max* deterministically picks the same sample, overlapping semitransparent regions can get filtered out by one more dense region. Despite allowing some regions to steal visibility from others, it merits intuitive usage. Figure 4 shows a visualization of *max* blending.

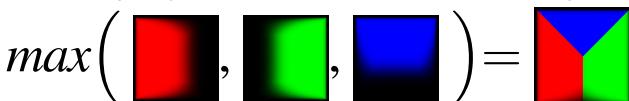


Figure 4: Visualization of *max* blending function from Equation 8 using the same three fields from Figure 3.

Mixing with user-defined weights (i.e., *mix*) allows users to adjust weights for each channel to allow color mixing among them. Using weights, w_n , *mix* blending can be achieved by using the same $\sigma_a(x, \theta_n)$ as Equation 8 and re-defining the radiance as:

$$L_e(x, \omega, w_n) = \frac{\sum_{n=1}^N f_{cm}(\theta_n(x)) \cdot \theta_n(x) \cdot w_n}{\max(\theta_1(x), \theta_2(x), \dots, \theta_N(x))}, \text{ where } \sum_{n=1}^N (w_n) = 1 \quad (10)$$

Although the user interface complexity for the *mix* blending scales

linearly with the number of channels, it can show correlations with mixed colors. Please see Figure 5 for a rendering with the *mix* function.

$$mix\left(\begin{array}{c} \text{Red} \\ \text{Green} \\ \text{Blue} \end{array}\right) = \text{Rainbow Color}$$

Figure 5: Visualization of *mix* blending function from Equation 10 using equal weights and the same three fields from Figure 3.

Another downside of blending modes like *mix* and *max* is that they require finding a maximum among N samples as a normalization factor, which has the algorithmic complexity of $\Theta(N)$. In contrast, our density-based blending function from Subsection 4.1 does not require a normalization factor (e.g., maximum opacity selection) that entails a linear cost. Instead, it probabilistically selects one sample among channels, resulting in better average algorithmic complexity.

We present definitions for some well-known functions but this approach is not constrained by them, and using the open-ended formulation we provided in this section, users can invent new blending functions. $\sigma_a(x, \theta_n)$ controls collision possibilities, and it can be parameterized to pick something other than the maximum opacity sample, or $L_e(x, \omega, \theta_n)$ can define a nonlinear combination of N channels.

There are also blending modalities completely disassociated from the rendering process that occurs in the image space, such as compositing images after rendering each channel (post-render compositing). Post-render compositing offers an occlusion-free view, but it loses the depth information. It is commonly used in tools like Fluorender.

Figure 2 illustrates some of these blending functions' benefits and visual distinctions on a real dataset.

4.3. Setup for the Woodcock Renderer

We developed a base renderer in CUDA that stores N structured grids of scalars in 3D textures, each with dimensions $k \times l \times m$. Each loaded channel utilizes their separate transfer function to determine their user-defined densities and colors. The renderer builds a coarser $k' \times l' \times m'$ *macrocell* grid. Within each macrocell, scalar ranges for each channel (a min and max pair) are stored, resulting in a memory consumption of $N \times k' \times l' \times m' \times 2$ floating-point values per channel. The construction of the macrocells can be considered a rasterization process, where the original grid of scalars is projected onto a grid with reduced resolution.

Every time a transfer function changes, we calculate that channel's majorants for each macrocell. The process involves determining the maximum opacity associated with the scalar value range of each

macrocell for the given channel. Therefore, for each macrocell, we iterate between the min and the max scalar values, fetching the opacity mapped to them from the respective channel's transfer function. The majorant for that macrocell is the maximum opacity found after the loop. We run these calculations using a CUDA kernel in parallel. The process yields a majorant grid with the exact resolution as the macrocell grid. Therefore, we store $k' \times l' \times m' \times N$ floating-point numbers for the majorants.

Listing 1: Implementation of Woodcock stepping function

```
float woodcockStep(float majorant) {
    return -(log(1.0f - randFloat()) / majorant);}
```

Using a 3D Digital Differential Analyzer (DDA) traversal, we employ ray tracing through the majorant grids. Woodcock steps are taken as shown in [Equation 3](#) (calculated with [Listing 1](#)) within each macrocell based on the active majorants. Upon identifying a collision at point p , we sample the 3D textures to retrieve the relevant channels' scalar value at p . The scalar is given to the respective channel's transfer function (TF) to get the color and opacity. We interpret the linear RGB returned from TF as an emissive color. We employ rejection sampling using the null collision probability in [Equation 4](#) to see if the sample is accepted (see [Listing 2](#)). Upon acceptance, we halt the traversal and return the color value (discarding the "w" component). We utilize an accumulation buffer to allow convergence over multiple frames.

Listing 2: Implementation of rejection scheme via null collisions

```
bool nullCollision(float majorant, float density) {
    return (density < randFloat() * majorant);}
```

We have identified several meaningful approaches to achieving this. They differ in implementation, however they ultimately converge to the same image when using the same parameters.

4.4. Method #1: N-DDA Traversals over Majorant Grids

One obvious way to approximate [Equation 6](#) is by simply applying Woodcock tracking for each channel as separate volumes and choosing the closest sample. This approach, though not the most computationally efficient, allows for serializing the rendering of individual channels by loading volumes and their corresponding transfer functions one at a time to mitigate memory constraints. It can be particularly useful for an out-of-core implementation. Moreover, this offers a good baseline for our comparisons and can be viewed as an initial step for understanding the multi-channel Woodcock tracking.

For this approach, we build majorant grids for each channel where the maximum density for a group of cells contained within a macrocell is stored. The majorants are recalculated every time a TF is edited.

The rendering process, illustrated in [Figure 6](#), involves applying Woodcock tracking to all N channels using their respective majorant grids. Specifically, for a given channel C_0 , we perform a DDA traversal over the majorant grid M_0 . Within each cell of M_0 , Woodcock steps are taken using the majorant $m_{0,i}$, where i corresponds to the 1D index of the current cell of M_0 , for the ray position. The DDA traversal and steps within are repeated until we encounter a real collision. Upon hitting a particle, we record the distance (t_0) and color for C_0 . This entire procedure is then repeated for C_1, C_2, \dots, C_N , utilizing their respective majorant grids M_1, M_2, \dots, M_N . Once we complete iterations through all channels, we select the closest sample — $\min(t_0, t_1, \dots, t_n, \dots, t_N)$. One obvious optimization we apply to this approach is early termination of the ray traversal for the channel C_{n+1} if t_{n+1} surpasses the closest hit distance, t_{hit} ,

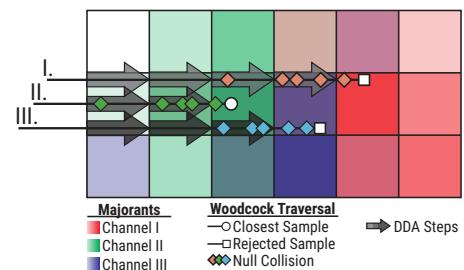


Figure 6: The Woodcock tracking process with three DDA traversals across three majorant grids from a perspective of one ray: Starting from channel I, a DDA traversal is initiated, and between each DDA step, Woodcock steps are taken where null collisions are discarded. These processes are repeated until a sample for each channel is found. The closest sample among the three channels (channel II) is accepted, and the others (I and III) are rejected.

for channels between C_0-C_n , as it would have been discarded in the final stage anyway. C-style pseudo code for this method is shown in [Listing 3](#).

Listing 3: Multi-channel Woodcock tracking using N DDA traversals

```
void nDDAMultiChannelWoodcock(
    Ray ray, HitRecord& hit, int N) {
    hit.t = FLT_MAX;
    for(int n = 0; n < N; n++) { //do N DDA traversals
        DDA dda(ray);
        int cellID = dda.curCell();
        do { //DDA traversal
            float t = dda.cellEntryT();
            float maj = majorants[cellID * N + n];
            while(true) { //Woodcock steps in the cell
                t += woodcockStep(maj);
                if(t > hit.t) { //a closer sample exists?
                    dda.stop();
                    break; //No need to sample
                }
                if(!dda.InCurCell(t)) //bounds check
                    break; //go to the next cell
                float scalar = volumeAt(ray.org
                    + ray.dir * t, n);
                //scalar's color (r,g,b) and density (a)
                float4 sample = trFunc(scalar, n);
                //null collision check n-th channel
                if(!nullCollision(maj, sample.a)) {
                    //Return the color, discard the "a"
                    hit.color = float3(sample);
                    hit.t = t;
                    dda.stop(); break; //stop this traversal
                }
            }
            cellID = dda.nextCell();
        } while(!dda.shouldStop()); } }
```

4.5. Method #2: One DDA Traversal over Majorant Grids

The next logical step on top of the method presented in [Subsection 4.4](#) is using a single synchronized DDA traversal over N channels to prune to-be-discarded samples earlier.

The data structures employed remain consistent with those outlined in [Subsection 4.4](#), featuring majorant grids for each channel, subject to updates upon transfer function modification.

The rendering commences with a singular DDA traversal over the

majorant grid. Multiple Woodcock steps are taken within each majorant cell for each channel until either a sample is accepted or all of the rays leave the current majorant grid cell. Similarly to the method in [Listing 3](#), we pick the closest sample to the ray origin if there are multiple hits within a majorant cell. In other words, the outer for-loop iterating over N channels in [Listing 3](#) is placed inside the DDA traversal.

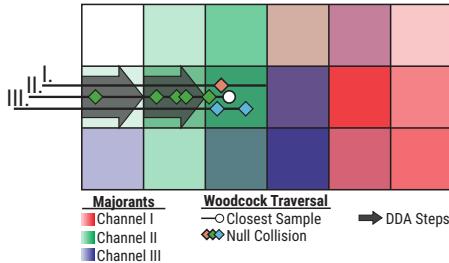


Figure 7: The Woodcock tracking process with a single DDA traversal across three majorant grids from a perspective of one ray: A DDA traversal is initiated, and between each DDA step, three Woodcock traversals are initiated in the order of channels I, II, and III. The processes are repeated until a sample is taken (by channel II in this case). We let a final round of Woodcock traversal run for channel III as it could find a closer sample. Only the closest sample from channel II is accepted.

By synchronizing the N Woodcock tracking steps into one DDA traversal, we potentially reach the closest sample sooner and with fewer iterations (compared to [Subsection 4.4](#)). Also refer to [Figure 7](#) for an illustration of this method and [Listing 4](#) for a streamlined kernel code.

Listing 4: Multi-channel Woodcock tracking using a synchronized DDA

```
void syncDDAMultiChannelWoodcock (
    Ray ray, HitRecord& hit, int N) {
    DDA dda(ray);
    int cellID = dda.curCell();
    hit.t = FLT_MAX;
    do{ //do a DDA traversal
        for (int n = 0; n < N; n++) { //for channels...
            float maj = majorants[cellID * N + n];
            float t = dda.cellEntryT();
            while(true){ //Woodcock steps in the cell
                t += woodcockStep(maj);
                if(t > hit.t || //a closer sample exists?
                    !dda.InCurCell(t)) //bounds check
                    break; //go to the next channel/cell
                float scalar = volumeAt(ray.org
                    + ray.dir * t, n);
                //scalar's color (r,g,b) and density (a)
                float4 sample = trFunc(scalar, n);
                //null collision check for n-th channel
                if(!nullCollision(majs[n], sample.a)){
                    //return the color, discard "a"
                    hit.color = float3(sample);
                    hit.t= t;
                    //regardless of there is a closer
                    //sample or not we need to stop DDA
                    dda.stop();
                    break; //go to the next channel/stop
                }
            }
        } // no channel was selected: null collision
        cellID = dda.nextCell();
    }while(!dda.shouldStop()); }
```

5. Results and Evaluation

In this section, we evaluate the performance impacts of parameters such as macrocell resolution ([Subsection 5.1](#)), number of channels ([Subsection 5.2](#)), and visual effect of using our, and some of the previously used, blending functions ([Subsection 5.4](#)).

We utilize diverse datasets to drive more extensive conclusions. *NYX* is a cosmological simulation that is in $512 \times 512 \times 512$ resolution, *Hurricane* is a weather simulation that is flatter in the z dimension with $500 \times 500 \times 100$ resolution, *Miranda* is a hydrodynamics simulation with highly overlapping features in $256 \times 384 \times 384$ resolution and, *Zebrafish* is a well-known multi-channel microscopy dataset in $640 \times 640 \times 121$ resolution [[ZDL*20](#)].

We run our experiments on an NVIDIA RTX4090 GPU using the program we implemented using CUDA. We take the average timing of 500 frames for each data point after rendering 50 warm-up frames, which are excluded from the average.

5.1. Impact of Macrocell Resolution

The resolution of the macrocell and majorant grids directly impacts our method's performance, as these grids approximate density at specific volume points. This approximation influences the Woodcock stepping size and collision probabilities. A finer grid allows for better stepping and fewer null collisions but is more memory-intensive, while a coarser grid results in more null collisions, increasing the time spent per rays [[MLB*23](#)]. In this section, we measure the performance of our two multi-density Woodcock tracking methods from [section 4](#) over macrocell size.

We measure the rendering times of each algorithm for direct volume rendering using the emission+absorption (E+A) model alone and the E+A model with shadows.

Four channels of *NYX*, *Hurricane*, and *Miranda* and three channels of *Zebrafish* are used for this experiment. [Figure 8](#) shows the render times vs. decreasing macrocell resolution. For each data point going forward, we insert twice as many cells per dimension within a macrocell, i.e., $8 \times$ less resolution. We record 291, 504, 453, and 686 frames per second (FPS) for *NYX*, *Hurricane*, *Miranda*, and *Zebrafish* datasets. Our fastest timings with volumetric shadows are 178, 286, 299, and 577 FPS, respectively. We also report the memory consumption of our application using the most performant data macrocell resolution in [Table 1](#).

We observe that method #2 from [Subsection 4.5](#) with 1-DDA traversal achieves the highest frame rates. For both methods, the performance peaks around 64 cells per macrocell ($4 \times 4 \times 4$) for almost all datasets. We record the memory consumption of our data structures to be between 0.58-4.69% of the original data sizes at the peak performance.

We also tested a third method that uses a single majorant buffer, summing all majorants into one value and testing each sample against the cumulative majorant. The idea was to potentially reduce the number of steps, but the results were negative. This method achieved only 59.88-84.03% of the performance of our fastest method and required about $8 \times$ the macrocell size, consuming more memory. Consequently, we omit this method and its results. Notably, this approach converges to a technique described in the overlapping volumes chapter of Fong et al. Production Volume Rendering course [[FWKH17](#)].

5.2. Performance Impact of Multiple Channels

The number of channels is one of the factors that may add a significant cost to algorithms such as ray marchers. Therefore, we document the performance of our methods against the increasing number of channels.

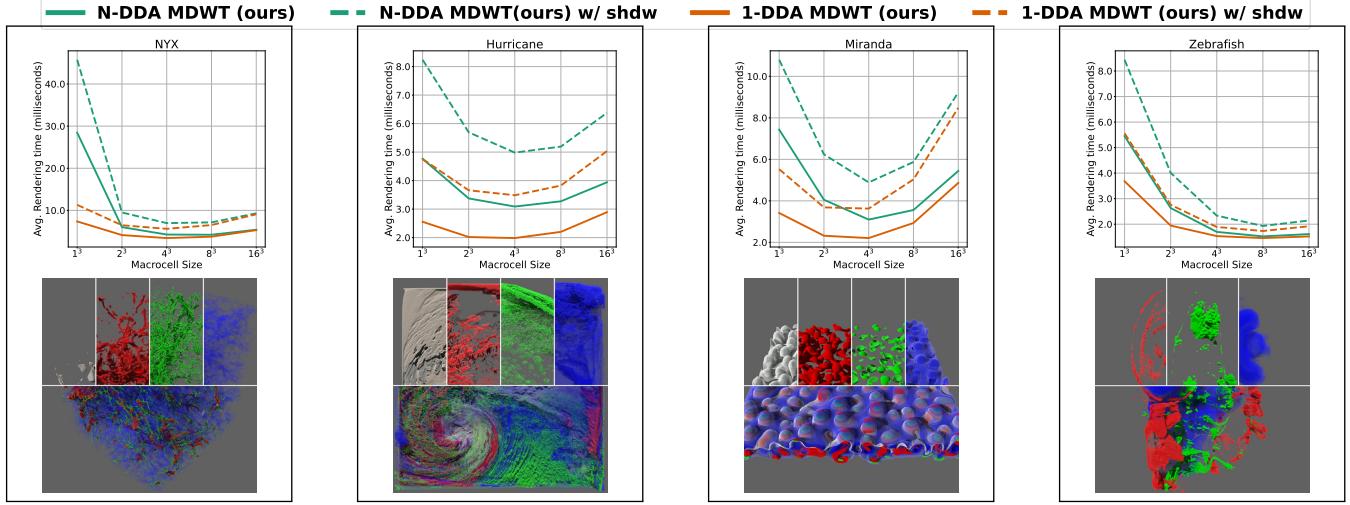


Figure 8: The plots of average rendering times in seconds (lower is better) against increasing macrocell sizes (# of cells within a macrocell) for four datasets. We compare our two methods: “N-DDA MDWT (Multi-density Woodcock Tracking)” from [Subsection 4.4](#), and “1-DDA MDWT” from [Subsection 4.5](#). Solid lines are for the Emission+Absorption (E+A) model, while dashed lines represent timings using E+A with shadows. The split images in the second row show individual channels on the top halves and multi-channel visualization in the bottom halves (with shadows).

We use a standard ray marcher as a baseline in these experiments. To mirror the behavior of out-of-the-box visualization solutions closely, we use the common mix blending function with equal weights for the ray marcher.

[Figure 9](#) depicts the performance of our two methods and the aforementioned ray marcher against the increasing number of channels. The ray marcher and our methods are configured to produce similar-quality images using the same transfer function and scene configurations.

We artificially amplify three of our datasets’ channel counts by duplicating the same channel data with different transfer functions to allow testing for higher scalability. Unlike the others, the Hurricane dataset contains 12 unique channels, so we use it as is.

Our most efficient method ([Subsection 4.5](#)) can render 12 channels simultaneously, equivalent to rendering approximately $\approx 4-8$ channels sequentially with volumetric shadows. Adding more channels creates occlusion, hiding previously visible layers; our multi-density Woodcock tracking algorithm exploits this fact and terminates on the first hit, yielding amortized timings. In contrast, ray marching collects partial opacities along the ray direction, leading to a linear traversal cost. The mix operator must also blend N different colors, introducing another linear cost, causing ray marching to scale linearly with the number of channels.

We observe some rapid changes in the performance trend, especially observable in NYX between 8-10 channels (in [Figure 9](#)). This spike is upward, indicating a loss of performance. Although channels 8-10 are more occlusive, they do not improve performance because each volume occupies a large portion of the space without overlapping significantly. This non-overlapping space consumption causes the majorant grids to become more dense, and less adaptive which in-turn causes our algorithm to explore and stop more often to check for collisions.

Our N-DDA multi-density Woodcock tracking (method #1 from [Subsection 4.4](#)) does not obtain as much amortization from occlusion as the 1-DDA multi-density Woodcock tracking (method #2 from [Subsection 4.5](#)) does despite using the same data structure. From our observations, this outcome can be explained by the increasing number of redundant traversals since this method does not prune some of the redundant traversals that are pruned in method #2 as explained in [Subsection 4.5](#) and seen in [Figure 7](#).

Table 1: Memory consumption of our implementation. The columns report the consumption for the specified number of channels. “Base” shows the memory allocated for the dataset; for the macrocells and the majorant grid, the consumption for the best-performing resolution is reported.

Dataset\Memory(MB)	Base	Macrocells	Majorants	Total
NYX (4 channels)	2048.00	64.00	32.00	2144.00
Hurricane (4 channels)	381.47	11.92	5.96	400.35
Miranda (4 channels)	576.00	18.00	9.00	603.00
Zebrafish (3 channels)	567.15	2.19	1.01	570.35

5.3. Quality vs. Performance

Woodcock tracking enables one sample per ray, allowing frames to accumulate and gradually reduce variance. In contrast, ray marching uses an analytic solution that takes multiple partial samples per ray, resulting in images with virtually no variance but at a significantly higher time cost (as shown in [Figure 9](#)).

This is indeed a trade-off between per-frame performance and noise. To evaluate the cost-efficiency of our method, we conduct “similar performance” and “similar quality” experiments using the NYX and Miranda datasets, as seen in [Figure 10](#). First, we compare the ray marcher’s performance to our 1-DDA MDWT in similar quality setups. We find 16 samples per pixel(spp) to yield visuals of similar quality to the ray marcher sampling at the Nyquist rate (to avoid aliasing artifacts). Next, we fix our method’s parameters and compare its performance to the ray marcher by linearly increasing the sampling interval from the artifact-free Nyquist rate to the point where both methods perform similarly.

Our results show that the ray marcher can achieve high-quality images in less time, particularly with the NYX dataset, which reaches a similar quality image $\approx 3 \times$ faster. However, similar quality is achieved at around the same frame rate in datasets with more empty space, like Miranda. In the “similar performance” benchmarks, ray marcher images rendered with sampling intervals larger than the Nyquist rate exhibit significant aliasing artifacts, obscuring features. This explains the drastic variations seen in RM-Mix images in [Figure 10](#), where undersampling leads to severe drops in quality. Unlike our Monte Carlo-based solution, ray marching cannot improve image quality incrementally. Additionally, increasing spp in our method incurs a linear cost while reducing the ray marcher’s

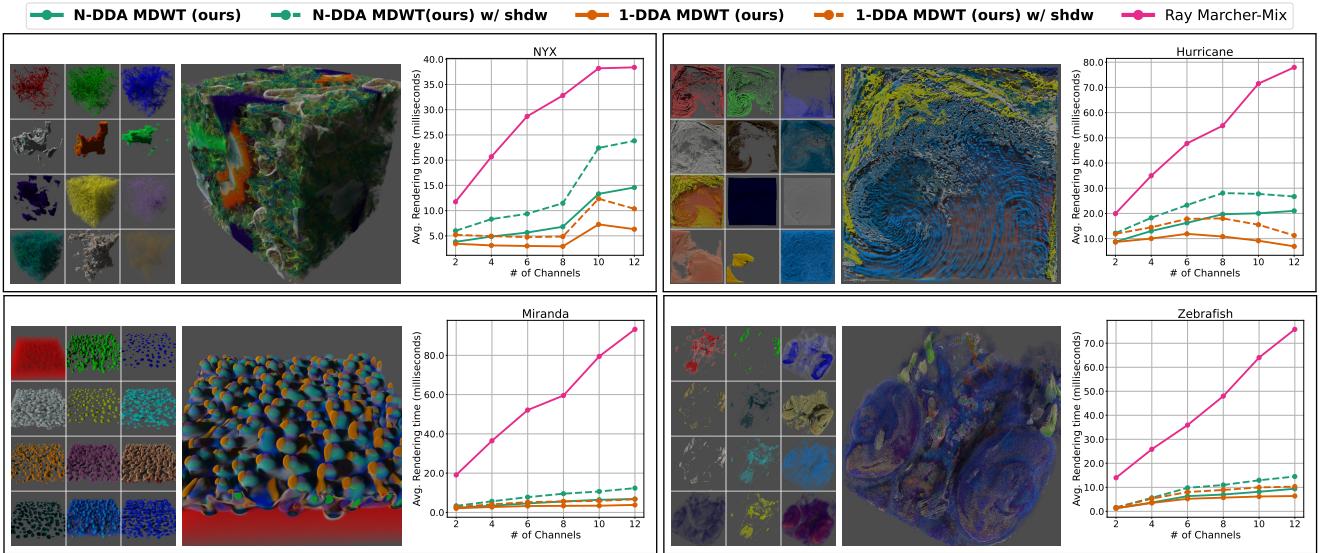


Figure 9: The average rendering times in milliseconds (lower is better) versus the increasing number of channels across four datasets. For each dataset, we present the rendering of individual channels on the left, all channels combined with multi-channel Woodcock tracking on the middle, and a line plot on the right. The line plots compare “N-DDA MDWT (Multi-density Woodcock Tracking)” from Figure 6, our “1-DDA MDWT” method from Subsection 4.5, and “Ray Marcher-Mix,” a ray marcher using mix blending. Solid lines represent DVR timings with only the Emission+Absorption (E+A) model, and dashed lines represent the performance of our methods using E+A plus volumetric shadows.

sampling interval scales nearly exponentially. Our approach maintains higher interactivity without requiring the highest spp from the start.

5.4. Visual Comparisons for Blending Functions

We compare the visual results of various inter-channel blending operators over three datasets in Table 2, using F1IP and Peak Signal to Noise Ratio(PSNR). We use four flat colormaps: red, green, blue, and white. In these experiments, we compare an image-space method, two not physically based blending functions, to our physically-based density blending function. These are post-render compositing and mixing with equal weights from Fluorender [WOH¹⁷] and max opacity selection from Rice and Schulze [RS04].

Our results indicate the post-render compositing offers the least physically correct result and the farthest image in terms of difference. Although this mode is intended to offer another way to examine the data, it makes denser datasets such as NYX, and Hurricane harder to distinguish as both the colors and depth relationships are altered.

The mixing function has the unique property of blending colors into a mixture of two colors. However, this can be counterintuitive to untrained users as these colors can be out of the current set of colormaps, or they can be confused with other fields if the color is present in a colormap. This mixing is more obvious in the Hurricane dataset, where reds and blues turn into magenta (which is not a color in the set of colormaps).

Upon reviewing the experiments in Subsection 5.2 and Subsection 5.4, we notice that visual clutter becomes more pronounced as the number of channels increases. Beyond five channels, the blending colors may start appearing counterintuitive. Additionally, employing blending functions that generate out-of-colormap colors like mix can exacerbate this problem, potentially leading to confusion between a field and overlap of the other two fields.

The results closest to our blending function are from max opacity selection, as it often selects the same maximum opacity sample as ours.

However, this approach can lead to the loss of semitransparent regions due to the sharp cut-off of $\max()$.

When tested on the Zebrafish dataset, the mix and max functions show minimal differences from our approach. In contrast to simulation data, where features cluster in specific regions, the microscopy data exhibit less overlap between channels. This reduced overlap leads to less color blending, resulting in similar outcomes across blending modes.

Density-based blending function offers an intuitive way to blend colors as it can blend the colors through more neutral tones while not losing information of more transparent regions due to sharp cut-offs. It remains physically based as it stems from the direct generalization of Woodcock tracking’s sampling to multiple channels.

6. Discussion and Conclusion

In this paper, we presented an efficient method for rendering multiple volume channels by extending the Woodcock tracking algorithm to handle multiple densities. We explored two approaches: an intuitive method (Subsection 4.4) using N serial traversals to find the closest sample and a more performant method (Subsection 4.5) that combines N traversals into one for fewer steps. Additionally, we introduced a Monte Carlo estimator for a physically motivated inter-channel blending function. We generalized the framework to support other blending functions, including user-weighted blending and max opacity selection.

Our most efficient method achieved real-time frame rates for multi-channel rendering by leveraging occlusion to terminate traversal upon sampling. Unlike ray marching, which struggles with interactivity when casting shadow rays, our approach samples and casts shadow rays for a single sample without performance penalties. While initial frames exhibited variance, the results quickly converged to high-quality visuals within 20-30 ms. Despite producing initial noise-free images, taking fewer samples to increase the performance is not an option for ray marcher, which results in aliasing artifacts.

We proposed a density-based blending approach that avoided the

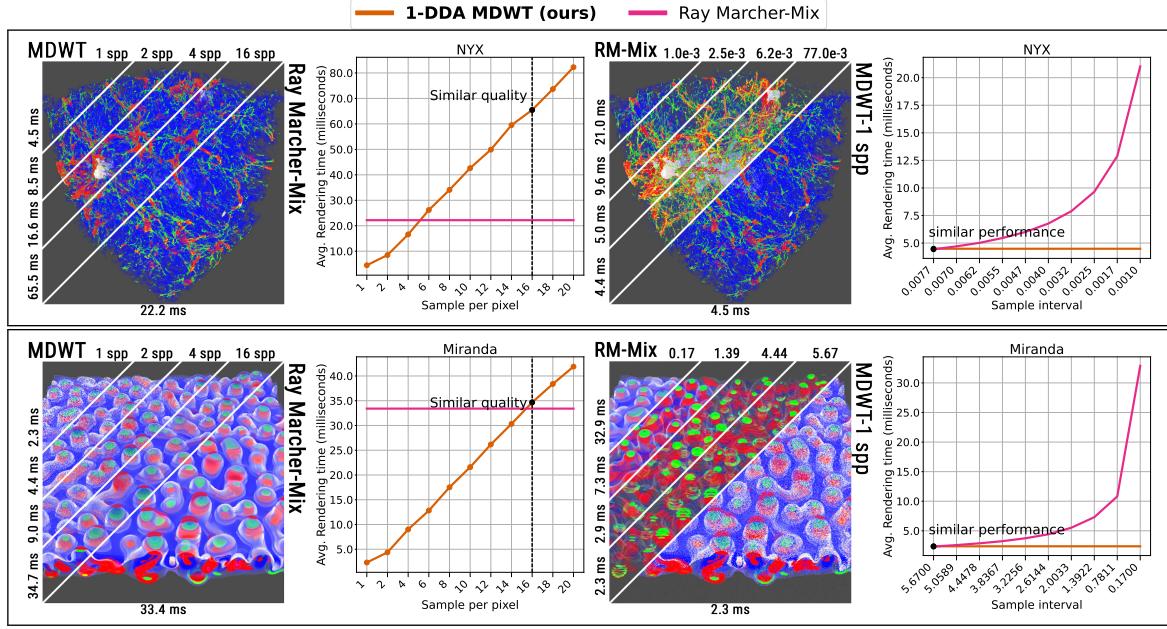


Figure 10: Similar quality (left) and performance (right) benchmarks for two datasets. The similar quality images display slices rendered with our Multi-Density Woodcock Tracking (MDWT) at increasing samples per pixel (spp), while the bottom-right half shows the ray marcher using mix blending with half of the minimum cell size as the sampling interval (Nyquist rate). The similar performance images show slices rendered with the ray marcher and mix blending (RM-Mix) for increasing sampling intervals, with the bottom-right half showing the same scene rendered by MDWT at 1 spp. The rendering time for each slice is indicated on the left side of each image, and each image is accompanied by a line plot of the average rendering times in milliseconds (lower is better).

pitfalls of opacity and compositing methods, producing distinct results without out-of-transfer-function colors. It required no complex interfaces, handled occlusion with many opaque channels, and unified prior blending functions, allowing users to choose blending based on their needs, such as weighted mixing for correlation discovery or max opacity for emphasizing dominant features.

A vital discussion is the impracticality of residual ratio tracking methods [NSJ14, KHLN17] in the sci-vis context due to how color information is derived. Unlike cinematic rendering—where RGB-classified data can be treated as separate volumes with dedicated majorants—sci-vis applies a transfer function to a single volume. Storing per-color-channel majorants is inefficient using uniform grids—which often outperform hierarchical structures in traversal and update performance [ZWS*24]—as this would incur significant memory overhead if replicated for each transfer function channel. Moreover, integrating minorant-based optimizations is also unsuitable as the analytically solvable control component of the extinction (alpha) and emission spectra (the RGB) could differ for the same transfer function.

This work suggests several potential avenues for future research. In Subsection 4.2, we introduced a framework for defining new blending functions, which could be expanded through a user study and investigating novel functions. Attribute-aware radial basis functions [MZS*23] offer a promising multi-channel blending strategy. The potential of the N -DDA method (Subsection 4.4) to reduce memory contention in out-of-core use cases could also be explored. Additionally, extending the approach to unstructured meshes using techniques like those from Morrical et al. [MSG*23], as well as research into ensemble and uncertainty visualization, would be valuable.

Ultimately, our multi-density Woodcock tracking method offers a performant and versatile solution for multi-channel volume visualization. Adapting contemporary rendering research, it delivers high-fidelity images obtained in real-time and provides more interactivity than previous

multi-channel sci-vis methods. Our physically driven blending function preserves colormaps and depth-ordering without complex user interfaces, enabling intuitive, user-driven visualization.

Acknowledgments

This work was supported by NSF[†] (Awards OAC 2138811, OISE 2330582), ARPA-H[‡] (Grant D24AC00338-00), the Intel oneAPI Center of Excellence at the University of Utah, NASA AMES (cooperative agreement 80NSSC23M0013), JPL (Subcontract 1685389), and the DFG[§] (Grant 456842964). We thank Lukic et al., Wang et al., Peter Lindstrom, and Yong Wan for providing the NYX, Hurricane, Miranda, and Zebrafish datasets, respectively, and NVIDIA Corporation for providing the hardware.

References

- [ANA21] ANDERSSON P., NILSSON J., AKENINE-MÖLLER T.: Visualizing and Communicating Errors in Rendered Images. In *Ray Tracing Gems II*, Marrs A., Shirley P., Wald I., (Eds.). 2021, ch. 19, pp. 301–320. [10](#)
- [CCF94] CABRAL B., CAM N., FORAN J.: Accelerated Volume Rendering and Tomographic Reconstruction Using Texture Mapping Hardware. In *Proceedings of the 1994 Symposium on Volume Visualization* (1994), VVS '94. [doi:10.1145/197938.197972](#). 2
- [CN93] CULLIP T., NEUMANN U.: *Accelerating Volume Reconstruction with 3D Texture Hardware*. Tech. rep., University of North Carolina at Chapel Hill, 1993. 2
- [CNLE09] CRASSIN C., NEYRET F., LEFEBVRE S., EISEMANN E.: Gigavoxels: ray-guided streaming for efficient and detailed voxel rendering. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games* (2009), I3D '09, Association for Computing Machinery. [doi:10.1145/1507149.1507152](#). 2

[†] National Science Foundation

[‡] Advanced Research Projects Agency for Health

[§] German Science Foundation

Table 2: Visual comparisons of inter-channel blending functions over test datasets: Direct volume renderings using Emission+Absorption model with identical parameters for various blending functions, accompanied by heat map images depicting the difference to our physically motivated density-based blending function using FLIP metric [ANA21]. The overall mean difference and PSNR to density-based blending are also reported under the heat maps.

Dataset \ Blending	Post-Render Composite	Equal Mix	Max Opacity	Density-Based (ours)	
NYX	Direct Volume Rendering				
	Difference to ours				
	Mean FLIP Diff., PSNR	47%, 16.56	21%, 23.03	13%, 22.85	
Hurricane	Direct Volume Rendering				
	Difference to ours				
	Mean FLIP Diff., PSNR	77%, 11.61	33%, 19.82	24%, 23.02	
Zebrafish	Direct Volume Rendering				
	Difference to ours				
	Mean FLIP Diff., PSNR	34%, 16.89	6%, 34.62	6%, 34.06	

[CR08] CABAN J., RHEINGANS P.: Texture-Based Transfer Functions for Direct Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008). [doi:10.1109/TVCG.2008.169](https://doi.org/10.1109/TVCG.2008.169). 2

[CS99] CAI W., SAKAS G.: Data Intermixing and Multi-volume Rendering. *Computer Graphics Forum* 18, 3 (1999). [doi:10.1111/1467-8659.00356](https://doi.org/10.1111/1467-8659.00356). 2, 3

[EKE01] ENGEL K., KRAUS M., ERTL T.: High-Quality Pre-Integrated Volume Rendering Using Hardware-Accelerated Pixel Shading. In *Proceedings of the ACM SIGGRAPH/Eurographics Workshop on Graphics Hardware* (2001), HWWS '01. [doi:10.1145/383507.383515](https://doi.org/10.1145/383507.383515). 1, 2

[FWKH17] FONG J., WRENNINGE M., KULLA C., HABEL R.: Production Volume Rendering: SIGGRAPH 2017 Course. In *ACM SIGGRAPH 2017 Courses* (2017), SIGGRAPH '17. [doi:10.1145/3084873.3084907](https://doi.org/10.1145/3084873.3084907). 2, 3, 6

[GBC*13] GALTIER M., BLANCO S., CALIOT C., COUSTET C., DAUCHET J., EL HAFI M., EYMET V., FOURNIER R., GAUTRAIS J., KHUONG A., PIAUD B., TERRÉE G.: Integral formulation of null-collision monte carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (2013), 57–68. [doi:10.1016/j.jqsrt.2013.04.001](https://doi.org/10.1016/j.jqsrt.2013.04.001). 2

[GKT16] GÜNTHER T., KUHN A., THEISEL H.: MCFTLE: Monte Carlo Rendering of Finite-Time Lyapunov Exponent Fields. *Computer Graphics Forum* 35 (2016). [doi:10.1111/cgf.12914](https://doi.org/10.1111/cgf.12914). 2

[GMH*19] GEORGIEV I., MISSO Z., HACHISUKA T., NOWROUZEZAHRAI D., KŘIVÁNEK J., JAROSZ W.: Integral formulations of volumetric transmittance. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (2019). [doi:10.1145/3355089.3356559](https://doi.org/10.1145/3355089.3356559). 2

[HHK*23] HERZBERGER L., HADWIGER M., KRÜGER R., SORGER P., PFISTER H., GROELLER E., BEYER J.: Residency Octree: A Hybrid Approach for Scalable Web-Based Multi-Volume Rendering. *IEEE Transactions on Visualization and Computer Graphics* 30 (2023). [doi:10.1109/TVCG.2023.3327193](https://doi.org/10.1109/TVCG.2023.3327193). 2

[HMES20] HOFMANN N., MARTSCHINKE J., ENGEL K., STAMMINGER M.: Neural Denoising for Path Tracing of Medical Volumetric Data. *ACM Transactions on Graphics (Proceedings of SIGGRAPH '20)* 3, 2 (2020). [doi:10.1145/3406181](https://doi.org/10.1145/3406181). 2

[HSS*05] HADWIGER M., SIGG C., SCHARSACH H., BÜHLER K., GROSS M.: Real-Time Ray-Casting and Advanced Shading of Discrete Isosurfaces. *Computer Graphics Forum* 24 (2005). [doi:10.1111/j.1467-8659.2005.00855.x](https://doi.org/10.1111/j.1467-8659.2005.00855.x). 2

[IKLH04a] IKITS M., KNISS J., LEFOHN A., HANSEN C.: *Chapter 39. Volume Rendering Techniques*. Addison-Wesley, 2004, p. 352–365. 2

[IKLH04b] IKITS M., KNISS J., LEFOHN A., HANSEN C.: Volume Rendering Techniques. In *GPU Gems*. 2004. Available at <https://developer.nvidia.com/gpumemsamples>.

nvidia.com/sites/all/modules/custom/gpugems/books/GPUGems/gpugems_ch39.html, Accessed: 24 June 2022. 1

[KHLN17] KUTZ P., HABEL R., LI Y. K., NOVÁK J.: Spectral and decomposition tracking for rendering heterogeneous volumes. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2017)* 36, 4 (2017). [doi:10.1145/3072959.3073665](https://doi.org/10.1145/3072959.3073665). 2, 9

[Kim11] KIM H. S.: *Visual Exploration in Volume Rendering for Multi-Channel Data*. PhD thesis, University of California, San Diego, 2011. 2

[KKSS13] KHLEBNIKOV R., KAINZ B., STEINBERGER M., SCHMALSTIEG D.: Noise-Based Volume Rendering for the Visualization of Multivariate Volumetric Data. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013). [doi:10.1109/TVCG.2013.180](https://doi.org/10.1109/TVCG.2013.180). 2

[KSC*10a] KIM H. S., SCHULZE J. P., CONE A. C., SOSINSKY G. E., MARTONE M. E.: Dimensionality Reduction on Multi-Dimensional Transfer Functions for Multi-Channel Volume Data Sets. *Information Visualization* 9, 3 (2010). [doi:10.1057/ivs.2010.6](https://doi.org/10.1057/ivs.2010.6). 2

[KSC*10b] KIM H. S., SCHULZE J. P., CONE A. C., SOSINSKY G. E., MARTONE M. E.: Multichannel Transfer Function with Dimensionality Reduction. *SPIE Proceedings* (2010). [doi:10.1117/12.839526](https://doi.org/10.1117/12.839526). 2

[Kvh84] KAJIYA J. T., VON HERZEN B. P.: Ray Tracing Volume Densities. *SIGGRAPH Comput. Graph.* 18, 3 (1984). [doi:10.1145/964965.808594](https://doi.org/10.1145/964965.808594). 3

[KW03] KRÜGER J., WESTERMANN R.: Acceleration Techniques for GPU-based Volume Rendering. In *Proceedings of IEEE Visualization* (2003), IEEE Visualization Conference. [doi:10.1109/VISUAL.2003.1250384](https://doi.org/10.1109/VISUAL.2003.1250384). 2

[KZX*23] KUMAR A., ZHANG X., XIN H. L., YAN H., HUANG X., XU W., MUELLER K.: RadVolViz: An Information Display-Inspired Transfer Function Editor for Multivariate Volume Visualization. *IEEE Transactions on Visualization and Computer Graphics* (2023). [doi:10.1109/TVCG.2023.3263856](https://doi.org/10.1109/TVCG.2023.3263856). 2

[Lia08] LIANG C. K.: *Bridging the Resolution Gap: Superimposition of Multiple Multi-Channel Volumes*. Master's thesis, University of California, San Diego, 2008. 2, 3, 4

[LKC05] LEE T.-H., KIM Y. J., CHANG J.: High Quality Volume Rendering for Large Medical Datasets Using GPUs. In *Systems Modeling and Simulation: Theory and Applications* (2005). [doi:10.1007/978-3-540-30585-9_74](https://doi.org/10.1007/978-3-540-30585-9_74). 2

[LLY06] LJUNG P., LUNDSTRÖM C., YNNERMAN A.: Multiresolution Interblock Interpolation in Direct Volume Rendering. In *EUROVIS - Eurographics/IEEE VGTC Symposium on Visualization* (2006). [doi:10.2312/VisSym/EuroVis06/259-266](https://doi.org/10.2312/VisSym/EuroVis06/259-266). 2

[ME11] MORAN P., ELLSWORTH D.: Visualization of AMR Data With Multi-Level Dual-Mesh Interpolation. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011). [doi:10.1109/TVCG.2011.252](https://doi.org/10.1109/TVCG.2011.252). 2

[MHK*19] MARTSCHINKIE J., HARTNAGEL S., KEINERT B., ENGEL K., STAMMINGER M.: Adaptive Temporal Sampling for Volumetric Path Tracing of Medical Data. *Computer Graphics Forum* (2019). 2

[MJW*13] MACIEJEWSKI R., JANG Y., WOO I., JÄNICKE H., GAITHER K. P., EBERT D. S.: Abstracting Attribute Space for Transfer Function Exploration and Design. *IEEE Transactions on Visualization and Computer Graphics* 19, 1 (2013). [doi:10.1109/TVCG.2012.105](https://doi.org/10.1109/TVCG.2012.105). 2

[MLB*23] MISSO Z., LI Y. K., BURLEY B., TEECE D., JAROSZ W.: Progressive null-tracking for volumetric rendering. In *ACM SIGGRAPH Conference Papers* (2023). [doi:10.1145/3588432.3591557](https://doi.org/10.1145/3588432.3591557). 6

[MSG*23] MORRICAL N., SAHISTAN A., GÜDÜKBAY U., WALD I., PASCUCCI V.: Quick Clusters: A GPU-Parallel Partitioning for Efficient Path Tracing of Unstructured Volumetric Grids. *IEEE Transactions on Visualization and Computer Graphics* 29, 01 (2023). [doi:10.1109/TVCG.2022.3209418](https://doi.org/10.1109/TVCG.2022.3209418). 2, 9

[MUWP19] MORRICAL N., USHER W., WALD I., PASCUCCI V.: Efficient Space Skipping and Adaptive Sampling of Unstructured Volumes Using Hardware Accelerated Ray Tracing. In *2019 IEEE Visualization Conference* (2019). [doi:10.1109/VISUAL.2019.8933539](https://doi.org/10.1109/VISUAL.2019.8933539). 2

[MZS*23] MORRICAL N., ZELLMANN S., SAHISTAN A., SHRIWISE P., PASCUCCI V.: Attribute-Aware RBFs: Interactive Visualization of Time Series Particle Volumes Using RT Core Range Queries. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2023). [doi:10.1109/TVCG.2023.3327366](https://doi.org/10.1109/TVCG.2023.3327366). 2, 9

[NSJ14] NOVÁK J., SELLE A., JAROSZ W.: Residual ratio tracking for estimating attenuation in participating media. *ACM Trans. Graph.* 33, 6 (2014). [doi:10.1145/2661229.2661292](https://doi.org/10.1145/2661229.2661292). 2, 9

[PH89] PERLIN K., HOFFERT E. M.: Hypertexture. *SIGGRAPH Comput. Graph.* 23, 3 (1989). [doi:10.1145/74334.74359](https://doi.org/10.1145/74334.74359). 2

[PJH23] PHARR M., JAKOB W., HUMPHREYS G.: 11 Volume Scattering. In *Physically based rendering: From theory to implementation*, 4 ed. The MIT Press, 2023, pp. 697–736. 3

[PLL*24] PAN B., LU J., LI H., CHEN W., WANG Y., ZHU M., YU C., CHEN W.: Differentiable Design Galleries: A Differentiable Approach to Explore the Design Space of Transfer Functions. *IEEE Transactions on Visualization and Computer Graphics* 30, 1 (2024). [doi:10.1109/TVCG.2023.3327371](https://doi.org/10.1109/TVCG.2023.3327371). 2

[RGW*03] RÖTTGER S., GUTHÉ S., WEISKOPF D., ERTL T., STRASSER W.: Smart Hardware-Accelerated Volume Rendering. In *Proceedings of the Symposium on Data Visualisation* 2003 (2003), vol. 3 of *VISSYM '03*. 2

[RS04] RICE A., SCHULZE J. P.: Real-Time Volume Rendering of Four Channel Data Sets. In *Visualization Conference, IEEE* (2004). [doi:10.1109/VISUAL.2004.89](https://doi.org/10.1109/VISUAL.2004.89). 2, 3, 4, 8

[RSK08] RAAB M., SEIBERT D., KELLER A.: Unbiased Global Illumination with Participating Media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006* (2008). [doi:10.1007/978-3-540-74496-2_35](https://doi.org/10.1007/978-3-540-74496-2_35). 2

[SDM*21] SAHISTAN A., DEMIRCI S., MORRICAL N., ZELLMANN S., AMAN A., WALD I., GÜDÜKBAY U.: Ray-traced Shell Traversal of Tetrahedral Meshes for Direct Volume Visualization. In *Proceedings of the IEEE Visualization Conference-Short Papers* (2021), VIS '21. [doi:10.1109/VIS49827.2021.9623298](https://doi.org/10.1109/VIS49827.2021.9623298). 2

[SKTM11] SZIRMAY-KALOS L., TÓTH B., MAGDICS M.: Free Path Sampling in High Resolution Inhomogeneous Participating Media. *Computer Graphics Forum* 30, 1 (2011). [doi:10.1111/j.1467-8659.2010.01831.x](https://doi.org/10.1111/j.1467-8659.2010.01831.x). 2

[TT84] TUY H. K., TUY L. T.: Direct 2-d display of 3-d objects. *IEEE Computer Graphics and Applications* 4, 10 (1984), 29–34. [doi:10.1109/MCG.1984.6429333](https://doi.org/10.1109/MCG.1984.6429333). 1

[WMPT65] WOODOCK E., MURPHY T., P. H., T. C. L.: *Techniques used in the GEM Code for Monte Carlo Neutronics Calculation in Reactors and Other Systems of Complex Geometry*. Tech. rep., Argonne National Laboratory, 1965. 1, 3

[WOCH09] WAN Y., OTSUNA H., CHIEN C.-B., HANSEN C.: An Interactive Visualization Tool for Multi-Channel Confocal Microscopy Data in Neurobiology Research. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009). [doi:10.1109/TVCG.2009.118](https://doi.org/10.1109/TVCG.2009.118). 2

[WOCH12] WAN Y., OTSUNA H., CHIEN C.-B., HANSEN C.: FluorRender: An Application of 2D Image Space Methods for 3D and 4D Confocal Microscopy Data Visualization in Neurobiology Research. In *2012 IEEE Pacific Visualization Symposium* (2012). [doi:10.1109/pacificvis.2012.6183592](https://doi.org/10.1109/pacificvis.2012.6183592). 2

[WOH*17] WAN Y., OTSUNA H., HOLMAN H. A., BAGLEY B., ITO M., LEWIS A. K., COLASANTO M., KARDON G., ITO K., HANSEN C.: Fluorender: Joint Freehand Segmentation and Visualization for Many-Channel Fluorescence Data Analysis. *BMC Bioinformatics* 18, 1 (2017). [doi:10.1186/s12859-017-1694-9](https://doi.org/10.1186/s12859-017-1694-9). 2, 8

[WZU*21] WALD I., ZELLMANN S., USHER W., MORRICAL N., LANG U., PASCUCCI V.: Ray Tracing Structured AMR Data Using ExaBricks. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2021). [doi:10.1109/TVCG.2020.3030470](https://doi.org/10.1109/TVCG.2020.3030470). 2

[YIC*10] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Unbiased, Adaptive Stochastic Sampling for Rendering Inhomogeneous Participating Media. *ACM Transactions on Graphics* 29, 6 (2010). [doi:10.1145/1882261.1866199](https://doi.org/10.1145/1882261.1866199). 2

[ZDL*20] ZHAO K., DI S., LIAN X., LI S., TAO D., BESSAC J., CHEN Z., CAPPELLO F.: Sdrbench: Scientific data reduction benchmark for lossy compressors. In *2020 IEEE International Conference on Big Data* (2020), IEEE. URL: <https://sdrbench.github.io>, [doi:10.1109/bigdata50022.2020.9378449](https://doi.org/10.1109/bigdata50022.2020.9378449). 6

[ZWS*24] ZELLMANN S., WU Q., SAHISTAN A., MA K.-L., WALD I.: Beyond ExaBricks: GPU Volume Path Tracing of AMR Data. *Computer Graphics Forum* 43, 3 (2024). [doi:10.1111/cgf.15095](https://doi.org/10.1111/cgf.15095). 2, 9