

Volume Encoding Gaussians: Transfer Function-Agnostic 3D Gaussians for Volume Rendering

Landon Dyken, Andres Sewell, Will Usher, Nathan Debardeleben, Steve Petruzza, Sidharth Kumar

Abstract—Visualizing the large-scale datasets output by HPC resources presents a difficult challenge, as the memory and compute power required become prohibitively expensive for end user systems. Novel view synthesis techniques can address this by producing a small, interactive model of the data, requiring only a set of training images to learn from. While these models allow accessible visualization of large data and complex scenes, they do not provide the interactions needed for scientific volumes, as they do not support interactive selection of transfer functions and lighting parameters. To address this, we introduce *Volume Encoding Gaussians* (VEG), a 3D Gaussian-based representation for volume visualization that supports arbitrary color and opacity mappings. Unlike prior 3D Gaussian Splatting (3DGS) methods that store color and opacity for each Gaussian, VEG decouple the visual appearance from the data representation by encoding only scalar values, enabling transfer function-agnostic rendering of 3DGS models. To ensure complete scalar field coverage, we introduce an opacity-guided training strategy, using differentiable rendering with multiple transfer functions to optimize our data representation. This allows VEG to preserve fine features across a dataset’s full scalar range while remaining independent of any specific transfer function. Across a diverse set of volume datasets, we demonstrate that our method outperforms the state-of-the-art on transfer functions unseen during training, while requiring a fraction of the memory and training time.

Index Terms—Volume Rendering, 3D Gaussian Splatting

1 INTRODUCTION

As high-performance computing (HPC) systems continue to grow in scale, scientific simulations are generating ever-larger and more complex datasets. These simulations frequently use adaptive mesh refinement (AMR) [2, 3, 5, 8, 9, 21] or unstructured meshes [2, 11, 34, 35] to reduce the memory footprint of output volumes. However, visualizing these volumes remains a significant challenge: datasets often exceed GPU memory, and visualization itself may require HPC resources, limiting interactivity and accessibility.

To address these challenges, recent work has explored machine learning (ML) to enable large-scale volume visualization, both for direct compression of volumetric datasets [24, 57, 60, 61] and for reducing the compute requirements during rendering [4, 10, 56, 58]. At the same time, novel view synthesis (NVS) techniques have emerged as a powerful method for representing 3D scenes within small memory and compute budgets. Requiring only a sparse set of training images depicting the original scene, a NVS model can learn to infer viewpoints at unseen angles, allowing for interactive rendering of a dataset with a memory footprint independent of the original size. Additionally, image-based training allows NVS techniques to natively support arbitrary data formats, as long as they can be rendered to a set of images. This presents an opportunity for volume visualization, where existing ML techniques have largely focused on structured volumes rather than AMR or unstructured data.

Within the field of NVS, 3D Gaussian Splatting (3DGS) has shown itself as an impressive technique for real-time rendering, and has enjoyed an explosion of follow-up work since its initial debut [13, 19]. In 3DGS, a scene is represented by a set of 3D Gaussians, each defined by a position, scale, rotation, view-dependent color, and opacity. These Gaussians are projected into screen space and alpha-blended to form the final image. Each Gaussian’s parameters are optimized by an ML framework using differentiable rendering, and the training pipeline dynamically densifies and prunes the total number of Gaussians. Once optimized, 3DGS enables fast and high-quality rendering of 3D scenes.

However, base 3DGS is fundamentally designed for surface-based, photogrammetry-derived datasets, not scientific volumes, as it embeds view-dependent color and opacity into each Gaussian. This makes 3DGS incompatible with the workflows that are typical for volume visualization, where users must be able to adjust color and opacity mappings interactively by changing transfer functions (TFs). To address this, previous work has extended 3DGS for volume visualization tasks [1, 49] using composition of multiple sets of *editable Gaussians* that allow for adjusting color, opacity, and lighting after training. Each set of editable Gaussians is trained using a collection of multi-view images rendered from a single TF. By using many TFs that highlight disjoint areas of the volume’s range, the composed set can represent the original data. While this method is accurate for the TFs the model was trained on, it is unable to robustly render unseen TFs. In addition, because the method involves composing multiple sets of 3DGS models to produce a full model, its training time and memory footprint quickly grow as more sets are required to represent larger datasets.

To address these issues, we present *Volume Encoding Gaussians* (VEG), a novel 3DGS representation tailored for scientific volume visualization. Rather than storing color and opacity for each Gaussian, VEG store only scalar field values, decoupling the data representation from its visual properties. This allows VEG to support fully transfer function-agnostic rendering, where arbitrary color and opacity mappings can be applied dynamically at render time. Additionally, the entire function range of a volume can be accurately represented with only one set of VEG, greatly reducing memory and training time compared to methods which rely on composing multiple models.

Our work presents the following contributions:

- Landon Dyken and Sidharth Kumar are with the University of Illinois Chicago. E-mail: ldyke|sidharth@uic.edu.
- Andres Sewell and Steve Petruzza are with Utah State University. E-mail: a02024444|steve.petruzza@usu.edu.
- Will Usher is with Luminary Cloud. E-mail: will@willusher.io.
- Nathan Debardeleben is with Los Alamos National Laboratory. E-mail: ndebard@lanl.gov.

- A method for novel view synthesis of volume data that can faithfully visualize color and opacity mappings unseen during training by approximating the underlying scalar field.
- An opacity-guided training technique that allows representing the full function range of volume data without separate training and composition for multiple transfer functions.
- An evaluation demonstrating that VEG achieve higher reconstruction quality on unseen transfer functions than the state-of-the-art, while having greatly improved training time and memory footprint (average $6.8 \times$ faster training time and $17.0 \times$ smaller files).

2 RELATED WORK

Interactive rendering of large-scale volumes has long been a focus area in scientific visualization [7, 39, 43, 55]. Ever larger data sets and the desire for faster rendering and higher visual quality pose a continuous challenge, while new hardware capabilities provide opportunities for new techniques [28–30, 37, 38]. For a current overview of the field, we refer the reader to the survey by Sarton et al. [39] for more details.

In this work, we propose an image-based machine learning method for volume visualization. While our work does not produce a neural representation, it is similar to methods that do in that we use machine learning to optimize our VEG to represent volume data. Because of this, we review neural representations for compressing and rendering volume data in Section 2.1. Next, we review image-based methods for volume visualization in Section 2.2. Finally, since our work draws heavily from 3DGS, we give background for this work and related methods in Section 2.3.

2.1 Neural Representations for Volume Data

Neural representations for volume data, also known as scene representation networks (SRNs), implicitly model a volumetric function from input positions to scalar values using neural networks. Because neural networks can be sampled quickly and represent data without scaling linearly with dataset size, they offer a compressed representation that can be used for rendering. While there has been much work in this field (see the recent survey by Wang and Han [52]), here we will focus on papers that apply to scientific volume data. Weiss et al. [59] showed the possibility of reconstructing volume data from ground truth images using differentiable volume rendering. Jain et al. [16] presented a neural method for compressed volume rendering of multivariate time-varying volumes. Kim et al. created NeuralVDB [20], which used neural networks to further improve the compression ability of the VDB [32] data structure for sparse volumes. Lu et al. created *Neurcomp* [24], which adapted SIREN’s [44] implicit neural representation for efficient volume data compression. Weiss et al. developed fV-SRN [57], which used dense-grid encoding [48] and fast tensor core inference to build a neural representation that supports interactive rendering. Wu et al. then improved on fV-SRN’s performance (in terms of reconstruction quality, training time, and rendering speed) with InstantVNR [61], which utilizes multi-resolution hash encoding [31], along with out-of-core training and batch inference during rendering.

Although our method is similar to these representations in its goal of enabling visualization of large-scale volumes, there are multiple advantages arising from our method being image-based. First, image-based methods easily fit into simulation workflows, where output images can be written at each timestep and later used for training without transfer of the prohibitively large output volumes. Next, image-based models naturally support any data type, while only a small set of neural representations have been demonstrated to work directly with AMR or unstructured volume data [23, 45]. Finally, image-based methods offer performance that is completely decoupled from the original data size, only depending on image resolution and scene complexity. This allows training of image-based models to be done even on consumer hardware, rather than requiring HPC resources.

2.2 Image-Based Methods for Volume Visualization

As image-based rendering and modeling methods are not tied to dataset size, they are important in enabling volume visualization for use cases where traditional rendering methods have prohibitively high compute and memory requirements. Early work in this area from Tikhonova et al. [50, 51] involved creating exploratory images, which allowed for changing color and opacity mappings of a rendered image without accessing the original volume data. Since then, image-based work has been extended to support synthesis of images from unseen view points using various ML approaches. Berger et al. [6] proposed pre-trained generative adversarial networks to synthesize volume-rendered images and explore the space of transfer functions while guiding user selection. Yao et al. [63] used a NeRF-based approach to train a network to represent pre-shaded volumetric scenes from a small number of rendered

images. VisNerf [62] presented a multidimensional radiance field representation that allows for novel view synthesis of volumetric scenes with interactive control over transfer functions, isovalue, timesteps, and simulation parameters. Most importantly to our work, iVR-GS [49] showed how 3D Gaussian splatting can be extended for volume visualization by implementing Blinn-Phong lighting and editable Gaussian colors after training. Instead of directly storing color for each Gaussian, they store a palette color common to a set of Gaussians, train an offset color for each Gaussian, then produce each Gaussian’s final color during rendering by adding the offset color to the palette color. This allows them to edit color after training by modifying the palette color of the set. We differ from iVR-GS in that we fully remove color and opacity attributes for each Gaussian, and instead allow our model to learn the scalar field represented by volume data directly. In doing so, we remove the need to train multiple models for each desired TF, achieving a transfer-function independent model after training.

2.3 3D Gaussian Splatting

Splatting methods involve projecting a set of 3-dimensional primitives directly onto a 2-dimensional image, then α -blending the primitives in order of their depth. Zwicker et al. [68] first proposed the use of Gaussian filters as a primitive for splatting, showing their effectiveness for rendering both surface and volume data. Since the landmark work on utilizing neural radiance fields (NeRF) [26], novel view synthesis has seen an explosion of work. While neural radiance fields have been shown to produce highly accurate renderings, they regularly suffer from high training times and low rendering frame rates, reducing their applicability for real-time interaction. Kerbl and Kopanas et al. [19] noted that, while the rendering algorithm for Gaussian splatting is very different, its image formation model is similar to traditional volume rendering methods such as those in neural radiance field rendering (NeRF) [26]. At the same time, the explicit nature of 3D Gaussians enables fast rendering compared to the expensive sampling necessary for rendering with the continuous implicit neural representation of NeRFs. 3D Gaussian splatting [19] was developed to optimize an explicit representation for novel view synthesis, allowing for greatly improved rendering performance while maintaining reconstruction quality.

Originally developed for photogrammetry-based scene reconstruction, in 3DGS each Gaussian stores an opacity along with view-dependent color using spherical harmonics, similar to previous work [14, 31]. In addition, initializing a set of 3D Gaussians relies on point clouds derived from NeRFs, COLMAP [40, 41] or similar structure-from-motion pipelines. These features of 3DGS are inherently designed for real-world image datasets, making them unsuitable for scientific volume data, where direct volume rendering [25] techniques are commonly used to map scalar field values to optical properties using transfer functions consisting of color and opacity maps. Arbitrary transfer functions are needed for data exploration, so volume rendering must support interactive transfer function adjustments, meaning that hard-coded color and opacities in the data representation are not suitable. While recent work [17] has used distributed 3DGS [67] for visualization of large scientific datasets, this method supports only a single static transfer function per 3DGS model. In our work, we remove this limitation through our development of Volume Encoding Gaussians which represent the underlying data field of a volume directly. Moreover, we replace the use of scene reconstruction for initialization, instead using the given volume dataset or random initialization.

Our work also takes advantage of recent extensions to 3DGS with regards to lighting effects and reduced memory footprint. Gao et al. [15] introduced relightable Gaussians, which augmented 3DGS with extra parameters to support Bidirectional Reflectance Distribution Function (BRDF) lighting. iVR-GS [49] adapted this method to the Blinn-Phong reflection model, which is more common to volume rendering use cases. Our method adopts their technique of Blinn-Phong lighting for producing shaded images with our own model. The original 3DGS formulation had an unnecessarily high memory requirement, and several works have been created to address this limitation [12, 33, 36]. We adopt these methods by utilizing vector quantization to compress our models while maintaining reconstruction quality, clustering most model

parameters into a codebook for storage.

3 DESIGN OF VEG

In this section, we give background on volume data types (Section 3.1), then describe the VEG representation (Section 3.2), an extension of 3DGS tailored to scientific volume visualization. Finally, we present how VEG are used for differentiable rendering with arbitrary transfer functions (Section 3.3).

3.1 Volume Types

A structured volume is defined by a set of scalar field values on a regular grid, with neighboring points assumed to be connected grid cells. The geometry of a structured grid is defined implicitly; only the data values need to be stored along with the grid dimensions and spacing. Any point in space can then be interpolated directly from the surrounding grid points. In contrast, unstructured volume data requires explicit storage of the geometry over which the scalar field is defined. Given a position in space, a test is performed to determine what cell in the volume, if any, contains it, before interpolating between that cell’s vertices to find the scalar value at that position.

3.2 The VEG Representation

In our work, we propose an explicit, point-based representation to encode a volume’s geometry and underlying scalar field. To do this, we use a 3D Gaussian formulation, where each point is defined as an ellipsoid centered at position μ with 3-dimensional scaling and rotation matrices, S and R . This allows the points to translate, stretch, and rotate independently in each dimension, giving the representation freedom to adapt to the arbitrary geometries of structured and unstructured volumes. Each Gaussian stores a value, allowing it to record the scalar field over the covered area of the volume’s geometry. In this way, our Volume Encoding Gaussians (VEG) give a compact representation that can encode the full information of the volume being rendered without storing any connectivity information for unstructured volumes.

3.3 Rendering of VEG

The goal of our representation is to approximate direct volume rendering [25] of the ground truth dataset. Specifically, we target raymarching with the emission-absorption model, where the color of each pixel is given by accumulating the contribution of N samples along the ray using α -blending with the discretized equation:

$$P = \sum_{i=0}^N c_i \cdot \alpha_i \cdot \prod_{j=0}^{i-1} (1 - \alpha_j) \quad (1)$$

where samples i of the represented scalar field produce each opacity α_i (emission) and color c_i (absorption) through transfer function application and lighting effects. We use 3D Gaussian splatting to model Equation 1 approximately, where the Volume Encoding Gaussians representing a scalar field are sampled through splatting. To render each Gaussian, the scaling and rotation matrices are used to create the Gaussians’ 3-dimensional covariance matrix, which is then projected to a 2-dimensional covariance matrix using the method of Zwicker et al. [68]. This converts the 3D Gaussian in object space into a 2D Gaussian in image space. The list of 2D Gaussians intersecting the screen can then be used to create the rendered image using the efficient tile-based rasterizer of Kerbl and Kopanas et al. [19]. This rasterizer employs α -blending of splatted Gaussians in front-to-back order, where a pixel accumulates each Gaussian intersecting it by evaluating its 2D covariance matrix at the pixel’s position (using the method of Yifan et al. [65]). This solves Equation 1, effectively approximating the emission-absorption model for raymarching-based volume rendering.

To determine how each VEG contributes to the color and opacity of a pixel, we first use a given transfer function to map each Gaussian to opacity and color. Rather than evaluating a transfer function when accumulating samples, it is more efficient to apply the transfer function as the first step of our rendering algorithm. This allows us to

immediately cull Gaussians with values mapped to low opacities, saving computation time similarly to empty space skipping in traditional volume rendering.

We found that, although this transfer function mapping follows the principles of volume rendering, it leads to poor training performance due to limiting the model’s ability to optimize opacity directly. This is important, as 3DGS effectively uses opacity to weigh the importance of different Gaussians, increasing the visibility of those that are useful and removing those with low opacity for density control. To reintroduce this ability into VEG, we augment our Gaussians with a trainable weight parameter analogous to the original opacity, and output the final opacity for each VEG as the product of its transfer function-dependent opacity and its weight, i.e. $\alpha_i = O(v_i) \cdot w_i$ where $O(v)$ is the opacity mapping of the transfer function, v_i is the Gaussian’s scalar value and w_i is the Gaussian’s weight.

Additionally, our model needs to support lighting effects, which are essential for 3D perception during rendering. For volume visualization, the most common lighting model used is Blinn-Phong. The Blinn-Phong model produces lighting effects by summing ambient (c^a), diffuse (c^d), and specular (c^s) colors output by samples, rather than just the transfer function-determined color $C(v)$. We adopt the technique of iVR-GS [49] for Blinn-Phong lighting in our model by attaching ambient, diffuse, and specular coefficients (k_i^a , k_i^d , and k_i^s), a normal vector (n_i), and a shininess term (β_i) to each Gaussian. Then, given a light direction l , we produce the final color for each Gaussian as:

$$c_i = k_i^a \cdot C(v_i) + k_i^d \cdot C(v_i) \cdot |n_i \cdot l| + c_i^s, \quad (2)$$

$$c_i^s = \begin{cases} k_i^s \cdot c_{white} \cdot |n_i \cdot \frac{v+l}{|v+l|} \beta_i|, & \text{if } |n \cdot l| > 0 \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

This gives the final parameter list for each VEG (G_i) in our model as $(\mu_i, S_i, R_i, v_i, w_i, n_i, k_i^a, k_i^d, k_i^s, \beta_i)$. In the backwards pass of rendering, gradients are computed manually for each Gaussian parameter using a modified version of Kerbl and Kopanas et al.’s [19] rasterizer. In order to optimize the attached scalar values for each Gaussian, we backpropagate gradients with respect to color and opacity through the lighting equations and transfer function used to render that frame. We assume linear interpolation of color and opacity between each control point of a given transfer function, but allow for an arbitrary number of control points to be given. To constrain values within $[0, 1]$ and ensure smooth gradients during training, we utilize a sigmoid activation function for storing scalar values and weights.

4 OPTIMIZATION OF VEG

The essence of our method is our optimization of VEG to accurately represent volume data. In this section we outline our strategy, starting with how we initialize our models in Section 4.1. We then detail our method for training these models to completion in Section 4.2.

4.1 Initialization

As shown in previous work [19], proper initialization can be essential for learning explicit scene representations. Because the input to our method is a volume dataset, it is important to initialize our VEG with some knowledge of the scalar field they will be trained to represent. To do this, we experimented with initializing VEG from either cell centers or vertices of a volume, but found both methods to give around equal performance, so present here only the vertex-centered method. In order to initialize our VEG, we first drop the connectivity information and treat the scalar-valued vertices as a point cloud with each point storing a value. We then transform each point into a VEG centered at the same position with the same scalar value. We initialize each VEG’s scale depending on the proximity of its closest neighbors, with larger Gaussians in areas of empty space and smaller Gaussians in areas of high density. This process is shown in Figure 2.

In practice, we randomly drop out all but a small number of data points before initializing the VEG in order to create compact models, as each VEG can represent many data points. This also makes our method easier to fit into HPC workflows, as only a small sampling of data is

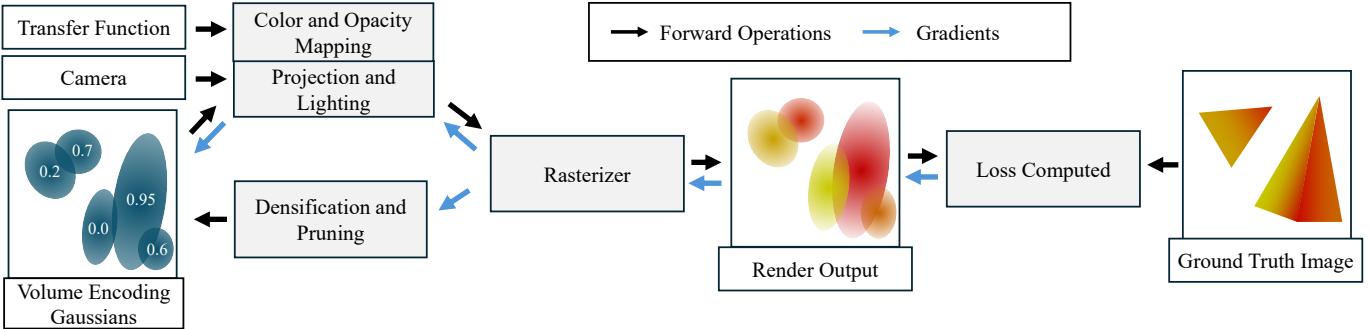


Fig. 1: Diagram of our training pipeline, where operation flow is depicted with black arrows and gradient backpropagation is depicted with blue. We conduct image-space training, in which our differentiable rendering algorithm is used to optimize Volume Encoding Gaussians from loss computed against ground truth images. Each training step, we render our VEG by opacity/color mapping with a given transfer function, applying lighting effects with the Blinn-Phong model, projection into 2D with a given camera view, then rasterization and blending of the 2D colored primitives. We then compare the rendered output against shaded reference images created from volume rendering with an emission-absorption model. After computing loss, gradients are backpropagated through to the scaling, rotation, position, scalar values, weight and lighting parameters of the VEG for optimization. Gradients are also used for densification and pruning of the VEG, in which the density of 3D Gaussians in space is adaptively controlled.

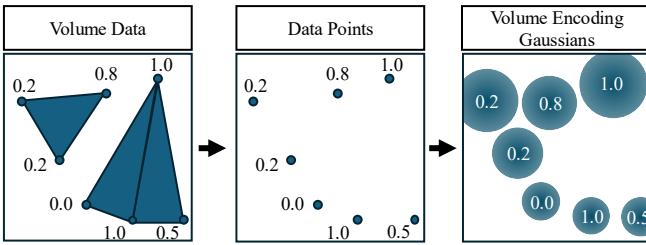


Fig. 2: Diagram of our initialization strategy to create VEG from volume data. Our approach begins by deleting all connectivity information, leaving only scalar-valued data points. Next, we extend these data points with 3D scaling and rotation matrices, so that each point is a 3D Gaussian with attached data value. These are initialized to make each initial 3D Gaussian a sphere with size relative to the density of other points around it, with lower density creating larger Gaussians.

required to be stored for later training. For use cases where this is not possible, our method can also perform well with random initialization, as shown in Section 5.3. The exact number of Gaussians needed for a VEG model depends on a dataset’s complexity, and so we rely on adaptive density control during training (see Section 4.2) and only use a very rough estimate for initialization. For unstructured volumes, we also experimented with selecting data points uniformly in space rather than randomly from the unstructured mesh, but found this to have little effect on training; in our evaluation, we use random selection of data points for initialization.

4.2 Training

To optimize our VEG, we conduct image-space training using our fast differentiable rendering algorithm described in Section 3.3. The complete training pipeline is shown in Figure 1. The model is iteratively optimized to match training images via backpropagation to each Gaussians’ parameters. Through this, the VEG are able to learn both the geometry of the volume and value of the scalar field.

In order to accurately recreate a dataset, the model must have training data that can allow it to learn all of the dataset’s information. This process requires creating a robust set of training images with corresponding transfer functions and camera information. For our evaluation, we accomplish this by rendering a set of ground truth images using PyVista’s [46] VTK backend [42]. For these images, PyVista ray-marches volumes with an emission-absorption model and shades them with Blinn-Phong lighting with a headlight from the camera. Images are rendered in a full camera orbit to ensure complete coverage of the given dataset.

To allow our model to render images for transfer functions never seen during training, we must ensure the VEG are able to reconstruct

the full function range of the dataset. For this, the transfer functions used to create the training set of images must be carefully chosen. Because optimization is done from loss computed between images, transfer functions which highlight certain ranges of the scalar field effectively weight the set of VEG to reconstructing the regions of the dataset containing those ranges. We found that the mapping of values to color did not have large effect on the resulting trained VEG, as long as the colormap used was not diverging. Opacity mapping, on the other hand, was especially significant to deciding which areas of the volume were reconstructed. This is intuitive, since areas of value mapped to high opacity will be rendered with high opacity in resulting images, making them more impactful on loss, which leads optimization to assign more relevance towards reducing their error.

We take advantage of this to improve the reconstruction quality of our models with *opacity-guided training*, where we use the opacity applied by transfer functions as an importance mapping to reconstruct all areas of the given dataset. To create our training set of images, we render each camera angle with multiple opacity maps highlighting different areas of the scalar field’s range. We design these opacity maps so that, in combination, every area of the range is weighted as equally important. Because these opacity maps will be used for optimizing scalar values, it is also important that there is slope between each data point, in order for gradients to move values up or down during training. We build the set of opacity maps by first dividing the function range into a number of steps, which we take as an input parameter. We then create an opacity map centered in each step, with linear drop-off around this center having slope equal to the number of steps. A visual comparison of trained Gaussian density resulting from our opacity-guided training technique against naive approaches is given in Figure 3. We evaluate the quality of models trained with our method against the naive approach, as well as investigate the number of opacity steps needed for accurate training in Section 5.3.

For each training iteration, a random camera angle and transfer function are chosen from the training set. The current set of VEG are then rendered and compared against the ground truth image for the same settings in order to compute loss. We compute image loss as a mixture of L1 and SSIM [53] loss, i.e. $(1 - \lambda)L_1 + \lambda L_{SSIM}$, where λ weighs them against each other. We set λ to 0.2 for all our tests, borrowing this convention from other work in 3D Gaussian splatting [19, 49]. Additionally, we also optimize Gaussian normals directly using the normal consistency loss from Gao et al. [15], and Gaussian lighting attributes (k_i^a, k_i^d, k_i^s , and β_i) using a small bilateral smoothness regularization term [49, 64, 66].

For best resulting models, training needs some way to adaptively control the density of Gaussians in space. This allows for removing Gaussians from areas where they are unneeded, saving memory and computation, and adding them to areas where the current number of

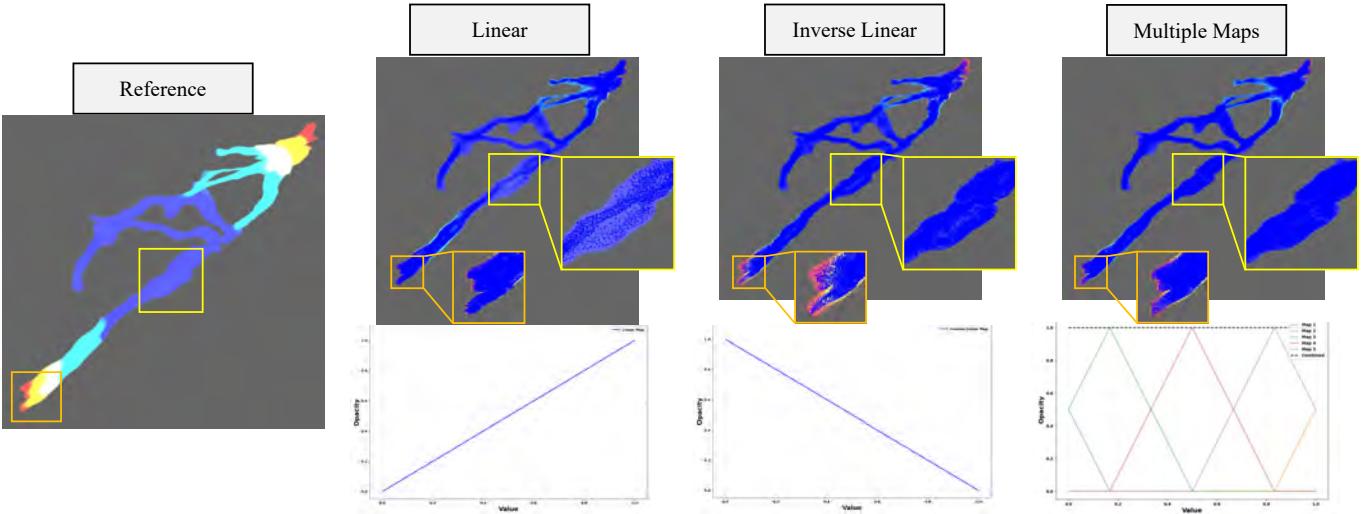


Fig. 3: Gaussian density comparison of VEG models trained using either a linear opacity map, an inverse linear opacity map, or our multiple opacity map method using three steps on the RBL dataset. Reference shows volume rendering with a pale rainbow colormap, where high values are red and low values are pale blue. Gaussian centers are painted using deep blue for each of the trained models, on top of the reference image, using SuperSplat’s [47] 3D Gaussian editor. Selected areas of high value are highlighted with an orange box, and areas of low value are highlighted with a yellow box. Notice training with the linear transfer function (which gives high opacity to high values) leads to greater Gaussian density in areas of high values and very low density in areas of low values, while the model trained with the inverse linear map does the opposite. Our multiple opacity map method leads to equal density for both areas of the function range.

Gaussians cannot accurately represent the volume’s complexity. This is accomplished via densification and pruning heuristics during training. For both, we follow the method of Kerbl and Kopanas et al. [19]. They show that, if a Gaussian has large view-space positional gradients, it is likely to be in a region that is not well reconstructed. Therefore, new Gaussians need to be added to that area, which is done by either cloning or splitting the Gaussian. For pruning, their work removes Gaussians with low attached opacity, so we follow this heuristic by using each VEG’s attached weight parameter w_i . We find these heuristics sufficient for creating accurately compact models depending on the complexity of the volume being reconstructed.

5 EVALUATION

In this section, we perform an extensive evaluation of our method both quantitatively and qualitatively. We list our datasets and system setup in Section 5.1. Our experimental study is to compare our method with the state-of-the-art image-based method for volume visualization, iVR-GS [49]. We evaluate both methods with regards to reconstruction quality, memory footprint, training time, rendering performance, and generalizability to unseen transfer functions in Section 5.2. We then perform an ablation study of the design decisions of our method in Section 5.3. Finally, we present and briefly describe a qualitative showcase of our method on more complex transfer functions in Section 5.4.

In addition, we also adapted the interactive renderer of Tang et al. [49] to our use case, as presented in Figure 4. Please see the supplemental video for a demo of interaction with our renderer and comparison with iVR-GS. We additionally created a web application port of this renderer, accelerated with WebGPU [54], to illustrate the potential of our method to enable lightweight client devices within a low memory footprint.

Dataset	Dimensions	Dataset	Tetrahedra
Vortex	(512, 512, 512)	RBL	3.89M
Chameleon	(1024, 1024, 1080)	Mito	5.54M

Table 1: Chosen datasets and their size. Separated into structured volumes (left) and unstructured volumes (right).

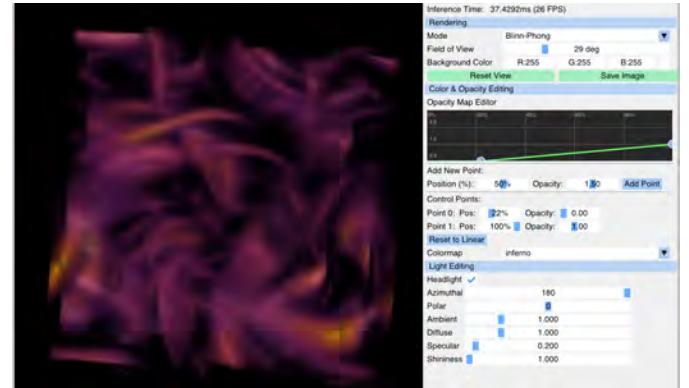


Fig. 4: Interactive VEG renderer adapted from Tang et al. [49]. Because VEG models are trained to generalize to arbitrary transfer functions, we replace their per-basic-set color and opacity sliders with selection of colormaps and an interactive opacity map creation tool.

5.1 Setup

Datasets To showcase the potential of VEG across different volume types, we benchmark our method on 4 datasets, consisting of 2 structured volumes (Vortex [49] and Chameleon [22]) and 2 unstructured volumes (Mito and RBL [37]). The size of each is given in Table 1.

Systems All training was performed on a workstation with an Nvidia A40 GPU and an AMD EPYC 9124 16-Core Processor. To illustrate the applicability of both methods to less powerful end-user systems, all rendering performance results were achieved on an XPS 17 laptop with an RTX 4070 Laptop GPU and an i9-13900H CPU.

5.2 Evaluation of Metrics

For our evaluation, we create training and testing sets of 800×800 resolution images using PyVista’s [46] VTK backend [42]. We perform volume rendering with Blinn-Phong shading using a headlight from the camera for all images. We use 10 training opacity maps as described in Section 4.2 for all datasets except RBL, where we use only 5 steps due to the dataset’s sparsity. Multiple colormaps were experimented with, but since all led to similar results we choose to use only the Matplotlib [18] “viridis” colormap during training for simplicity. For every transfer function used for training or testing,

Dataset	Method	Training TF		Unseen Colormaps		Broad Opacity		Narrow Opacity		Training Time (min)	# of Gaussians	File Size (MB)	Avg. FPS
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM				
Vortex	iVR-GS	37.83	0.99	28.41	0.96	23.20	0.92	21.69	0.90	219.8	3,848,171	90.2	3.9
	VEG	35.19	0.98	33.98	0.97	25.85	0.92	22.22	0.89	29.7	216,328	4.9	40.9
Chameleon	iVR-GS	34.41	0.96	29.27	0.94	24.42	0.90	24.99	0.91	144.8	1,327,237	31.1	17.5
	VEG	32.60	0.95	31.44	0.94	28.93	0.93	24.45	0.91	20.1	70,742	1.6	51.3
RBL	iVR-GS	37.38	0.99	34.03	0.99	27.83	0.98	23.49	0.98	49.8	195,384	4.6	47.1
	VEG	39.83	0.99	40.65	0.99	32.81	0.99	31.08	0.99	11.3	30,218	0.7	61.2
Mito	iVR-GS	37.23	0.99	25.55	0.98	23.20	0.96	23.57	0.97	137.0	1,105,228	25.9	19.8
	VEG	36.36	0.99	33.75	0.99	28.00	0.97	28.08	0.98	17.1	49,969	1.1	57.8

Table 2: Evaluation metrics for novel view synthesis across different datasets for our method (VEG) and the state of the art for volume visualization (iVR-GS). Results show that our models have drastically improved training times and memory footprints. Additionally, our models outperform iVR-GS when testing on color and opacity mappings not seen during training (Unseen Colormaps, Broad Opacity, Narrow Opacity), although iVR-GS has greater reconstruction quality when testing on the transfer functions used for training. Best results for each metric are in bold.

we render 160 images where camera positions are sampled from a spherical orbit around the volume iterating through 16 azimuthal and 10 elevation angles. The camera zoom was selected as the minimum which shows the entire volume’s bounding box in each image, and this design decision is motivated in Section 5.3. All testing images use camera views not seen during training, so that all results are for novel view synthesis. We create VEG models for each dataset by initializing Gaussians from 500,000 data points of the original volume. While the number of Gaussians necessary for a model depends on the represented dataset’s complexity, our densification and pruning methods allow the model to find this number during training, so an initial guess of 500,000 is sufficient. As seen in Table 2, models regularly end with drastically different numbers of Gaussians after training. All models are trained for 30,000 iterations with settings matching the defaults in 3DGS and iVR-GS wherever possible. Densification and pruning were enabled from iteration 500 until iteration 20,000. We explored using iVR-GS’s two-stage training approach for Gaussians to learn geometry before scalar values and lighting effects, but found best performance with training all parameters from the first iteration. We use a learning rate of 0.00025 for scalar values and 0.025 for weights, along with a pruning threshold of 0.005 for removing Gaussians with low weight. These values were found by conducting a hyperparameter search on the Mito dataset, and choosing the values that led to the highest reconstruction quality to number of Gaussians ratio.

We also create composed iVR-GS models for each dataset, following their approach of combining multiple basic sets of editable Gaussians trained from transfer functions with disjoint opacity bumps. For example, to train 10 basic sets, one will be trained using renderings where a transfer function left all values transparent except $[0, 0.1]$, the next with all values transparent except $[0.1, 0.2]$, and so on. We tested using flat opacity bumps for each basic set’s transfer function, but found equal performance when training with the ramped opacity maps used by our method, so here we present results using ramped opacity maps for consistency with our method. Choosing an appropriate number of basic sets to train for each dataset is not trivial, as composing more sets improves reconstruction quality while worsening training time, memory footprint, and rendering performance significantly. The original paper [49] does not provide an easy heuristic to decide where to make this tradeoff, so for our study we choose to compose 10 basic sets for all datasets except RBL, where we compose 5 due to its sparsity. As with our models, all training images use the “viridis” colormap. We use the default settings for their method in all cases, and follow their two-stage training approach for a total of 40,000 training iterations per basic model. The training time reported for iVR-GS models in Table 2 is the sum of training times for each of the basic sets, not including time to compose the sets together. The file size, number of Gaussians, and FPS reported are those of the composed models. All quality metrics are computed using renderings from the composed models. For testing of transfer functions, we apply colormaps by setting the palette color for each basic set to the colormaps’ value at the corresponding portion of the function range, taken at the midpoint of the value range that set represents. We apply opacity mappings in the same way, assigning the opacity for each basic set according to the given opacity function.

In order to measure the generalizability of both approaches for arbitrary transfer functions, we evaluate reconstruction quality on three use cases. First, performing novel view synthesis on the transfer functions seen during training. Next, rendering unseen views with colormaps that were unseen during training. We select five colormaps for variance with regards to color assigned to a specific part of the function range, along with the maps’ popularity. Since all models were trained on Matplotlib “viridis”, for testing we use “rainbow”, “rainbow reversed”, “cool to warm”, “warm to cool”, and “red blue yellow.” Finally, we render unseen views with opacity mappings not seen during training. For this evaluation, we want to test each approach’s ability to both infer volume renderings where large portions of the function range are visible, and surface-type renderings where only narrow portions of the function range are visible. To do this, we create sets of opacity maps where the slope of the opacity functions were either double the slope of the training transfer functions, leading to volume rendering-type images, or half the slope of the training transfer functions, leading to surface rendering-type images. We call these sets of opacity mappings “Broad Opacity” and “Narrow Opacity” in our results.

We present these results, along with training time, number of Gaussians in resulting models, file size, and rendering FPS in Table 2. We find several conclusions from this study. First, we find that our method drastically improves on training time, number of Gaussians, and file size compared to iVR-GS, with average improvements of $6.76\times$, $16.28\times$, and $16.99\times$ respectively. These improvements can mostly be explained by our method directly training a single set of 3D Gaussians to represent a volume, while iVR-GS requires training multiple basic sets individually, then composing them to create their full volume model. Even so, the number of Gaussians in our trained VEG models is less than the number of Gaussians in a single basic set of the corresponding iVR-GS models. We find that VEG models are capable of offering a high-quality interactive rendering of the 4.2GB chameleon dataset while requiring only 2.7MB of memory and 1.6MB for storage. Interestingly, while we do find large frame rate speedups with our method, they are far from linearly correlated with the number of Gaussians, likely due to fixed costs per frame and higher GPU utilization with larger models. All models are able to achieve interactive framerates except for iVR-GS’s Vortex model, showing the advantage of 3DGS-based methods for real-time rendering, even with a laptop GPU.

Next, we find that while iVR-GS outperforms our method for novel view synthesis using the transfer functions that the models were trained on, our method can more accurately infer novel views for transfer functions unseen during training. When generalizing to unseen colormaps, VEG models consistently outperform iVR-GS models, with an average increase of 5.64 PSNR. With regards to broad opacity maps, VEG models have an average increase of 4.24 PSNR. Finally, for narrow opacity maps, while iVR-GS actually outperforms VEG on Chameleon, VEG models still produce an average increase of 3.02 PSNR. These results show that with our training of scalar field values directly, we are able to create a transfer function-agnostic model in only 30,000 iterations of training, while iVR-GS’s color storing approach requires a total of 400,000 (200,000 for RBL) iterations of training while creating models with worse accuracy for transfer functions unseen during training.

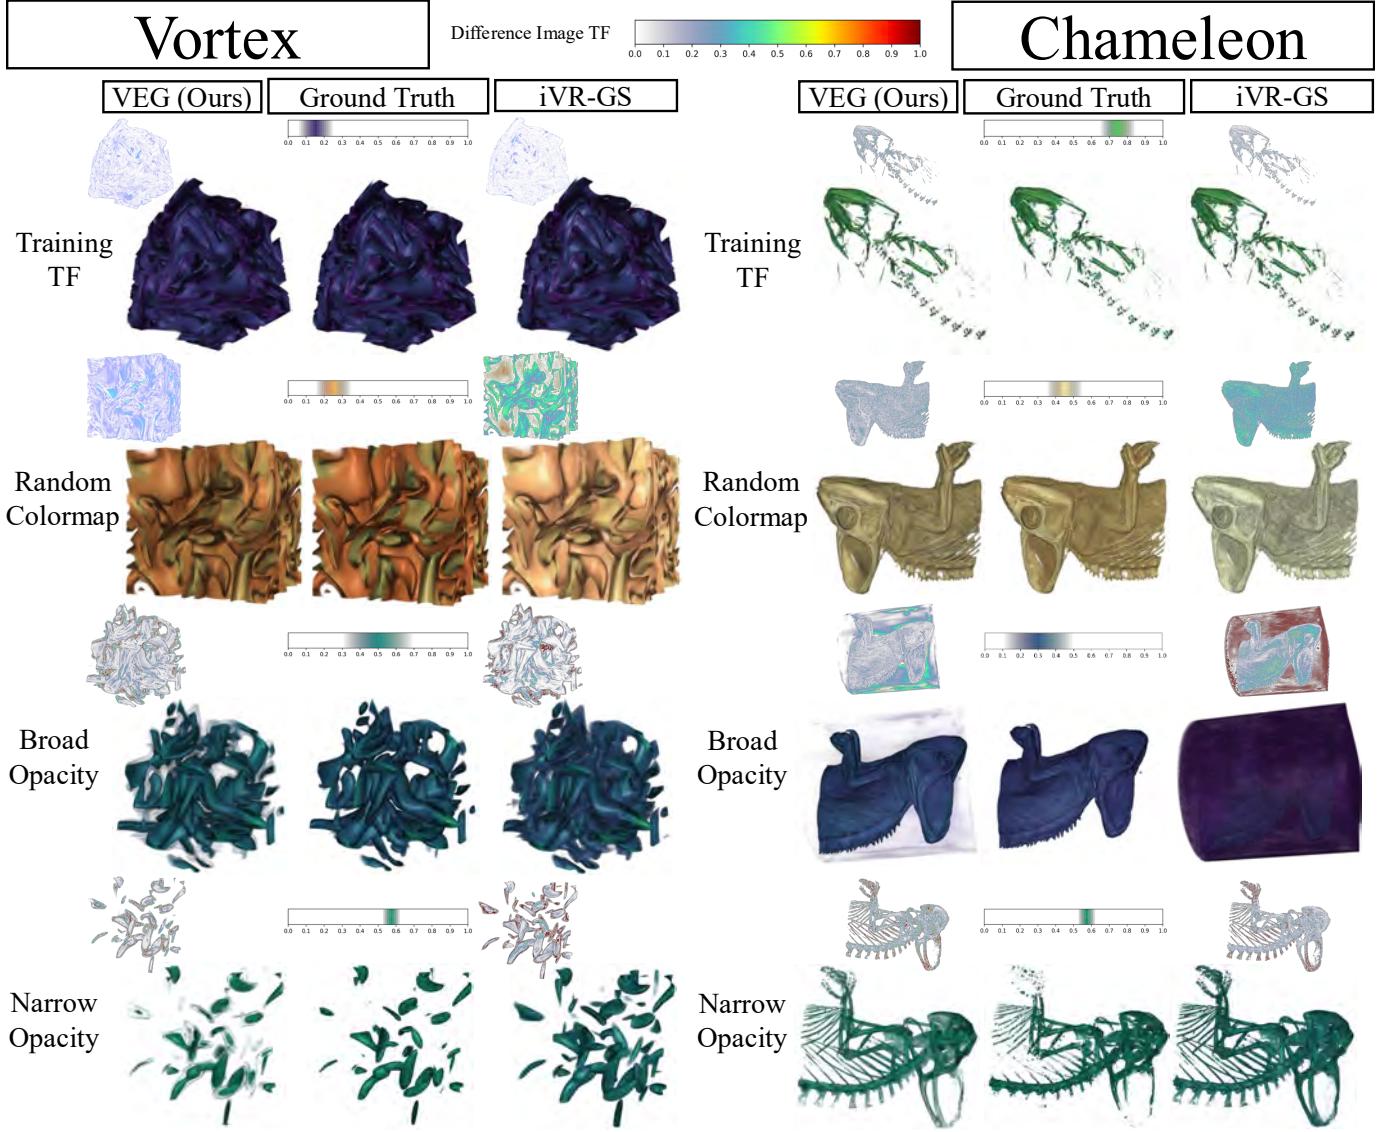


Fig. 5: Qualitative comparison of our method with iVR-GS on the Vortex and Chameleon datasets. For each dataset and method, we show results from the tests run in Table 2, along with difference images showing perceivable difference in the CIE LUV color space between reconstructions and ground truth. Difference images use the Matplotlib 'jet' colormap with linear opacity map, shown at top, to highlight areas of low error (purple to blue), medium error (cyan to green), and high error (yellow to red). These results showcase how VEG models outperform iVR-GS composed models for unseen color and opacity mappings. For example, see "Broad Opacity" for chameleon, where the iVR-GS model incorrectly assigns opacity to large portions of the volume that are supposed to be transparent.

We present a qualitative evaluation of our method in Figures 5 and 6, for structured and unstructured datasets respectively. We include a rendering for both our method and iVR-GS for every dataset, for each of the reconstruction quality tests assessed in Table 2. Ground truth renderings are given as well, along with images showing the pixel-wise perceivable difference between each reconstructed rendering and the ground truth. By analyzing these images, it's possible to see what causes the difference in reconstruction quality between methods for each type of test. Starting with the tests on training TFs, iVR-GS models outperform VEG for two reasons. First, by training color directly instead of optimizing scalar values through a transfer function, iVR-GS models can more accurately reconstruct the exact rendered colors of the ground truth. Second, iVR-GS models contain many more Gaussians than VEG models, so they are better able to approximate fine-grained volume geometry for the collection of surfaces that are seen during training. For the random colormap tests, iVR-GS's method of editing Gaussian color by changing the palette color for each basic set does not generalize well when the color being applied has high perceptual difference to the original palette color. In contrast, VEG models handle

recoloring through direct application of transfer functions. In both of the opacity mapping tests, the main reason for higher reconstruction quality with VEG models appears to be avoiding assigning opacity to areas that should not be rendered, illustrated by the Chameleon results for "Broad Opacity" in Figure 5. Storing scalar values allows VEG models to cull Gaussians more precisely where the opacity function determines, rather than iVR-GS which can only control opacity set-wide.

5.3 Ablation Study

In this section, we conduct several tests for validation of design decisions made for VEG. Specifically, we test the effect of our initialization strategy, our use of additional weight parameters per Gaussian, our opacity-guided training strategy, the generalizability of training with different numbers of opacity maps, and optimal zoom level for training image selection. To do this, we train models on the Mito dataset while changing various aspects of our algorithm. Unless otherwise stated, training and testing follow the procedure from Section 5.2 exactly. We present our results in Table 3. The "Default" row represents usage of the default parameters and method of our algorithm, and the metrics

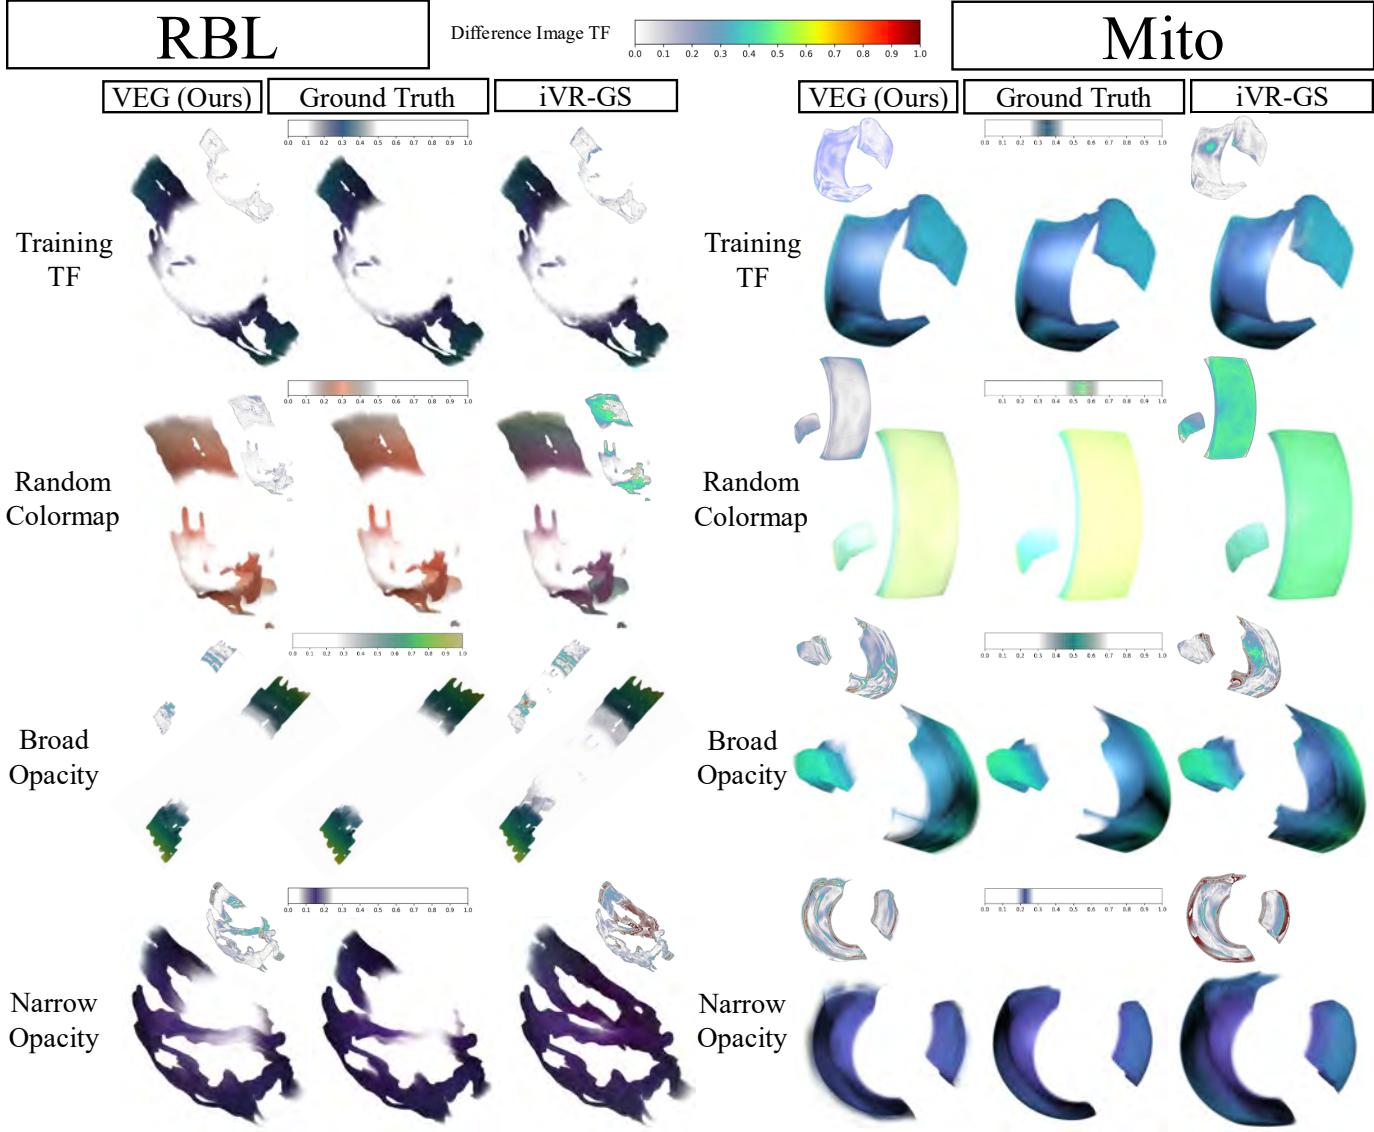


Fig. 6: Qualitative comparison of our method with iVR-GS on the RBL and Mito datasets. We show results from each of the tests run in Table 2, along with difference images computed in the same way as Figure 5. These results showcase how VEG models outperform iVR-GS composed models for unseen color and opacity mappings. See both of the "Random Colormap" sections, where iVR-GS models are incapable of applying the transfer function correctly, as it causes the basic sets to have palette colors perceptually different than during training. Also see "Broad Opacity" and "Narrow Opacity" sections, where the VEG models more closely follow the ground truth while the iVR-GS models incorrectly render more of the volume.

are taken directly from Table 2.

Initialization To test our volume initialization strategy, we compare our default initialization against a model trained starting from 500k data points with random positions and values. These results are given in the "Random Init" row of Table 3. Surprisingly, we find that the performance of this model is around equal to the default model in all cases except the training TF, where the default's quality is higher. This shows that our method is robust to random initialization, likely due to the densification and pruning strategies that allow creating and destroying new geometry as needed.

Weight Parameter Training As mentioned in Section 3.3, our VEG models include a trainable weight parameter per Gaussian that is multiplied with the transfer function-determined opacity when rendering. This allows for more direct model optimization, along with the ability to remove ineffective Gaussians as their weight parameters are pushed below our pruning threshold during training. To test this design decision, we present results from training a VEG model without the use of this parameter in Table 3, labeled "No Weights". As can be seen, the reconstruction quality of this model dips significantly from the default model. In addition, because there is no way to prune based on the weight

threshold, the model can only grow from its starting 500,000 Gaussians, ending with 563,223 compared to the default model's 49,969.

Multiple Opacity Maps An important contribution of our work is our opacity-guided training, in which we use chosen opacity maps to ensure that the entire function range of a volume is accurately represented. To test this, we compare training with our default 10 steps of opacity maps against naive training with a single linear opacity map. These results are given in the "Linear Opacity Map" row of Table 3. This model achieved low quality for all other sets of opacity maps, illustrating that the model did not learn the dataset's full function range.

In addition, we experiment with the number of opacity maps necessary for accurate volume reconstruction. We conduct training using our multiple opacity map method from Section 4.2 with number of steps 5 ("5 Opac Steps"), 10 ("Default"), and 20 ("20 Opac Steps"), then evaluate how well each model generalizes to each of the other sets of opacity maps. We present the results in Table 3, with "Training TF" testing the opacity maps from the set with 10 steps, "Broad Opacity" testing from the set with 5 steps, and "Narrow Opacity" testing from the set with 20 steps. As expected, each model outperformed the others when tested on the set of opacity maps it was trained on. We also find, as the number

Ablation	Training TF		Unseen Colormaps		Broad Opacity		Narrow Opacity		Training Time (min)	# of Gaussians	File Size (MB)
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM			
Default	36.36	0.99	33.75	0.99	28.00	0.97	28.08	0.98	17.1	49,969	1.1
Random Init	35.53	0.99	33.91	0.98	27.76	0.97	28.19	0.98	20.5	40,906	0.9
No Weights	33.65	0.99	33.72	0.99	24.76	0.95	25.92	0.97	37.8	563,223	13.3
Linear Opac	18.28	0.93	17.45	0.93	17.37	0.92	18.66	0.93	27.0	70,492	1.6
5 Opac Steps	26.23	0.97	26.92	0.97	36.08	0.99	25.82	0.97	18.2	88,678	2.1
20 Opac Steps	29.28	0.98	28.18	0.97	28.18	0.97	34.74	0.99	12.4	39,001	0.9

Table 3: Evaluation metrics for novel view synthesis with different ablations of our method, evaluated on the Mito dataset. Best results in bold.

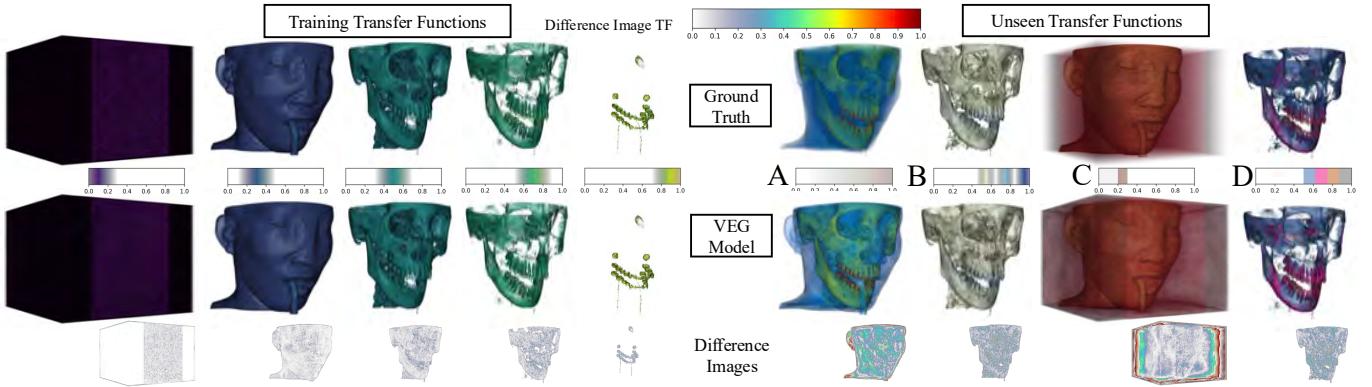


Fig. 7: Qualitative showcase of novel view synthesis for a VEG model on complex unseen transfer functions. We train a model for the Paraview head dataset using the five transfer functions on the left, and test with those and the four functions on the right (A-D), where the opacity and color maps were unseen during training. The top, middle, and bottom rows of images show ground truth, inference, and difference images (as in Figure 5).

Training Set	Default		Zoomed In		Training Time (min)	# of Gaussians
	PSNR	SSIM	PSNR	SSIM		
Default	36.36	0.99	33.34	0.98	17.1	49,969
Zoom	34.87	0.99	33.43	0.98	22.7	119,821

Table 4: Evaluation metrics for novel view synthesis for default and zoomed in testing sets of our method, evaluated on the Mito dataset. We compare our model trained on the default training set and training set including zoomed images. The default model outperforms on the default test set, while maintaining near equal quality on the zoomed in test set.

of training opacity maps increases, the model seems to perform better on opacity maps not seen during training. This is likely because wider opacity maps can prevent the model from learning features for regions of the volume that are overlapping for that map’s portion of the function range. Interestingly, the training time and number of Gaussians in the resulting models correlates inversely with the number of opacity maps used for training, likely because narrower opacity maps lead to sparser training images.

Optimal Zoom Level We also validate our selection of zoom level for training images. We create a set of test images identical to our Training TF testing set from Section 5.2 for the Mito dataset, but with varying zoom levels at either 1 (default), 2, or 3 \times . We name this testing set "Zoomed In," and report the performance of our default model on this and the default testing set in Table 4. Additionally, we train another model "Zoom," which is identical to the default except that its training image set also included images at 2 and 3 \times zoom, and report its results on both testing sets. We find that, while our model does perform worse when rendering zoomed in views, including similarly zoomed views in the training set does not improve model quality, validating our camera sampling strategy for training images. In fact, the "Zoom" model shows noticeably worse performance on the default views, while using more than 2 \times the number of Gaussians as the default model.

5.4 Complex TF Showcase

The final evaluation we present is to test our method on more complex transfer functions than those in Section 5.2. To do this, we use a Paraview example dataset of a CT scan of a boy’s head made up of a structured volume of 256x256x94 cells. We choose this dataset as it

consists of easily highlightable anatomical features, allowing for examples such as the opaque skull rendered underneath semi-transparent skin. We train a single model for this dataset in the same way as previously, using 5 opacity maps. We then manually select 4 transfer functions for testing, where the opacity and color maps were both unseen during training. From A to D in Figure 7, these were selected for demonstrating a linear opacity map over the full function range, multiple narrow opacity bumps, a semi-transparent opacity bump along with a high opacity bump, and a constant opacity as well as a qualitative colormap. Results of the VEG model for novel view synthesis on the training and unseen transfer functions are given in Figure 7. While the model performs more poorly on our custom transfer functions (especially in overfilling semi-transparent areas), these results show that VEG models can perform high-quality novel view synthesis even with arbitrary transfer functions disparate from those seen during training.

6 CONCLUSION

In this paper, we introduced *Volume Encoding Gaussians* (VEG), a novel 3D Gaussian splatting (3DGS) approach tailored specifically for structured and unstructured scientific volume rendering. VEG departs from conventional 3DGS techniques by removing color and opacity encodings and instead directly encoding scalar valued fields from volume datasets. This decouples the visual appearance from data representation, enabling lightweight, transfer function-agnostic models suitable for interactive visualization. Our evaluation demonstrates that VEG can achieve high reconstruction quality with robust generalization to previously unseen transfer functions, while maintaining tiny storage sizes, requiring less than 5 megabytes for all datasets studied.

For future work, we will explore support for time-varying volumetric data and distributed training schemes in order to improve scalability for very large datasets and enable training in HPC contexts. Additionally, we are interested in investigating the possibility of VEG to fully reconstruct volumes, rather than only render images. Future work will look into creating training pipelines for compact, high-quality VEG models to use as compressed representations for volume sampling directly. These could then be rendered through 3D Gaussian raytracing, similarly to recent work [27].

7 ACKNOWLEDGMENTS

This work was partly funded by NSF Collaborative Research Awards 2401274, 2221812, and 2106461, and NSF PPoSS Planning and Large awards 2217036 and 2316157.

REFERENCES

- [1] K. Ai, K. Tang, and C. Wang. NLI4VolVis: Natural language interaction for volume visualization via llm multi-agents and editable 3d gaussian splatting. *IEEE Transactions on Visualization and Computer Graphics (IEEE VIS 2025)*, 32(1), 2026. Preprint available at <https://arxiv.org/abs/2507.12621>. 1
- [2] W. K. Anderson, R. T. Biedron, J.-R. Carlson, J. M. Derlaga, B. Diskin, C. T. Druyor Jr, P. A. Gnoffo, D. P. Hammond, K. E. Jacobson, W. T. Jones, et al. Fun3d manual: 14.1, 2024. 1
- [3] D. Arndt, W. Bangerth, D. Davydov, T. Heister, L. Heltai, M. Kronbichler, M. Maier, J.-P. Pelteret, B. Turcksin, and D. Wells. The deal.II finite element library: Design, features, and insights. *Computers & Mathematics with Applications*, 2021. 1
- [4] D. Bauer, Q. Wu, and K.-L. Ma. FoVolNet: Fast Volume Rendering using Foveated Deep Neural Networks. *IEEE Transactions on Visualization and Computer Graphics*, 2023. doi: 10.1109/TVCG.2022.3209498 1
- [5] M. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82(1):64–84, 1989. doi: 10.1016/0021-9991(89)90035-1 1
- [6] M. Berger, J. Li, and J. A. Levine. A Generative Model for Volume Rendering. *IEEE Transactions on Visualization & Computer Graphics*, 25(04):1636–1650, Apr. 2019. doi: 10.1109/TVCG.2018.2816059 2
- [7] P. Bunyk, A. Kaufman, and C. T. Silva. Simple, Fast, and Robust Ray Casting of Irregular Grids. In *Scientific Visualization Conference (Dagstuhl '97)*, 1997. 2
- [8] C. Burstedde, L. C. Wilcox, and O. Ghattas. P4est : Scalable Algorithms for Parallel Adaptive Mesh Refinement on Forests of Octrees. *SIAM Journal on Scientific Computing*, 2011. 1
- [9] A. Dubey, K. Antypas, A. C. Calder, C. Daley, B. Fryxell, J. B. Gallagher, D. Q. Lamb, D. Lee, K. Olson, L. B. Reid, P. Rich, P. M. Ricker, K. M. Riley, R. Rosner, A. Siegel, N. T. Taylor, K. Weide, F. X. Timmes, N. Vladimirova, and J. ZuHone. Evolution of flash, a multi-physics scientific simulation code for high-performance computing. *The International Journal of High Performance Computing Applications*, 28(2):225–237, 2014. doi: 10.1177/1094342013505656 1
- [10] L. Dyken, W. Usher, and S. Kumar. Interactive Isosurface Visualization in Memory Constrained Environments Using Deep Learning and Speculative Raycasting. *IEEE Transactions on Visualization & Computer Graphics*, 31(02):1582–1597, Feb. 2025. doi: 10.1109/TVCG.2024.3420225 1
- [11] T. D. Economou, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso. SU2: An Open-Source Suite for Multiphysics Simulation and Design. *AIAA Journal*, 2016. 1
- [12] Z. Fan, K. Wang, K. Wen, Z. Zhu, D. Xu, and Z. Wang. Lightgaussian: unbounded 3d gaussian compression with 15x reduction and 200+ fps. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, article no. 4447, 21 pages. Curran Associates Inc., Red Hook, NY, USA, 2025. 2
- [13] B. Fei, J. Xu, R. Zhang, Q. Zhou, W. Yang, and Y. He. 3d gaussian splatting as new era: A survey. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–20, 2024. doi: 10.1109/TVCG.2024.3397828 1
- [14] Fridovich-Keil and Yu, M. Tancik, Q. Chen, B. Recht, and A. Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2
- [15] J. Gao, C. Gu, Y. Lin, Z. Li, H. Zhu, X. Cao, L. Zhang, and Y. Yao. Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XLV*, 17 pages, p. 73–89. Springer-Verlag, Berlin, Heidelberg, 2024. doi: 10.1007/978-3-031-72995-9_5 2, 4
- [16] W. Griffin, A. Godil, J. Bullard, J. Terrill, A. Varshney, and S. Jain. Compressed volume rendering using deep learning, 2017-10-02 2017. 2
- [17] M. Han, A. Sewell, J. Insley, J. Knowles, V. A. Mateevitsi, M. E. Papka, S. Petruzza, and S. Rizzi. Toward distributed 3d gaussian splatting for high-resolution isosurface visualization. In *2025 IEEE International Conference on eScience (eScience)*, pp. 331–332, 2025. doi: 10.1109/eScience65000.2025.00056 2
- [18] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55 5
- [19] B. Kerbl, G. Kopanas, T. Leimkuehler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4), article no. 139, 14 pages, July 2023. doi: 10.1145/3592433 1, 2, 3, 4, 5
- [20] D. Kim, M. Lee, and K. Museth. Neuralvldb: High-resolution sparse volume representation using hierarchical neural networks. *ACM Trans. Graph.*, 43(2), article no. 20, 21 pages, Feb. 2024. doi: 10.1145/3641817 2
- [21] C. C. Kiris, J. A. Housman, M. F. Barad, C. Brehm, E. Sozer, and S. Moini-Yekta. Computational framework for launch, ascent, and vehicle aerodynamics (lava). *Aerospace Science and Technology*, 55:189–219, 2016. doi: 10.1016/j.ast.2016.05.008 1
- [22] P. Klacansky. Open scivis datasets, December 2017. <https://klacansky.com/open-scivis-datasets/>. 5
- [23] K. Liu, C. Jiao, X. Gao, and C. Bi. Uginr: large-scale unstructured grid reduction via implicit neural representation. *J. Vis.*, 27(5):983–996, 14 pages, June 2024. doi: 10.1007/s12650-024-01003-y 2
- [24] Y. Lu, K. Jiang, J. A. Levine, and M. Berger. Compressive neural representations of volumetric scalar fields. *Computer Graphics Forum*, 40(3):135–146, 2021. doi: 10.1111/cgf.14295 1, 2
- [25] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, June 1995. doi: 10.1109/2945.468400 2, 3
- [26] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: representing scenes as neural radiance fields for view synthesis. *Commun. ACM*, 65(1):99–106, 8 pages, Dec. 2021. doi: 10.1145/3503250 2
- [27] N. Moenne-Loccoz, A. Mirzaei, O. Perel, R. de Lutio, J. M. Esturo, G. State, S. Fidler, N. Sharp, and Z. Gojcic. 3d gaussian ray tracing: Fast tracing of particle scenes. *ACM Transactions on Graphics and SIGGRAPH Asia*, 2024. 9
- [28] N. Morrical, A. Sahistan, U. Güdükbay, I. Wald, and V. Pascucci. Quick clusters: A gpu-parallel partitioning for efficient path tracing of unstructured volumetric grids. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):537–547, 2023. doi: 10.1109/TVCG.2022.3209418 2
- [29] N. Morrical, W. Usher, I. Wald, and V. Pascucci. Efficient Space Skipping and Adaptive Sampling of Unstructured Volumes Using Hardware Accelerated Ray Tracing. In *IEEE VIS Short Papers*, 2019. doi: 10.1109/VISL.2019.893539 2
- [30] N. Morrical, I. Wald, W. Usher, and V. Pascucci. Accelerating Unstructured Mesh Point Location with RT Cores. *IEEE Transactions on Visualization and Computer Graphics*, 2020. doi: 10.1109/TVCG.2020.3042930 2
- [31] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4), article no. 102, 15 pages, July 2022. doi: 10.1145/3528223.3530127 2
- [32] K. Museth. Vdb: High-resolution sparse volumes with dynamic topology. *ACM Trans. Graph.*, 32(3), article no. 27, 22 pages, July 2013. doi: 10.1145/2487235 2
- [33] S. Niedermayr, J. Stumpfegger, and R. Westermann. Compressed 3D Gaussian Splatting for Accelerated Novel View Synthesis. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10349–10358. IEEE Computer Society, Los Alamitos, CA, USA, June 2024. doi: 10.1109/CVPR52733.2024.00985 2
- [34] N. Offermans, A. Peplinski, O. Marin, and P. Schlatter. Adaptive mesh refinement for steady flows in Nek5000. *Computers & Fluids*, 2020. doi: 10.1016/j.compfluid.2019.104352 1
- [35] F. Palacios, J. Alonso, K. Duraisamy, M. Colonna, J. Hicken, A. Aranake, A. Campos, S. Copeland, T. Economou, A. Lonkar, T. Lukaczyk, and T. Taylor. Stanford University Unstructured (SU²): An open-source integrated computational environment for multi-physics simulation and design. In *51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. American Institute of Aeronautics and Astronautics, 2013. 1
- [36] P. Papantoniakis, G. Kopanas, B. Kerbl, A. Lanvin, and G. Drettakis. Reducing the memory footprint of 3d gaussian splatting. *Proc. ACM Comput. Graph. Interact. Tech.*, 7(1), article no. 16, 17 pages, May 2024. doi: 10.1145/3651282 2
- [37] B. Rathke, I. Wald, K. Chiu, and C. Brownlee. SIMD Parallel Ray Tracing of Homogeneous Polyhedral Grids. In *Eurographics Symposium on Parallel Graphics and Visualization*, 2015. doi: 10.2312/pgv.20151153 2, 5
- [38] A. Sahistan, S. Demirci, N. Morrical, S. Zellmann, A. Aman, I. Wald, and U. Güdükbay. Ray-traced Shell Traversal of Tetrahedral Meshes for Direct Volume Visualization. In *2021 IEEE Visualization Conference (VIS)*. IEEE, New Orleans, LA, USA, 2021. doi: 10.1109/VIS49827.2021.9623298 2
- [39] J. Sarton, S. Zellmann, S. Demirci, U. Güdükbay, W. Alexandre-Barff, L. Lucas, J. M. Dischler, S. Wesner, and I. Wald. State-of-the-art in Large-Scale Volume Visualization Beyond Structured Data. *Computer Graphics Forum*, 2023. doi: 10.1111/cgf.14857 2
- [40] J. L. Schönberger and J.-M. Frahm. Structure-from-motion revisited. In *Conference*

on Computer Vision and Pattern Recognition (CVPR), 2016. 2

[41] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 2

[42] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit (4th ed.)*. Kitware, 2006. 4, 5

[43] C. Silva, J. Comba, S. Callahan, and F. Bernardon. A survey of GPU-Based volume rendering of unstructured grids. *Brazilian Journal of Theoretic and Applied Computing (RITA)*, 2005. 2

[44] V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions, 2020. 2

[45] H. Son, J. Noh, S. Jeon, C. Wang, and W.-K. Jeong. Mc-inr: Efficient encoding of multivariate scientific simulation data using meta-learning and clustered implicit neural representations, 2025. 2

[46] C. B. Sullivan and A. Kaszynski. PyVista: 3d plotting and mesh analysis through a streamlined interface for the visualization toolkit (VTK). *Journal of Open Source Software*, 4(37):1450, may 2019. doi: 10.21105/joss.01450 4, 5

[47] SuperSplat <https://github.com/playcanvas/supersplat>. 5

[48] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezehrai, A. Jacobson, M. McGuire, and S. Fidler. Neural Geometric Level of Detail: Real-time Rendering with Implicit 3D Shapes . In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11353–11362. IEEE Computer Society, Los Alamitos, CA, USA, June 2021. doi: 10.1109/CVPR46437.2021.01120 2

[49] K. Tang, S. Yao, and C. Wang. ivr-gs: Inverse volume rendering for exploratory visualization via editable 3d gaussian splatting. *IEEE Transactions on Visualization and Computer Graphics*, 31(6):3783–3795, 2025. doi: 10.1109/TVCG.2025.3567121 1, 2, 3, 4, 5, 6

[50] A. Tikhonova, C. D. Correa, and K.-L. Ma. Explorable images for visualizing volume data. In *2010 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 177–184, 2010. doi: 10.1109/PACIFICVIS.2010.5429592

[51] A. Tikhonova, C. D. Correa, and K.-L. Ma. An exploratory technique for coherent visualization of time-varying volume data. In *Proceedings of the 12th Eurographics / IEEE - VGTC Conference on Visualization*, EuroVis'10, 10 pages, p. 783–792. The Eurographics Association & John Wiley & Sons, Ltd, Chichester, GBR, 2010. doi: 10.1111/j.1467-8659.2009.01690.x 2

[52] C. Wang and J. Han. DL4scvis: A state-of-the-art survey on deep learning for scientific visualization. *IEEE Transactions on Visualization and Computer Graphics*, 29(8):3714–3733, 20 pages, Aug. 2023. doi: 10.1109/TVCG.2022.3167896 2

[53] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004. 4

[54] WebGPU. <https://gpuweb.github.io/gpuweb/>. 5

[55] M. Weiler, M. Kraus, M. Merz, and T. Ertl. Hardware-based ray casting for tetrahedral meshes. In *IEEE Visualization*. IEEE Computer Society, 2003. doi: 10.1109/VISUAL.2003.1250390 2

[56] S. Weiss, M. Chu, N. Thuerey, and R. Westermann. Volumetric Isosurface Rendering with Deep Learning-Based Super-Resolution. *IEEE Transactions on Visualization and Computer Graphics*, 2021. doi: 10.1109/TVCG.2019.2956697 1

[57] S. Weiss, P. Hermüller, and R. Westermann. Fast neural representations for direct volume rendering. *Computer Graphics Forum*, 41(6):196–211, June 2022. doi: 10.1111/cgf.14578 1, 2

[58] S. Weiss, M. İslık, J. Thies, and R. Westermann. Learning adaptive sampling and reconstruction for volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 28(7):2654–2667, 2022. doi: 10.1109/TVCG.2020.3039340 1

[59] S. Weiss and R. Westermann. Differentiable direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):562–572, 2022. doi: 10.1109/TVCG.2021.3114769 2

[60] Q. Wu, D. Bauer, M. J. Doyle, and K.-L. Ma. Interactive Volume Visualization via Multi-Resolution Hash Encoding based Neural Representation. *IEEE Transactions on Visualization and Computer Graphics*, 2023. doi: 10.1109/TVCG.2023.3293121 1

[61] Q. Wu, D. Bauer, M. J. Doyle, and K.-L. Ma. Interactive volume visualization via multi-resolution hash encoding based neural representation. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5404–5418, 2024. doi: 10.1109/TVCG.2023.3293121 1, 2

[62] S. Yao, Y. Lu, and C. Wang. ViSNeRF: Efficient Multidimensional Neural Radiance Field Representation for Visualization Synthesis of Dynamic Volumetric Scenes . In *2025 IEEE 18th Pacific Visualization Conference (PacificVis)*, pp. 235–245. IEEE Computer Society, Los Alamitos, CA, USA, Apr. 2025. doi: 10.1109/PacificVis64226.2025.00029 2

[63] S. Yao and C. Wang. Revolve: Neural reconstruction of volumes for visualization enhancement of direct volume rendering. *Computers and Graphics*, 131:104295, 2025. doi: 10.1016/j.cag.2025.104295 2

[64] Y. Yao, J. Zhang, J. Liu, Y. Qu, T. Fang, D. McKinnon, Y. Tsin, and L. Quan. Neilf: Neural incident light field for physically-based material estimation. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, 17 pages, p. 700–716. Springer-Verlag, Berlin, Heidelberg, 2022. doi: 10.1007/978-3-031-19821-2_40 4

[65] W. Yifan, F. Serena, S. Wu, C. Öztïrel, and O. Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Trans. Graph.*, 38(6), article no. 230, 14 pages, Nov. 2019. doi: 10.1145/3355089.3356513 3

[66] J. Zhang, Y. Yao, S. Li, J. Liu, T. Fang, D. McKinnon, Y. Tsin, and L. Quan. NeILF++: Inter-Reflectable Light Fields for Geometry and Material Estimation . In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3578–3587. IEEE Computer Society, Los Alamitos, CA, USA, Oct. 2023. doi: 10.1109/ICCV51070.2023.00333 4

[67] H. Zhao, H. Weng, D. Lu, A. Li, J. Li, A. Panda, and S. Xie. On scaling up 3d gaussian splatting training, 2024. 2

[68] M. Zwicker, H. Pfister, J. van Baar, and M. Gross. Ewa volume splatting. In *Proceedings Visualization, 2001. VIS '01.*, pp. 29–538, 2001. doi: 10.1109/VISUAL.2001.964490 2, 3