

Efficient Neural Representation of Volumetric Data using Coordinate-Based Networks.

S. Devkota and S. Pattanaik

University of Central Florida, Orlando FL, USA

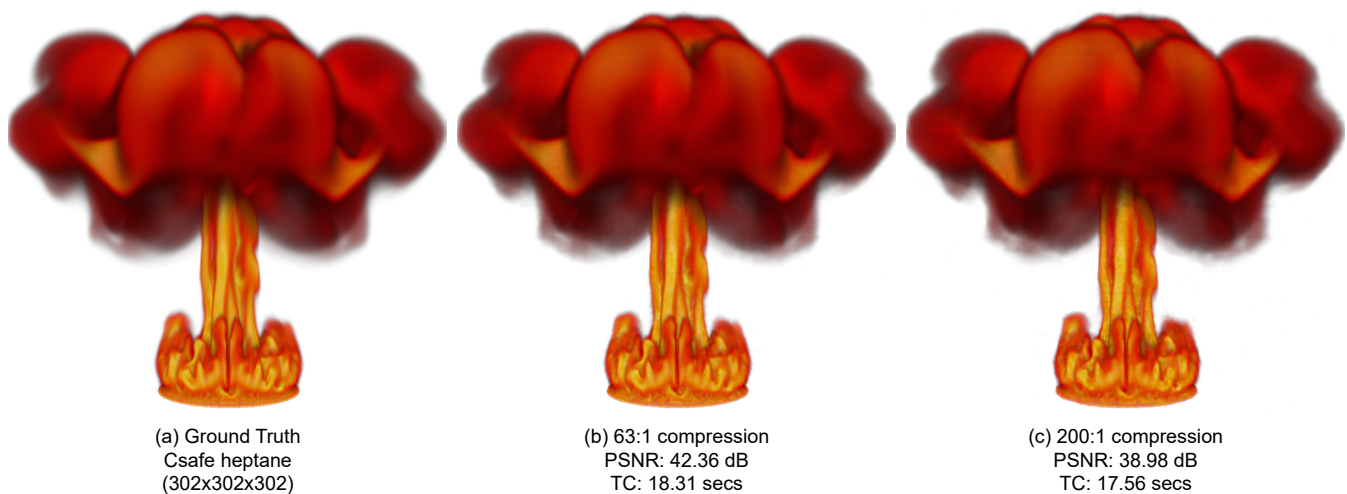


Figure 1: Neural representation of volumetric data using our proposed method. (PSNR: Peak Signal to Noise Ratio, TC: Time to compress)

Abstract

In this paper, we propose an efficient approach for the compression and representation of volumetric data utilizing coordinate-based networks and multi-resolution hash encoding. Efficient compression of volumetric data is crucial for various applications, such as medical imaging and scientific simulations. Our approach enables effective compression by learning a mapping between spatial coordinates and intensity values. We compare different encoding schemes and demonstrate the superiority of multi-resolution hash encoding in terms of compression quality and training efficiency. Furthermore, we leverage optimization-based meta-learning, specifically using the Reptile algorithm, to learn weight initialization for neural representations tailored to volumetric data, enabling faster convergence during optimization. Additionally, we compare our approach with state-of-the-art methods to showcase improved image quality and compression ratios. These findings highlight the potential of coordinate-based networks and multi-resolution hash encoding for an efficient and accurate representation of volumetric data, paving the way for advancements in large-scale data visualization and other applications.

CCS Concepts

• Human-centered computing → Visualization; • Computing methodologies → Image compression;

1. Introduction

Visualization of large-scale data can be a challenging problem. First, large-scale data sets can require significant amounts of storage space, making it difficult to store and access the data efficiently. Second, large-scale data sets can take a long time to transfer over

a network, which can impact the performance of visualization. Another challenge with large-scale data visualization is that the sheer size of the data can overwhelm visualization tools and systems, leading to slow rendering times, unresponsive interfaces, and difficulty in exploring the data.

Data compression can help address some of these challenges by reducing the amount of data that needs to be stored, transferred, and processed. There are many different techniques that can be used for volume data compression, ranging from traditional methods such as lossless and lossy compression to more recent approaches based on deep learning and neural networks.

One of the earliest methods for lossy compression was introduced in the 1990s by Ning and Hesselink [NH92], who proposed vector quantization (VQ) as a way to compress 3D scalar data. In their method, a codebook of representative 3D vectors is constructed by clustering similar data points. The original data is then compressed by replacing each data point with the closest codebook vector. Other early research on lossy compression for volume rendering also includes using discrete cosine transform (DCT) based compression [CLCM91] [YL95] and Fourier domain-based volume rendering [Mal93] [Lev92] [TL93]. These techniques are particularly useful for compressing smooth data, where most of the energy is concentrated in low-frequency components. Since the human visual system is more perceptive to low-frequency components, the high-frequency components from the image are safely discarded while still maintaining most of the information contained in the image.

Recently, scene representation using neural networks have gained a lot of popularity in the visualization community where the network encodes a field function that maps input 3D coordinates (coupled with direction vectors in some cases) to a field value, such as density or radiance, using neural networks. This approach has enabled a range of applications, such as novel view reconstruction [MST*20] and compression [SPY*22] [TET*22]. Thanks to the flexibility and differentiability of neural networks, this new approach has opened up many possibilities for volume data compression and visualization [LJLB21] [WHW21].

In our work, inspired by scene representation networks (SCN), we plan to represent volume data by directly approximating the mapping from spatial coordinates to volume values using a multi-layer perceptron (MLP). The trained MLP is then considered a compressed version of the original data. This representation is efficient because the memory footprint of a neural network is often orders of magnitude smaller than the original data, and sampling the representation is flexible as one can arbitrarily query volume values without explicit decompression and interpolation.

Recent research has demonstrated the efficacy of optimization-based meta-learning in reducing the number of gradient descent steps required for optimizing coordinate based neural networks in various domains, including images [SPY*22], signed distance fields [SCT*20], and radiance fields [TMW*20]. In our work, we focus on learning the weight initialization for neural representations specifically tailored to volumetric data. By utilizing learned values as the initial network weights, we establish a strong prior that facilitates faster convergence during optimization. To achieve this, we employ Reptile [NAS18], an optimization-based meta-learning algorithm, to generate optimized initial weights for representing a specific signal class, such as medical volume datasets. We opt for Reptile over alternative meta-learning algorithms like MAML [FAL17] due to its simplicity and computational efficiency. While MAML differentiates through the computation graph of the

gradient descent algorithm, Reptile employs gradient descent individually on each task, eliminating the need for graph unrolling or second derivative calculations [NAS18]. This key distinction allows for lower memory consumption and better computational efficiency.

To summarize, the main contributions from our work are as follows:

- A fast neural compression approach for volume data based on multiresolution hash encoding and a study on performance comparison with other input encoding techniques.
- Evaluation on the effectiveness of representation transfer with meta-learned initialization for neural compression of volume data.
- Experimental results on the efficacy of our volume compression approach against other state of the art techniques

2. Related Work

Volume data compression has been extensively researched in the field of computer science and data compression. Various lossy and lossless techniques have been proposed to achieve high compression ratios while preserving the quality of data. In this section, we provide an overview of some of the key works in three different research areas that are closely tied with each other: volume compression, neural scene representation, and neural volume representation.

2.1. Volume Compression

Recent years have seen a growing interest in lossy compression of large-scale volumes, driven by the expanding capabilities of simulation and data acquisition. Early research focused primarily on discrete cosine transforms [CLCM91] [YL95], and wavelet-based compression [IP99] [Mur93]. These techniques involve removing high-frequency information to achieve sparser data representation, which can then be quantized to reduce the number of bits required to store the data. Quantization-based volume compression techniques [NH92] [VGK96] [LH96] [TCRS00] involve dividing the volume into non-overlapping regions and mapping each region to a discrete value a set of representative values, which can be either pre-defined or learned from the data [SW03] [FM07] [GG16].

Inspired by fixed-rate texture compression methods, ZFP [Lin14] provides a compression method for floating point data in multi-dimensional arrays. This method uses a fixed-rate compression approach that maps small blocks of values in multiple dimensions to a user-specified number of bits per block. In recent years, tensor decomposition-based compression techniques [BRP15] [SMP13] are also being used for compactly encoding large multidimensional arrays. A notable mention in this category is TTHRESH [BRLP20] which utilizes Higher Order Singular Value Decomposition (HOSVD) with bit-plane, run-length, and arithmetic coding for volume data compression.

2.2. Scene Representation using Neural Networks

Scene representation networks (SRN) are a class of deep neural networks that can be used to encode occupancy fields [MON*19]

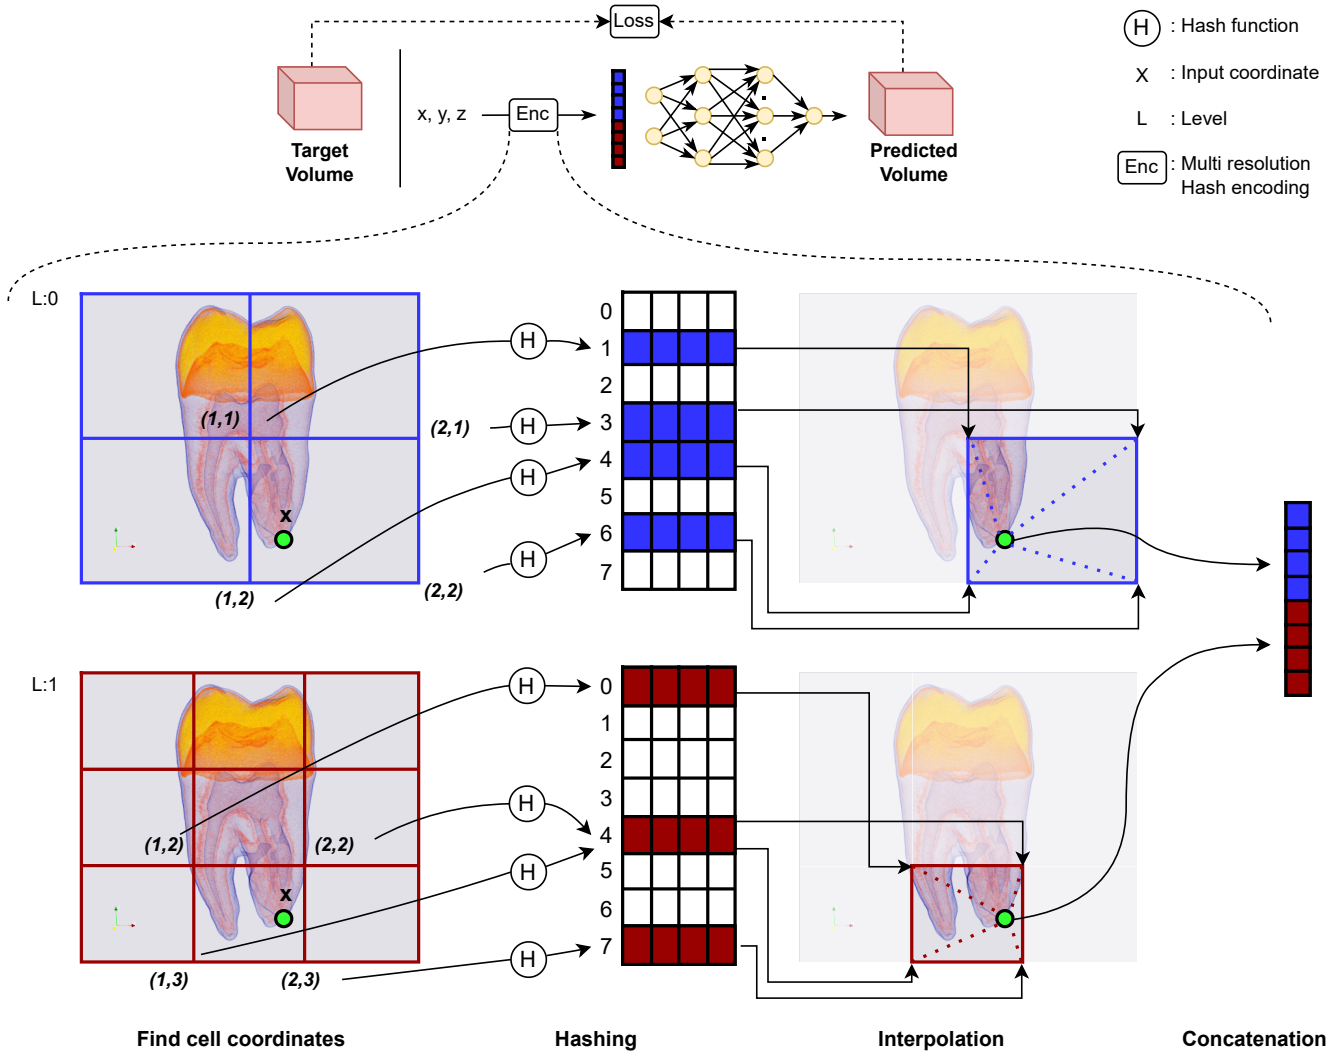


Figure 2: Overview of Multi-resolution hash encoding. In this two-dimensional example, we first break the image in L resolution levels (grids). Figure shows an example with two levels, $L:0$ and $L:1$. For a given normalized input coordinate x , the integer coordinates of the surrounding corners are hashed to obtain an index to a hash table with size T . Every entry in the hash table is a trainable feature vector of size W . In the example shown above, we have $L=2$, $T=8$, and $W=4$. The feature vectors from the surrounding corners are linearly interpolated to obtain the feature vector at coordinate x . Then all the feature vectors for x from each level are concatenated with each other which forms the final encoded vector for input x .

[PNM*20], implicit surfaces like SDF [MLL*21] [MPJ*19] [PFS*19], or radiance fields [LGL*20] [MST*20] [MRNK21] of complex 3D scenes in a compact and efficient manner. Most of these approaches use an encoding method that maps the input coordinates to a higher dimensional space before passing them to the network. For instance, to synthesize novel views of complex scenes using a sparse set of input images, Mildenhall *et al.* proposed to encode coordinates as a multi-resolution sequence of sine and cosine functions for the NeRF algorithm [MST*20]. Later, it was shown that using sinusoids with logarithmically-spaced axis-aligned frequencies further improves the reconstruction ability of coordinate-

based networks [TSM*20]. These encoding schemes are referred to as frequency-based encoding.

Recently, parametric encodings have been introduced which achieve state-of-the-art results. This encoding scheme involves arranging additional trainable parameters in auxiliary task-specific data structures such as a tree [TLY*21] or grid [JSM*20] [SSC22] [YFKT*21]. Although the total number of trainable parameters is higher for parametric encoding, it enables the use of much smaller coordinate-based networks and can be trained to converge much faster without sacrificing approximation quality. Recently, Müller *et al.* proposed the use of a multiresolution hash table of trainable parameters for input encoding which is task-independent and

outperforms all the previous approaches [MESK22]. Motivated by this, we plan to implement this encoding scheme in our volume representation.

While our approach draws inspiration from the principles of Müller *et al.*'s work [MESK22], there are notable differences that distinguish our approach. Unlike their work, we focus on neural representation of volumetric data where the network maps input coordinates to voxel values. Additionally, our study incorporates meta-learning to facilitate more efficient parameter initialization for the neural network representation. This inclusion of meta-learning allows us to harness domain-specific knowledge and further optimize the compression process for volumetric data.

2.3. Volume Representation using Neural Networks

Several approaches have been proposed to use deep learning in representing volume data for visualization. Some of the previous works introduced a new technique for volume visualization, which does not rely on the traditional rendering pipeline. Instead, they used either Generative Adversarial Networks [BLL19] or encoder-decoder [HLY19]. He *et al.* [HWG*20] used a convolutional regression model to learn the mapping from the simulation and visualization parameters to the final visualization. This allowed flexible exploration of parameter space for large-scale ensemble simulations. These networks can render the volume data directly using the compact representations stored in the network weights. However, the networks may not perform well if the training data does not contain the specific combination of views and transfer functions that are used in the test data.

The field of super-resolution is closely related to volume compression and can be used to enhance the visual quality of low-resolution volume data or rendered frames. Instead of storing data in higher resolution, super-resolution networks can efficiently up-scale the resolution of the data. One approach by Weiss *et al.* [WCTW21] used a deep learning-based architecture for isosurface rendering. Another method by Devkota *et al.* [DP22] implemented temporal reprojection for volumetric cases to perform super-resolution for direct volume rendering. Although these methods were applied to volumetric scalar fields, other works have focused on temporal [HW20] and spatial [HW22b] super-resolution for time-varying vector field data. Additionally, recent advancements in neural representation have led to methods that handle diverse scientific visualization tasks in a single framework. For instance, Han *et al.* [HW22a] proposed a unified coordinate-based neural network architecture capable of tackling both super-resolution and visualization tasks relevant to time-varying volumetric data visualization.

Recent contributions have brought forward numerous SCN-based approaches for volume representation. Neurcomp [LJLB21] demonstrated the effectiveness of coordinate-based networks for volume compression for scalar field data. Using implicit neural representation, their approach frames compression as function approximation, achieving highly compact representations that outperform existing volume compression methods. Their utilization of neural networks introduces a novel way of handling scalar field compression, and its performance benchmarks will be used for direct comparisons in our evaluation. Weiss *et al.* [WHW21] proposed to use

SCNs with tensor cores for faster decoding time and lower memory consumption during data reconstruction. Kim *et al.* introduced NeuralVDB [KLM22], a hybrid storage scheme with hierarchical neural network and wide VDB tree structure for efficient storage of sparse volumetric data. These studies have shown that SCN-based compression method achieves high compression ratios while maintaining a high level of visual fidelity compared to traditional compression methods. This makes it a promising approach for compressing large volumetric datasets in scientific visualization and medical imaging applications.

3. Methodology

We are interested in a function $\phi_\theta(x)$ that satisfies the objective function of the form

$$\arg \min_{\theta} F(\delta_x, \phi_\theta(x)) \quad (1)$$

The function $\phi_\theta(x)$ is a coordinate-based network with parameters θ that maps the spatial coordinates $x \in \mathbb{R}^d$ to some value that is as close as possible to δ_x , which is the intensity at coordinate x . Thus, our goal is to define such a volume representation network $\phi_\theta(x)$ that takes spatial coordinates as input to learn a mapping between the input coordinates and a target output.

In coordinate-based scene representation networks, the input coordinates undergo an encoding process to transform them to a higher-dimensional space before being inputted into the neural network. The encoding of the input plays a significant role in capturing the spatial information of the input data. If the input is not encoded, the network can only learn a relatively smooth function of position, which results in an inadequate representation of the intensity field [MESK22].

Recently, state-of-the-art results have been achieved by parametric encodings [JSM*20] [SSC22] [YFKT*21] for coordinate based networks. Motivated by these works, we propose to use multi-resolution hash encoding [MESK22] which we explain in the following section.

3.1. Multi-resolution Hash Encoding

Multi-resolution Hash Encoding introduces a multi-step process to effectively encode input coordinates. This subsection details each step of the encoding process.

Step 1 - Find cell coordinates

The encoding starts with breaking up a d dimensional array into L different levels (grids) with increasing resolution. Figure 2 shows an example domain with two dimensions (an image). The resolution for each level is a constant multiple of the previous level. The constant multiple is given by a growth factor b , which is computed based on three hyper-parameters: number of levels L , resolution of the coarsest level N_0 and resolution of the finest level N_{L-1}

$$b = e^{\frac{\ln N_{L-1} - \ln N_0}{L-1}} \quad (2)$$

Therefore, the resolution for each grid level l is computed as

$$N_l = \lfloor N_{l-1} * b \rfloor = \lfloor N_0 * b^l \rfloor. \quad (3)$$

Consider a normalized input coordinate $x \in \mathbb{R}^d$ in $[0, 1]$. For each level l , x is first scaled by the resolution of that level and is rounded up and down to find the integer coordinates of each corner of the cell containing x .

Step 2 - Hashing

Each of the 2^d integer coordinates surrounding the coordinate x is hashed using a hash function $H(x)$ [THM*03].

$$H(x) = \bigoplus_{i=1}^d x_i \pi_i \bmod T \quad (4)$$

Here, π_i denotes large prime numbers for each dimension i and \bigoplus denotes bit-wise XOR operation. The output from the hashing function is the index to a trainable hash table of size T . Each entry in the hash table is an array of trainable weights of size W . Thus, each grid level l has a corresponding hash table which is described by two hyper-parameters W and T . During training, the gradients are back-propagated all the way back to the hash table entries, dynamically optimizing them to learn a good input encoding.

Step 3 - Interpolation

After all the surrounding coordinates are mapped to index values which is used to query the hash table, the trainable weights from the corresponding slots are linearly interpolated in d dimensions to obtain the feature vector at coordinate x .

Step 4 - Concatenation

These feature vectors of size W from each of the L levels are concatenated together to form the input to a multi-layer perceptron. The output from the MLP is the predicted intensity value at coordinate x .

3.2. Meta Learned initialization

When we overfit any model to a volume data sample, it learns their parameters from scratch during training. With weights typically initiated at random, these models do not carry any domain knowledge about the volume data they are assigned to compress. Simply put, these models suffer from a lack of inductive biases. Inductive biases are assumptions that a model makes about the data that can help the model to learn more effectively.

To address this, we propose to apply the Reptile algorithm [NAS18]. This first-order meta-learning algorithm works by consistently selecting a task, conducting stochastic gradient descent on the chosen task, and subsequently shifting the initial parameters in the direction of the final parameters that were learned during that task. By building on inductive biases, it aims to enrich the model with a greater understanding of the data domain before the actual learning process begins.

In our case, we start with multiple volume datasets, sourced from a specific distribution V_d (for example: medical imaging or scientific simulation). The goal here is to identify initial weights, represented by θ , which would yield the lowest possible loss when

we optimize a network to represent a novel and previously unseen volume from the same distribution.

For our purpose, the algorithm works as shown in 1

Algorithm 1 Reptile Meta-Learning [NAS18]

- 1: **Initialize** θ , the initial parameter vector
 - 2: **for** iteration $1, 2, 3, \dots$ **do**
 - 3: Select a volume V at random from distribution V_d
 - 4: Perform $k > 1$ steps of gradient descent on V , starting with parameters θ , resulting in parameters W
 - 5: **Update:** $\theta \leftarrow \theta + \epsilon(W - \theta)$
 - 6: **end for**
 - 7: **return** θ
-

4. Experiments

In this section, we perform an array of experiments, ranging from investigating different encoding schemes and hyperparameters to performance comparisons against state-of-the-art techniques. In all of our experiments, we adopt mixed precision training [MNA*18], where the neural network weights are stored as *float16*. During training, in order to match the accuracy of the *float32* networks, a *float32* master copy of weights is maintained for parameter update. Unless otherwise stated, our hash encoding network in all of our experiments employs an MLP with two hidden layers and 64 neurons. Additionally, we use ReLU (rectified linear unit) as the activation function and L_2 -loss to guide the training process.

4.1. Encoding Schemes

In the context of coordinate-based networks, the encoding of inputs holds significant importance, and we compare the performance of different encoding schemes to determine their effectiveness. These encoding schemes are essential for mapping the coordinate inputs to a higher-dimensional space. One such scheme is frequency encoding, where each coordinate $x \in \mathbb{R}$ is represented as a sequence of sine and cosine functions :

$$E(x) = (\sin(2^0 \pi x), \cos(2^0 \pi x) \dots \sin(2^M \pi x), \cos(2^M \pi x)),$$

In our implementation, we select a value of $M = 10$. Another encoding scheme, called triangle wave encoding, replaces the sine function with a more computationally efficient triangle wave and omits the cosine function. Additionally, we explore the one-blob encoding scheme, a generalized version of one-hot encoding, where a Gaussian kernel is applied over the normalized input coordinate and discretized into multiple bins. In our case, we use $k = 64$ bins for this encoding scheme.

To showcase why we choose hash encoding over other schemes, we conduct a performance comparison of various encoding schemes for compressing the skull dataset. Figure 3 illustrates this comparison with similar compression ratios ranging from 81:1 to 91:1. We ensure that all networks have a similar number of trainable parameters. For the hash encoding network, we use a small MLP with 2 hidden layers and 64 neurons each, along with trainable multi-resolution hashtables ($L = 6, W = 8, T = 12$). In contrast, the other networks consist of 12 hidden layers with 128 neurons

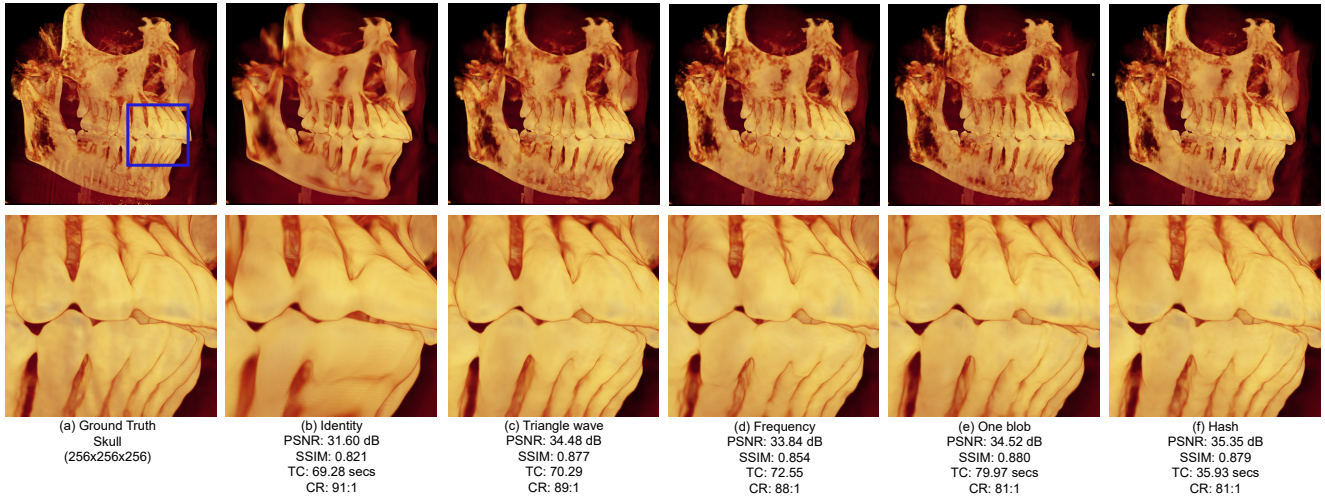


Figure 3: A comparison of various encoding schemes for compressing the skull dataset is presented. Each configuration was trained for 50 epochs, where each epoch is one complete pass over the entire volume. The top left image represents the ground truth rendering (a). The reconstruction quality is assessed using metrics such as PSNR and SSIM, while additional information including time to compress (TC) and compression ratio (CR) is provided beneath each image. Notably, the reconstructed rendering using Hash encoding demonstrates comparable quality to Triangle wave (c) and One blob encoding (e), but exhibits faster compression time, by at least 2x

each. We train the networks for 50 epochs where each epoch is one complete pass over the entire volume. When the network employed no encoding, specifically identity encoding, it struggled to preserve the high-frequency details present in the volume data, resulting in a smoothed approximation of the field data. Conversely, schemes such as frequency encoding, triangle wave encoding, and one-blob encoding enabled a more accurate representation of the field data within the same-sized network. However, these schemes still take substantial training times. In the case of multi-resolution hash encoding, the inclusion of trainable parameters within the hash table facilitated training with a smaller MLP architecture. Moreover, since the hash tables for all resolutions are queried in parallel, we achieved faster training times for compression. Notably, for our skull dataset, we observed a minimum 2x speedup with hash encoding compared to the alternative encoding schemes.

4.2. Hyperparameter study

Multi-resolution hash encoding offers flexibility in determining the number of encoding parameters, which is given by the product of hyperparameters L (number of levels), W (number of weights in each entry of the hash table), and T (size of the hash table). The choice of the hash table size T involves a trade-off between compression performance, memory usage, and compression quality. For T , we rely on the findings reported by Müller *et al.* [MESK22]. Based on their results, higher values of T lead to improved reconstruction quality, but at the expense of increased memory usage and slower training and inference. The memory footprint is linear in T , whereas quality and performance tend to scale sub-linearly. In our experiments, we select T values ranging from 2^8 to 2^{12} for different volume datasets, to lower the training and inference time while still having acceptable compression quality.

Table 1: Datasets used in our evaluation.

S.N.	Name	Dimension	Type	Source
1	Tooth	103x94x161	uint8	Open scivis [Ope]
2	MRI ventricles	256x256x124	uint8	[Bar]
3	MRHead	256x256x130	uint16	Slicer [Wik]
4	Aneurism	256x256x256	uint8	Philips Research [Phi]
5	Skull	256x256x256	uint8	Siemens [Sie]
6	Foot	256x256x256	uint16	Philips Research [Phi]
7	Mrt-angio	416x512x112	uint16	Institute for Neuroradiology [Gür]
8	Panoramix	441x321x215	int16	Slicer [Wik]
9	CT-chest	512x512x139	int32	Slicer [Wik]
10	CTA-cardio	512x512x321	int16	Slicer [Wik]
11	Manix	512x512x460	int16	[OSI]
12	Boston teapot	256x256x178	uint8	Terarecon Inc [Ter]
13	Backpack	512*512*373	uint16	Viatronix Inc [Kre]
14	Engine	256x256x128	uint8	General Electric [Gen]
15	Tacc turbulence	256x256x256	float32	Open scivis [Ope]
16	Csafe heptane	302x302x302	uint16	[Cen]
17	Vorticity magnitude	480x720x120	float32	[Vor]
18	Magnetic reconstruction	512x512x512	float32	[GLDL14]

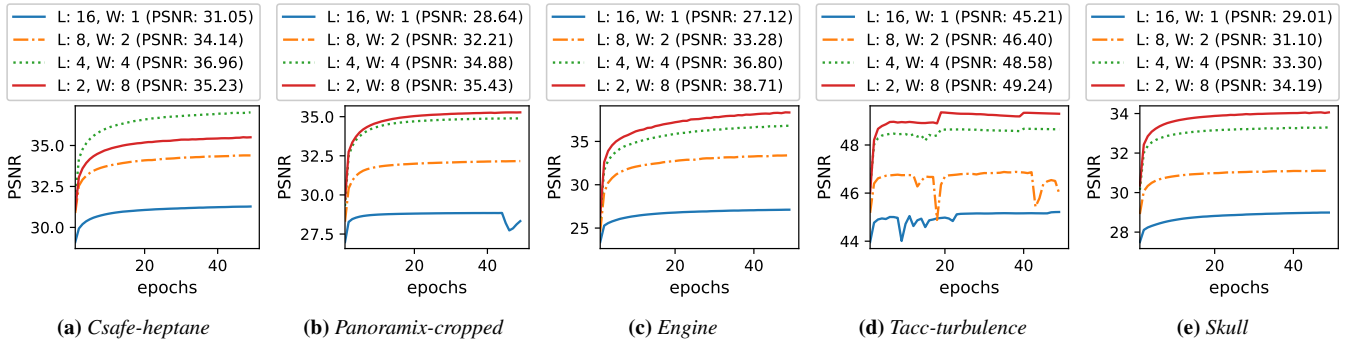


Figure 4: Figure shows compression error against the number of epochs for various combinations of L and W while keeping the total encoding parameters as constant. Notably, we find that configurations with W values ranging from 4 to 8 consistently achieve higher PSNR values across all datasets.

The hyperparameters L and W also influence the trade-off between compression quality and performance. To determine the optimal range of L and W for our specific use case, we plot the compression error against the number of epochs for various combinations of L and W , while keeping the total number of encoding parameters fixed (i.e., maintaining a constant compression ratio). As depicted in Figure 4, a configuration with W between 4 and 8 appears to yield favorable results across most datasets, and thus, we adopt this configuration for the majority of our evaluations. We vary L to achieve different compression levels while keeping the total encoding parameters constant.

4.3. Performance Comparison with Meta-Initialization

In this section, we investigate the potential benefits of incorporating meta-learned initialization into our volume representation network. Initially, we evaluate the impact of meta-initialization within a specific distribution of volume data. We perform meta-learning with a particular set of medical volume data and then use the learned parameters to compress another medical volume data that was not included in the meta-learning stage. Additionally, we also investigate whether meta-learned initialization from one domain, such as medical volume data, can enhance the performance of the network when applied to compress volume data from a different domain, such as scientific simulations.

4.3.1. Intra-domain Weight Transfer

For the first experiment, we apply meta-learning to optimize coordinate-based networks for representing medical datasets. The underlying dataset for meta-learning consists of medical volume data # 1 to 10 from Table 1. We use Reptile learning to meta-learn the initial weights for each medical volume data shown in Figure 5. For each experiment, we hold out one of the volume data as testing data and perform meta-learning using the remaining datasets. For every iteration of the meta-learning stage, we randomly sample a volume data from the dataset pool and perform k gradient updates on that sample before updating the initial parameters using the update rule outlined in Algorithm 1. Here, k gradient updates correspond to completing one full pass over the entire volume.

At testing time, we optimize a similar-sized network initialized with meta-learned parameters to compress the test dataset. For comparison, we also optimize another network with random initialization for the same test dataset. The underlying MLP architecture consists of 2 hidden layers with 64 neurons each, and we apply multi-resolution hash encoding with $L = 6, W = 4$, and $T = 12$. As seen in Figure 5, using the learned initial weights enables faster convergence, which is particularly evident at the initial training phase. While the random initialization approach eventually achieves a similar PSNR to the meta-learned approach, the latter surpasses random initialization after only a few gradient updates.

4.3.2. Inter-domain Weight Transfer

In the second experiment, we utilize the same set of medical volume datasets (datasets # 1 to 10 from Table 1) to generate meta-learned parameters. Subsequently, during the testing phase, we use these meta-learned parameters to optimize a coordinate-based network for compressing different volume datasets that are not from the medical domain. The comparison with random initialization for these datasets is presented in Figure 6. We observe that, with the exception of Tacc turbulence and Boston teapot datasets where the benefits of meta-learned initialization were slight or insignificant, our method typically resulted in faster convergence across the majority of the datasets. These two evaluations suggest that leveraging domain-specific knowledge through meta-learned initialization can bring significant improvements to the efficiency of volumetric data compression and representation.

4.4. Comparison with State of the Art Volume Compression Methods

In this section, we primarily benchmark our approach against Neurcomp, which represents a state-of-the-art in neural compression technique. Additionally, we perform comparisons with Tthresh, an advanced CPU-based volume compression technique. To assess the effectiveness of our volumetric neural representation in encoding the ground truth volume data, we sequentially decode the volume at its original resolution from the compressed representation, and measure the similarity between the original volume data and the volume predicted by our methodology, Neurcomp, and Tthresh.

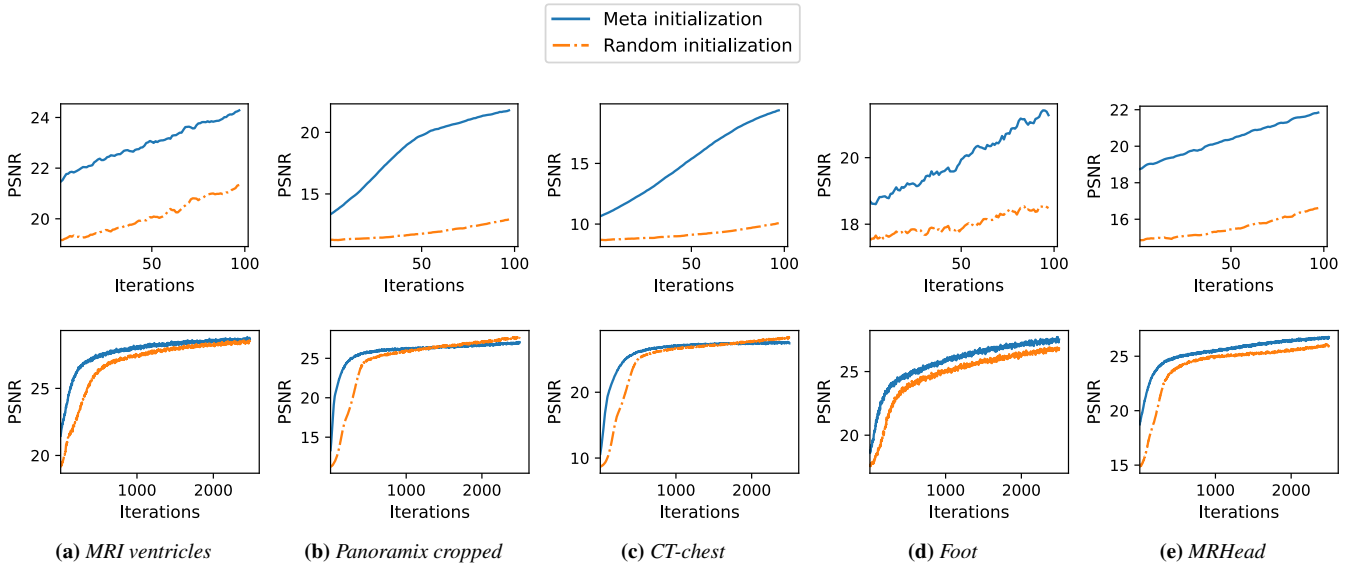


Figure 5: Comparison of convergence speed between meta-learned initialization and random initialization for intra-domain weight transfer. The reconstruction PSNR is reported for the first 100 iterations (top row) and 2500 iterations (bottom row) for each dataset. The number of iterations corresponds to the number of gradient updates performed during the training process. The meta-learned approach exhibits faster convergence, particularly evident at the initial training phase. While the random initialization approach eventually achieves a similar PSNR to the meta-learned approach, the latter surpasses random initialization after only 100 iterations in terms of PSNR.

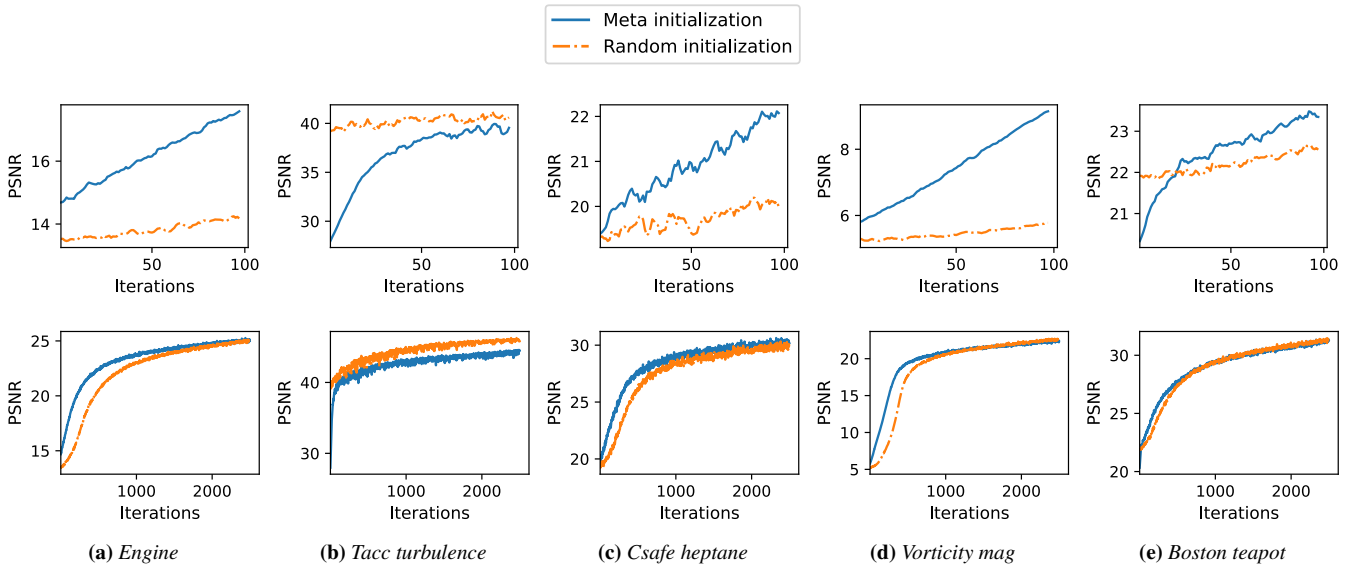


Figure 6: Comparison of convergence speed between meta-learned initialization and random initialization for inter-domain weight transfer. The reconstruction PSNR is reported for the first 100 iterations (top row) and 2500 iterations (bottom row) for each dataset. The number of iterations corresponds to the number of gradient updates performed during the training process. The meta-learned approach demonstrates slightly faster convergence for datasets a) engine, c) Csafe heptane, and d) Vorticity magnitude. However, it provides minimal to no advantage for dataset e) Boston teapot and performs poorly for dataset b) Tacc turbulence.

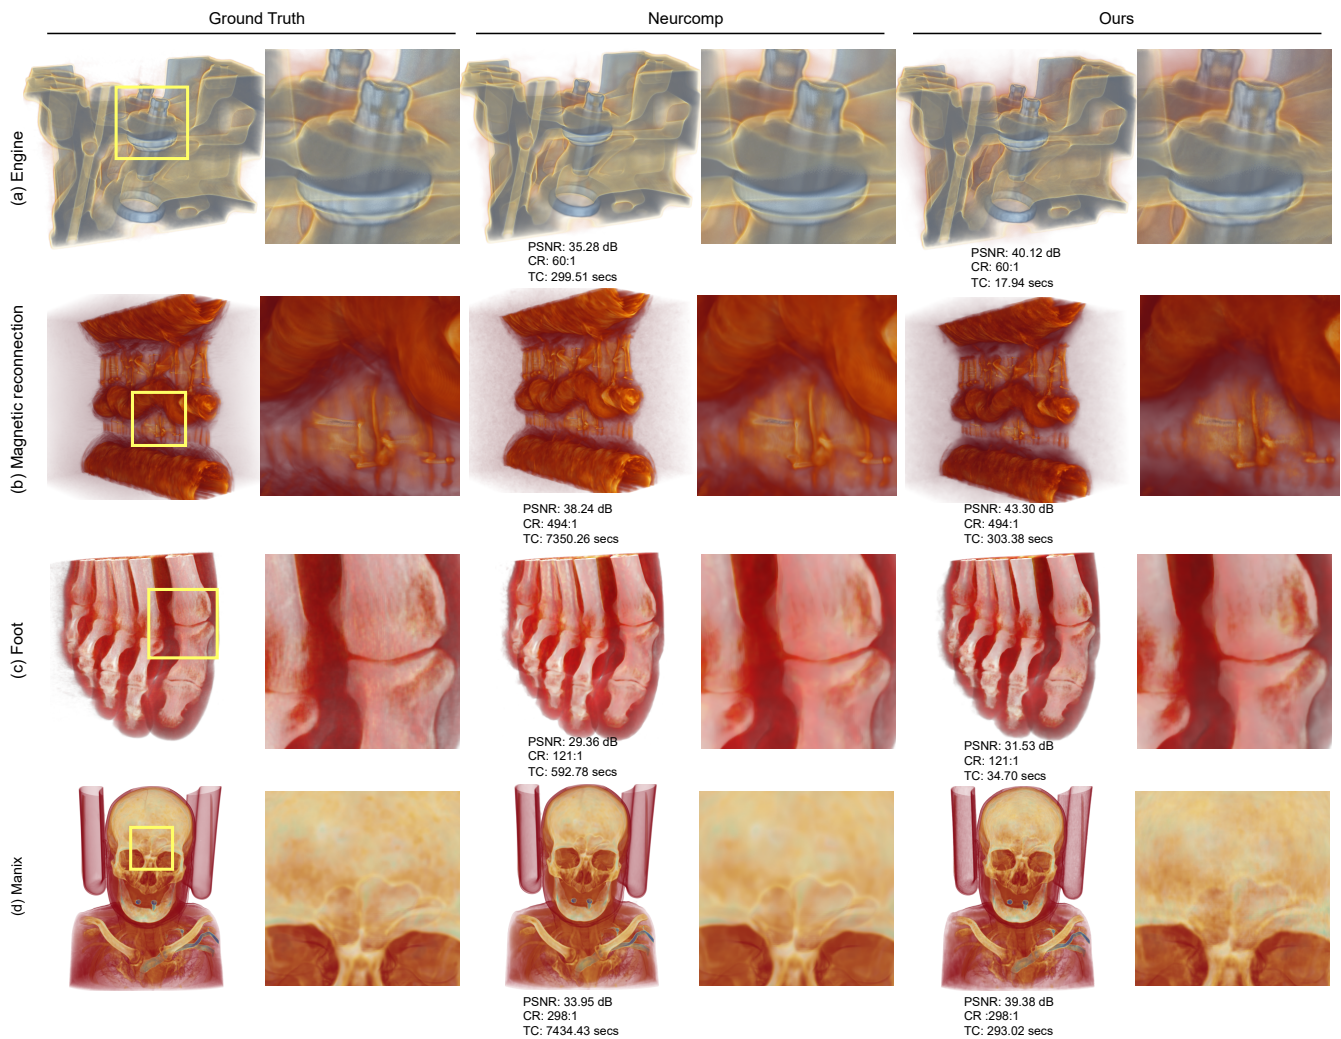


Figure 7: Comparative Compression Analysis - Our Approach vs Neurcomp. The first column presents the ground truth renderings of four distinct datasets: a) Engine, b) Magnetic Reconnection, c) Foot, and d) Manix. The second column showcases the renderings using Neurcomp, while the third column highlights the results achieved with our proposed method. This visual comparison shows the effectiveness of our approach in compressing diverse datasets.

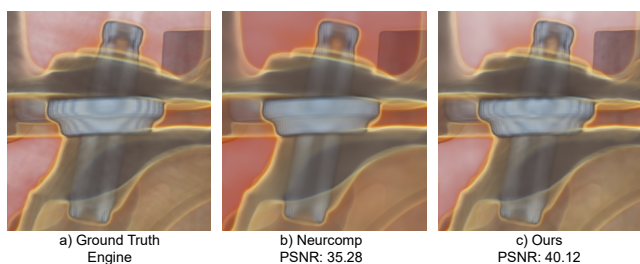


Figure 8: Visual comparison between Neurcomp and Ours for the engine dataset.

For the comparisons we make with the two baselines: Neurcomp and Tthresh, we opt to not use meta-learning. This decision is based on our observations from sections 4.3.1 and 4.3.2, which indicated that the benefits of meta-learned initialization might be minimal or insignificant for certain scenarios, particularly for inter-domain weight transfer. Since the datasets we use for evaluation in this section come from different domains, we choose to make the comparisons without using meta-learning. Thus, the metric "TC" (time to compress) solely represents the time required for the actual compression process, excluding the meta-learning stage.

Furthermore, we use different transfer functions for different datasets because the optimal transfer function for a particular dataset may vary depending on the characteristics of the data. However, to ensure a fair comparison, we use the same transfer func-

Table 2: Comparison of compression results for different datasets using *Neurcomp*, *Tthresh* and our method. (CR: Compression ratio, TC: Time to compress in seconds, GT: Ground Truth, FPS: Frames per second). Here GT FPS means the frames per second we observe from rendering the ground truth volume data.

Dataset	Neurcomp			Tthresh			Ours				GT
	PSNR	CR	TC	PSNR	CR	TC	PSNR	CR	TC	FPS	FPS
Tooth	34.52	40:1	36.58	34.13	133:1	4.04	33.92	40:1	3.89	24	53
MRI ventricles	22.19	58:1	289.05	25.10	309:1	1.48	24.69	58:1	19.84	23	47
MRHead	22.78	61:1	302.72	26.21	230:1	1.53	25.75	61:1	21.02	19	46
Aneurism	31.64	121:1	596.20	38.83	35:1	3.43	39.46	121:1	40.43	21	52
Skull	30.43	121:1	593.18	34.41	80:1	3.33	34.88	121:1	40.27	19	40
Foot	29.36	121:1	592.78	31.43	56:1	3.30	31.53	121:1	34.70	17	50
Mrt-angio	27.39	172:1	850.60	31.16	64:1	5.90	31.80	172:1	48.04	17	51
Panoramix	27.35	220:1	1095.98	35.60	162:1	8.60	35.07	220:1	107.63	22	57
CT-chest	40.78	67:1	4849.20	40.10	75:1	8.52	39.95	67:1	95.02	13	42
CTA-cardio	35.99	208:1	5220.44	38.61	221:1	19.61	38.36	208:1	205.45	16	48
Manix	33.95	298:1	7434.43	39.69	341:1	20.446	39.38	298:1	293.02	18	53
Boston teapot	35.16	42:1	622.73	40.19	396:1	1.95	39.68	42:1	25.74	18	44
Backpack	36.62	55:1	12918.76	37.15	55:1	16.13	38.77	55:1	394.41	17	47
Engine	35.28	60:1	299.51	40.70	128:1	1.34	40.12	60:1	17.94	18	51
Tacc turbulence	37.72	528:1	421.91	44.37	464:1	2.59	45.73	528:1	10.93	19	49
Csafe-heptane	37.11	134:1	1199.07	40.78	203:1	5.5	40.83	134:1	60.77	18	55
Vorticity magnitude	29.10	152:1	2079.05	29.64	85:1	38.87	30.14	152:1	92.14	16	52
Magnetic reconnection	38.24	494:1	7350.26	43.25	470:1	20.91	43.3	494:1	303.38	31	55

tion for comparing our method with ground truth, *neurcomp*, and *tthresh*. We also represent the ground truth data in single precision floating point format for all comparisons. This guarantees that our approach accurately reflects the performance compared to the baseline techniques.

4.4.1. Comparison with *Neurcomp*

In order to provide a comprehensive comparison between our method and *Neurcomp* [LJLB21], we execute our method on different datasets for varying compression ratios, as shown in Figure 7 and table 2. Both methods were run under the same conditions, utilizing a batch size of 2^{14} and the networks were trained for a total of 50 epochs, with each epoch representing a complete pass over the entire volume of data. The underlying architecture we employed for the comparison encompasses two hidden layers, each containing 64 neurons. For input encoding, we use multi-resolution hash tables with parameters $W=8$, and $T=12$, and we vary L between 4 and 12 to reflect different compression ratios. To ensure equivalent compression ratios across all datasets for both methods, we use an 8-layer network for *Neurcomp*, while adjusting the number of neurons accordingly. This adjustment ensured that both networks achieved the same level of compression.

A qualitative analysis reveals that the rendered images produced by both *Neurcomp* and our method look similar for the same compression ratios. However, our approach outperforms *Neurcomp* in terms of PSNR, indicating superior image quality. Upon closer examination, we observe that *Neurcomp* has a tendency to generate smoother surfaces in the resulting renderings, which may at times lead to an under-representation of the high-frequency noise inherent in the original, ground truth data. This phenomenon is particularly noticeable in the case of the engine dataset, as depicted in

Figure 8. Our method, conversely, exhibits better capability to preserve and represent this high-frequency noise, leading to a more accurate representation of the ground truth data,

4.4.2. Comparison with *Tthresh*

In order to assess our method in comparison to *Tthresh*, we modify the input parameters of *Tthresh*. As it accepts an error parameter (for instance, PSNR) rather than a specific compression ratio, we adjust the PSNR values for *Tthresh* roughly equivalent to those obtained from our approach and compare the compression ratios.

The quantitative outcomes of our approach versus *Tthresh* are presented in Table 2. Our analysis reveals that the performance of the two methods is data-dependent: for certain datasets, our method outperforms *Tthresh*, while for others, *Tthresh* achieves superior results. While the PSNR values obtained with *Tthresh* are comparable to those achieved by our method, a key distinction exists in the post-compression handling of data. Unlike *Tthresh*, our method does not require data decompression prior to rendering; instead, we can directly perform ray-traced direct volume rendering from the compressed representation. The observed frames per second (FPS) for rendering different volume data with our method is shown in Table 2. Our approach treats the neural representation of data as a compressed version of the original dataset, which significantly reduces memory usage. Moreover, this method permits flexible sampling without the necessity for explicit interpolation.

Figure 9 and Figure 10, present a comparison between *Tthresh* and our method for the backpack and magnetic reconnection datasets. In these examples, our approach demonstrates a slight advantage over *Tthresh*, achieving similar PSNR values but with better compression ratios.

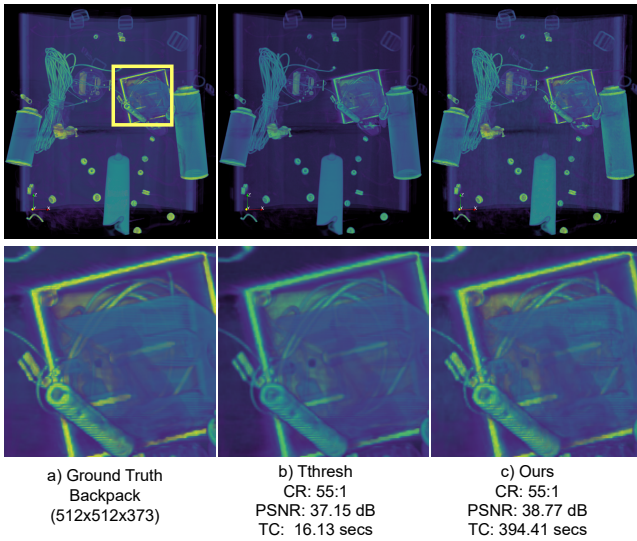


Figure 9: Comparison with Thresh for lower compression ratio for Backpack dataset

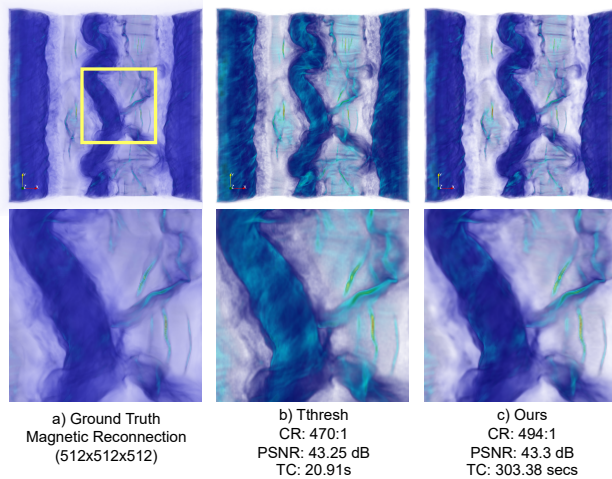


Figure 10: Comparison with Thresh for higher compression ratio for Magnetic reconnection dataset

Table 3: Results for CTA-cardio dataset: with and without gradient loss.

	Without grad. loss	With grad. loss
PSNR	38.36 dB	37.45 dB
Grad. loss	0.3517	0.3728

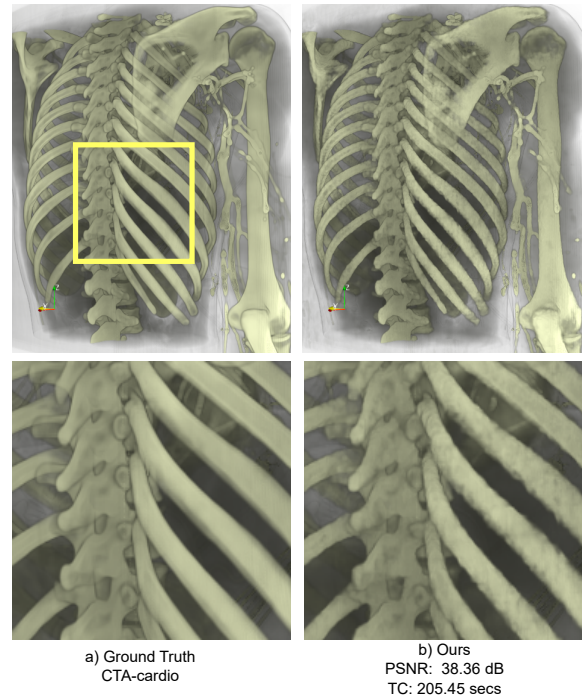


Figure 11: Ground truth compared to ours for CTA cardio dataset. Unwanted artifacts can be seen on the bone structure. Compression ratio: 208:1

4.5. Limitations

While our approach shows promise, there are some limitations to consider. For some datasets, particularly at higher compression ratios, we observed the introduction of noise and artifacts in the compressed volume. An example of this can be seen in Figure 11 for the CTA cardio dataset at a high compression ratio of 208:1, where visually undesirable surface roughness is present on the bone structure.

To address this issue, we explored the incorporation of a small percentage of gradient mean squared error (MSE) loss in our total loss function, as gradient-based volume rendering techniques play a crucial role in accurate visualization of complex structures. In particular, we implement the central difference method to compute the gradients for the predicted volume and the ground truth volume. While we were able to notice an improvement in the overall gradient MSE loss, it came at the expense of lower PSNR, as shown in Table 3. This also aligns with the findings reported by Lu *et al.* [LJLB21]. Further exploration is needed to fully understand and mitigate these trade-offs.

5. Discussion and Conclusion

In this work, we proposed a volume representation network based on coordinate-based networks and multi-resolution hash encoding. Our method achieved efficient and high-quality compression of volumetric data by leveraging spatial encoding and trainable hash tables. Our comparison to the existing neural compression method

showcased our method's superior performance in terms of both quality of representation and time to convergence.

We also introduced the concept of meta-learned initialization for volume compression, which leverages domain-specific knowledge to enhance the efficiency of neural volume representation. While we found that the benefits of meta-learned initialization are highly data-dependent, we observed consistent improvements across a broad range of datasets, suggesting that our approach can bring general benefits in a variety of contexts.

Looking ahead, there are several potential directions for future work. While our current method treats the entire volume as a single entity, future work could look into developing more sophisticated models that are capable of recognizing and separately handling different regions within the volume data. This would allow for region-specific compression that could potentially yield even better compression ratios and visual quality. We hope our work will pave the way for further research into the use of meta-learning for volume data compression and opens new opportunities for efficient volume data management.

Another interesting avenue for future work is the exploration of hybrid compression techniques that combine the strengths of neural network-based compression with traditional compression algorithms. Integrating neural network-based methods with existing compression techniques, such as lossless or lossy compression algorithms, could potentially yield even better compression ratios and preservation of fine details.

References

- [Bar] BARTZ D.: Vcm, university of tübingen, germany. URL: <http://www.volvis.org>. 6
- [BLL19] BERGER M., LI J., LEVINE J. A.: A generative model for volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 25, 4 (2019), 1636–1650. 4
- [BRLP20] BALLESTER-RIPOLL R., LINDSTROM P., PAJAROLA R.: Tthresh: Tensor compression for multidimensional visual data. *IEEE Transactions on Visualization and Computer Graphics* 26, 9 (2020), 2891–2903. doi:10.1109/TVCG.2019.2904063. 2
- [BRP15] BALLESTER-RIPOLL R., PAJAROLA R.: Lossy volume compression using tucker truncation and thresholding. *The Visual Computer* 32 (05 2015). doi:10.1007/s00371-015-1130-y. 2
- [Cen] Center for simulation of accidental fires and explosions. URL: <http://uintah.utah.edu/>. 6
- [CLCM91] CHAN K. K., LAU C. C., CHUANG K.-S., MORIOKA C. A.: Visualization and volumetric compression. In *Medical Imaging* (1991). 2
- [DP22] DEVKOTA S., PATTANAİK S.: Deep learning based super-resolution for medical volume visualization with direct volume rendering. In *Advances in Visual Computing* (Cham, 2022), Bebis G., Li B., Yao A., Liu Y., Duan Y., Lau M., Khadka R., Crisan A., Chang R., (Eds.), Springer International Publishing, pp. 103–114. 4
- [FAL17] FINN C., ABBEEL P., LEVINE S.: Model-agnostic meta-learning for fast adaptation of deep networks, 2017. arXiv:1703.03400. 2
- [FM07] FOUT N., MA K.-L.: Transform coding for hardware-accelerated volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1600–1607. doi:10.1109/TVCG.2007.70516. 2
- [Gen] General electric. URL: <http://www.volvis.org>. 6
- [GG16] GUTHE S., GOESELE M.: Variable length coding for gpu-based direct volume rendering. In *Vision, Modeling and Visualization* (2016), Hullin M., Stamminger M., Weinkauff T., (Eds.), The Eurographics Association. doi:10.2312/vmv.20161345. 2
- [GLDL14] GUO F., LI H., DAUGHTON W., LIU Y.-H.: Formation of hard power laws in the energetic particle spectra resulting from relativistic magnetic reconnection. *Phys. Rev. Lett.* 113 (oct 2014), 155005. doi:10.1103/PhysRevLett.113.155005. 6
- [Gür] GÜRVIIT Ö.: Institute for neuroradiology, frankfurt. URL: <http://www.volvis.org>. 6
- [HLY19] HONG F., LIU C., YUAN X.: Dnn-volvis: Interactive volume visualization supported by deep neural network. In *2019 IEEE Pacific Visualization Symposium (PacificVis)* (2019), pp. 282–291. doi:10.1109/PacificVis.2019.00041. 4
- [HW20] HAN J., WANG C.: Tsr-tvd: Temporal super-resolution for time-varying data analysis and visualization. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 205–215. doi:10.1109/TVCG.2019.2934255. 4
- [HW22a] HAN J., WANG C.: Coordnet: Data generation and visualization for time-varying volumes via a coordinate-based neural network. *IEEE Transactions on Visualization and Computer Graphics* (2022), 1–12. doi:10.1109/TVCG.2022.3197203. 4
- [HW22b] HAN J., WANG C.: Ssr-tvd: Spatial super-resolution for time-varying data analysis and visualization. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2022), 2445–2456. doi:10.1109/TVCG.2020.3032123. 4
- [HWG*20] HE W., WANG J., GUO H., WANG K., SHEN H., RAJ M., NASHED Y. G., PETERKA T.: Insitunet: Deep image synthesis for parameter space exploration of ensemble simulations. *IEEE Transactions on Visualization and Computer Graphics* 26, 01 (jan 2020), 23–33. doi:10.1109/TVCG.2019.2934312. 4
- [IP99] IHM I., PARK S.: Wavelet-based 3d compression scheme for interactive visualization of very large volume data. *Computer Graphics Forum* 18, 1 (1999), 3–15. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.00298>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.00298>, doi:<https://doi.org/10.1111/1467-8659.00298>. 2
- [JSM*20] JIANG C. M., SUD A., MAKADIA A., HUANG J., NIESSNER M., FUNKHOUSER T. A.: Local implicit grid representations for 3d scenes. *CoRR abs/2003.08981* (2020). URL: <https://arxiv.org/abs/2003.08981>, arXiv:2003.08981. 3, 4
- [KLM22] KIM D., LEE M., MUSETH K.: Neuralvdb: High-resolution sparse volume representation using hierarchical neural networks, 2022. arXiv:2208.04448. 4
- [Kre] KREEGER K.: Viatronix inc. URL: <http://www.volvis.org>. 6
- [Lev92] LEVOY M.: Volume rendering using the fourier projection-slice theorem. In *Proceedings of the Conference on Graphics Interface '92* (San Francisco, CA, USA, 1992), Morgan Kaufmann Publishers Inc., p. 61–69. 2
- [LGL*20] LIU L., GU J., LIN K. Z., CHUA T., THEOBALT C.: Neural sparse voxel fields. *CoRR abs/2007.11571* (2020). URL: <https://arxiv.org/abs/2007.11571>, arXiv:2007.11571. 3
- [LH96] LEVOY M., HANRAHAN P.: Light field rendering. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), SIGGRAPH '96, Association for Computing Machinery, p. 31–42. URL: <https://doi.org/10.1145/237170.237199>, doi:10.1145/237170.237199. 2
- [Lin14] LINDSTROM P.: Fixed-rate compressed floating-point arrays. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2674–2683. doi:10.1109/TVCG.2014.2346458. 2

- [LJLB21] LU Y., JIANG K., LEVINE J. A., BERGER M.: Compressive neural representations of volumetric scalar fields. *Computer Graphics Forum* 40, 3 (2021), 135–146. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14295>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14295>, doi:<https://doi.org/10.1111/cgf.14295>. 2, 4, 10, 11
- [Mal93] MALZBENDER T.: Fourier volume rendering. *ACM Trans. Graph.* 12, 3 (jul 1993), 233–250. URL: <https://doi.org/10.1145/169711.169705>, doi:10.1145/169711.169705. 2
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *CoRR abs/2201.05989* (2022). URL: <https://arxiv.org/abs/2201.05989>, arXiv:2201.05989. 4, 6
- [MLL*21] MARTEL J. N. P., LINDELL D. B., LIN C. Z., CHAN E. R., MONTEIRO M., WETZSTEIN G.: ACORN: adaptive coordinate networks for neural scene representation. *CoRR abs/2105.02788* (2021). URL: <https://arxiv.org/abs/2105.02788>, arXiv:2105.02788. 3
- [MNA*18] MICIKEVICIUS P., NARANG S., ALBEN J., DIAMOS G., ELSEN E., GARCIA D., GINSBURG B., HOUSTON M., KUCHAIEV O., VENKATESH G., WU H.: Mixed precision training, 2018. arXiv:1710.03740. 5
- [MON*19] MESCHEDER L., OECHSLE M., NIEMEYER M., NOWOZIN S., GEIGER A.: Occupancy networks: Learning 3d reconstruction in function space, 2019. arXiv:1812.03828. 2
- [MPJ*19] MICHALKIEWICZ M., PONTES J. K., JACK D., BAKTASH-MOTLAGH M., ERIKSSON A.: Implicit surface representations as layers in neural networks. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)* (2019), pp. 4742–4751. doi:10.1109/ICCV.2019.00484. 3
- [MRNK21] MÜLLER T., ROUSSELLE F., NOVÁK J., KELLER A.: Real-time neural radiance caching for path tracing. *ACM Trans. Graph.* 40, 4 (jul 2021). URL: <https://doi.org/10.1145/3450626.3459812>, doi:10.1145/3450626.3459812. 3
- [MST*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHI R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *CoRR abs/2003.08934* (2020). URL: <https://arxiv.org/abs/2003.08934>, arXiv:2003.08934. 2, 3
- [Mur93] MURAKI S.: Volume data and wavelet transforms. *IEEE Computer Graphics and Applications* 13, 4 (1993), 50–56. doi:10.1109/38.219451. 2
- [NAS18] NICHOL A., ACHIAM J., SCHULMAN J.: On first-order meta-learning algorithms. *CoRR abs/1803.02999* (2018). URL: <http://arxiv.org/abs/1803.02999>, arXiv:1803.02999. 2, 5
- [NH92] NING P., HESSELINK L.: Vector quantization for volume rendering. In *Proceedings of the 1992 Workshop on Volume Visualization* (New York, NY, USA, 1992), VVS '92, Association for Computing Machinery, p. 69–74. URL: <https://doi.org/10.1145/147130.147152>, doi:10.1145/147130.147152. 2
- [Ope] Open scivis dataset. URL: <https://klacansky.com/open-scivis-datasets/>. 6
- [OSI] OSIRIX: Dicom image library. URL: <https://www.osirix-viewer.com/resources/dicomimage-library/>. 6
- [PFS*19] PARK J. J., FLORENCE P., STRAUB J., NEWCOMBE R. A., LOVEGROVE S.: DeepSDF: Learning continuous signed distance functions for shape representation. *CoRR abs/1901.05103* (2019). URL: <http://arxiv.org/abs/1901.05103>, arXiv:1901.05103. 3
- [Phi] Philips research. URL: <https://www.philips.com/a-w/research/locations/hamburg.cs>. 6
- [PNM*20] PENG S., NIEMEYER M., MESCHEDER L., POLLEFEYS M., GEIGER A.: Convolutional occupancy networks. In *Computer Vision – ECCV 2020* (Cham, 2020), Vedaldi A., Bischof H., Brox T., Frahm J.-M., (Eds.), Springer International Publishing, pp. 523–540. 3
- [SCT*20] SITZMANN V., CHAN E. R., TUCKER R., SNAVELY N., WETZSTEIN G.: MetaSDF: Meta-learning signed distance functions. *CoRR abs/2006.09662* (2020). URL: <https://arxiv.org/abs/2006.09662>, arXiv:2006.09662. 2
- [Sie] Siemens medical solutions. URL: <http://www.volvis.org>. 6
- [SMP13] SUTER S., MAKHYNIA M., PAJAROLA R.: Tamresh – tensor approximation multiresolution hierarchy for interactive volume visualization. *Computer Graphics Forum* 32, 3pt2 (2013), 151–160. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12102>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12102>, doi:<https://doi.org/10.1111/cgf.12102>. 2
- [SPY*22] STRÜMLER Y., POSTELS J., YANG R., VAN GOOL L., TOMBARI F.: Implicit neural representations for image compression, 2022. arXiv:2112.04267. 2
- [SSC22] SUN C., SUN M., CHEN H.-T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction, 2022. arXiv:2111.11215. 3, 4
- [SW03] SCHNEIDER J., WESTERMANN R.: Compression domain volume rendering. In *IEEE Visualization, 2003. VIS 2003.* (2003), pp. 293–300. doi:10.1109/VISUAL.2003.1250385. 2
- [TCRS00] TARINI M., CIGNONI P., ROCCHINI C., SCOPIGNO R.: Real time, accurate, multi-featured rendering of bump mapped surfaces. *Computer Graphics Forum* 19, 3 (2000), 119–130. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.00404>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.00404>, doi:<https://doi.org/10.1111/1467-8659.00404>. 2
- [Ter] Terarecon inc, merl, brigham and women’s hospital. URL: <http://www.volvis.org>. 6
- [TET*22] TAKIKAWA T., EVANS A., TREMBLAY J., MÜLLER T., MCGUIRE M., JACOBSON A., FIDLER S.: Variable bitrate neural fields, 2022. arXiv:2206.07707. 2
- [THM*03] TESCHNER M., HEIDELBERGER B., MÜLLER M., POMERANETS D., GROSS M.: Optimized spatial hashing for collision detection of deformable objects. *VMV’03: Proceedings of the Vision, Modeling, Visualization 3* (12 2003). 5
- [TL93] TOTSUKA T., LEVOY M.: Frequency domain volume rendering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1993), SIGGRAPH '93, Association for Computing Machinery, p. 271–278. URL: <https://doi.org/10.1145/166117.166152>, doi:10.1145/166117.166152. 2
- [TLY*21] TAKIKAWA T., LITALIEN J., YIN K., KREIS K., LOOP C., NOWROUZSAHRAI D., JACOBSON A., MCGUIRE M., FIDLER S.: Neural geometric level of detail: Real-time rendering with implicit 3d shapes, 2021. arXiv:2101.10994. 3
- [TMW*20] TANCIK M., MILDENHALL B., WANG T., SCHMIDT D., SRINIVASAN P. P., BARRON J. T., NG R.: Learned initializations for optimizing coordinate-based neural representations. *CoRR abs/2012.02189* (2020). URL: <https://arxiv.org/abs/2012.02189>, arXiv:2012.02189. 2
- [TSM*20] TANCIK M., SRINIVASAN P. P., MILDENHALL B., FRIDOVICH-KEIL S., RAGHAVAN N., SINGHAL U., RAMAMOORTHI R., BARRON J. T., NG R.: Fourier features let networks learn high frequency functions in low dimensional domains. *CoRR abs/2006.10739* (2020). URL: <https://arxiv.org/abs/2006.10739>, arXiv:2006.10739. 3

- [VGK96] VAN GELDER A., KIM K.: Direct volume rendering with shading via three-dimensional textures. In *Proceedings of 1996 Symposium on Volume Visualization* (1996), pp. 23–30. doi:10.1109/SVV.1996.558039. 2
- [Vor] Volume rendering. URL: <https://www.cs.purdue.edu/homes/cs530/projects/project3.html>. 6
- [WCTW21] WEISS S., CHU M., THUEREY N., WESTERMANN R.: Volumetric isosurface rendering with deep learning-based super-resolution. *IEEE Transactions on Visualization and Computer Graphics* 27, 6 (2021), 3064–3078. doi:10.1109/TVCG.2019.2956697. 4
- [WHW21] WEISS S., HERMÜLLER P., WESTERMANN R.: Fast neural representations for direct volume rendering. *CoRR abs/2112.01579* (2021). URL: <https://arxiv.org/abs/2112.01579>, arXiv:2112.01579. 2, 4
- [Wik] WIKI S.: Sampledata — slicer wiki,. [Online; accessed 5-June-2023]. URL: <https://www.slicer.org/w/index.php?title=SampleData&oldid=62556>. 6
- [YFKT*21] YU A., FRIDOVICH-KEIL S., TANCIK M., CHEN Q., RECHT B., KANAZAWA A.: Plenoxels: Radiance fields without neural networks, 2021. arXiv:2112.05131. 3, 4
- [YL95] YEO B.-L., LIU B.: Volume rendering of dct-based compressed 3d scalar data. *IEEE Transactions on Visualization and Computer Graphics* 1, 1 (1995), 29–43. doi:10.1109/2945.468390. 2