

# GaussianImage: 1000 FPS Image Representation and Compression by 2D Gaussian Splatting

Xinjie Zhang<sup>1,3\*†</sup>, Xingtong Ge<sup>2,3\*†</sup>, Tongda Xu<sup>4</sup>, Dailan He<sup>5</sup>, Yan Wang<sup>4</sup>, Hongwei Qin<sup>3</sup>, Guo Lu<sup>6</sup>, Jing Geng<sup>2†</sup>, and Jun Zhang<sup>1†</sup>

<sup>1</sup> The Hong Kong University of Science and Technology

<sup>2</sup> Beijing Institute of Technology    <sup>3</sup> SenseTime Research

<sup>4</sup> Institute for AI Industry Research (AIR), Tsinghua University

<sup>5</sup> The Chinese University of Hong Kong    <sup>6</sup> Shanghai Jiaotong University

xzhangga@connect.ust.hk, xingtong.ge@gmail.com, x.tongda@nyu.edu  
hedailan@link.cuhk.edu.hk, wangyan202199@163.com, qinhongwei@sensetime.com

luguoc2014@sjtu.edu.cn, janegeng@bit.edu.cn, eejzhang@ust.hk

**Abstract.** Implicit neural representations (INRs) recently achieved great success in image representation and compression, offering high visual quality and fast rendering speeds with 10-1000 FPS, assuming sufficient GPU resources are available. However, this requirement often hinders their use on low-end devices with limited memory. In response, we propose a groundbreaking paradigm of image representation and compression by 2D Gaussian Splatting, named GaussianImage. We first introduce 2D Gaussian to represent the image, where each Gaussian has 8 parameters including position, covariance and color. Subsequently, we unveil a novel rendering algorithm based on accumulated summation. Remarkably, our method with a minimum of  $3\times$  lower GPU memory usage and  $5\times$  faster fitting time not only rivals INRs (e.g., WIRE, I-NGP) in representation performance, but also delivers a faster rendering speed of 1500-2000 FPS regardless of parameter size. Furthermore, we integrate existing vector quantization technique to build an image codec. Experimental results demonstrate that our codec attains rate-distortion performance comparable to compression-based INRs such as COIN and COIN++, while facilitating decoding speeds of approximately 2000 FPS. Additionally, preliminary proof of concept shows that our codec surpasses COIN and COIN++ in performance when using partial bits-back coding. Code is available at <https://github.com/Xinjie-Q/GaussianImage>.


**Keywords:** 2D Gaussian Splatting · Image Representation · Image Compression

## 1 Introduction

Image representation is a fundamental tasks in signal processing and computer vision. Traditional image representation methods, including grid graphics, wavelet

---

\* Equal Contribution. <sup>†</sup> This work was done when Xinjie Zhang and Xingtong Ge interned at SenseTime Research. <sup>†</sup> Corresponding Authors.



**Fig. 1:** Image representation (left) and compression (right) results with different decoding time on the Kodak and DIV2K dataset, respectively. The radius of each point indicates the parameter size (left) or bits per pixel (right). Our method enjoys the fastest decoding speed regardless of parameter size or bpp.

transform [26], and discrete cosine transform [3], have been extensively applied across a broad spectrum of applications, from image compression to vision task analysis. However, these techniques encounter significant obstacles when processing large-scale datasets and striving for highly efficient storage solutions.

The advent of implicit neural representations (INRs) [48, 50] marks a significant paradigm shift in image representation techniques. Typically, INRs employ a compact neural network to derive an implicit continuous mapping from input coordinates to the corresponding output values. This allows INRs to capture and retain image details with greater efficiency, which provides considerable benefits across various applications, including image compression [20, 21, 24, 51], deblurring [44, 59, 61], and super-resolution [14, 38, 43]. However, most state-of-the-art INR methods [22, 45, 47, 48, 53] rely on a large high-dimensional multi-layer perceptron (MLP) network to accurately represent high-resolution images. This dependency leads to prolonged training times, increased GPU memory requirements, and slow decoding speed. While recent innovations [17, 39, 42, 52] have introduced multi-resolution feature grids coupled with a compact MLP to accelerate training and inference, they still require enough GPU memory to support their fast training and inference, which is difficult to meet when resources are limited. Consequently, these challenges substantially hinder the practical deployment of INRs in real-world scenarios.

In light of these challenges, our research aims to develop an advanced image representation technique that enables efficient training, friendly GPU memory usage, and fast decoding. To achieve this goal, we resort to Gaussian Splatting (GS) [30] that is recently developed for 3D scene reconstruction. By leveraging explicit 3D Gaussian representations and differentiable tile-based rasterization, 3D GS not only enjoys high visual quality with competitive training times, but also achieves real-time rendering capabilities.

Nevertheless, it is non-trivial to directly adapt 3D GS for efficient single image representation. **Firstly**, considering that existing 3D GS methods [11, 30]

depend on varying camera transformation matrices to render images from different perspectives, a straightforward adaptation for single image is fixing the camera transformation matrix to render an image from a single viewing angle. Unfortunately, each 3D Gaussian usually includes 59 learnable parameters [30] and thousands of 3D Gaussians are required for representing a single image. This naive approach substantially increases the storage and communication demands. As can be inferred from Table 1, the storage footprint for a single image with tens of kilobytes can escalate to dozens of megabytes, which makes rendering difficult on low-end devices with limited memory. **Secondly**, the rasterization algorithm [30] in 3D GS, designed for  $\alpha$ -blending approximation, necessitates pre-sorted Gaussians based on depth information derived from camera parameters. This poses a challenge for single images because detailed camera parameters are often not known in natural individual image, while non-natural images, including screenshots and AI-generated content, are not captured by cameras. Without accurate depth information, the Gaussian sorting might be impaired, diminishing the final fitting performance. Moreover, the current rasterization process skips the remaining Gaussians once the accumulated opacity surpasses the given threshold, which results in underutilization of Gaussian data, thereby requiring more Gaussians for high-quality rendering.

To address these issues, we propose a new paradigm of image representation and compression, namely GaussianImage, using 2D Gaussian Splatting. **Firstly**, we adopt 2D Gaussians in lieu of 3D for a compact and expressive representation. Each 2D Gaussian is defined by 4 attributes (9 parameters in total): position, anisotropic covariance, color coefficients, and opacity. This modification results in a  $6.5\times$  compression over 3D Gaussians with equivalent Gaussian points, significantly mitigating storage demands of Gaussian representation. **Subsequently**, we advocate a unique rasterization algorithm that replaces depth-based Gaussian sorting and  $\alpha$ -blending with an accumulated summation process. This novel approach directly computes each pixel’s color from the weighted sum of 2D Gaussians, which not only fully utilizes the information of all Gaussian points covering the current pixel to improve fitting performance, but also avoids the tedious calculation of accumulated transparency to accelerate training and inference speed. More important, this summation mechanism allows us to merge color coefficients and opacity into a singular set of weighted color coefficients, reducing parameter count to 8 and further improving the compression ratio to  $7.375\times$ . **Finally**, we transfer our 2D Gaussian representation into a practical image codec. Framing image compression as a Gaussian attribute compression task, we employ a two-step compression strategy: attribute quantization-aware fine-tuning and encoding. By applying 16-bit float quantization, 6-bit integer quantization [10], and residual vector quantization (RVQ) [66] to positions, covariance parameters, and weighted color coefficients, respectively, we successfully develop the first image codec based on 2D Gaussian Splatting. As a preliminary proof of concept, the partial bits-back coding [46, 54] is optionally used to further improve the compression performance of our codec. Overall, our contributions are threefold:

- We present a pioneering paradigm of image representation and compression by 2D Gaussian Splatting. With compact 2D Gaussian representation and a novel accumulated blending-based rasterization method, our approach achieves high representation performance with short training duration, minimal GPU memory overhead and remarkably, 2000 FPS rendering speed.
- We develop a low-complexity neural image codec using vector quantization. Furthermore, a partial bits-back coding technique is optionally used to reduce the bitrate.
- Experimental results show that when compared with existing INR methods, our approach achieves a remarkable training and inference acceleration with less GPU memory usage while maintaining similar visual quality. When used as an efficient image codec, our approach offers competitive compression performance comparable to COIN [21] and COIN++ [20]. Comprehensive ablations and analyses demonstrate the effectiveness of each proposed component.

## 2 Related Works

### 2.1 Implicit Neural Representation

Recently, implicit neural representation has gained increasing attention for its wide-ranging potential applications, such as 3D scene rendering [7, 8, 40, 60], image [17, 42, 47, 48] and video [12, 13, 35, 67] representations. We roughly classified existing image INRs into two categories: (i) MLP-based INRs [22, 45, 47, 48, 53] take position encoding of spatial coordinates as input of an MLP network to learn the RGB values of images, while they only rely on the neural network to encode all the image information, resulting in inefficient training and inference especially for high-resolution image. (ii) Feature grid-based INRs [17, 39, 42, 52] adopt a large-scale multi-resolution grid, such as quadtree and hash table, to provide prior information for a compact MLP. This reduces the learning difficulty of MLP to a certain extent and accelerates the training process, making INRs more practical. Unfortunately, they still consume large GPU memory, which is difficult to accommodate on low-end devices. Instead of following existing INR methods, we aim to propose a brand-new image representation paradigm based on 2D Gaussian Splatting, which enables us to enjoy swifter training, faster rendering, and less GPU resource consumption.

### 2.2 Gaussian Splatting

Gaussian Splatting [30] has recently gained tremendous traction as a promising paradigm to 3D view synthesis. With explicit 3D Gaussian representations and differentiable tile-based rasterization, GS not only brings unprecedented control and editability but also facilitates high-quality and real-time rendering in 3D scene reconstruction. This versatility has opened up new avenues in various domains, including simultaneous localization and mapping (SLAM) [28, 29, 62],



**Fig. 2:** Our proposed GaussianImage framework. 2D Gaussians are first formatted and then rasterized to generate the output image. The rasterizer uses our proposed accumulated blending for efficient 2D image representation.

dynamic scene modeling [37, 57, 64], AI-generated content [15, 18, 69], and autonomous driving [63, 68]. Despite its great success in 3D scenarios, the application of GS to single image representation remains unexplored. Our work pioneers the adaptation of GS for 2D image representation, leveraging the strengths of GS in highly parallelized workflow and real-time rendering to outperform INR-based methods in terms of training efficiency and decoding speed.

### 2.3 Image Compression

Traditional image compression techniques, such as JPEG [55], JPEG2000 [49] and BPG [9], follow a transformation, quantization, and entropy coding procedure to achieve good compression efficiency with decent decompression speed. Recently, learning-based image compression methods based on variational auto-encoder (VAE) have re-imagined this pipeline, integrating complex nonlinear transformations [4, 19, 25, 36] and advanced entropy models [5, 6, 31, 41]. Despite these methods surpassing traditional codecs in rate-distortion (RD) performance, their extremely high computational complexity and very slow decoding speed severely limit their practical deployment. To tackle the computational inefficiency of existing art, some works have explored INR-based compression methods [20, 21, 24, 33, 34, 51]. However, as image resolutions climb, their decoding speeds falter dramatically, challenging their real-world applicability. In this paper, our approach diverges from VAE and INR paradigms, utilizing 2D Gaussian Splatting to forge a neural image codec with unprecedented decoding efficiency. This marks an important milestone for neural image codecs.

## 3 Method

Fig. 2 delineates the overall processing pipeline of our GaussianImage. Our approach begins by forming 2D Gaussians on an image plane, a process mainly calculating the 2D covariance matrix  $\Sigma$ . Afterwards, we employ an accumulated blending mechanism to compute the value for each pixel. In what follows, we begin with the formation of a 2D Gaussian in Section 3.1. Next, we describe how to adapt the rasterization process of 3D GS to align the unique characteristics of 2D image representation and upgrade 2D Gaussian with less parameters in Section 3.2. Then, we present a two-step compression pipeline to convert our GaussianImage into a neural image codec in Section 3.3. Finally, we state the training process on the image representation and compression tasks in Section 3.4.

### 3.1 2D Gaussian Formation

In 3D Gaussian Splatting, each 3D Gaussian is initially mapped into a 2D plane through viewing and projection transformation. Then a differentiable rasterizer is used to render the current view image from these projected Gaussians. Since our application is no longer oriented to 3D scenes, but to 2D image representation, we discard many bloated operations and redundant parameters in 3D GS, such as project transformation, spherical harmonics, etc.

In our framework, the image representation unit is a 2D Gaussian. The basic 2D Gaussian is described by its position  $\boldsymbol{\mu} \in \mathbb{R}^2$ , 2D covariance matrix  $\boldsymbol{\Sigma} \in \mathbb{R}^{2 \times 2}$ , color coefficients  $\boldsymbol{c} \in \mathbb{R}^3$  and opacity  $o \in \mathbb{R}$ . Note that covariance matrix  $\boldsymbol{\Sigma}$  of a Gaussian distribution requires positive semi-definite. Typically, it is difficult to constrain the learnable parameters using gradient descent to generate such valid matrices. To avoid producing invalid matrix during training, we choose to optimize the factorized form of the covariance matrix. Here, we present two decomposition ways to cover all the information of the original covariance matrix. One intuitive decomposition is the Cholesky factorization [27], which breaks down  $\boldsymbol{\Sigma}$  into the product of a lower triangular matrix  $\boldsymbol{L} \in \mathbb{R}^{2 \times 2}$  and its conjugate transpose  $\boldsymbol{L}^T$ :

$$\boldsymbol{\Sigma} = \boldsymbol{L}\boldsymbol{L}^T. \quad (1)$$

For the sake of writing, we use a Choleksy vector  $\boldsymbol{l} = \{l_1, l_2, l_3\}$  to represent the lower triangular elements in matrix  $\boldsymbol{L}$ . When compared with 3D Gaussian having 59 learnable parameters, our 2D Gaussian only require 9 parameters, making it more lightweight and suitable for image representation.

Another decomposition follows 3D GS [30] to factorize the covariance matrix into a rotation matrix  $\boldsymbol{R} \in \mathbb{R}^{2 \times 2}$  and scaling matrix  $\boldsymbol{S} \in \mathbb{R}^{2 \times 2}$ :

$$\boldsymbol{\Sigma} = (\boldsymbol{R}\boldsymbol{S})(\boldsymbol{R}\boldsymbol{S})^T, \quad (2)$$

where the rotation matrix  $\boldsymbol{R}$  and the scaling matrix  $\boldsymbol{S}$  are expressed as

$$\boldsymbol{R} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}, \quad \boldsymbol{S} = \begin{bmatrix} s_1 & 0 \\ 0 & s_2 \end{bmatrix}. \quad (3)$$

Here,  $\theta$  represents the rotation angle.  $s_1$  and  $s_2$  are scaling factors in different eigenvector directions. While the decomposition of the covariance matrix is not unique, they have equivalent capabilities to represent the image. However, the robustness to compression of different decomposition forms is inconsistent, which is explained in detail in the appendix. Therefore, we need to carefully choose the decomposition form of the covariance matrix when facing different image tasks.

### 3.2 Accumulated Blending-based Rasterization

During the rasterization phase, 3D GS first forms a sorted list of Gaussians  $\mathcal{N}$  based on the projected depth information. Then the  $\alpha$ -blending is adopted to

render pixel  $i$ :

$$\mathbf{C}_i = \sum_{n \in \mathcal{N}} \mathbf{c}_n \cdot \alpha_n \cdot T_n, \quad T_n = \prod_{m=1}^{n-1} (1 - \alpha_m), \quad (4)$$

where  $T_n$  denotes the accumulated transparency. The  $\alpha_n$  is computed with projected 2D covariance  $\Sigma$  and opacity  $o_n$ :

$$\alpha_n = o_n \cdot \exp(-\sigma_n), \quad \sigma_n = \frac{1}{2} \mathbf{d}_n^T \Sigma^{-1} \mathbf{d}_n, \quad (5)$$

where  $\mathbf{d} \in \mathbb{R}^2$  is the displacement between the pixel center and the projected 2D Gaussian center.

Since the acquisition of depth information involves viewing transformation, it requires us to know the intrinsic and extrinsic parameters of the camera in advance. However, it is difficult for natural individual image to access the detailed camera parameters, while non-natural images, such as screenshots and AI-generated content, are not captured by the camera. In this case, retaining the  $\alpha$ -blending of the 3D GS without depth cues would result in arbitrary blending sequences, compromising the rendering quality. Moreover, 3D GS only maintains Gaussians with a 99% confidence interval in order to solve the problem of numerical instability in computing the projected 2D covariance, but this makes only part of Gaussians covering pixel  $i$  contribute to the rendering of pixel  $i$ , leading to inferior fitting performance.

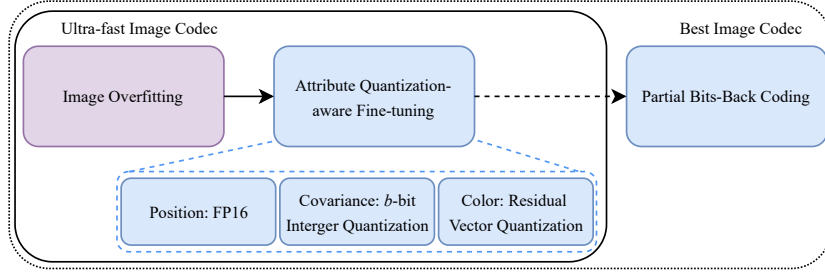
To overcome these limitations, we propose an accumulated summation mechanism to unleash the potential of our 2D Gaussian representation. Since there is no viewpoint influence when rendering an image, the rays we observe from each element are determined, and so as all the  $\alpha$  values. Therefore, we merge the  $T_n$  part in Equation 4 into the  $o_n$  term, and simplify the computation consuming  $\alpha$ -blending to a weighted sum:

$$\mathbf{C}_i = \sum_{n \in \mathcal{N}} \mathbf{c}_n \cdot \alpha_n = \sum_{n \in \mathcal{N}} \mathbf{c}_n \cdot o_n \cdot \exp(-\sigma_n). \quad (6)$$

This removes the necessity of Gaussian sequence order, so that we can remove the sorting from rasterization.

This novel rasterization algorithm brings multiple benefits. First, our accumulated blending process is insensitive to the order of Gaussian points. This property allows us to avoid the impact of the random order of Gaussian points on rendering, achieving robustness to any order of Gaussian points. Second, when compared with Equation 4, our rendering skips the tedious sequential calculation of accumulated transparency  $T_n$ , improving our training efficiency and rendering speed. Third, since the color coefficients  $\mathbf{c}_n$  and the opacity  $o_n$  are learnable parameters, they can be merged to further simplify Equation 6:

$$\mathbf{C}_i = \sum_{n \in \mathcal{N}} \mathbf{c}'_n \cdot \exp(-\sigma_n), \quad (7)$$



**Fig. 3:** Compression pipeline of our proposed GaussianImage. After overfitting image, we apply attribute quantization-aware fine-tuning to build an ultra-fast image codec. Partial bits-back coding is used to achieve the best compression performance.

where the weighted color coefficients  $\mathbf{c}'_n \in \mathbb{R}^3$  is no longer limited in the range of  $[0, 1]$ . In this way, instead of the basic 2D Gaussian that requires 4 attributes in Section 3.1, our upgraded 2D Gaussian is described by only 3 attributes (i.e., position, covariance, and weighted color coefficients) with a total of 8 parameters. This further improves the compression ratio to  $7.375\times$  when compared with 3D Gaussian under equivalent Gaussian points.

### 3.3 Compression Pipeline

After overfitting the image, we propose a compression pipeline for image compression with GaussianImage. As shown in Fig. 3, our standard compression pipeline is composed of two steps: image overfitting and attribute quantization-aware fine-tuning. To achieve the best compression performance, partial bits-back coding [46, 54] is an optional strategy. Herein, we elucidate the compression process using our GaussianImage based on Cholesky factorization as an example. **Attribute Quantization-aware Fine-tuning.** Given a set of 2D Gaussian points fit on an image, we apply distinct quantization strategies to various attributes. Since the Gaussian location is sensitive to quantization, we adopt 16-bit float precision for position parameters to preserve reconstruction fidelity. For Cholesky vector  $\mathbf{l}_n$  in the  $n$ -th Gaussian, we incorporate a  $b$ -bit asymmetric quantization technique [10], where both the scaling factor  $\gamma_i$  and the offset factor  $\beta_i$  are learned during fine-tuning:

$$\hat{l}_i^n = \left\lfloor \text{clamp} \left( \frac{l_i^n - \beta_i}{\gamma_i}, 0, 2^b - 1 \right) \right\rfloor, \quad \bar{l}_i^n = \hat{l}_i^n \times \gamma_i + \beta_i, \quad (8)$$

where  $i \in \{0, 1, 2\}$ . Note that we share the same scaling and offset factors at all Gaussians in order to reduce metadata overhead. After fine-tuning, the covariance parameters are encoded with  $b$ -bit precision, while the scaling and offset values required for re-scaling are stored in 32-bit float precision.

As for weighted color coefficients, a codebook enables representative color attribute encoding via vector quantization (VQ) [23]. While naively applying vector quantization leads to inferior rendering quality, we employ residual vector



quantization (RVQ) [66] that cascades  $M$  stages of VQ with codebook size  $B$  to mitigate performance degradation:

$$\begin{aligned}\hat{\mathbf{c}}_n^m &= \sum_{k=1}^m \mathcal{C}^k[i^k], \quad m \in \{1, \dots, M\}, \\ i_n^m &= \arg \min_m \|\mathcal{C}^m[k] - (\mathbf{c}'_n - \hat{\mathbf{c}}_n^{m-1})\|_2^2, \quad \hat{\mathbf{c}}_n^0 = 0,\end{aligned}\tag{9}$$

where  $\hat{\mathbf{c}}_n^m$  denotes the output color vector after  $m$  quantization stages,  $\mathcal{C}^m \in \mathbb{R}^{B \times 3}$  represents the codebook at the stage  $m$ ,  $i^m \in \{0, \dots, B-1\}^N$  is the codebook indices at the stage  $m$ , and  $\mathcal{C}[i] \in \mathbb{R}^3$  is the vector at index  $i$  of the codebook  $\mathcal{C}$ . To train the codebooks, we apply the commitment loss  $\mathcal{L}_c$  as follows:

$$\mathcal{L}_c = \frac{1}{N \times B} \sum_{k=1}^M \sum_{n=1}^N \|\text{sg}[\mathbf{c}'_n - \hat{\mathbf{c}}_n^{k-1}] - \mathcal{C}^k[i_n^k]\|_2^2,\tag{10}$$

where  $N$  is the number of Gaussians and  $\text{sg}[\cdot]$  is the stop-gradient operation.

**Partial Bits-Back Coding.** As we have not adopted any auto-regressive context [41] to encode 2D Gaussian parameters, any permutation of 2D Gaussian points can be seen as an equivariant graph without edge. Therefore, we can adopt bits-back coding [54] for equivariant graph described by [32] to save bitrate. More specifically, [32] show that an unordered set with  $N$  elements has  $N!$  equivariant, and bits-back coding can save a bitrate of

$$\log N! - \log N,\tag{11}$$

compared with directly store those unordered elements.

However, the vanilla bits-back coding requires initial bits [54] of  $\log N!$ , which means that it can only work on a dataset, not on a single image. To tackle this challenge, [46] introduces a partial bits-back coding strategy that segments the image data, applying vanilla entropy coding to a fraction of the image as the initial bit allocation, with the remainder encoded via bits-back coding.

In our case, we reuse the idea of [46]. Specifically, we encode the initial  $K$  Gaussians by vanilla entropy coding, and the subsequent  $N - K$  Gaussians by bits-back coding. This segmented approach is applicable to single image compression, contingent upon the bitrate of the initial  $K$  Gaussian exceeding the initial bits  $\log(N - K)!$ . Let  $R_k$  denotes the bitrate of  $k$ -th Gaussian, the final bitrate saving can be formalized as:

$$\log(N - K^*)! - \log(N - K^*),\tag{12}$$

$$\text{where } K^* = \inf K, \text{ s.t. } \sum_{k=1}^K R_k - \log(N - K^*)! \geq 0.\tag{13}$$

Despite its theoretical efficacy, bits-back coding may not align with the objective of developing an ultra-fast codec due to its slow processing latency [32]. Consequently, we leave this part as a preliminary proof of concept on the best rate-distortion performance our codec can achieve, instead of a final result of our codec can achieve with 2000 FPS.

**Table 1:** Quantitative comparison with various baselines in PSNR, MS-SSIM, training time, rendering speed, GPU memory usage and parameter size.

(a) Kodak dataset						
Methods	PSNR↑	MS-SSIM↑	Training Time(s)↓	FPS↑	GPU Mem(MiB)↓	Params(K)↓
WIRE [47]	41.47	0.9939	14338.78	11.14	2619	136.74
SIREN [48]	40.83	0.9960	6582.36	29.15	1809	272.70
I-NGP [42]	43.88	0.9976	490.61	1296.82	1525	300.09
NeuRBF [17]	43.78	0.9964	991.83	663.01	2091	337.29
3D GS [30]	43.69	0.9991	339.78	859.44	557	3540.00
Ours	44.08	0.9985	106.59	2092.17	419	560.00

(b) DIV2K dataset						
Methods	PSNR↑	MS-SSIM↑	Training Time(s)↓	FPS↑	GPU Mem(MiB)↓	Params(K)↓
WIRE [47]	35.64	0.9511	25684.23	14.25	2619	136.74
SIREN [48]	39.08	0.9958	15125.11	11.07	2053	483.60
I-NGP [42]	37.06	0.9950	676.29	1331.54	1906	525.40
NeuRBF [17]	38.60	0.9913	1715.44	706.40	2893	383.65
3D GS [30]	39.36	0.9979	481.27	640.33	709	4130.00
Ours	39.53	0.9975	120.76	1737.60	439	560.00

### 3.4 Training

For image representation, our objective is to minimize the distortion between the original image  $x$  and reconstructed image  $\hat{x}$ . To this end, we employ the L2 loss function to optimize the Gaussian parameters. It is worth noting that previous GS method [30] introduces adaptive density control to split and clone Gaussians when optimizing 3D scenes. Since there exists many empty areas in the 3D space, they need to consider avoiding populating these areas. By contrast, there is no so-called empty area in the 2D image space. Therefore, we discard adaptive density control, which greatly simplifies the optimization process of 2D image representation.

As for image compression task, the overall loss  $\mathcal{L}$  consists of the reconstruction loss  $\mathcal{L}_{rec}$  and the commitment loss  $\mathcal{L}_c$ :

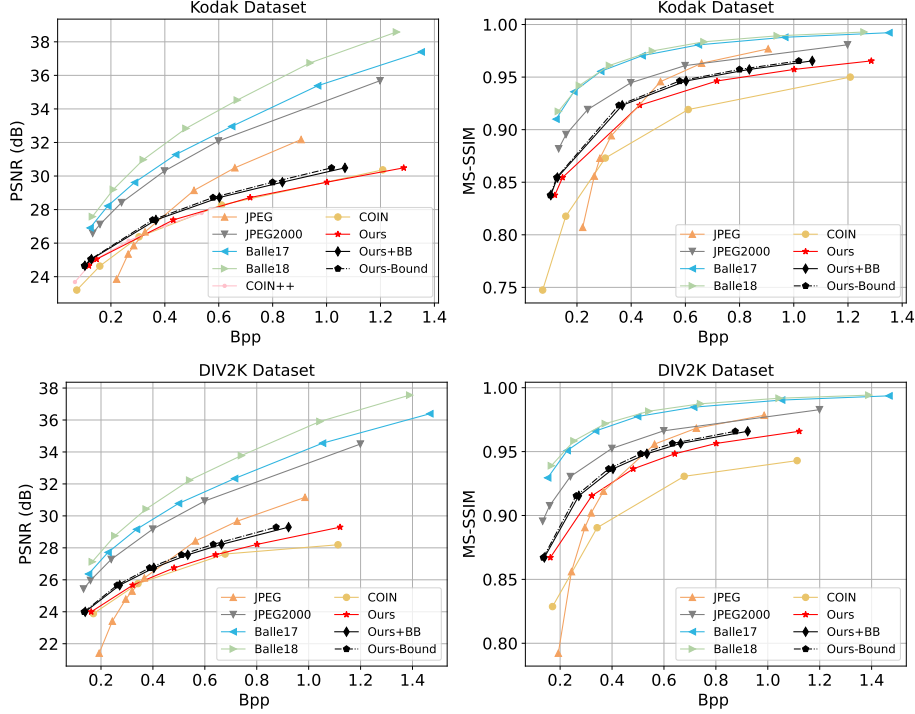
$$\mathcal{L} = \mathcal{L}_{rec} + \lambda \mathcal{L}_c, \quad (14)$$

where  $\lambda$  serves as a hyper-parameter, balancing the weight of each loss component. The color codebooks are initialized using the K-means algorithm, providing a robust starting point for subsequent optimization. During fine-tuning, we adopt the exponential moving average mode to update the codebook.

## 4 Experiments

### 4.1 Experimental Setup

**Dataset.** Our evaluation in image representation and compression is conducted on two popular datasets. We use the Kodak dataset [1], which consists of 24 images with a resolution of  $768 \times 512$ , and the DIV2K validation set [2] with  $2 \times$  bicubic downscaling, featuring 100 images with dimensions varying from  $408 \times 1020$  to  $1020 \times 1020$ .




**Fig. 4:** Rate-distortion curves of our approach and different baselines on Kodak and DIV2K datasets in PSNR and MS-SSIM. BB denotes partial bits-back coding. Bound denotes the theoretical rate of our codec.

**Evaluation Metrics.** To assess image quality, we employ two esteemed metrics: PSNR and MS-SSIM [56], which measure the distortion between reconstructed images and their originals. The bitrate for image compression is quantified in bits per pixel (bpp).

**Implementation Details.** Our GaussianImage, developed on top of gsplat [65], incorporates custom CUDA kernels for rasterization based on accumulated blending. We represent the covariance of 2D Gaussians using Cholesky factorization unless otherwise stated. The Gaussian parameters are optimized over 50000 steps using the Adan optimizer [58], starting with an initial learning rate of  $1e^{-3}$ , halved every 20000 steps. During attribute quantization-aware fine-tuning, the quantization precision  $b$  of covariance parameters is set to 6 bits, with the RVQ color vectors' codebook size  $B$  and the number of quantization stages  $M$  fixed at 8 and 2, respectively. The iterations of K-means algorithm are set to 5. Experiments are performed using NVIDIA V100 GPUs and PyTorch, with further details available in the supplementary material.

**Benchmarks.** For image representation comparisons, GaussianImage is benchmarked against competitive INR methods like SIREN [48], WIRE [47], I-NGP [42], and NeuRBF [17]. As for image compression, baselines span traditional codecs (JPEG [55], JPEG2000 [49]), VAE-based codecs (Ballé17 [5], Ballé18 [6]),

**Fig. 5:** Subjective comparison of various codecs on Kodak at low Bpp.**Table 2:** Computational complexity of traditional and learning-based image codecs on DIV2K Dataset at low and high Bpp.

Methods	Bpp↓	PSNR↑	MS-SSIM↑	Encoding FPS↑	Decoding FPS↑
JPEG [55]	0.3197/0.5638	25.2920/28.4299	0.9020/0.9559	608.61/557.35	614.68/545.59
JPEG2000 [49]	0.2394/0.5993	27.2792/30.9294	0.9305/0.9663	3.46/3.40	4.32/3.93
Ballé17 [5]	0.2271/0.4987	27.7168/30.7759	0.9508/0.9775	21.23/16.53	18.83/17.87
Ballé18 [6]	0.2533/0.5415	28.7548/32.2351	0.9584/0.9816	16.53/13.56	15.87/15.20
COIN [21]	0.3419/0.6780	25.8012/27.6126	0.8905/0.9306	$5.30e^{-4}/3.51e^{-4}$	166.31/93.74
Ours	0.3221/0.6417	25.6631/27.5656	0.9154/0.9483	$4.11e^{-3}/4.73e^{-3}$	1970.76/1980.54

INR-based codecs (COIN [21], COIN++ [20]). We utilize the open-source PyTorch implementation [16] of NeuRBF for I-NGP. These INR methods maintain consistent training steps with GaussianImage. Detailed implementation notes for baselines are found in the appendix.

## 4.2 Image Representation

Fig. 1 (left) and Table 1 show the representation performance of various methods on the Kodak and DIV2K datasets under the same training steps. Although MLP-based INR methods (SIREN [48], WIRE [47]) utilize fewer parameters to fit an image, they suffer from enormous training time and hyperslow rendering speed. Recent feature grid-based INR methods (I-NGP [42], NeuRBF [17]) accelerate training and inference, but they demand substantially more GPU memory compared to GS-based methods. Since the original 3D GS uses 3D Gaussian as the representation unit, it face the challenge of giant parameter count, which decelerates training and restricts inference speed. By choosing 2D Gaussian as the representation unit, our method secures pronounced advantages in training duration, rendering velocity, and GPU memory usage, while substantially reducing the number of stored parameters yet preserving comparable fitting quality.

## 4.3 Image Compression

**Coding performance.** Fig. 4 presents the RD curves of various codecs on the Kodak and DIV2K datasets. Notably, our method achieves comparable compression performance with COIN [21] and COIN++ [20] in PSNR. With the help of the partial bits-back coding, our codec can outperform COIN and COIN++. Furthermore, when measured by MS-SSIM, our method surpasses COIN by a large margin. Fig. 5 provides a qualitative comparison between our method, JPEG [55], and COIN, revealing that our method restores image details more effectively and delivers superior reconstruction quality by consuming lower bits.

**Table 3:** Ablation study of image representation on Kodak dataset with 30000 Gaussian points over 50000 training steps. AR means accumulated blending-based rasterization, M indicates merging color coefficients  $\mathbf{c}$  and opacity  $o$ . RS denotes decomposing the covariance matrix into rotation and scaling matrices. The final row in each subclass represents our default solution.

Methods	PSNR $\uparrow$	MS-SSIM $\uparrow$	Training Time(s) $\downarrow$	FPS $\uparrow$	Params(K) $\downarrow$
3D GS (w/ L1+SSIM)	37.75	0.9961	285.26	1067	1770
3D GS (w/ L2)	37.41	0.9947	197.90	1190	1770
Ours (w/ L2+w/o AR+w/o M)	37.89	0.9961	104.76	2340	270
Ours (w/ L2+w/ AR+w/o M)	38.69	0.9963	98.54	2555	270
Ours(w/ L2+w/ AR+w/ M)	38.57	0.9961	91.06	2565	240
Ours (w/ L1)	36.46	0.9937	92.68	2438	240
Ours (w/ SSIM)	35.65	0.9952	183.20	2515	240
Ours (w/ L1+SSIM)	36.57	0.9945	188.22	2576	240
Ours (w/ L2+SSIM)	34.73	0.9932	189.17	2481	240
Ours (w/ L2)	38.57	0.9961	91.06	2565	240
Ours-RS	38.83	0.9964	98.55	2321	240
Ours-Cholesky	38.57	0.9961	91.06	2565	240

**Computational complexity.** Table 2 reports the computational complexity of several image codecs on the DIV2K dataset, with learning-based codecs operating on an NVIDIA V100 GPU and traditional codecs running on an Intel Core(TM) i9-10920X processor at a base frequency of 3.50GHz in single-thread mode. Impressively, the decoding speed of our codec reaches around 2000 FPS, outpacing traditional codecs like JPEG, while also providing enhanced compression performance at lower bitrates. This establishes a significant advancement in the field of neural image codecs.

#### 4.4 Ablation Study

**Effect of different components.** To highlight the contributions of the key components in GaussianImage, we conduct a comprehensive set of ablation studies, as detailed in Table 3. Initially, the original 3D GS [30] method, which employs a combination of L1 and SSIM loss, is adapted to use L2 loss. This modification halves the training time at a minor cost to performance. Then, we replace the 3D Gaussian with the basic 2D Gaussian in Section 3.1, which not only improves the fitting performance and decreases training time by  $\frac{1}{3}$ , but also doubles the inference speed and reduces parameter count by  $6.5\times$ . By simplifying alpha blending to accumulated blending, we eliminate the effects of random 2D Gaussian ordering and bypasses the complex calculations for the accumulated transparency  $T$ , resulting in a significant 0.8dB improvement in PSNR alongside notable training and inference speed gains. This underscores the efficiency of our proposed accumulated blending approach. Furthermore, by merging the color vector  $\mathbf{c}$  and opacity  $o$  to form our upgraded 2D Gaussian, we observe a 10% reduction in parameter count with a negligible 0.1dB decrease in PSNR.

**Loss function.** We evaluate various combinations of L2, L1, and SSIM losses, with findings presented in Table 3. These results confirm that L2 loss is optimally

**Table 4:** Ablation study of quantization schemes on Kodak dataset. The first row denotes our final solution and is set as the anchor.

Variants	BD-PSNR (dB) $\uparrow$ BD-rate (%) $\downarrow$	BD-MS-SSIM $\uparrow$ BD-rate (%) $\downarrow$
Ours	0	0
(V1) w/o $\mathcal{L}_c$ +w/ RVQ + 6bit	-3.145 333.16	-0.0824 337.84
(V2) w/o $\mathcal{L}_c$ +w/o RVQ + 6bit	-0.159 7.02	-0.0030 6.14
(V3) w/o $\mathcal{L}_c$ +w/o RVQ + 8bit	-0.195 11.69	-0.0127 62.77

suites for our approach, significantly improving image reconstruction quality while facilitating rapid training.

**Factorized form of covariance matrix.** As outlined in Section 3.1, we optimize the factorized form of the covariance matrix through decomposition. The findings detailed in Table 3 demonstrate that various factorized forms possess similar capabilities in representing images, despite the decomposition’s inherent non-uniqueness. The appendix provides additional analysis on the compression robustness of different factorized forms.

**Quantization strategies.** Table 4 investigates the effect of different quantization schemes on image compression. Without the commitment loss  $\mathcal{L}_c$  (V1), the absence of supervision for the RVQ codebook leads to significant deviations of the codebook vector from the original vector, adversely affecting reconstruction quality. Moreover, eliminating RVQ in favor of 6-bit integer quantization for color parameters (V2) resulted in a 7.02% increase in bitrate consumption when compared with our default solution. This outcome suggests that the color vectors across different Gaussians share similarities, making them more suitable for RVQ. Further exploration into the use of higher bit quantization (V3) reveals a deterioration in compression efficiency.

## 5 Conclusion

In this work, we introduce GaussianImage, an innovative paradigm for image representation that leverages 2D Gaussian Splatting. This approach diverges significantly from the commonly utilized implicit neural networks, offering a discrete and explicit representation of images. When compared to 3D Gaussian Splatting, employing 2D Gaussian kernels brings forth two notable benefits for image representation. Firstly, the computationally intensive alpha blending is simplified to an efficient and permutation-invariant accumulated summation blending. Secondly, the quantity of parameters required for each Gaussian diminishes drastically from 59 to just 8, marking a substantial reduction in complexity. Consequently, GaussianImage emerges as a highly efficient and compact technique for image coding. Experimental results confirm that this explicit representation strategy enhances training and inference efficiency substantially. Moreover, it delivers a competitive rate-distortion performance after adopting vector quantization on parameters, compared to methods adopting implicit neural representation. These findings suggest promising avenues for further exploration in non-end-to-end image compression and representation strategies.

**Acknowledgments.** This work was supported by the National Natural Science Fund of China (Project No. 42201461, 62102024, 62331014) and the General Research Fund (Project No. 16209622) from the Hong Kong Research Grants Council.

## References

1. Kodak lossless true color image suite. <https://r0k.us/graphics/kodak/> (1999) 10
2. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (July 2017) 10
3. Ahmed, N., Natarajan, T., Rao, K.R.: Discrete cosine transform. *IEEE transactions on Computers* **100**(1), 90–93 (1974) 2
4. Ballé, J., Laparra, V., Simoncelli, E.P.: Density modeling of images using a generalized normalization transformation. *arXiv preprint arXiv:1511.06281* (2015) 5
5. Ballé, J., Laparra, V., Simoncelli, E.P.: End-to-end optimized image compression. In: *International Conference on Learning Representations* (2017) 5, 11, 12
6. Ballé, J., Minnen, D., Singh, S., Hwang, S.J., Johnston, N.: Variational image compression with a scale hyperprior. In: *International Conference on Learning Representations* (2018) 5, 11, 12
7. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5855–5864 (2021) 4
8. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5470–5479 (2022) 4
9. Bellard, F.: Bpg image format. <https://bellard.org/bpg/> (2014) 5
10. Bhalgat, Y., Lee, J., Nagel, M., Blankevoort, T., Kwak, N.: Lsq+: Improving low-bit quantization through learnable offsets and better initialization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 696–697 (2020) 3, 8
11. Chen, G., Wang, W.: A survey on 3d gaussian splatting. *arXiv preprint arXiv:2401.03890* (2024) 2
12. Chen, H., Gwilliam, M., Lim, S.N., Shrivastava, A.: Hnerv: A hybrid neural representation for videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10270–10279 (2023) 4
13. Chen, H., He, B., Wang, H., Ren, Y., Lim, S.N., Shrivastava, A.: Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems* **34**, 21557–21568 (2021) 4
14. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8628–8638 (2021) 2
15. Chen, Y., Chen, Z., Zhang, C., Wang, F., Yang, X., Wang, Y., Cai, Z., Yang, L., Liu, H., Lin, G.: Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2024) 5



16. Chen, Z., Li, Z., Song, L., Chen, L., Yu, J., Yuan, J., Xu, Y.: <https://github.com/oppo-us-research/NeuRBF> 12
17. Chen, Z., Li, Z., Song, L., Chen, L., Yu, J., Yuan, J., Xu, Y.: Neurbf: A neural fields representation with adaptive radial basis functions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4182–4194 (2023) 2, 4, 10, 11, 12
18. Chen, Z., Wang, F., Liu, H.: Text-to-3d using gaussian splatting. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2024) 5
19. Cheng, Z., Sun, H., Takeuchi, M., Katto, J.: Learned image compression with discretized gaussian mixture likelihoods and attention modules. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7939–7948 (2020) 5
20. Dupont, E., Loya, H., Alizadeh, M., Golinski, A., Teh, Y., Doucet, A.: Coin++: neural compression across modalities. Transactions on Machine Learning Research 2022(11) (2022) 2, 4, 5, 12
21. Dupont, E., Golinski, A., Alizadeh, M., Teh, Y.W., Doucet, A.: Coin: Compression with implicit neural representations. In: Neural Compression: From Information Theory to Applications–Workshop@ ICLR 2021 (2021) 2, 4, 5, 12
22. Fathony, R., Sahu, A.K., Willmott, D., Kolter, J.Z.: Multiplicative filter networks. In: International Conference on Learning Representations (2020) 2, 4
23. Gray, R.: Vector quantization. IEEE Assp Magazine 1(2), 4–29 (1984) 8
24. Guo, Z., Flamich, G., He, J., Chen, Z., Hernández-Lobato, J.M.: Compression with bayesian implicit neural representations. Advances in Neural Information Processing Systems 36 (2024) 2, 5
25. He, D., Yang, Z., Peng, W., Ma, R., Qin, H., Wang, Y.: Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5718–5727 (2022) 5
26. Heil, C.E., Walnut, D.F.: Continuous and discrete wavelet transforms. SIAM review 31(4), 628–666 (1989) 2
27. Higham, N.J.: Cholesky factorization. Wiley interdisciplinary reviews: computational statistics 1(2), 251–254 (2009) 6
28. Huang, H., Li, L., Cheng, H., Yeung, S.K.: Photo-slam: Real-time simultaneous localization and photorealistic mapping for monocular, stereo, and rgb-d cameras. arXiv preprint arXiv:2311.16728 (2023) 4
29. Keetha, N., Karhade, J., Jatavallabhula, K.M., Yang, G., Scherer, S., Ramanan, D., Luiten, J.: Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. arXiv preprint arXiv:2312.02126 (2023) 4
30. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics 42(4) (2023) 2, 3, 4, 6, 10, 13
31. Koyuncu, A.B., Gao, H., Boev, A., Gaikov, G., Alshina, E., Steinbach, E.: Contextformer: A transformer with spatio-channel attention for context modeling in learned image compression. In: European Conference on Computer Vision. pp. 447–463. Springer (2022) 5
32. Kunze, J., Severo, D., Zani, G., van de Meent, J.W., Townsend, J.: Entropy coding of unordered data structures. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=afQuNt3Ruh> 9
33. Ladune, T., Philippe, P., Henry, F., Clare, G., Leguay, T.: Cool-chic: Coordinate-based low complexity hierarchical image codec. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13515–13522 (2023) 5



34. Leguay, T., Ladune, T., Philippe, P., Clare, G., Henry, F., Déforges, O.: Low-complexity overfitted neural image codec. In: 2023 IEEE 25th International Workshop on Multimedia Signal Processing (MMSP). pp. 1–6. IEEE (2023) 5
35. Li, Z., Wang, M., Pi, H., Xu, K., Mei, J., Liu, Y.: E-nerv: Expedite neural video representation with disentangled spatial-temporal context. In: European Conference on Computer Vision. pp. 267–284. Springer (2022) 4
36. Liu, J., Sun, H., Katto, J.: Learned image compression with mixed transformer-cnn architectures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14388–14397 (2023) 5
37. Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. arXiv preprint arXiv:2308.09713 (2023) 5
38. Ma, C., Yu, P., Lu, J., Zhou, J.: Recovering realistic details for magnification-arbitrary image super-resolution. IEEE Transactions on Image Processing 31, 3669–3683 (2022) 2
39. Martel, J.N., Lindell, D.B., Lin, C.Z., Chan, E.R., Monteiro, M., Wetzstein, G.: Acorn: adaptive coordinate networks for neural scene representation. ACM Transactions on Graphics (TOG) 40(4), 1–13 (2021) 2, 4
40. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European Conference on Computer Vision. pp. 405–421. Springer (2020) 4
41. Minnen, D., Ballé, J., Toderici, G.D.: Joint autoregressive and hierarchical priors for learned image compression. In: Advances in neural information processing systems (2018) 5, 9
42. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) 41(4), 1–15 (2022) 2, 4, 10, 11, 12
43. Nguyen, Q.H., Beks, W.J.: Single image super-resolution via a dual interactive implicit neural network. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 4936–4945 (2023) 2
44. Quan, Y., Yao, X., Ji, H.: Single image defocus deblurring via implicit neural inverse kernels. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12600–12610 (2023) 2
45. Ramasinghe, S., Lucey, S.: Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. In: European Conference on Computer Vision. pp. 142–158. Springer (2022) 2, 4
46. Ryder, T., Zhang, C., Kang, N., Zhang, S.: Split hierarchical variational compression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 386–395 (2022) 3, 8, 9
47. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: Wire: Wavelet implicit neural representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18507–18516 (2023) 2, 4, 10, 11, 12
48. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. Advances in neural information processing systems 33, 7462–7473 (2020) 2, 4, 10, 11, 12
49. Skodras, A., Christopoulos, C., Ebrahimi, T.: The jpeg 2000 still image compression standard. IEEE Signal processing magazine 18(5), 36–58 (2001) 5, 11, 12
50. Stanley, K.O.: Compositional pattern producing networks: A novel abstraction of development. Genetic programming and evolvable machines 8, 131–162 (2007) 2

51. Strümpfer, Y., Postels, J., Yang, R., Gool, L.V., Tombari, F.: Implicit neural representations for image compression. In: European Conference on Computer Vision. pp. 74–91. Springer (2022) [2](#), [5](#)
52. Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11358–11367 (2021) [2](#), [4](#)
53. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems* **33**, 7537–7547 (2020) [2](#), [4](#)
54. Townsend, J., Bird, T., Barber, D.: Practical lossless compression with latent variables using bits back coding. *arXiv preprint arXiv:1901.04866* (2019) [3](#), [8](#), [9](#)
55. Wallace, G.K.: The jpeg still picture compression standard. *Communications of the ACM* **34**(4), 30–44 (1991) [5](#), [11](#), [12](#)
56. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003. vol. 2, pp. 1398–1402. Ieee (2003) [11](#)
57. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528* (2023) [5](#)
58. Xie, X., Zhou, P., Li, H., Lin, Z., Shuicheng, Y.: Adan: Adaptive nesterov momentum algorithm for faster optimizing deep models. In: Has it Trained Yet? NeurIPS 2022 Workshop (2022) [11](#)
59. Xu, D., Wang, P., Jiang, Y., Fan, Z., Wang, Z.: Signal processing for implicit neural representations. *Advances in Neural Information Processing Systems* **35**, 13404–13418 (2022) [2](#)
60. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5438–5448 (2022) [4](#)
61. Xu, W., Jiao, J.: Revisiting implicit neural representations in low-level vision. In: International Conference on Learning Representations Workshop (2023) [2](#)
62. Yan, C., Qu, D., Wang, D., Xu, D., Wang, Z., Zhao, B., Li, X.: Gs-slam: Dense visual slam with 3d gaussian splatting. *arXiv preprint arXiv:2311.11700* (2023) [4](#)
63. Yan, Y., Lin, H., Zhou, C., Wang, W., Sun, H., Zhan, K., Lang, X., Zhou, X., Peng, S.: Street gaussians for modeling dynamic urban scenes. *arXiv preprint arXiv:2401.01339* (2024) [5](#)
64. Yang, Z., Yang, H., Pan, Z., Zhu, X., Zhang, L.: Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *arXiv preprint arXiv:2310.10642* (2023) [5](#)
65. Ye, V., Turkulainen, M., the Nerfstudio team: gsplat, <https://github.com/nerfstudio-project/gsplat> [11](#)
66. Zeghidour, N., Luebs, A., Omran, A., Skoglund, J., Tagliasacchi, M.: Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **30**, 495–507 (2021) [3](#), [9](#)
67. Zhang, X., Yang, R., He, D., Ge, X., Xu, T., Wang, Y., Qin, H., Zhang, J.: Boosting neural representations for videos with a conditional decoder. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024) [4](#)
68. Zhou, X., Lin, Z., Shan, X., Wang, Y., Sun, D., Yang, M.H.: Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. *arXiv preprint arXiv:2312.07920* (2023) [5](#)

69. Zielonka, W., Bagautdinov, T., Saito, S., Zollhöfer, M., Thies, J., Romero, J.: Drivable 3d gaussian avatars. arXiv preprint arXiv:2311.08581 (2023) [5](#)