Towards Edge Holography via Implicit Neural Representation and Compression

Hyunmin Ban, Wenbin Zhou, and Yifan Peng

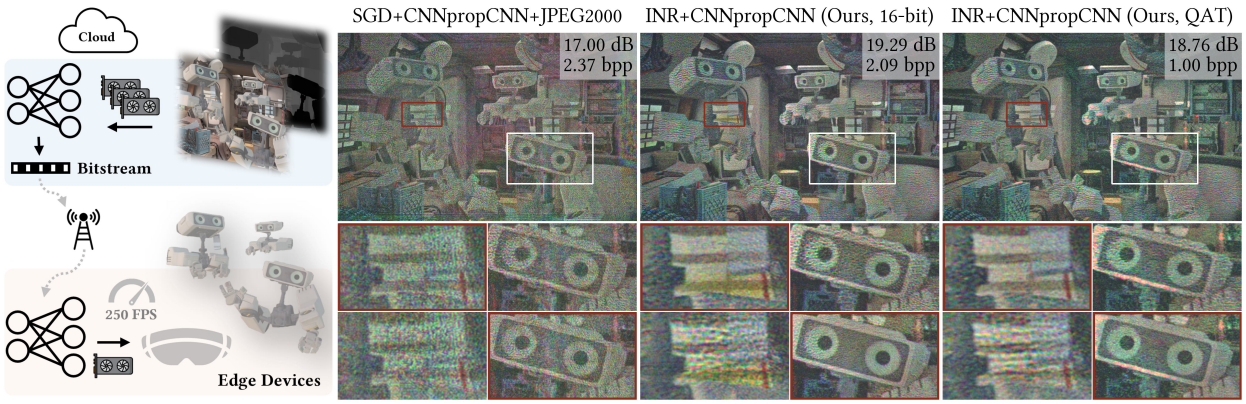

Fig. 1: The emerging cloud-edge computing mechanism offers great potential for deploying edge holography in the near future. In light of this trend, our edge computer-generated holography (CGH) exploration employs an *implicit neural representation (INR)* to compactly encode hologram data, facilitating efficient compression and fast decoding (left diagram). Further, by integrating camera-calibrated wave propagation, our framework achieves high-fidelity, unfiltered, 3D holographic displays in experiments (second and third columns), outperforming state-of-the-art iterative methods in compression. Metrics presented at the corners denote the reconstructed PSNR and compression rate.

Abstract—Holographic displays offer the promise of realistic 3D visualization for virtual and augmented wearable solutions. Nevertheless, existing computer-generated holography (CGH) methods often struggle with either a high computational burden or limited display realism. While the emerging cloud-edge computing mechanism can enable the real-time streaming of holograms, classic image compression techniques struggle to efficiently encode and decode the substantial high-frequency information inherent in hologram data. In light of these challenges, we present a display-aware and lightweight CGH framework, leveraging implicit neural representations (INRs) and camera-calibrated wave propagation, to generate and compress high-fidelity phase-only holograms. Specifically, our approach interprets hologram generation as a continuous function approximation problem, enabling the network, with reduced parameters, to effectively learn the inherent periodicity and high-frequency components of 2D and 3D hologram data. To enable efficient deployment, we further incorporate quantization-aware training, followed by entropy coding. Experimental results evaluated on an unfiltered holographic display prototype demonstrate that the proposed INR-CGH retains image quality comparable to that of existing optimization-based methods in both 2D and 3D scenarios. In addition, our compact INR representation achieves up to 11$\times$ compression rate with minimal quality degradation and can be further reduced via quantization-aware training. The resulting model enables $\geq$250 fps in decoding speed, paving the way towards edge holography.

Index Terms—CGH, INR, compression, learned propagation.

1 Introduction

Holographic displays have long been recognized for their potential to deliver realistic 3D perception in a compact system form factor [7, 32], facilitating various virtual and augmented reality applications [34, 38]. Yet, a critical challenge for the practical deployment of holography lies in the enduring trade-off between computational load and image quality [54]. In the current landscape, cloud computing has emerged as a promising solution to deliver enhanced visual experiences to general users on edge devices with constrained processing capabilities and limited battery life [42], aligning with the advancement of next-generation holographic wearables [26].

Recent advances in computer-generated holography (CGH) have led to substantial enhancements in image quality and large etendue expansion [44], typically through pixel-wise iterative optimization approaches such as stochastic gradient descent (SGD) [9, 49]. Recently, a vast number of deep neural network models [52, 67, 68] have also been utilized. Nonetheless, both methodologies may exhibit challenges when launched on edge devices within cloud computing environments [7]. While SGD can achieve high-fidelity reconstructions, its iterative process is computationally intensive and unsuitable for edge-side decoding. Likewise, although neural networks facilitate real-time hologram synthesis, their inherent high resolution and large receptive field requirements in CGH can be problematic. Existing network architectures, notably those based on convolutional neural networks (CNNs), tend to produce phase holograms with checkerboard artifacts or spurious fringes [70], necessitating additional optomechanical components to filter out unwanted light.

It is worth noting that holograms exhibit very distinct statistical properties compared to natural images, rendering holographic re-

- *Hyunmin Ban, Wenbin Zhou, and Yifan Peng are with the Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong SAR, China (E-mail: hmban@connect.hku.hk; zhouwb@connect.hku.hk; evanpeng@hku.hk).*
- *Corresponding author: Yifan Peng.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

constructions highly susceptible to distortions in the hologram data. Consequently, classical lossy codecs tailored for natural images or videos [46, 71] may prove less effective for compressing holograms. The conventional sequential processing practice — hologram generation first and then compression to devices — could be suboptimal in saving data bandwidth while preserving effective details for holography. As such, the end-to-end optimization concept has inspired emerging hologram compression methods [3, 54], realizing significant bitrate reduction while maintaining reconstruction quality in simulation. Among them, DPRC [66] stands out, leveraging a variational autoencoder to jointly generate and compress phase holograms. The compression performance of existing methods remains modest, reporting typically 1.3∼4.4 bits per pixel (bpp). While effective in simulation, these CNN-based models still rely on band-limited assumptions, such as Fourier apertures or optical filtering, to suppress high-frequency artifacts, and they do not explicitly account for hardware imperfections during generation and compression.

In parallel, implicit neural representations (INRs) have emerged as compact, signal-adaptive models that encode high-frequency variations through global coordinate-based mappings [28, 56, 58], offering a promising alternative to convolution-based pipelines. In this work, we explore an INR-empowered CGH framework that can generate and compress phase-only holograms efficiently while retaining high reconstruction fidelity (Fig. 1). Unlike traditional CNN architectures, INRs represent signals as continuous functions over spatial coordinates, enabling the network to naturally capture high-frequency variations and periodic structures inherent in holographic data. To the practicality of our pipeline for real-time decoding, we incorporate a neural compression framework that combines quantization-aware training and entropy coding. Consequently, the learned MLP parameters can be stored and transmitted in a compact format, suitable for embedded hardware setups. Importantly, our experimental results reveal approximately 32× compression relative to pixel-wise optimization (e.g., SGD) methods at comparable visual quality, while decoding holograms at over 250 frames per second. In summary, our technical contributions are:

- We explore INRs for phase-only 3D CGH, demonstrating their superior capability to capture high-frequency details over classic CNN-based methods.
- The lightweight nature of our proposed INR pipeline enables real-time hologram generation on edge devices with limited computational resources.
- By integrating quantization-aware training and entropy coding for the INR parameters, we reduce memory while retaining reconstruction fidelity, aiming for future deployment.
- We build an unfiltered holographic display prototype and realize reasonable optical reconstruction fidelity, evidenced by quantitative and visual assessments at multiple depths.

Overview of Limitations. Although the proposed edge INR-CGH pipeline delivers real-time decoding and enhanced image quality over state-of-the-art iterative CGH methods by using lower bit rates, it is currently demonstrated only in the setting of pre-encoded holographic video streaming, due to the encoding process not being real-time. Yet, this pipeline naturally aligns with a cloud-edge computing mechanism, wherein computationally intensive encoding is performed offline on powerful cloud servers, while lightweight INR decoding enables real-time hologram synthesis on resource-constrained edge devices. In addition, while our reconstructions appear robust under unfiltered optical configurations, there exist residual color artifacts, particularly for uniform or low-contrast inputs.

2 Related Work

Neural Network-empowered CGH. Previous works already show deep neural networks can accelerate the phase-only hologram generation [30, 31, 40]. Notably, Peng et al. [49] accounted for the optical aberrations into the CNN model, thereby enhancing image fidelity while reducing computational demands compared to iterative counterparts. In a similar vein, Shi et al. [52] employed a lightweight ResNet architecture to generate complex wavefronts, which were subsequently transformed into phase-only representations via the double phase amplitude encoding technique [59]. While such direct methods are capable of generating phase-only holograms in near real time on edge devices, they can be prone to artifacts such as regular stripe or checkerboard patterns, typically requiring additional physical filtering [53]. Recent studies have also explored neural rendering paradigms for CGH [12, 13], primarily synthesizing complex-valued holograms and likewise require a subsequent phase-only conversion, with the assumption of optical filtering. In contrast, we study an implicit neural representation that produces phase-only holograms directly and is optimized for unfiltered conditions.

Implicit Neural Representation. Recent advancements of INRs [56, 61] have demonstrated significant potential in addressing limitations inherent in traditional CNN-based approaches. Their recent innovations have further expanded to capture high-frequency and complex patterns through refined activation functions [41, 50] and multi-scale frameworks [45, 63]. Accordingly, INRs have proven effective across various domains, including image compression [21], video representation [11], 3D scenes [43], underscoring its potential to represent high-frequency information and complex spatial variations — characteristics essential for precise 3D holographic reconstruction. Most recently, Choi et al. [16] leveraged INRs to model partially coherent wavefront propagation through waveguides. While these results indicate the promise of INRs in holography, we explore INRs as a compact representation of phase-only holograms for compression and fast edge-side decoding.

Neural Image Compression. Conventional image codecs, including JPEG and JPEG2000 [17, 65], rely on hand-crafted transformations (e.g., discrete cosine transform, wavelet transform), quantization, and entropy coding, which limit adaptability and joint optimization with modern neural pipelines. Recent neural image compression advances based on variational autoencoders [2, 29, 36] replace hand-crafted stages with end-to-end learned frameworks, achieving improved rate-distortion performance. More recent INR-based compression approaches [8, 20, 57] overfit compact MLPs to individual images, thereby reducing decoding overhead compared to large-scale autoencoders. These models further leverage entropy modeling per image [33, 39] to optimize compressibility. Although originally designed for natural images, we explore the applicability of INR-based compression techniques in the holography domain, which poses distinct challenges due to its phase-only representation nature and rich high-frequency content.

Hologram Compression. The inherently high-frequency, complex-valued nature of holographic data poses challenges for efficient compression, where traditional image and video codecs struggle to preserve the subtle phase relationships required for high-fidelity holographic reconstructions [7, 54]. Neural network-based approaches have been proposed, leveraging variational autoencoders [66] that jointly learns to generate and compress holograms in the latent space and variants extending this for phase hologram video compression [3]. However, these methods typically rely on filtered holography or impose constraints that compromise reconstruction fidelity in practical display scenarios. In addition, prior codec-based works have been introduced to compress complex holograms [35, 51] and object planes [5]. In contrast, our work explores this problem via INR-based compression through quantization-aware training and entropy coding, tailored for unfiltered, phase-only holograms.

Camera-calibrated Wave Propagation. To bridge the gap between simulation and physical reconstructions caused by hardware imperfections [26], such as SLM lookup tables (LUTs) inaccuracies, optical aberrations, and nonuniform light source, various calibration techniques are employed. Explicit calibration methods involve capturing the SLM surface while displaying a series of structured patterns, for instance, uniform phase distributions or checkerboard patterns, so as to characterize spatial variance of LUTs and/or optical aberrations [53], which is often a tedious process. Researchers also calibrate the system in an end-to-end manner. Online calibration involves updating the phase

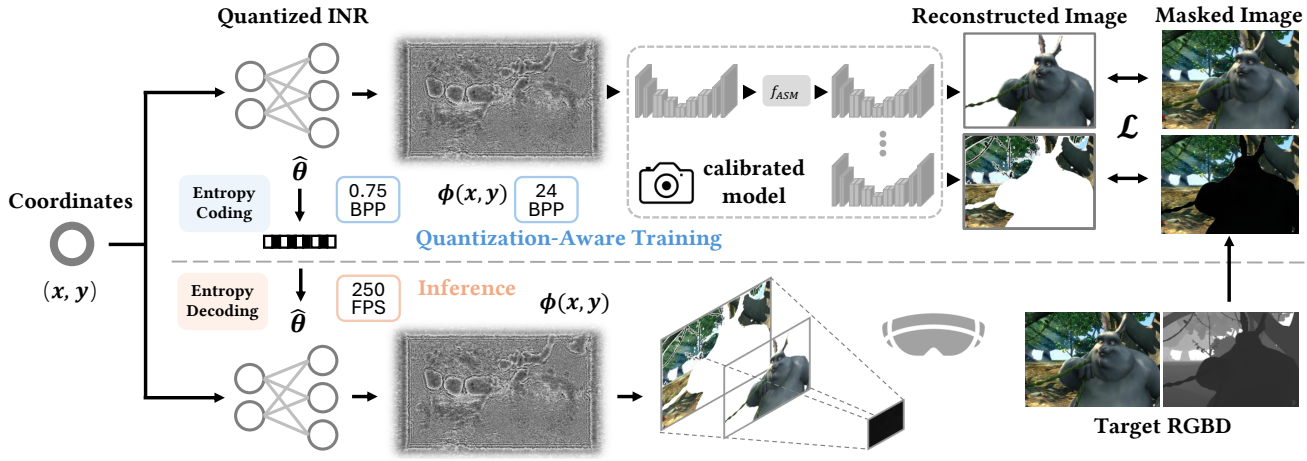

Fig. 2: Pipeline overview of our proposed INR-based hologram generation and compression, dubbed *INR-CGH* for brevity. On the server side, coordinates are input to the INR model, producing a compact and continuous representation of a phase-only hologram optimized through quantization-aware training with a camera-calibrated, learned propagation model and the RGBD supervision. INR parameters are subsequently compressed through quantization and entropy coding for efficient transmission. On the edge side, the compressed INR parameters are decoded and loaded, enabling rapid direct inference of phase-only holograms, which are then displayed on the spatial light modulator (SLM) for high-fidelity, unfiltered 3D reconstruction.

patterns using live feedback from a camera, typically achieving high fidelity results but requiring calibration for each individual image [10,48]. In contrast, offline calibration focuses on explicitly or implicitly learning a *compensation* model [14], using a collection of captured images on the display, so as to rectify the wave propagation model and improve generalization. We also leverage a camera-calibrated, learned wave propagation model, akin to the strategy studied by Choi et al. [15] and its variant [72]. Such pre-computation addresses hardware mismatches in the training stage and thus facilitates real-time decoding on resource-constrained edge devices.

3 Method

The objective of our proposed INR-based CGH framework is to generate and compress phase-only holograms that reconstruct the target field at multiple depth planes, as illustrated in Fig. 2. The pipeline begins with an INR, where an MLP maps coordinates to phase values on the SLM plane. This phase-only hologram is then propagated to the target plane via either a wave propagation model, such as the angular spectrum method (ASM), or a camera-calibrated neural network, which accounts for actual hardware imperfections. Once the MLP is trained to overfit the phase-only hologram, its weights and biases can either be retained as an efficient representation or further compressed via quantization-aware training followed by entropy coding to realize additional storage and transmission efficiency. We delve into individual parts below.

3.1 INR for Phase Holograms

We seek to represent a phase-only hologram as a continuous function $\phi(x,y)$ defined over spatial coordinates through a fully connected neural network. Specifically, our INR-CGH employs a global optimization framework, inherently suitable for modeling the Fourier-based propagation inherent in holography. Unlike CNNs, which operate on discrete pixel grids and optimize phase holograms using local convolution kernels that likely introduce checkerboard artifacts and localized inaccuracies, an INR, being a fully connected network, naturally optimizes the hologram globally, directly aligning with the global Fourier transforms utilized in standard wave propagation mechanisms like ASM.

Formally, we define our INR, F_θ, as a fully connected multilayer perception (MLP) parameterized by θ, which maps normalized spatial coordinates (x,y) on the SLM plane to phase values, as follows:

$$(x,y) \mapsto \phi(x,y) = F_\theta(x,y), \tag{1}$$

where this F_θ involves multiple hidden layers, each performing a nonlinear transformation using variable-periodic activation functions as introduced by FINER [41]. Given input spatial coordinates $\mathbf{h}_0 = [x,y]^\top$, each hidden layer can be defined as:

$$\mathbf{y}_i = \mathbf{W}_i\mathbf{h}_i + b_i \text{ and } \mathbf{h}_{i+1} = \sin\big(\omega_0\,(|\mathbf{y}_i|+1)\,\mathbf{y}_i\big), \tag{2}$$

where $\mathbf{h}_i$ represents the output of the i^{th} layer, $\mathbf{W}_i$ and b_i denote its weights and biases, and ω_0 is a frequency scaling factor. The final layer outputs the phase at each spatial coordinate $\phi(x,y)$.

To further improve phase representation quality, we leverage the spectral adaptivity of FINER activations. This activation function in Eq. (2) has been shown to enhance high-frequency modeling by adjusting the network's frequency response. In light of this, we investigate the optimal state-of-the-art INR representation suitable for our phase-only hologram generation task, as illustrated in Fig. 3. We observe that the use of FINER activation function appears to enable comparable reconstruction quality compared to the conventional sinusoidal activations (SIREN), without increasing the number of parameters. More detailed analysis is presented in Sec. 5.3.

3.2 Free-space/Learned Propagation and Loss

The output of the INR, $\phi(x,y)$, directly modulates the incident coherent wavefront (assuming unit amplitude) as $e^{j\phi(x,y)}$. This modulated wavefront then undergoes wave propagation using the well-established angular spectrum method [24] to reconstruct amplitudes at target planes. The wave propagation can be efficiently modeled using the Fourier transform as:

$$f^z_{ASM}(\phi(x,y)) = \mathscr{F}^{-1}\left(\mathscr{F}\left(e^{j\phi(x,y)}\right)\cdot H(f_x,f_y;z)\right), \tag{3}$$

where $\mathscr{F}$ and $\mathscr{F}^{-1}$ denote the forward and inverse Fourier Transforms, and $H(f_x,f_y)$ is a transfer function modeling free-space propagation over distance z. The predicted amplitude $a_{\text{pred}}(x,y;z)$ at the target plane is described by $|f^z(\phi(x,y))|$. The detailed derivation of propagation could be found in the supplementary material.

To determine the optimal phase $\phi(x,y)$, one can pose the CGH problem as an optimization task in which the phase hologram is refined to minimize discrepancies between the predicted amplitude and the ground-truth amplitude $a_{\text{gt}}(x,y;z)$. In the single-depth case, this distortion can be quantified using a loss function such as:

$$\mathscr{L}_z = \|s\cdot a_{\text{pred}}(x,y;z) - a_{\text{gt}}(x,y;z)\|_2^2, \tag{4}$$

where s is an optimizable scale factor accounting for the power of the laser. To extend the single-plane phase representation to multi-depth holograms, we can compute the wave propagation $f^z(\cdot)$ at multiple

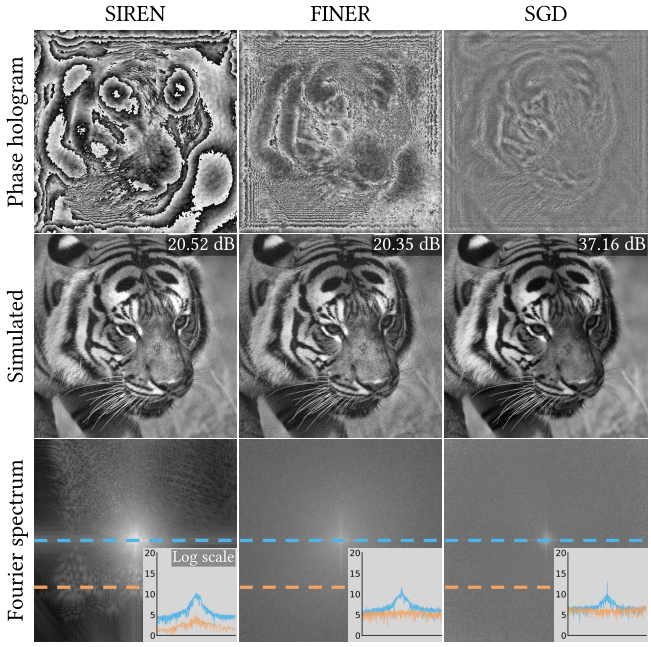

Fig. 3: Comparison of simulation results for studied INR-based (SIREN and FINER) and iterative (e.g., SGD) CGH methods. From top to bottom: phase holograms, reconstructed amplitudes, and corresponding magnitudes of the Fourier spectrum (log scale). FINER and SGD maintain broader spectral coverage and energy in high-frequency regions, while SIREN shows low-frequency components and exhibits sparse high-frequency activation. Dashed plots at the bottom indicate representative regions sampled to visualize the spectral distribution.

planes ($z \in z_1, z_2, \ldots$). In this scenario, each plane z_k can include an in-focus mask $M(u,v;z_k)$ specifying regions of interest for holographic reconstruction. The joint loss function for multi-depth holograms can be defined as:

$$\mathscr{L} = \sum_k \sum_{(u,v)} M(u,v;z_k) \left\| (s \cdot a_{\text{pred}}(x,y;z_k) - a_{\text{gt}}(x,y;z_k) \right\|_2^2. \quad (5)$$

Hereby, minimizing $\mathscr{L}$ drives the phase function $\phi(x,y)$ to reconstruct accurate amplitudes across multiple depths, selectively emphasizing each plane's in-focus region

To address the hardware imperfections (e.g., SLM imperfections and lens distortions) causing mismatches between ASM simulations and physical displays, we replace $f^z_{ASM}(\cdot)$ with $f^z_{model}(\cdot)$, a camera-calibrated wave propagation network [15]. This learnable model replicates the physical capture process, bridging the simulation-reality gap. Our INR can then back-propagate through this calibrated model to adapt and compress the phase hologram $\phi(x,y)$ in an end-to-end framework, with least impacts to the display fidelity. Detailed structure and training strategies of this model are presented in the following implementation section.

3.3 Implicit Neural Compression

While vanilla MLP-based INRs can be relatively compact, deploying them in resource-constrained or high-resolution settings can still be challenging. We opt to extend our overfitted phase hologram representation network with a lightweight compression pipeline, involving an effective *two-step* procedure: quantization followed by entropy coding. In practice, this can reduce the INR's size by an order of magnitude or more, while incurring only a minimal impact on the phase hologram reconstruction fidelity.

Weight Quantization. Notably, quantizing all weights in the INR model from 32-bit to 16-bit floating point already yields high performance in hologram reconstruction. To further reduce the model size, we adopt quantization-aware training to minimize accuracy degradation under lower bit precision. This quantization is informed by the statistical properties of our INR weights, akin to the theoretical analysis presented in SIREN [56], which demonstrates that repeated linear projections of sinusoidal inputs yield Gaussian-distributed outputs at each layer. We observe that the hidden-layer weights in our model empirically exhibit near-zero-centered Gaussian distributions with similar variance across layers. Accordingly, we apply symmetric uniform quantization without offsets, leveraging a single global scale factor α across all layers, effectively reducing metadata size and ensuring consistent quantization behavior. Specifically, this global scale factor is optimized through the learned step size quantization (LSQ) framework [6, 22], which dynamically adjusts the quantization range during training:

$$\hat{\mathbf{W}}_i = \left\lfloor clamp(\frac{\mathbf{W}_i}{\alpha}), -2^{b-1}, 2^{b-1}-1 \right\rceil, \bar{\mathbf{W}}_i = \hat{\mathbf{W}}_i \times \alpha, \quad (6)$$

where $\mathbf{W}_i$ denotes the weight matrix of the i^{th} hidden layer. We quantize all hidden-layer weights to b bits, and the scale factor α is stored in 32-bit floating-point precision. The remaining parameters, including input/output weights and all biases, are also kept in 32-bit precision, since they constitute only a small fraction of the model yet are highly sensitive to quantization. To handle the non-differentiability of the rounding operation, we use the straight-through estimator (STE) [4], which enables gradients to be backpropagated through quantization during training.

Entropy Coding. Following quantization, we further compress the model parameters using asymmetric numeral systems (ANS) [19], which is an efficient entropy coding technique that exploits the symbol frequency distribution and produces a tightly packed bitstream optimized for storage and transmission. At inference time, the bitstream is entropy decoded to recover the quantized weights, and then the learned scale α is used to rescale and approximate float parameters, enabling efficient storage and transmission of the hologram representation.

4 Implementation Details

Training Details. We train our INR model on the NVIDIA RTX 4090 GPU with the Adam optimizer using a learning rate of $5 \times e^{-4}$. Each phase hologram is fitted for 4,000 steps for single depth and 6,000 steps for multiple depths. The network architecture contains 3 layers with 180 channels for single depth and 172 channels for multiple depth, due to the GPU memory limit, and the ω_0 factor is set to 30. The biases of the first layer are initialized uniformly in the range $[-5,5]$ to expand the range of high-frequency coverage. We use images from DIV2K [1] for training the camera-calibrated wave propagation model and apply spatial masking to constrain in-focus regions per depth plane. For multi-depth hologram training, we propagate the hologram to five depth planes, uniformly spaced at 0.07 cm intervals, ranging from 10 cm to 10.28 cm. The wavelengths are set to 639.0 nm (red), 524.9 nm (green), and 445.8 nm (blue), with a pixel pitch of 8 µm. All these configurations are aligned with our experimental setup.

Display Prototype. We have built a benchtop holographic display prototype to validate our method, as shown in Fig. 4. Synthesized holograms were displayed on a phase-only SLM (HOLOEYE Pluto 2). In our experiment, the propagation distance is chosen to be around 10.00 cm to achieve substantial modulation capability while avoiding an increase in system form factor based on prior empirical studies [47, 49]. Illumination was provided by an RGB laser source (Fisba ReadyBeam). The beam was directed through a sequence of optical elements, including a collimating lens, a neutral density filter, a linear polarizer, and a beam splitter. The shaped beam was incident onto the SLM. The modulated wavefronts were reflected, propagated, and transmitted through an eyepiece (Nikon 50 mm SLR lens) and a lens (Canon 35 mm SLR lens) before being recorded by a sensor (FLIR Grasshopper 3). The focused distance was adjusted using an Arduino micro-controller to support re-focusing on multiple pre-defined planes. At this proof-of-concept stage, color results are obtained via post-image processing. Detailed specifications of the devices employed are provided in supplementary Sec. S3.

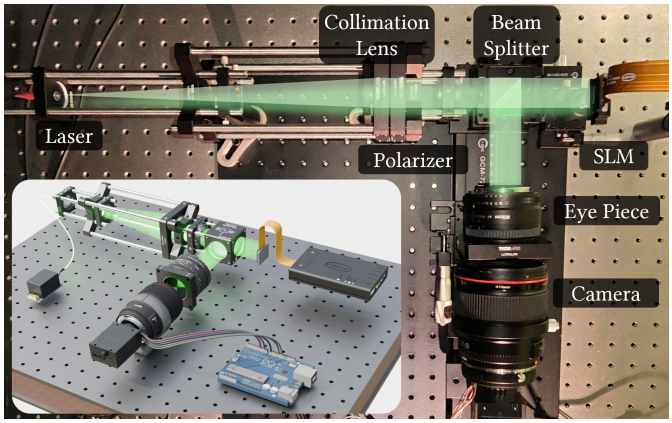

Fig. 4: Photograph and schematic diagram of our holographic display prototype. An Arduino is used to automatically control the focus of the camera. Note that our experimental setup doesn't involve any optical filtering parts that are often placed in-between the eye-piece and the camera.

Notably, our prototype was implemented without any conventional optical filter placed along the reconstruction path to largely mitigate unwanted diffraction orders and ambient artifacts, leading to more challenging scenarios. Likewise, the wave propagation model was trained without the inclusion of optical filters. Prior to data acquisition, we calibrated the SLM's phase response using standard vortex phase patterns, with manual voltage tuning to account for severe phase nonlinearity. To enhance dataset stability, we applied an established homography-based affine alignment to captured images, greatly correcting for camera-SLM misalignment and jitter.

CNNpropCNN-based Calibration. The camera-calibrated, learned wave propagation model is trained by sequentially displaying phase-only holograms on the SLM and capturing the resulting images at five target planes with a focus-tunable camera. Our dataset comprises 10,500 phase patterns and 52,500 captured images, derived from 2D images in the DIV2K dataset [1] at various depths, utilizing the proposed INR-CGH method with ASM as the wave propagator.

Our CNNpropCNN model is operating in the following manner: initially, an SLM UNet predicts the complex field to address crosstalk and imperfections on the SLM plane. This complex field is then propagated to 5 target planes using the ASM. Subsequently, the outputs are fed into a target UNet to account for optical aberrations, with supervision from the captured dataset. In particular, each UNet features 2 input/output channels (real and imaginary components). The first UNet comprises 8 down-/up-sampling layers, starting with 32 channels and doubling to 512 channels. The second UNet consists of 5 down-/up-sampling layers, starting with 8 and doubling to 128 channels. Both networks incorporate instance normalization and utilize Leaky ReLU during downsampling and ReLU during upsampling.

5 Results and Discussion

5.1 Implicit Neural Phase Representation

We first evaluate our proposed INR-CGH against the traditional SGD-based pixel-wise optimization, both integrated with a camera-calibrated, learned wave propagation model, as illustrated in Fig. 6. Both methods generate holograms for visually appealing reconstruction. However, distinct differences emerge when assessing hardware-display performance without optical filtering. Specifically, our INR-generated holograms consistently demonstrate high contrast and higher fidelity in optically captured results, illustrating robustness to high-frequency representation demands. While the SGD-based method achieves reasonable fidelity, it requires iterative optimization at the pixel level, frequently converges to local minima, and is challenging to compress or deploy in practice. In many cases, its pixel-wise optimization can yield sharper local details, for example, the house in Fig. 7. This illustrates that even

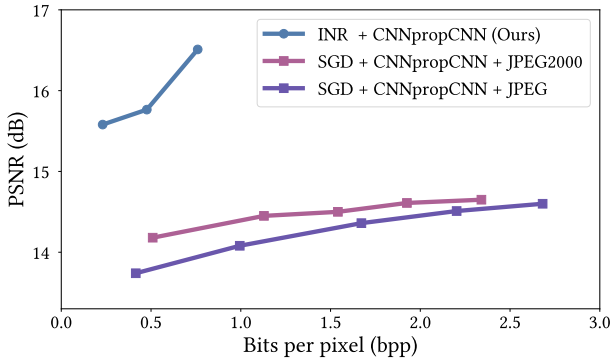

Fig. 5: Rate-distortion comparison on captured holograms employing three compression pipelines: INR + Model (ours), SGD + Model + JPEG, and SGD + Model + JPEG2000. The x-axis indicates bits per pixel (bpp), and the y-axis shows computed PSNR (dB). Each entity is averaged over 10 captured test images (green channel only), selected from the DIV2K validation set, shown in Supplementary Fig. S2.

Table 1: Per-channel decoding time under various compression schemes. Our quantized INR-CGH yields efficient runtime suitable for real-time holographic display applications. JPEG uses quality factors $Q = 5$–60; JPEG2000 is configured in PSNR-based mode with target qualities ranging 18–27 dB. JPEG and JPEG2000 are implemented on the CPU, while DPRC and Ours are on the GPU.

Method	JPEG	JPEG2000	DPRC	Ours (INR 6-bit)
Decoding Time (ms)	3–7	66–160	35	**4**

though our method achieves higher PSNR values, the metric does not always fully reflect perceptual quality. In contrast, our INR approach converges towards a globally coherent phase representation, with faster inference and requiring significantly less memory.

We next present 3D, aka., RGBD, results in Fig. 7. By leveraging the camera-calibrated, learned model, our INR incorporates hardware-specific propagation characteristics across multiple depth planes, thereby improving optical reconstruction quality. In particular, INR-based holograms exhibit fewer artifacts, improved focus consistency, and reduced background noise across depths spanning 10∼10.28 cm. To our best knowledge, we are the first in this domain to realize such a systematic integration of hardware-aware, INR-represented, compressed, fast, 3D hologram generation.

5.2 Implicit Neural Phase Compression

We validate the efficiency of our INR-CGH in the context of compression. Specifically, we compare our INR-based, quantization-aware training framework with holograms generated via SGD [49], which are then compressed using image codecs, such as JPEG [64] and JPEG2000 [17].

As shown in Fig. 5, our method achieves higher quality at a significantly reduced bitrate. The high-rate endpoint corresponds to a 16-bit quantized INR model, yielding 0.76 bpp with high visual fidelity. Further compression using quantization-aware training (QAT) yields 10-bit and 6-bit models with approximately 0.48 bpp (16.7× compression) and 0.23 bpp (34.7× compression), respectively, both of which outperform SGD-based holograms compressed with JPEG or JPEG2000 at equivalent bitrates.

Conventional codecs explicitly attempt to compress the high-frequency data inherent in SGD-based holograms, often resulting in severe reconstruction artifacts even at moderate compression ratios. We also evaluate compressing our INR-generated holograms with JPEG2000, which exhibits larger degradation than compressing INR parameters at comparable bitrates. In contrast, our approach implicitly encodes this high-frequency data within the INR parameters, enabling a compact yet robust representation that achieves much higher compression ratios. Table 1 highlights the decoding efficiency of our

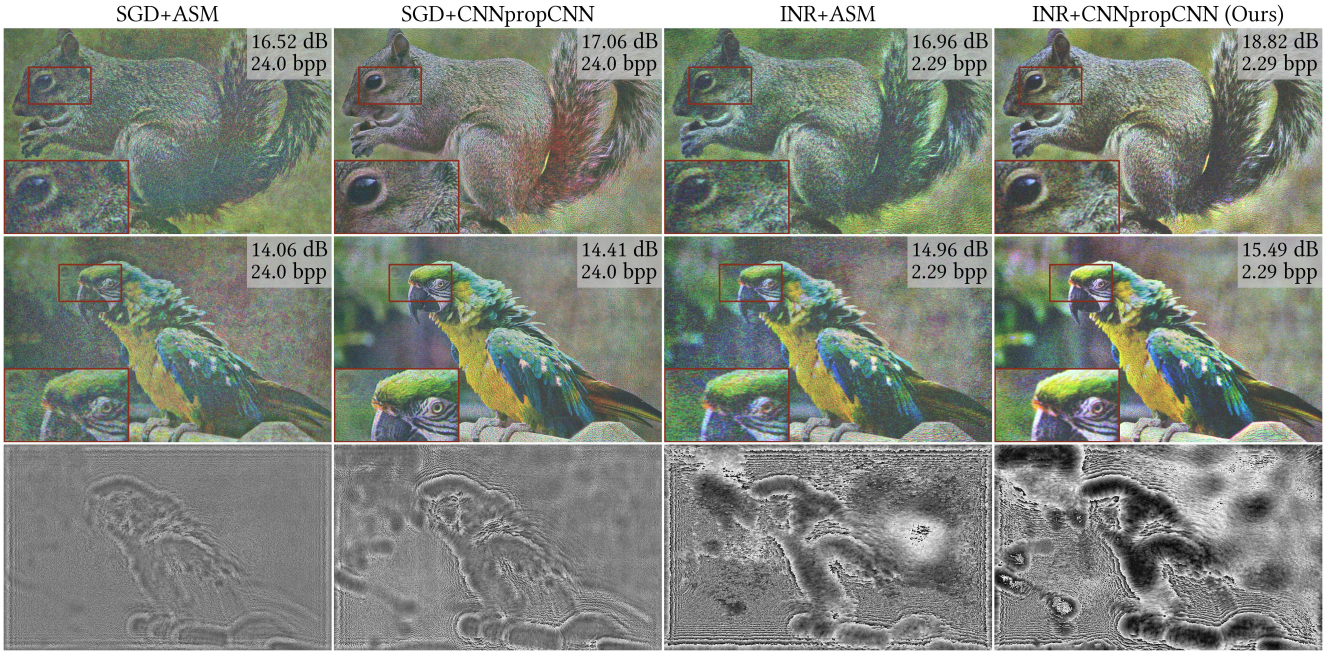

Fig. 6: Experimental reconstructions of 2D images using various CGH pipelines: SGD w/ ASM propagation, SGD w/ CNNpropCNN, INR w/ ASM, and INR w/ CNNpropCNN. This figure ablates the phase representation (SGD vs. INR) and the propagation model used during optimization (ASM vs. CNNpropCNN); the final configuration is INR w/ CNNpropCNN. All results are captured without a $4f$ system, akin to all other results in this manuscript. We present the achieved PSNR in dB and compression rates in bpp at their corners, as well as their corresponding phase-only holograms (green channel-only) for the second row example to display below.

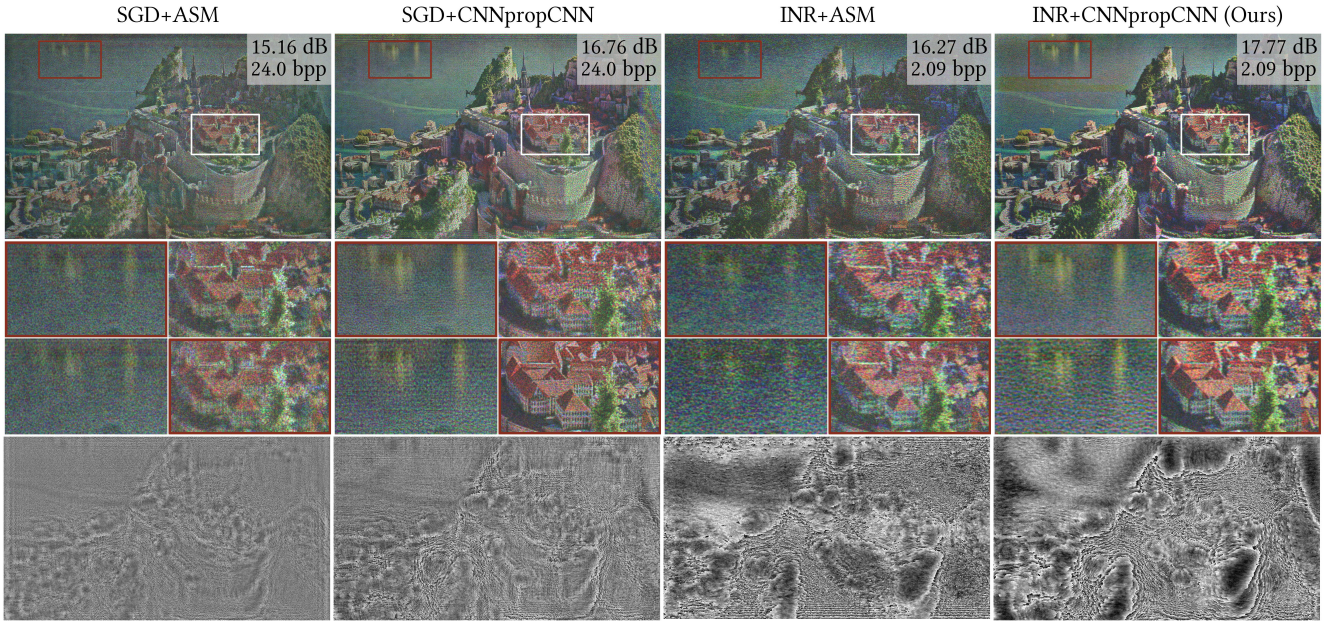

Fig. 7: Experimental reconstructions of 3D (RGBD) scenes using various CGH pipelines: SGD w/ ASM propagation, SGD w/ CNNpropCNN, INR w/ ASM, and INR w/ CNNpropCNN. We perform the same ablation in a 3D setting, varying the phase representation (SGD vs. INR) and the propagation model used during optimization (ASM vs. CNNpropCNN); the final configuration is INR w/ CNNpropCNN. Zoomed-in patches highlight the reconstruction at *near* and *far* depth planes. We present the corresponding phase-only holograms (green channel-only) below.

framework, where phase holograms can be decoded in 4 ms per frame, comparable to JPEG decoding speeds with low quality factors, whereas other baselines are not ready for real-time RGB decoding. Note that JPEG and JPEG2000 are implemented on the CPU, while DPRC and Ours are on the RTX 4090 GPU. Qualitative results are visualized in Fig. 8 and Fig. 9. In the 2D case, our INR-CGH produces cleaner images with higher contrast and reduced noise, despite operating at much lower bitrates. For 3D holography, INR-CGH further exhibits reduced noise across depth planes, particularly in the in-focus regions. However, under more aggressive quantization, fine details may become slightly smeared (see, for instance, the bamboo leaves in the cropped region of Fig. 9), reflecting a trade-off between sharpness and robustness. Refer to the **supplementary video** for side-by-side comparisons of experimental results.

5.3 Bells & Whistles

FINER vs. SIREN for CGH. While in the context of INR compression, fixed-frequency sinusoidal activations (e.g., SIREN) are often

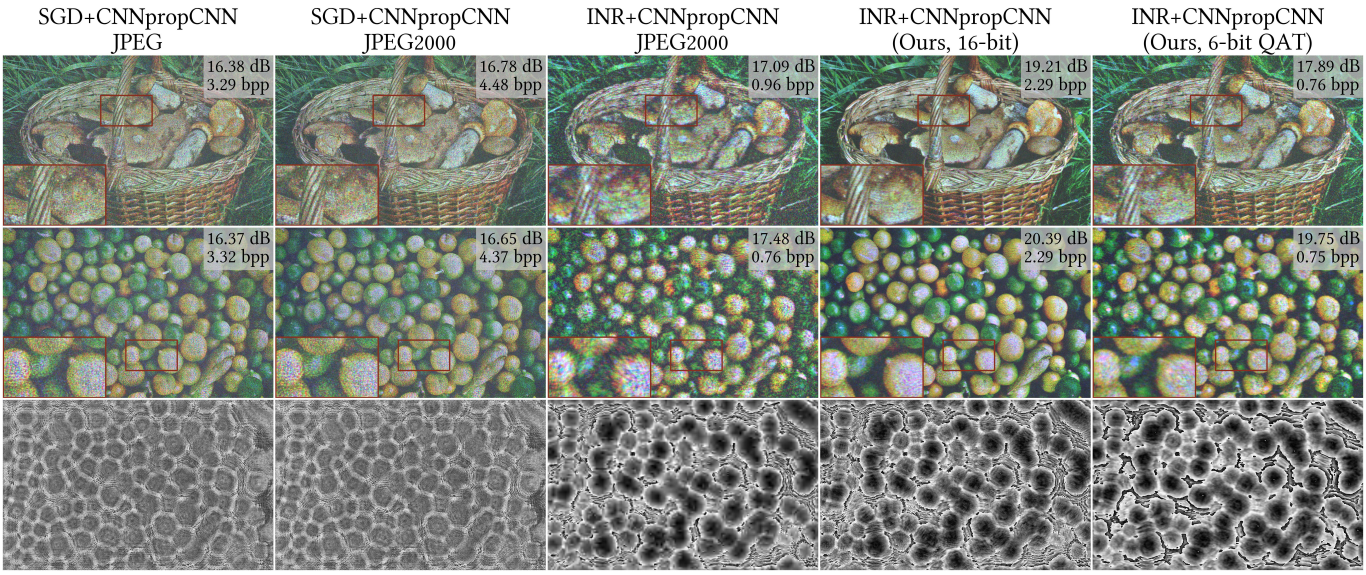

Fig. 8: Experimental reconstructions of 2D images using various compressed CGH pipelines: SGD w/ CNNpropCNN compressed via JPEG, SGD w/ CNNpropCNN compressed via JPEG2000, INR w/ CNNpropCNN compressed via JPEG2000, INR w/ CNNpropCNN using a 16-bit float representation, and INR w/ CNNpropCNN and a 6-bit quantization-aware training. We present the corresponding phase-only holograms (green channel-only) below.

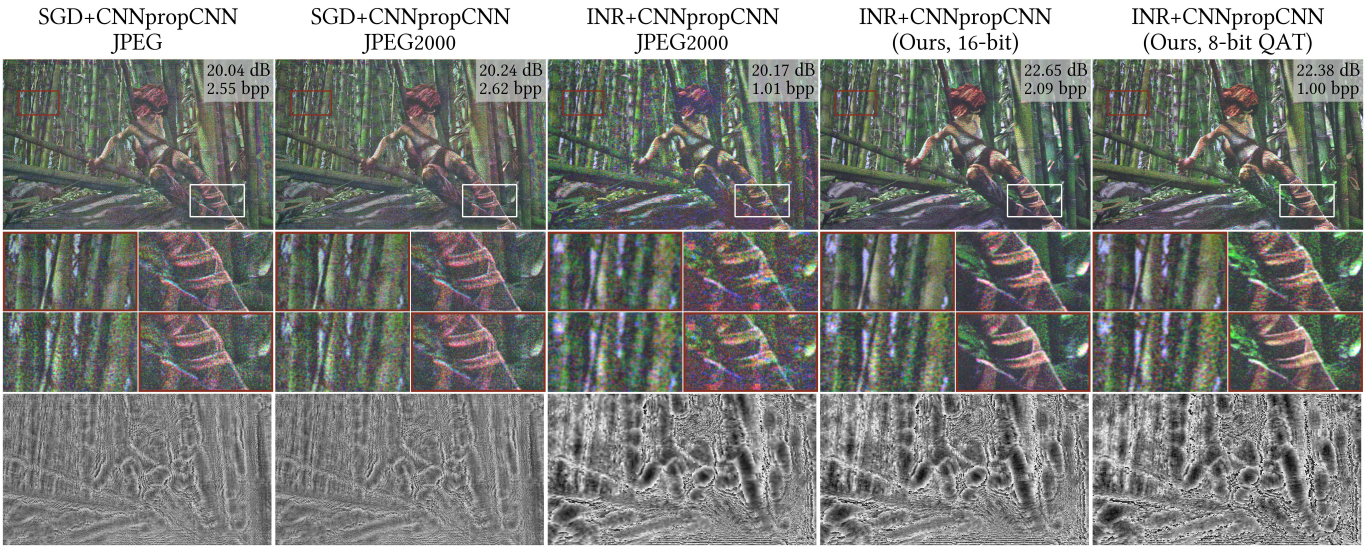

Fig. 9: Experimental reconstructions of 3D (RGBD) scenes using various compressed CGH pipelines: SGD w/ CNNpropCNN compressed via JPEG, SGD w/ CNNpropCNN compressed via JPEG2000, INR w/ CNNpropCNN compressed via JPEG2000, INR w/ CNNpropCNN using a 16-bit float representation, and INR w/ CNNpropCNN and an 8-bit quantization-aware training. Zoomed-in patches highlight the reconstruction at near and far depth planes. We present the corresponding phase-only holograms (green channel-only) below.

employed to overfit natural images due to their capacity in modeling continuous signals with sharp transitions, holographic data possess unique properties, notably structured high-frequency interference fringes. To assess the choice of activation functions for our CGH pipeline, we compare SIREN and FINER activations under identical training settings with 2,000 iterations. Both networks are trained to overfit phase-only holograms generated from a subset of the DIV2K dataset. As illustrated in Fig. 3, the FINER-based INR yields visibly sharper reconstructions and more spatially coherent phases, i.e., more "randomly-looking". Notably, spectral analysis of reconstructions indicates that FINER retains broader frequency content, which are often acknowledged as critical for holography [37, 55, 62] to preserve and deliver sufficient spectrum in 3D space. In contrast, SIREN exhibits energy concentration near the DC component and lacks support in higher spatial frequencies. These preliminary results suggest that FINER's adaptive periodicity may facilitate better alignment with CGH.

Comparison with CNN-based Hologram Generation. We compare our INR-based approach against two representative CNN-based CGH pipelines, HoloNet [49], which directly generates holograms, and DPRC [66], a joint generation and compression framework. As shown in Fig. 10, both CNN-based methods perform reasonably well in filtered simulations; However, their captured results exhibit strong artifacts and structured distortions without a physical filter. In unfiltered conditions, both methods experience notable degradation in both simulation and experiment. In contrast, our INR-CGH consistently maintains performance across all scenarios, exhibiting minimal simulation-to-capture discrepancy while delivering visually coherent results even in an unfiltered setup. This implies the robustness of our representation against optical imperfections and its adaptability to practical display scenarios.

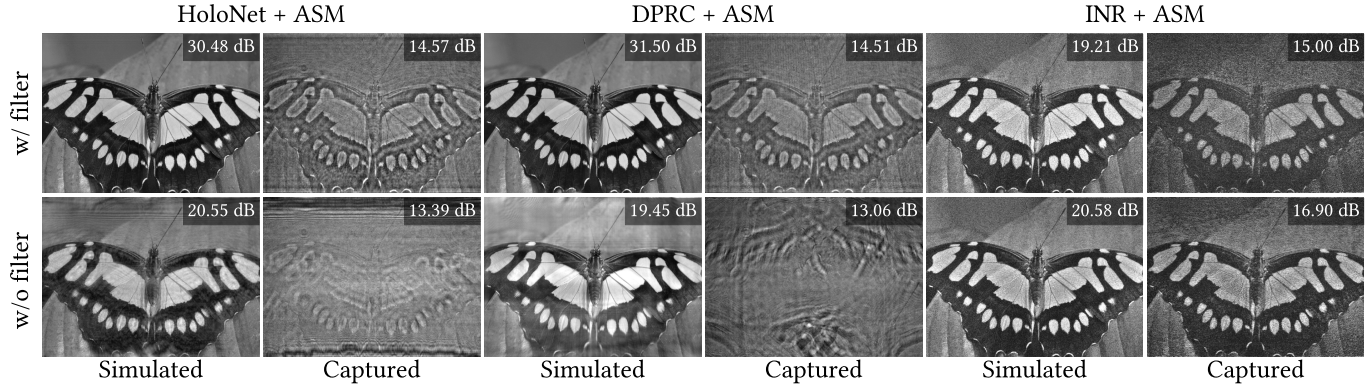

Fig. 10: Qualitative comparison for neural network-based CGH methods, including HoloNet, DPRC w/o compression, and our INR-CGH, each evaluated under two conditions: phase-only holograms generated with (top) and without (bottom) a Fourier aperture filter in simulation using ASM propagation. Within each group, the left column shows simulated reconstructions, while the right shows real captures with the same laser power and exposure time *without* a conventional $4f$ system, directly representing an unfiltered, compact display setup. Accordingly, the "with filter" simulation represents a reference configured with conventional Fourier aperture filtering. Note, we present green-channel acquisitions and visualize them as grayscale images.

Color-channel Joint Compression. At this proof-of-concept stage, our network generates a single-channel phase hologram per INR model. Nevertheless, this design overlooks the inter-channel redundancy that often exists in holographic scenes. Exploring jointly modeling RGB channels within a single INR can potentially reduce parameter count and improve compression efficiency [18]. As a preliminary test, we visualize holograms generated using a shared INR that regresses RGB phase channels simultaneously. While the model achieves moderate reconstruction quality, subtle color fringing indicates the need for more advanced coordination between channels, which we leave for future work. Refer to Supplementary Sec. S6 for detailed exploration in this domain.

6 Conclusion

In this work, we have explored an INR-empowered image processing framework tailored for unfiltered 3D hologram generation and compression to capture the high-frequency demands of hologram data. By integrating phase-only implicit representation with a multi-depth, camera-calibrated, learned propagation model, our method overcomes the typical simulation-to-display mismatch that likely affects CNN-based approaches without optical filtering. Furthermore, we compress the learned network parameters using quantization-aware training and entropy coding, enabling both efficient storage and fast decoding on edge devices. Experimental results indicate that our INR-CGH paradigm is a step towards practical holographic displays in pending cloud-edge computing scenarios, facilitating a good compromise between image fidelity and algorithm employability.

Follow-up Work. While our INR-CGH framework achieves efficient, aka., real-time, decoding, further improvements in encoding speed remain necessary. This can potentially leverage progressive training techniques or meta-learning strategies [23, 60], as well as the field-programmable gate arrays (FPGA) to accelerate the complex hologram generation process on edge devices. Another practical extension to improve performance and/or reduce model size with minimal overhead is using clustering-based bin placement (e.g., via k-means [27]), which may better match the Gaussian distribution of weights and reduce quantization error. In addition, as our framework is demonstrated to be compatible with *unfiltered* holographic display setups, integrating high-order gradient descent [25] algorithms can bootstrap image contrast. With respect to the 3D imaging capability, the optimization-based refinement using alternating direction method of multipliers [15] could be incorporated to regularize the in-focus/out-of-focus regions, potentially improving perceptual consistency in multi-depth holography. Last but not least, a natural extension of this work is to explore hardware-aware phase design objectives that can regulate phase complexity, for instance, random-phase-like characteristics [69], while accounting for current SLM imperfections such as pixel crosstalk. A systematic study of alternative random-phase formulations under these constraints remains an important research direction.

Acknowledgments

This research was supported by the National Science Foundation of China (62322217), the Research Grants Council of Hong Kong (GRF 17208023), and the Innovation and Technology Fund of Hong Kong (ITP/062/24AP and MHP/313/24).

References

[1] E. Agustsson and R. Timofte. 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. 4, 5

[2] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018. 2

[3] H. Ban, S. Choi, J. Y. Cha, Y. Kim, and H. Y. Kim. Nhvc: Neural holographic video compression with scalable architecture. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pp. 969–978. IEEE, 2024. 2

[4] Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 4

[5] M. V. Bernardo, P. Fernandes, A. Arrifano, M. Antonini, E. Fonseca, P. T. Fiadeiro, A. M. Pinheiro, and M. Pereira. Holographic representation: Hologram plane vs. object plane. *Signal Processing: Image Communication*, 68:193–206, 2018. 2

[6] Y. Bhalgat, J. Lee, M. Nagel, T. Blankevoort, and N. Kwak. 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. 4

[7] D. Blinder, T. Birnbaum, T. Ito, and T. Shimobaba. The state-of-the-art in computer generated holography for 3d display. *Light: Advanced Manufacturing*, 3(3):572–600, 2022. 1, 2

[8] L. Catania and D. Allegra. Nif: A fast implicit image compression with bottleneck layers and modulated sinusoidal activations. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 9022–9031, 2023. 2

[9] P. Chakravarthula, Y. Peng, J. Kollin, H. Fuchs, and F. Heide. Wirtinger holography for near-eye displays. *ACM Transactions on Graphics (TOG)*, 38(6):1–13, 2019. 1

[10] P. Chakravarthula, E. Tseng, T. Srivastava, H. Fuchs, and F. Heide. Learned hardware-in-the-loop phase retrieval for holographic near-eye displays. *ACM Transactions on Graphics (TOG)*, 39(6):1–18, 2020. 3

[11] H. Chen, B. He, H. Wang, Y. Ren, S. N. Lim, and A. Shrivastava. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 34:21557–21568, 2021. 2

[12] K. Chen, A. Wen, Y. Zhang, P. Chakravarthula, and Q. Sun. View synthesis for 3d computer-generated holograms using deep neural fields. *Optics Express*, 33(9):19399–19408, 2025. 2

[13] S. Choi, B. Chao, J. Yang, M. Gopakumar, and G. Wetzstein. Gaussian wave splatting for computer-generated holography. *ACM Transactions on Graphics (TOG)*, 44(4):1–13, 2025. 2

[14] S. Choi, M. Gopakumar, Y. Peng, J. Kim, M. O'Toole, and G. Wetzstein. Time-multiplexed neural holography: a flexible framework for holographic near-eye displays with fast heavily-quantized spatial light modulators. In *ACM SIGGRAPH 2022 Conference Proceedings*, pp. 1–9, 2022. 3

[15] S. Choi, M. Gopakumar, Y. Peng, J. Kim, and G. Wetzstein. Neural 3d holography: Learning accurate wave propagation models for 3d holographic virtual and augmented reality displays. *ACM Transactions on Graphics (TOG)*, 40(6):1–12, 2021. 3, 4, 8

[16] S. Choi, C. Jang, D. Lanman, and G. Wetzstein. Synthetic aperture waveguide holography for compact mixed-reality displays with large étendue. *Nature Photonics*, 19(8):854–863, 2025. 2

[17] C. Christopoulos, A. Skodras, and T. Ebrahimi. The jpeg2000 still image coding system: an overview. *IEEE transactions on consumer electronics*, 46(4):1103–1127, 2000. 2, 5

[18] K. Cui, A. Boev, E. Alshina, and E. Steinbach. Color image restoration exploiting inter-channel correlation with a 3-stage cnn. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):174–189, 2020. 8

[19] J. Duda. Asymmetric numeral systems: entropy coding combining speed of huffman coding with compression rate of arithmetic coding. *arXiv preprint arXiv:1311.2540*, 2013. 4

[20] E. Dupont, A. Goliński, M. Alizadeh, Y. W. Teh, and A. Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 2

[21] E. Dupont, H. Loya, M. Alizadeh, A. Goliński, Y. W. Teh, and A. Doucet. Coin++: Neural compression across modalities. *arXiv preprint arXiv:2201.12904*, 2022. 2

[22] S. K. Esser, J. L. McKinstry, D. Bablani, R. Appuswamy, and D. S. Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019. 4

[23] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017. 8

[24] J. W. Goodman. *Introduction to Fourier optics*. Roberts and Company publishers, 2005. 3

[25] M. Gopakumar, J. Kim, S. Choi, Y. Peng, and G. Wetzstein. Unfiltered holography: optimizing high diffraction orders without optical filtering for compact holographic displays. *Optics letters*, 46(23):5822–5825, 2021. 8

[26] M. Gopakumar, G.-Y. Lee, S. Choi, B. Chao, Y. Peng, J. Kim, and G. Wetzstein. Full-colour 3d holographic augmented-reality displays with metasurface waveguides. *Nature*, pp. 1–7, 2024. 1, 2

[27] C. Gordon, S.-F. Chng, L. MacDonald, and S. Lucey. On quantizing implicit neural representations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 341–350, 2023. 8

[28] Z. Guo, G. Flamich, J. He, Z. Chen, and J. M. Hernández-Lobato. Compression with bayesian implicit neural representations. *Advances in Neural Information Processing Systems*, 36:1938–1956, 2023. 2

[29] D. He, Z. Yang, W. Peng, R. Ma, H. Qin, and Y. Wang. 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. 2

[30] R. Horisaki, R. Takagi, and J. Tanida. Deep-learning-generated holography. *Applied optics*, 57(14):3859–3863, 2018. 2

[31] M. Hossein Eybposh, N. W. Caira, M. Atisa, P. Chakravarthula, and N. C. Pégard. Deepcgh: 3d computer-generated holography using deep learning. *Optics Express*, 28(18):26636–26650, 2020. 2

[32] D. Kim, S.-W. Nam, B. Lee, J.-M. Seo, and B. Lee. Accommodative holography: improving accommodation response for perceptually realistic holographic displays. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 1

[33] H. Kim, M. Bauer, L. Theis, J. R. Schwarz, and E. Dupont. C3: High-performance and low-complexity neural compression from a single image or video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9347–9358, 2024. 2

[34] J. Kim, M. Gopakumar, S. Choi, Y. Peng, W. Lopes, and G. Wetzstein. Holographic glasses for virtual reality. In *ACM SIGGRAPH 2022 Conference Proceedings*, pp. 1–9, 2022. 1

[35] R. Kizhakkumkara Muhamad, T. Birnbaum, D. Blinder, and P. Schelkens. Interfere, short-time fourier-transform-based compression of complex-valued holograms with bit depth and range-adaptive quantization. *Applied Optics*, 63(25):6740–6753, 2024. 2

[36] R. Kizhakkumkara Muhamad, C. Schretter, D. Blinder, and P. Schelkens. Autoregressive modeling for lossless compression of holograms. *Optics Express*, 31(23):38589–38609, 2023. 2

[37] G. Kuo, F. Schiffers, D. Lanman, O. Cossairt, and N. Matsuda. Multisource holography. *ACM Transactions on Graphics (Tog)*, 42(6):1–14, 2023. 7

[38] G. Kuo, L. Waller, R. Ng, and A. Maimone. High resolution étendue expansion for holographic displays. *ACM Transactions on Graphics (TOG)*, 39(4):66–1, 2020. 1

[39] T. Ladune, P. Philippe, F. Henry, G. Clare, and T. Leguay. Cool-chic: Coordinate-based low complexity hierarchical image codec. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13515–13522, 2023. 2

[40] J. Lee, J. Jeong, J. Cho, D. Yoo, B. Lee, and B. Lee. Deep neural network for multi-depth hologram generation and its training strategy. *Optics Express*, 28(18):27137–27154, 2020. 2

[41] Z. Liu, H. Zhu, Q. Zhang, J. Fu, W. Deng, Z. Ma, Y. Guo, and X. Cao. Finer: Flexible spectral-bias tuning in implicit neural representation by variable-periodic activation functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2713–2722, 2024. 2, 3

[42] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief. A survey on mobile edge computing: The communication perspective. *IEEE communications surveys & tutorials*, 19(4):2322–2358, 2017. 1

[43] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2

[44] S. Monin, A. C. Sankaranarayanan, and A. Levin. Analyzing phase masks for wide étendue holographic displays. In *2022 IEEE International Conference on Computational Photography (ICCP)*, pp. 1–12. IEEE, 2022. 1

[45] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. 2

[46] K.-J. Oh, H. Ban, S. Choi, H. Ko, and H. Y. Kim. Hevc extension for phase hologram compression. *Optics Express*, 31(6):9146–9164, 2023. 2

[47] N. Padmanaban, Y. Peng, and G. Wetzstein. Holographic near-eye displays based on overlap-add stereograms. *ACM Transactions on Graphics (TOG)*, 38(6):1–13, 2019. 4

[48] Y. Peng, S. Choi, J. Kim, and G. Wetzstein. Speckle-free holography with partially coherent light sources and camera-in-the-loop calibration. *Science advances*, 7(46):eabg5040, 2021. 3

[49] Y. Peng, S. Choi, N. Padmanaban, and G. Wetzstein. Neural holography with camera-in-the-loop training. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 1, 2, 4, 5, 7

[50] V. Saragadam, D. LeJeune, J. Tan, G. Balakrishnan, A. Veeraraghavan, and R. G. Baraniuk. Wire: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18507–18516, 2023. 2

[51] P. Schelkens, T. Ebrahimi, A. Gilles, P. Gioia, K.-J. Oh, F. Pereira, C. Perra, and A. M. Pinheiro. Jpeg pleno: Providing representation interoperability for holographic applications and devices. *ETRI journal*, 41(1):93–108, 2019. 2

[52] L. Shi, B. Li, C. Kim, P. Kellnhofer, and W. Matusik. Towards real-time photorealistic 3d holography with deep neural networks. *Nature*, 591(7849):234–239, 2021. 1, 2

[53] L. Shi, B. Li, and W. Matusik. End-to-end learning of 3d phase-only holograms for holographic display. *Light: Science & Applications*, 11(1):247, 2022. 2

[54] L. Shi, R. Webb, L. Xiao, C. Kim, and C. Jang. Neural compression for hologram images and videos. *Optics Letters*, 47(22):6013–6016, 2022. 1, 2

[55] T. Shimobaba and T. Ito. Random phase-free computer-generated hologram. *Optics express*, 23(7):9549–9554, 2015. 7

[56] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020. 2, 4

[57] Y. Strümpler, J. Postels, R. Yang, L. V. Gool, and F. Tombari. Implicit neural representations for image compression. In *European Conference on Computer Vision*, pp. 74–91. Springer, 2022. 2

[58] K. Su, M. Chen, and E. Shlizerman. Inras: Implicit neural representation for audio scenes. *Advances in Neural Information Processing Systems*,

35:8144–8158, 2022. 2

[59] X. Sui, Z. He, G. Jin, and L. Cao. Spectral-envelope modulated double-phase method for computer-generated holography. *Optics Express*, 30(17):30552–30563, 2022. 2

[60] M. Tancik, B. Mildenhall, T. Wang, D. Schmidt, P. P. Srinivasan, J. T. Barron, and R. Ng. Learned initializations for optimizing coordinate-based neural representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2846–2855, 2021. 8

[61] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020. 2

[62] E. Tseng, G. Kuo, S.-H. Baek, N. Matsuda, A. Maimone, F. Schiffers, P. Chakravarthula, Q. Fu, W. Heidrich, D. Lanman, et al. Neural étendue expander for ultra-wide-angle high-fidelity holographic display. *Nature communications*, 15(1):2907, 2024. 7

[63] K. Vaidyanathan, M. Salvi, B. Wronski, T. Akenine-Möller, P. Ebelin, and A. Lefohn. Random-access neural compression of material textures. *arXiv preprint arXiv:2305.17105*, 2023. 2

[64] G. K. Wallace. The jpeg still picture compression standard. *Communications of the ACM*, 34(4):30–44, 1991. 5

[65] G. K. Wallace. The jpeg still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992. 2

[66] Y. Wang, P. Chakravarthula, Q. Sun, and B. Chen. Joint neural phase retrieval and compression for energy-and computation-efficient holography on the edge. *ACM Transactions on Graphics*, 41(4), 2022. 2, 7

[67] J. Wu, K. Liu, X. Sui, and L. Cao. High-speed computer-generated holography using an autoencoder-based deep neural network. *Optics Letters*, 46(12):2908–2911, 2021. 1

[68] D. Yang, W. Seo, H. Yu, S. I. Kim, B. Shin, C.-K. Lee, S. Moon, J. An, J.-Y. Hong, G. Sung, et al. Diffraction-engineered holography: Beyond the depth representation limit of holographic displays. *Nature Communications*, 13(1):6012, 2022. 1

[69] D. Yoo, Y. Jo, S.-W. Nam, C. Chen, and B. Lee. Optimization of computer-generated holograms featuring phase randomness control. *Optics Letters*, 46(19):4769–4772, 2021. 8

[70] C. Zhong, X. Sang, B. Yan, H. Li, X. Xie, X. Qin, and S. Chen. Real-time 4k computer-generated hologram based on encoding conventional neural network with learned layered phase. *Scientific Reports*, 13(1):19372, 2023. 1

[71] M. Zhou, H. Zhang, S. Jiao, P. Chakravarthula, and Z. Geng. End-to-end compression-aware computer-generated holography. *Optics Express*, 31(26):43908–43919, 2023. 2

[72] W. Zhou, F. Qu, X. Meng, Z. Li, and Y. Peng. 3d-holonet: fast, unfiltered, 3d hologram generation with camera-calibrated network learning. *Optics Letters*, 50(4):1188–1191, 2025. 3