













Ultra-efficient physical field computing by complex-valued network quantization

Received: 14 January 2025

Accepted: 25 February 2026

Published online: 10 March 2026

 Check for updates

Zihan Geng ^{1,10} , Zhilin Li^{1,10}, Mi Zhou^{1,10}, Zijie Chen ¹, Ganzhangqin Yuan¹, Jaebum Noh², Seokho Lee ³, Cherry Park³, Hongya Geng ¹, Xiu Li ¹, Kui Jiang ^{4,5} , Mu Ku Chen ⁶  & Junsuk Rho ^{3,7,8,9} 

Neural network quantization is an established technique for compressing real-valued models, but its application to complex-valued networks—essential in electromagnetics, acoustics, and quantum physics—remains underdeveloped. Conventional quantization methods treat real and imaginary components as independent channels, thereby disrupting the algebraic structure of complex multiplication and distorting essential phase relationships. To address this problem, we propose a real-imaginary joint quantization method that minimizes error propagation in complex multiplication and maintains coherence in phase-sensitive tasks, thereby preserving amplitude-phase fidelity during complex-valued inference. Combined with physics-aware adaptive precision training, this approach demonstrates outstanding performance across hologram generation, audio classification, wireless signal classification, and synthetic aperture radar signal recognition tasks. Compared to the state-of-the-art hologram generation model HoloNet, our approach achieves a 3.9 dB improvement in peak signal-to-noise ratio while reducing computational load and memory consumption by 99.1% and 99.8%, respectively. This research establishes a pathway toward lightweight, high-fidelity complex-valued neural networks for scientific computing and coherent signal processing.

Neural network quantization is a technique for model compression through bit-width reduction of network parameters. It offers substantial memory footprint reduction, inference acceleration, and energy saving with minimal performance degradation¹. Its efficacy has attracted numerous researchers to innovate in architecture and training strategies, such as distribution rectification distillation², imbalanced quantization³, and zero-shot quantization⁴. Extensive research and engineering practice have demonstrated the

effectiveness of neural network quantization in accelerating real-valued tasks and compressing real-valued networks^{5,6}. However, for tasks involving complex-valued physical fields, the direct application of existing real-valued quantization operations often leads to a significant degradation in the quality of physical field reconstruction^{7–9}.

Complex-valued neural networks possess strong capabilities in fitting complex physical fields, serving as potent tools for resolving intricate physical problems. They align closely with the mathematical

¹Institute of Data and Information, Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen, Guangdong, China. ²Department of Optical Engineering, Kumoh National Institute of Technology, Gyeongsbuk, Republic of Korea. ³Department of Mechanical Engineering, Pohang University of Science and Technology (POSTECH), Pohang, Republic of Korea. ⁴School of Computer Science and Technology, Harbin Institute of Technology, Harbin, China. ⁵Harbin Institute of Technology Zhengzhou Research Institute, Zhengzhou, China. ⁶Department of Electrical Engineering, City University of Hong Kong, Hong Kong, China. ⁷Department of Chemical Engineering, Pohang University of Science and Technology (POSTECH), Pohang, Republic of Korea. ⁸Department of Electrical Engineering, Pohang University of Science and Technology (POSTECH), Pohang, Republic of Korea. ⁹POSTCO-POSTECH-RIST Convergence Research Center for Flat Optics and Metaphotonics, Pohang, Republic of Korea. ¹⁰These authors contributed equally: Zihan Geng, Zhilin Li, Mi Zhou. ✉ e-mail: geng.zihan@sz.tsinghua.edu.cn; jiangkui@hit.edu.cn; mkchen@cityu.edu.hk; jsrho@postech.ac.kr

essence of complex-valued physical fields. The complex number itself integrates the key information of a physical field—amplitude and phase—into a unified format. It encodes the relationship between phase and amplitude, thereby enabling the description of phenomena such as wave superposition and interference in electromagnetic waves, sound waves, quantum mechanics, and seismic waves^{10–15}. Complex values possess elegant differential properties, allowing differential equation problems to be transformed into simple algebraic equations, thereby effectively reducing computational costs. On the other hand, the exponential function form of complex numbers systematically captures oscillatory and periodic behaviors. Waves, vibrations, and modes within physical fields all exhibit periodicity, rendering their representation through complex numbers both highly natural and compact. For example, in fluid dynamics, complex-valued methods are highly valuable for the spectral analysis and boundary analysis of turbulence, laminar flow, and vortices^{16–18}. In thermodynamics, complex numbers simplify the periodic behaviors in heat conduction problems into algebraic operations¹⁹.

Although quantization for real-valued neural networks is a mature field, it is not directly applicable to complex-valued networks. Prevailing approaches treat the real and imaginary components as independent channels for quantization, a method we term independent component quantization. This paradigm, while straightforward, is mathematically suboptimal for coherent systems. It ignores the algebraic coupling between real and imaginary components, leading to uncorrelated quantization errors that severely disrupt the phase relationship during multiplication. Consequently, it introduces non-physical noise that degrades amplitude and—more critically—phase fidelity, resulting in artifacts that limit model utility in phase-sensitive tasks like holography or synthetic aperture radar (SAR) imaging. This limitation stems from a fundamental divergence in how physical information is mathematically represented.

To bridge this gap, we propose a foundational rethinking of quantization for complex-valued neural networks. We introduce a universal framework that respects the mathematical properties of complex operations, ensuring quantization noise is structured to preserve phase coherence. As shown in Fig. 1, our core theoretical innovation is a joint real-imaginary quantization scheme that explicitly models the error propagation in complex multiplication, minimizing corruption of the resultant vector angle and magnitude. To further enhance efficiency, the proposed quantization is augmented with an adaptive-precision strategy. We incorporate an adaptive-precision

mechanism that dynamically allocates layer-wise bit-widths, guided by the sensitivity of phase and amplitude inaccuracies to the physical task-specific objective.

The contributions of this work are threefold. (1) We formally reveal the limitation of independent component quantization and establish the necessity for algebraically consistent quantization in complex-valued deep learning. (2) We devise a holistic quantization framework that jointly optimizes real and imaginary parts for amplitude and phase fidelity, ensuring the quantized network remains a physically valid approximator. (3) We demonstrate that this framework enables unprecedented efficiency while preserving performance. Our method generalizes across diverse domains—hologram generation, audio classification, wireless signal recognition, and SAR processing—achieving superior accuracy while reducing computational load and memory footprint. Notably, we demonstrate inference speedup on a mobile device, proving the practical feasibility of deploying high-fidelity, complex-valued scientific models at the edge.

By establishing a principled approach to complex-valued network compression, this work enables the development of lightweight AI models for computationally demanding scientific fields, from electromagnetics and thermodynamics to quantum physics.

Results

Complex-valued mixed-precision quantization

Our complex-valued mixed-precision quantization method comprises two training stages: identifying the optimal quantization bit width for each network layer and quantizing the network to the optimal bit widths, as shown in Fig. 2.

Training stage 1: Identifying the optimal quantization bit width of the parameters in each network layer. As illustrated in Fig. 2, this stage aims to determine the best quantization bit width for the real and imaginary parts of activations and weights (a_{real} , a_{imag} , w_{real} , and w_{imag}) of each layer. In this study, full precision is employed for the input and output layers. For the intermediate layers, we determine the optimal bit width within the range of 1 to 4 bits. We initialize four probabilities corresponding to selecting 1 to 4 bits, represented by learnable parameters (p_0 , p_1 , p_2 , and p_3). During training, these probabilities are continuously adjusted via gradient descent to minimize the network's loss, which comprises task loss and complexity loss. Task loss assesses the quality of the quantized images, while complexity loss regulates the quantization bit width. Task loss encourages the selection of higher bit widths, whereas complexity loss favors lower bit widths. By

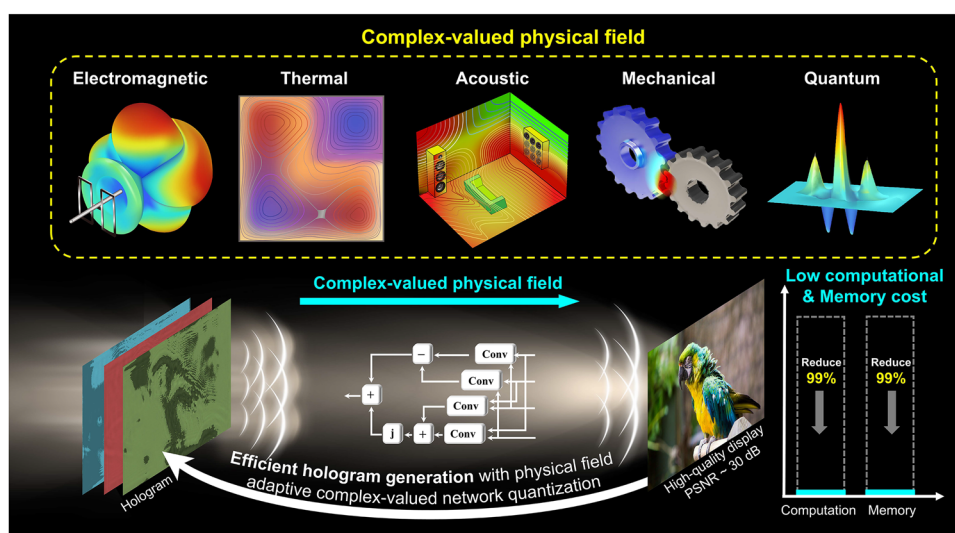


Fig. 1 | Overview. The proposed complex-valued network quantization for ultra-efficient physical field computing. The color image, “Green and yellow macaw bird” on the right-hand side, is originally posted to Unsplash by Andrew Li at https://unsplash.com/photos/gold-and-blue-macaw-on-brown-wooden-stick-iLVaLbfe9_g.

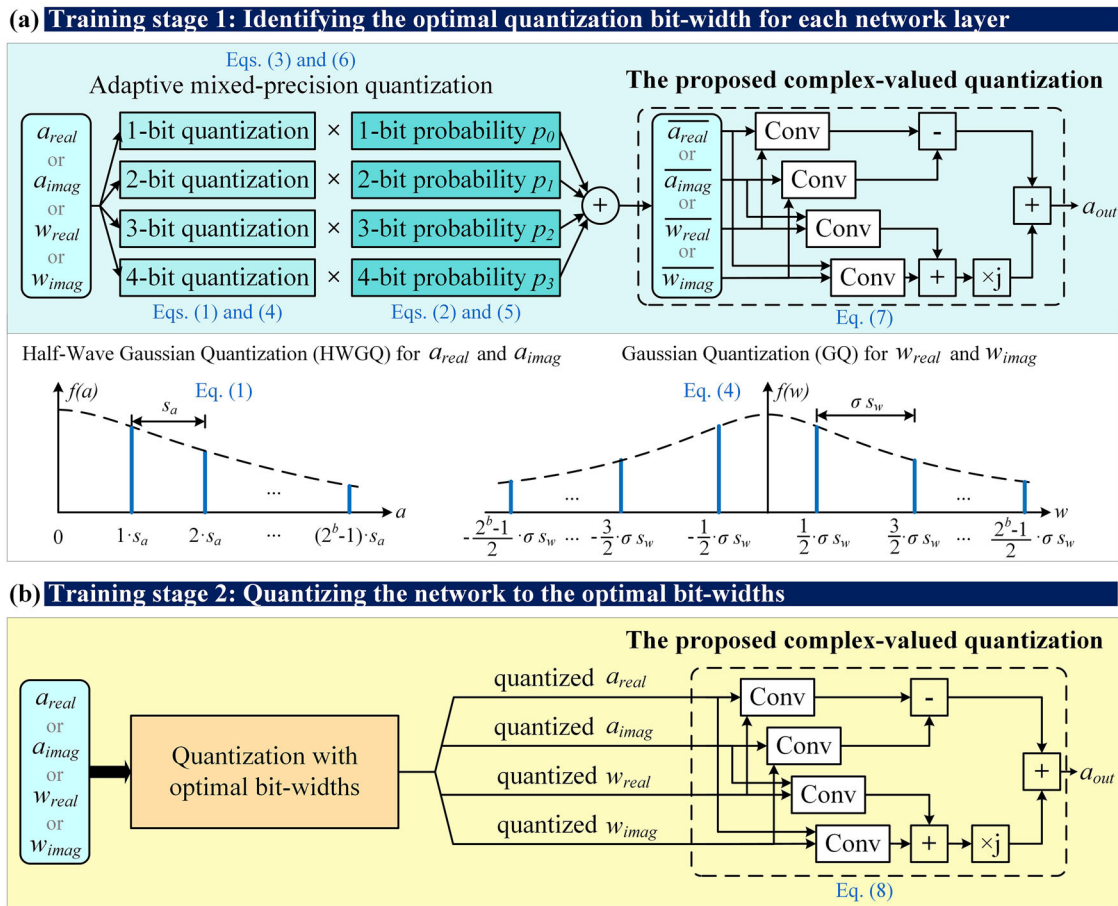


Fig. 2 | The training procedure of the proposed complex-valued mixed-precision quantization block. Our training strategy comprises two stages. **a** Training stage 1: Identifying the optimal quantization bit width of the parameters in each network layer. **b** Training stage 2: Quantizing the network to the optimal bit widths. $a_{real}, a_{imag}, w_{real},$ and w_{imag} are the real and imaginary parts of the input activations

and weights. $a_{real}^-, a_{imag}^-, w_{real}^-,$ and w_{imag}^- are the expectation of the quantized activations and weights. a_{out} is the output of the complex-valued mixed-precision block. s_a and s_w are the quantization spacing for activations and weights. σ is the variance of the distribution of w . $f(a)$ and $f(w)$ are the probability distributions of activations and weights. Conv denotes the real-valued convolution operator.

minimizing the network loss, we can achieve a balance between image quality and network complexity. After this training stage, the values of $(p_0, p_1, p_2,$ and $p_3)$ for each parameter are learnt. Among the four probabilities of one network parameter, the highest probability p_n indicates the optimal quantization bit width for this parameter is $n + 1$.

Training stage 2: Quantizing the network to the optimal bit widths. In the previous stage, the optimal quantization bit widths for each layer's parameters $a_{real}, a_{imag}, w_{real},$ and w_{imag} are identified. In this second stage, the network is reconstructed and retrained with these optimal bit widths. The network loss is minimized by tuning the network weights, w_{real} and w_{imag} . Here, the network loss only includes task loss to ensure high-quality image reconstruction. Ultimately, the quantized lightweight network is capable of generating high-quality holograms.

Complex-valued mixed-precision network for hologram generation

We take computer-generated holography (CGH)^{20–28} as a representative task for detailed methodological discussion and comparison. Hologram generation is highly sensitive to computational errors^{29–33}, making it a suitable benchmark for evaluating the impact of quantization errors on model performance. Furthermore, due to the complexity of inverse complex-valued light field reconstruction and the conversion from complex-valued fields to phase-only encoding, hologram generation is generally a computationally intensive task^{34–39}. Thus, it is necessary to develop corresponding compression

techniques to enable efficient execution on resource-constrained devices. In this work, we propose a light-field-aware ultra-low bit network (ULBN) with complex-valued mixed-precision quantization. This network design has two features:

Light-field-aware mixed-precision quantization: In holography, small perturbations, such as quantization noise, in the complex field can propagate and amplify during wavefront reconstruction. To mitigate this problem, we retain full-precision representations in shallow layers that directly interact with the light field, where preserving fine-grained amplitude and phase information is critical. As the signal moves deeper into the network, where operations become more abstract and less sensitive to high-frequency detail, we apply lower precision using an adaptive mixed-precision scheme. This sensitivity-aware bit allocation improves efficiency without sacrificing reconstruction quality.

Light field-aware loss function: Our reconstruction quality loss is computed on the complex field after light propagation. It implies that any quantization errors introduced during forward propagation are penalized according to their impact on the reconstructed light field, thereby establishing an optical feedback mechanism within the optimization process.

The proposed network architecture is depicted in Fig. 3. Our light-field-aware ULBN takes the amplitude of the target image as input. The output is a phase-only hologram (POH). The ULBN is composed of three subnetworks, namely a phase generator, a POH encoder, and a ringing artifacts compensator. The phase generator network takes the

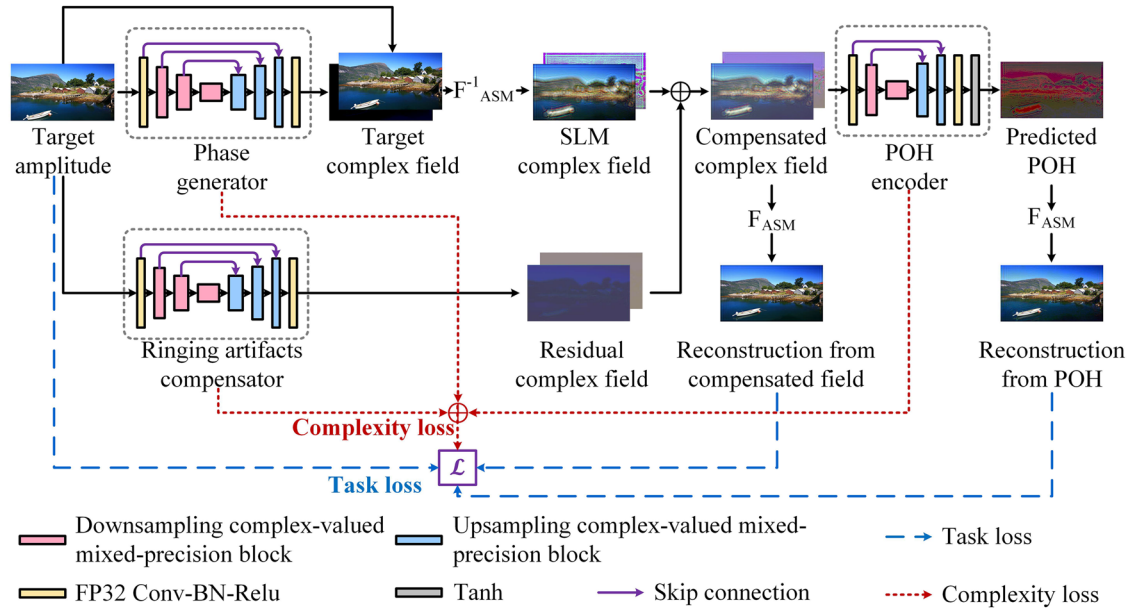


Fig. 3 | Structural diagram of our light-field-aware ultra-low bit network. The proposed model consists of three subnetworks, a phase generator, a phase-only hologram (POH) encoder, and a ringing artifacts compensator. The subnetworks have downsampling and upsampling complex-valued mixed-precision blocks.

F_{ASM}^{-1} and F_{ASM} correspond to opposite propagation distances and propagate optical fields using the ASM. The Target amplitude image is originally posted to Flickr by laszlo-photo at <https://www.flickr.com/photos/40467171@N00/4972189987>.

target amplitude distribution as input and predicts the phase of the target field. The resultant target complex field propagates backward from the target display plane to the spatial light modulator (SLM) plane via the angular spectrum method (ASM). The ringing artifacts compensator is proposed in our previous work⁴⁰. It is a plug-in neural network model that effectively reduces the ringing artifacts caused by the modeling error in the widely used forward and backward propagation CGH methods. The module generates a residual complex field to compensate for the two-fold diffraction propagation modeling error. The calculated residual complex field is added to the complex field at the SLM plane to generate the compensated complex field. The POH encoder transfers the complex-valued light field into a POH. The compensated complex light field and the predicted POH at the modulator plane propagate forward individually via the ASM method, generating two reconstructed amplitudes A_{rec_compen} and A_{rec_poh} at the target display plane.

Simulation results

The proposed ULBN network, U-Net³⁵, Holo-encoder³⁶, and HoloNet³⁷ are evaluated on the DIV2K dataset⁴¹ with bit operations, memory, peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), learned perceptual image patch similarity (LPIPS), natural image quality evaluator (NIQE), floating point operations (FLOPs), and parameters (Params) in Fig. 4a–h, respectively. Because our method effectively compresses the complex-valued model to 1–4 bits, it achieves substantially lower computational cost and memory usage than other CGH algorithms. Our ULBN requires only 88.352 gigabit operations to generate the phase data of the holography experiment, which is 99.1% less than HoloNet³⁷. The memory footprint is just 28.433 kilobytes, representing a 99.8% reduction compared to HoloNet. Despite these minimal resource demands, our method delivers the best overall performance across all evaluation metrics. For PSNR and SSIM, higher values indicate better quality, whereas for LPIPS and NIQE, lower values reflect greater perceptual similarity. With a PSNR of 30.75 dB, our ULBN surpasses HoloNet by about 4 dB. For the perceptual metrics, including SSIM, LPIPS, and NIQE, our approach still outperforms all other algorithms, delivering a better visual experience

for viewers. Additional test results on DIV2K (Table S4), Kodak (Table S5), and Flickr2K (Table S6) datasets are provided in Section 4.3 of the Supplementary Information.

The visualized results are shown in Fig. 5a. Our ULBN produces the best reconstructions, free from speckle noise and artifacts across the entire image. Gerchberg–Saxton (GS) algorithm⁴² requires iterative optimization for each input. The results of U-Net, Holo-encoder, and HoloNet contain noticeable artifacts that degrade the overall visual quality. Our quantized complex-valued network achieves improvements in both computational efficiency and reconstruction quality.

Experimental results

We conduct comprehensive optical experiments to evaluate the practical performance of our proposed methods on a benchtop prototype. Figure 5b–d illustrate the visualized reconstructions captured with the optical setup in Fig. 5e and f.

The 2D holographic image simulation results are shown in Fig. 5a. For comparison, we demonstrate the relevant experimental results in Fig. 5b. We compare the captured color images of GS, U-Net, Holo-encoder, HoloNet, and our proposed ULBN. For the GS algorithm, the speckle noise exists in the optical reconstruction results, leading to noticeable image quality degradation. Our proposed ULBN method models the properties of complex values in its quantization process and adopts an adaptive complex-valued mixed-precision quantization strategy, effectively reducing light field computation quantization errors and resulting in smoother POHs. This approach significantly reduces the speckle noise problem in the reconstructions. Consequently, compared with the simulations, the optical reconstructions of our method show minimal degradation, and the speckle noise is mostly suppressed, ensuring higher quality and clearer images.

To validate the generalizability of our method across different datasets, we conducted experiments on a binary image dataset in Fig. 5d with its enlarged results in Section 4.4 of Supplementary Information and Fig. S2. The results demonstrate that our approach consistently achieves high-resolution binary images. The checkboard image displays distinct alternating dark and light stripes. The digits on

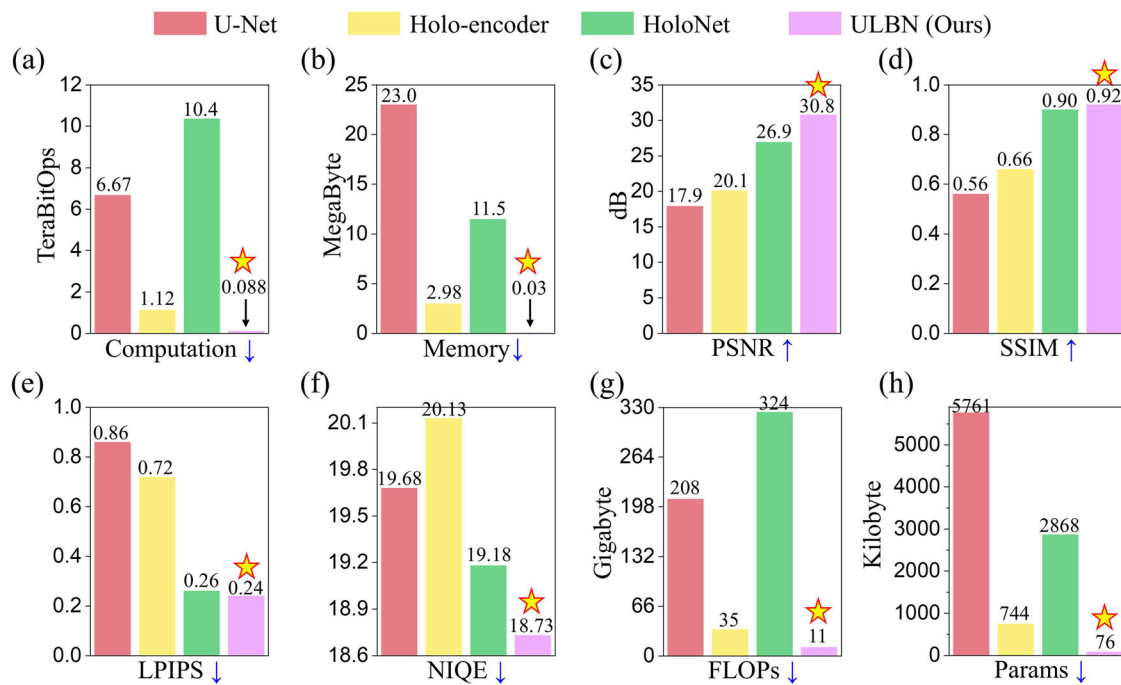


Fig. 4 | Performance and computational comparison of different methods. Comparison of U-Net, Holo-encoder, HoloNet, and our ULBN in terms of **a** bit operations, **b** memory, **c** PSNR, **d** SSIM, **e** LPIPS, **f** NIQE, **g** FLOPs, and **h** Params on the DIV2K dataset.

the Indian head image are clearly discernible. Both the numbers and stripes in the USAF-1951 image are distinctly visible.

To demonstrate the capabilities of our network in three-dimensional holographic displays, we present a 3D light-field-aware network architecture in the Supplementary Information Section 5.1 and Fig. S3. The 3D ULBN requires 104 giga FLOPs and 620 gigabits operations for computation, and 335 kilobytes for memory. As illustrated in Fig. 5c, our experimental results demonstrate the captured 3D high-quality images at different distances. Additional 3D holographic display results are presented in Figs. S4 and S5 of the Supplementary Information Sections 5.2 and 5.3. The camera is positioned at 19.5 cm, 20 cm, and 20.5 cm to capture pictures focusing on the hippopotamus, owl, and ring, respectively. The reconstructed images at varying distances demonstrate the ULBN's capabilities of generating 3D holograms.

Model deployment

The models are deployed on representative hardware platforms, including a high-performance desktop CPU and a resource-constrained Android phone. Given the inadequate support for extremely low bit-width mixed-precision quantization of complex-valued models in current deployment frameworks, we utilize standard uniform INT8 quantization to measure on-device latency of quantized models. Despite the lack of mixed-precision strategies, the results still affirm a key advantage: the effectiveness of our real-imaginary joint quantization approach. For the desktop evaluation, the INT8 models are produced using PyTorch post-training static quantization workflow (torch.ao.quantization) and executed in PyTorch 2.8 on an Intel Core i9-10980XE CPU. For the mobile evaluation, all the models are tested on an HONOR 70 smartphone with a Qualcomm Snapdragon 778 G Plus processor, 12 GB RAM, and Android 12. The models are exported and quantized using PyTorch 2.8. For mobile deployment, we convert them to the ExecuTorch 0.7.0 format, leveraging the XNNPACK backend - a highly optimized neural network inference engine for ARM CPUs. The resulting ExecuTorch models are packaged by an Android application package (APK) using the standard Java-based Android build tools for deployment and benchmarking. The latency results are

summarized in Table 1, with detailed results and analysis provided in Section 6 and Table S7 of the Supplementary Information.

On the desktop CPU, our 8-bit quantized model achieves more than a 2× speedup over its 32-bit (unquantized) counterpart and substantially outperforms prior CGH methods such as HoloNet, Holo-encoder, and U-Net.

As shown in Table 1, the latency trends on the mobile platform are consistent with those observed on the CPU. Our 8-bit joint real-imaginary quantization method achieves the lowest latency and delivers a 389× speedup compared with HoloNet. These results demonstrate that quantization accelerates hologram generation effectively, without introducing noticeable degradation in output quality.

Generalization capability of ULBN on other complex-valued physical signals

To evaluate the generalization capability of ULBN, we employ it to three representative complex-valued scenarios: acoustics, wireless modulation, and SAR, as illustrated in Fig. 6 and Section 7 of Supplementary Information.

For audio classification, it aims to identify the speaker of acoustic signals. In this study, we utilize a subset of the LibriSpeech dataset⁴³, which contains 28 speaker classes for training and testing. The structure of our baseline network is shown in Fig. 6a. Raw audio clips are transformed using the short-time Fourier transform (STFT) to obtain amplitude and phase representations, which are then processed by our ultra-low-bit and mixed-precision complex-valued network. The results are presented in Fig. 7a. Our quantized model achieves an accuracy of 98.93%, closely matching the full-precision complex-valued baseline (99.36%) and outperforming the real-valued model (97.65%). Crucially, the quantized network delivers substantial efficiency gains: the number of bit operations decreases by 85% relative to the complex-valued baseline, and memory usage decreases by more than 80%.

For wireless modulation, complex-valued wireless signals are typically degraded by noise, fading, and frequency shifts, requiring accurate classification of the modulation mode. We use the RadioML 2016.10a dataset⁴⁴, which provides simulated wireless signals and the

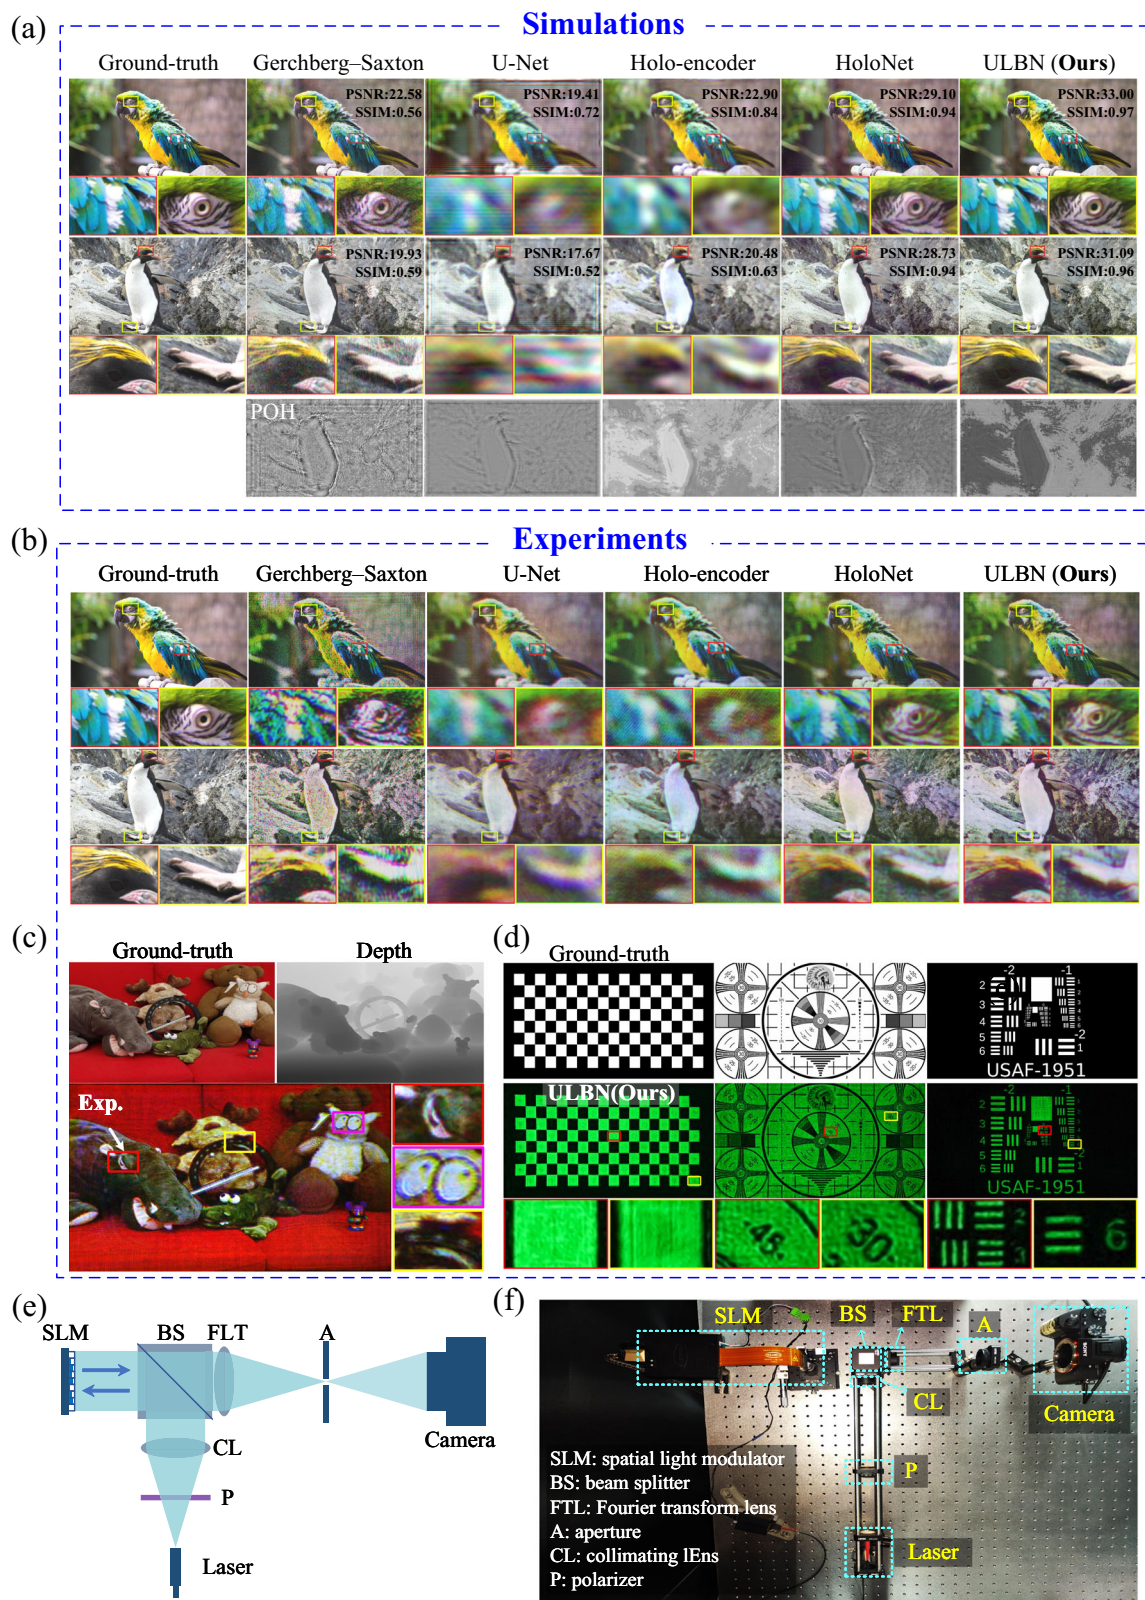


Fig. 5 | Comparison of holographic image reconstruction methods and experimental setup. **a** Simulated reconstruction of two-dimensional images. **b** Experimental optical reconstruction of two-dimensional images. **c** Experimental optical reconstruction of three-dimensional images. The image of the scenario is from the Couch dataset⁵⁰. **d** Experimental optical reconstruction of binary images. **e** Schematic diagram of the experimental setup. **f** Photograph of the experimental setup. The ground truth images in the first column of (a) and (b) are the “Green and

yellow macaw bird,” which is originally posted to Unsplash by Andrew Li at https://unsplash.com/photos/gold-and-blue-macaw-on-brown-wooden-stick-iLValbfe9_g. The other images in the first column of (a) and (b) are reproduced from the ground truth image. The images in (c) are reproduced from www.bigbuckbunny.org (© 2008, Blender Foundation) under a Creative Commons license (<https://creativecommons.org/licenses/by/3.0/>).

corresponding modulation mode. The backbone architecture is illustrated in Fig. 6b. With the proposed quantization scheme, the network efficiently classifies modulation modes from complex-valued noisy inputs. Our results are illustrated in Fig. 7b. Compared with the complex-valued method, our quantized complex-valued model achieves comparable accuracy. Meanwhile, our method substantially reduces bit operations by nearly 85%, and memory usage by about 81%, indicating substantial efficiency gains with minor compromise in performance. Compared with the real-valued method, our method achieves better classification accuracy, while with about a 41% reduction of bit operations and about 67% reduction of memory.

For SAR target recognition, complex-valued microwave signals carrying information about objects are reflected from the ground and captured by receivers. We use the MSTAR dataset⁴⁵ for training and testing. Figure 6c showcases the network for evaluation. Figure 7c summarizes the results. Our quantized complex-valued model achieves 97.98% accuracy, outperforming the real-valued baseline and closely approaching the full-precision complex-valued model. Meanwhile, our method substantially reduces bit operations by 87% and

memory consumption by 80% compared to the complex-valued baseline. Relative to the real-valued model, our approach reduces bit operations by about 47% and memory usage by 59%, while delivering superior accuracy.

Discussion

In this study, we propose a universal quantization framework for complex-valued neural networks, offering an efficient approach to capture the intrinsic coupling relations and physical laws in complex physical fields. By integrating physics-aware adaptive precision training, our approach achieves high-quality outputs with minimal computational and memory overhead, making it suitable for deployment on resource-constrained devices. Extensive evaluations across hologram generation, audio classification, wireless signal classification, and synthetic aperture radar tasks demonstrate that our method achieves superior performance and substantial reductions in computational load and memory usage. Compared to the state-of-the-art model HoloNet, our ULBN achieves approximately 4 dB improvement in PSNR while reducing computational and memory costs by 99.1% and 99.8%, respectively. Real-world edge device deployment further validates its practicality. These results highlight the potential of lightweight complex-valued neural networks for scientific computing, providing a broadly applicable solution with implications that extend beyond computational optics to the broader fields of machine learning and computational physics.

Methods

Two training stages for ULBN


Training stage 1: Identifying the optimal quantization bit width for each network layer. Each complex-valued mixed-precision block receives an input a_{in} , which is the output from the preceding network layer. This

Table 1 | Latency of hologram generation models on a desktop CPU and an Android device. The best performance in latency is marked in bold


Methods	Latency on CPU (ms)	Latency on Android (ms)
U-Net	982.876	70,221.286
Holo-encoder	493.238	13,758.553
HoloNet	1905.955	132,440.881
Ours, 32-bit	187.201	399.207
Ours, 8-bit	91.077	340.853

The best performance in latency is marked in bold.

(a) Complex-valued Audio Classification



(b) Complex-valued Wireless Modulation Mode Classification



(c) Complex-valued Synthetic Aperture Radar Recognition




Fig. 6 | Our methods applied to diverse physical signals. a Audio classification from the short-time Fourier transform (STFT) of raw audio clips. **b** Wireless modulation classification from complex-valued signals under noise. **c** SAR target recognition from complex-valued SAR data. In the subfigures, Conv, DSConv,

Linear, BN, ReLU, MPool, and APool represent convolution, depthwise separable convolution, linear layer, batch normalization, rectified linear unit, max pooling, and average pooling, respectively.

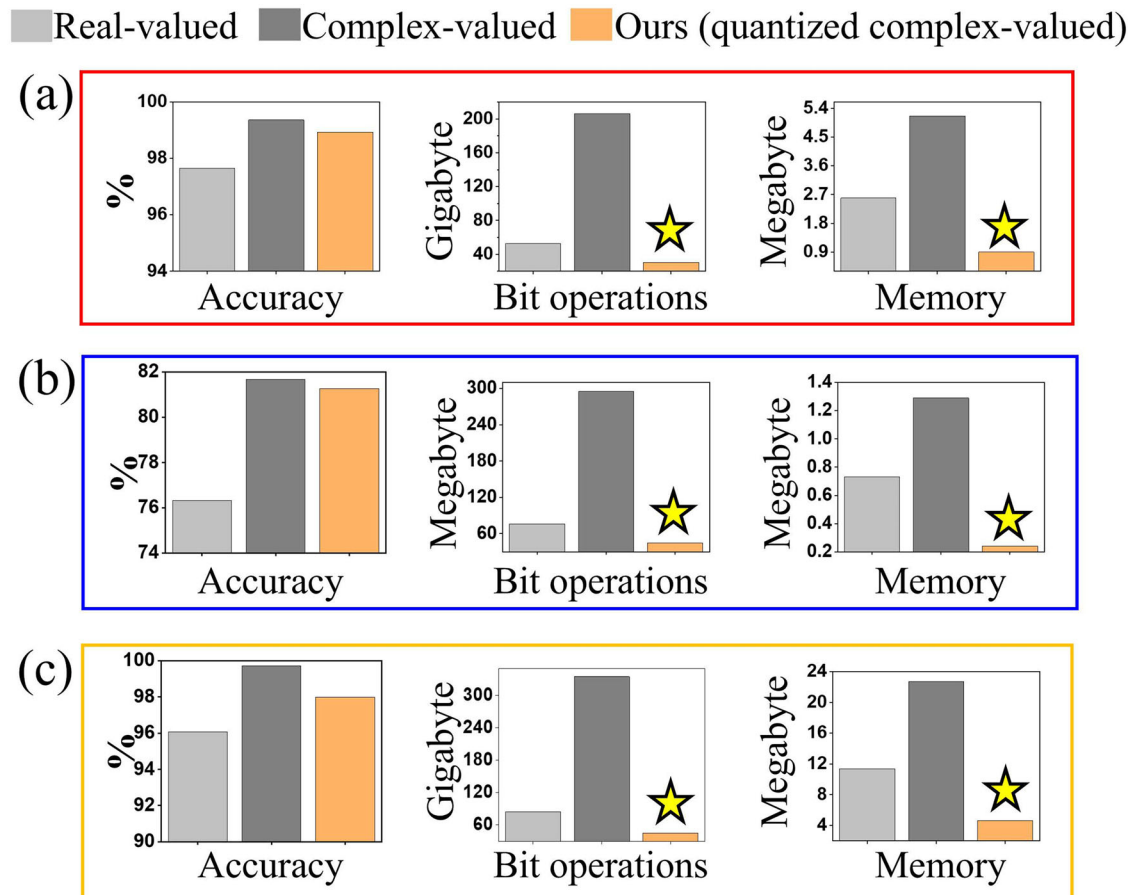


Fig. 7 | Comparison of real-valued, complex-valued, and our quantized complex-valued networks across diverse physical signal tasks. a Results on the audio classification task. **b** Results on the wireless modulation classification task. **c** Results on the SAR target recognition task.

input a_{in} comprises real and imaginary components, expressed as $a_{in} = a_{real} + ja_{imag}$. The identical quantization method is employed for both the real and imaginary parts. In this paper, a represents either a_{real} or a_{imag} . We assume that a follows the Gaussian distribution. The comparative results presented in Sections 1 and 2 of the Supplementary Information demonstrate that Gaussian quantization (GQ) outperforms learning step quantization (Table S1) and uniform quantization (Table S2), consistent with the Gaussian hypothesis. Firstly, as shown in Eq. (1) and visualized in Fig. 2, the Half-Wave Gaussian Quantization (HWGQ) method is utilized to perform ReLU activation and quantization simultaneously⁴⁶⁻⁴⁸. The use of HWGQ is motivated by the prior assumption that the distribution of activations follows a Gaussian distribution, combined with the ReLU activation function. The probability distribution of the negative part of a becomes zero after this operation, resulting in a half-wave Gaussian distribution for a . Because HWGQ includes activation and quantization processes, it is called activation quantization in this work and is calculated by:

$$HWGQ(a, b) = \text{clip} \left\{ \text{round} \left(\frac{a}{s_a} \right) \cdot s_a, 0, (2^b - 1) \cdot s_a \right\}, \quad (1)$$

where $\text{clip}\{\cdot, \cdot, \cdot\}$ is a truncation function. The first part refers to the quantization of activations to generate discrete values. The second and third parts are the lower and upper bounds of the quantization range. b denotes the bit width after quantization. Here, $b \in \{1, 2, 3, 4\}$. s_a is the quantization spacing. The value of s_a is given by the Lloyd's algorithm ($s_a \in \{0.799, 0.538, 0.3217, 0.185\}$, corresponding to the four-bit values)⁴⁶. The learnable parameters p_s^a represent the probability of

selecting the bit width $s + 1$ as the optimal activation quantization bit width, where $s \in \{0, 1, 2, 3\}$. For each activation parameter, the highest probability p_s^a indicates that its corresponding quantization bit width $s + 1$ is optimal. Since the values of probability p_s^a should lie between 0 and 1, it is given by normalizing another learnable parameter α . The probabilities are calculated by:

$$p_s^a = \frac{\exp(\alpha_s)}{\sum_{m=0}^{\|B^a\|-1} \exp(\alpha_m)}, \quad (2)$$

where B^a is the set of possible bit widths for activation. s and m are the indexes of α and p_s^a . Here, $B^a = \{1, 2, 3, 4\}$. s and $m \in \{0, 1, 2, 3\}$. Notably, $p_0^a + p_1^a + p_2^a + p_3^a = 1$. The expected value \bar{a} of the quantized activations $HWGQ(a, b_s^a)$ is calculated as a weighted sum across different quantization bit widths, with each weight given by the probability p_s^a .

$$\bar{a} = \sum_{s=0}^{\|B^a\|-1} p_s^a \cdot HWGQ(a, b_s^a), \quad (3)$$

where \bar{a} represents either \bar{a}_{real} or \bar{a}_{imag} . $b_s^a \in B^a$.

Secondly, quantization of weights w is performed. The symbol w represents either w_{real} or w_{imag} . As visualized in Fig. 2, unlike the activation quantization, the range of quantized weights extends from $-\infty$ to ∞ . The Gaussian Quantization method⁴⁶⁻⁴⁸ is adopted to quantize the network weights. Similar to HWGQ, the use of GQ is motivated by the assumption that the activations a follow a Gaussian distribution.

The weights after quantization are calculated by:

$$GQ(w, b) = \text{clip} \left\{ \left(\text{round} \left(\frac{w}{\sigma s_w} + \frac{1}{2} \right) - \frac{1}{2} \right) \cdot \sigma s_w, - \left(\frac{2^b - 1}{2} \right) \cdot \sigma s_w, \left(\frac{2^b - 1}{2} \right) \cdot \sigma s_w \right\}, \quad (4)$$

where b denotes the bit width. s_w is the quantization spacing. The value of s_w is given by the Lloyd's algorithm ($s_w \in \{1.596, 0.996, 0.586, 0.336\}$)⁴⁶. σ is the variance of the distribution of w . The probability associated with a quantized weight is:

$$p_t^w = \frac{\exp(\beta_t)}{\sum_{n=0}^{\|B^w\|-1} \exp(\beta_n)}, \quad (5)$$

where B^w is the set of possible bit widths for activation. t and n are the indices of learnable parameters β and p_t^w . $t, n \in \{0, 1, 2, 3\}$. The expected values \bar{w} of the quantized weights $GQ(w, b_t^w)$ are formulated by:

$$\bar{w} = \sum_{t=0}^{\|B^w\|-1} p_t^w \cdot GQ(w, b_t^w), \quad (6)$$

where \bar{w} represents either $\overline{w_{real}}$ or $\overline{w_{imag}}$. $b_t^w \in B^w$. The details of the learnt optimal quantization bit widths of each layer are presented in the Supplementary Information Section 1.

Thirdly, we propose a fusion method for quantizing complex-valued neural networks, which leverages the commutative and associative properties of complex multiplication. The output of this network layer a_{out} is represented as:

$$a_{out} = \overline{a_{complex} * w_{complex}} = \left(\overline{a_{real} * w_{real}} - \overline{a_{imag} * w_{imag}} \right) + j \left(\overline{a_{real} * w_{imag}} + \overline{a_{imag} * w_{real}} \right), \quad (7)$$

where $\overline{a_{complex}} = \overline{a_{real}} + j\overline{a_{imag}}$ and $\overline{w_{complex}} = \overline{w_{real}} + j\overline{w_{imag}}$. The operator $*$ indicates convolution.

Training stage 2: Quantizing the network to the optimal bit widths. Before deployment, the network is retrained and quantized to the optimal bit widths obtained from Stage 1. Similarly, this stage is divided into three parts, namely quantization of activations, quantization of weights, and complex-valued convolution of activations and weights. Firstly, the real and imaginary parts of the activations are quantized to the best bit widths b_{opt}^a with outputs $\text{HWGQ}(a_{real}, b_{opt}^a)$ and $\text{HWGQ}(a_{imag}, b_{opt}^a)$. Secondly, the real and imaginary parts of the weights are quantized to the optimal bit widths b_{opt}^w with the outputs $\text{GQ}(w_{real}, b_{opt}^w)$ and $\text{GQ}(w_{imag}, b_{opt}^w)$. Thirdly, the complex-valued quantized activations and the complex-valued quantized weights are involved in the convolution calculation. The output of the network layer a_{out}^{opt} is formulated as:

$$a_{out}^{opt} = \left\{ \text{HWGQ}(a_{real}, b_{opt}^a) * \text{GQ}(w_{real}, b_{opt}^w) - \text{HWGQ}(a_{imag}, b_{opt}^a) * \text{GQ}(w_{imag}, b_{opt}^w) \right\} + j \left\{ \text{HWGQ}(a_{real}, b_{opt}^a) * \text{GQ}(w_{imag}, b_{opt}^w) + \text{HWGQ}(a_{imag}, b_{opt}^a) * \text{GQ}(w_{real}, b_{opt}^w) \right\}, \quad (8)$$

Training configuration

The loss function is formulated with the objectives of enhancing hologram quality and determining the optimal bit width. Task loss promotes the selection of higher bit widths to enhance image quality, while complexity loss promotes lower bit widths to reduce network complexity. By minimizing the overall network loss, an optimal balance between image quality and network complexity is achieved. The

total loss function \mathcal{L}_{total} is expressed by:

$$\mathcal{L}_{total} = \mathcal{L}_{task} + \mu \mathcal{L}_{complexity}, \quad (9)$$

where \mathcal{L}_{task} and $\mathcal{L}_{complexity}$ denote the task loss and complexity loss, respectively. μ is the hyperparameter to control the complexity loss.

The task loss \mathcal{L}_{task} represents the discrepancy between the reconstructed image and the target image. It is given by:

$$\mathcal{L}_{total} = \mathcal{L}_{mse}(A_{rec_{poh}}, A_{target}) + \lambda_{com} \mathcal{L}_{mse}(A_{rec_{compen}}, A_{target}) + \lambda_{vgg} \mathcal{L}_{vgg}(A_{rec_{poh}}, A_{target}) + \lambda_{com} \mathcal{L}_{vgg}(A_{rec_{compen}}, A_{target}), \quad (10)$$

where \mathcal{L}_{mse} represents the pixel-wise mean squared error (MSE) between the target image and the reconstructed image. When the pixel-wise error approaches zero, the gradient of this loss is too small to properly cover high-frequency variations, resulting in blurred image reconstruction. Therefore, we also include a perceptual loss. It encourages generating natural and perceptually pleasing images based on high-level features extracted from pre-trained networks like the Visual Geometry Group (VGG) Net⁴⁹. \mathcal{L}_{vgg} is the Euclidean distance between the VGG-Net's hidden-layer feature representations of the reconstructed image and the target image. A_{target} is the amplitude of the target image. $A_{rec_{poh}}$ is the amplitude of the reconstructed image from POH. λ_{com} is a parameter to adjust the weight of the ringing artifacts compensation module. $\mathcal{L}_{mse}(A_{rec_{compen}}, A_{target})$ represents the loss between the target image A_{target} and the reconstructed amplitude from the compensated complex light field $A_{rec_{compen}}$. $\mathcal{L}_{mse}(A_{rec_{poh}}, A_{target})$ represents the loss between the target image A_{target} and the reconstructed amplitude from the predicted POH $A_{rec_{poh}}$. The coexistence of the two losses decouples the role of the ringing artifacts compensation network and the POH encoder, reducing the learning burden of the network.

The objective of the complexity loss $\mathcal{L}_{complexity}$ is to minimize the computational cost of network inference. It is proportional to the total bit operations across all network layers. The complexity loss of each layer is related to the expected value of the quantized parameter bit widths, the number of channels, as well as the output and convolution kernel sizes. The calculation of complexity loss is given as follows²⁰:

$$\mathcal{L}_{complexity} = \sum_{l=0}^{L-1} \left(\sum_{t=0}^{\|B^a\|-1} (p_t^l \cdot a \cdot b_t^a) \cdot \sum_{s=0}^{\|B^w\|-1} (p_s^{l,w} \cdot b_s^w) \right) \cdot comp^l, \quad (11)$$

$$comp^l = \frac{c_{in}^l \cdot c_{out}^l \cdot h_k^l \cdot w_k^l \cdot h_{out}^l \cdot w_{out}^l}{(s^l)^2}, \quad (12)$$

where L represents the total number of layers. l is the layer index. $p_t^l \cdot a$ and $p_s^{l,w}$ represent the probability of selecting the bit width $t+1$ or $s+1$ as the optimal quantization bit width for the activation a or weight parameter w of the layer l . The values of $p_t^l \cdot a$ and $p_s^{l,w}$ are obtained from Eqs. (2) and (5). The symbol b_t^a or b_s^w represents the candidate bit width from the set B^a or B^w for the activation or weight parameter. $comp^l$ is the complexity factor of the layer l . c_{in}^l and c_{out}^l are the number of input channels and the number of output channels of the layer l . h_k^l and w_k^l are the height and width of the convolution kernel in the layer l . h_{out}^l and w_{out}^l are the length and width of the outputs in the layer l . s^l is the filter stride of the layer l .

According to our ablation study of network structures in the Supplementary Information Section 3 and Table S3, our ULBN network is designed with full precision layers (FPL) at the input and output layers, along with quantization-aware (QAT) mixed-precision quantization (MPQ) of 1-4 bits in the intermediate layers for high-quality holography with reasonable computational cost. The model is trained using Python 3.8 on Ubuntu 20.04, with an NVIDIA A40 GPU and an

AMD EPYC 7543 32-Core Processor. In both stages, the models are trained with a learning rate of 10^{-3} and a batch size of 1. The weights for reconstruction loss, compensation loss, and complexity regularization are 1.0, 0.1, and 10^{-6} , respectively. The training on this hardware takes about one hour. The details of the quantization step size estimation method (Algorithm S1) and the bit widths of each network layer (Fig. S1) are in Sections 4.1 and 4.2 of the Supplementary Information.

Experimental details

In the experiment, a FISBA READYBeam emitting red light at 638 nm, green light at 520 nm, and blue light at 450 nm is utilized, as illustrated in Fig. 5e and f. The phase-only SLM employed is a HOLOEYE LETO-3-CF5-127 LCD modulator with a resolution of 1920×1080 and a pixel pitch of $6.4 \mu\text{m}$. Sony A7M3 camera is the receiver. The color holographic display is realized by time-multiplexing with 638 nm red, 520 nm green, and 450 nm blue laser sources.

Data availability

The data supporting the findings of this study are provided in the main text, the Supplementary Information, and GitHub, <https://github.com/THUIntelligentOpticsLab/ULBN>.

Code availability

The code supporting the findings of this study is available at the GitHub, <https://github.com/THUIntelligentOpticsLab/ULBN>.

References

- Rokh, B., Azarpeyvand, A. & Khanteymooi, A. A comprehensive survey on model quantization for deep neural networks in image classification. *ACM Trans. Intell. Syst. Technol.* **14**, 1–50 (2023).
- Xu, S. et al. Q-detr: An efficient low-bit quantized detection transformer. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.* (2023).
- Chen, T.-A., Yang, D.-N. & Chen, M.-S. Climhq: Class imbalanced quantization enabling robustness on efficient inferences. In *Proc. Adv. Neural Inf. Process. Syst.* **35**, 37134–37145 (2022).
- Cai, Y. et al. Zeroq: A novel zero-shot quantization framework. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 13169–13178 (2020).
- Zeng, C. et al. Abq-llm: Arbitrary-bit quantized inference acceleration for large language models. In *Proc. AAAI Conf. Artif. Intell.* **39**, 22299–22307 (2025).
- Lin, J. et al. AWQ: Activation-aware weight quantization for on-device LLM compression and acceleration. *Proc. Mach. Learn. Syst.* **6**, 87–100 (2024).
- Babikov, V. V. The phase-function method in quantum mechanics. *Sov. Phys. Usp.* **10**, 271 (1967).
- Carruthers, P. & Nieto, M. M. Phase and angle variables in quantum mechanics. *Rev. Mod. Phys.* **40**, 411 (1968).
- Luo, X. et al. Revolutionizing optical imaging: computational imaging via deep learning. *Photon. Insights* **4**, R03 (2025).
- Miksad, R. W., Jones, F. L., Powers, E. J., Kim, Y. C. & Khadra, L. Experiments on the role of amplitude and phase modulations during transition to turbulence. *J. Fluid Mech.* **123**, 1–29 (1982).
- Stuart, J. T. On finite amplitude oscillations in laminar mixing layers. *J. Fluid Mech.* **29**, 417–440 (1967).
- Xuan, A., Deng, B.-Q. & Shen, L. Study of wave effect on vorticity in Langmuir turbulence using wave-phase-resolved large-eddy simulation. *J. Fluid Mech.* **875**, 173–224 (2019).
- Curty, P. & Beck, H. Thermodynamics and phase diagram of high temperature superconductors. *Phys. Rev. Lett.* **91**, 257002 (2003).
- Shi, L., Li, B., Kim, C., Kellnhofer, P. & Matusik, W. Towards real-time photorealistic 3d holography with deep neural networks. *Nature* **591**, 234–239 (2021).
- Yang, D. et al. Diffraction-engineered holography: Beyond the depth representation limit of holographic displays. *Nat. Commun.* **13**, 6012 (2022).
- Zhong, C. et al. Real-time high-quality computer-generated hologram using complex-valued convolutional neural network. *IEEE Trans. Vis. Comput. Graph.* (2023).
- Nguyen, D. B. & Berntsen, S. The reconstruction of the relative phases and polarization of the electromagnetic field based on amplitude measurements. *IEEE Trans. Microw. Theory Tech.* **40**, 1805–1811 (2002).
- Brown, M. D. Phase and amplitude modulation with acoustic holograms. *Appl. Phys. Lett.* **115**, 051901 (2019).
- Li, Y. E. & Demanet, L. Phase and amplitude tracking for seismic event separation. *Geophysics* **80**, WD59–WD72 (2015).
- Hsueh, C.-K. & Sawchuk, A. A. Computer-generated double-phase holograms. *Appl. Opt.* **17**, 3874–3883 (1978).
- Maimone, A., Georgiou, A. & Kollin, J. S. Holographic near-eye displays for virtual and augmented reality. *ACM Trans. Graph.* **36**, 1–16 (2017).
- Chakravarthula, P., Peng, Y., Kollin, J., Fuchs, H. & Heide, F. Wirtinger holography for near-eye displays. *ACM Trans. Graph.* **38**, 1–13 (2019).
- Chang, C., Bang, K., Wetzstein, G., Lee, B. & Gao, L. Toward the next-generation VR/AR optics: a review of holographic near-eye displays from a human-centric perspective. *Optica* **7**, 1563–1578 (2020).
- Zhou, M., Zhang, H., Jiao, S., Chakravarthula, P. & Geng, Z. End-to-end compression-aware computer-generated holography. *Opt. Express* **31**, 43908–43919 (2023).
- Wei, Y. et al. Speckle-free holography with a diffraction-aware global perceptual model. *Photon. Res.* **12**, 2418–2423 (2024).
- Zhou, M. et al. Point spread function-inspired deformable convolutional network for holographic displays. *Adv. Fiber Laser Conf* **13104**, 667–678 (2024).
- Yuan, G. et al. Physics-aware cross-domain fusion aids learning-driven computer-generated holography. *Photon. Res.* **12**, 2747–2756 (2024).
- Choi, S., Gopakumar, M., Peng, Y., Kim, J. & Wetzstein, G. Neural 3d holography: learning accurate wave propagation models for 3d holographic virtual and augmented reality displays. *ACM Trans. Graph.* **40**, 1–12 (2021).
- Shimobaba, T. et al. Deep-learning computational holography: A review. *Front. Photonics* **3**, 8 (2022).
- Liu, K., Wu, J., He, Z. & Cao, L. 4k-dmdnet: diffraction model-driven network for 4k computer-generated holography. *Opto-Electron. Adv.* 220135 (2023).
- Zheng, H. et al. Diffraction model-driven neural network trained using hybrid domain loss for real-time and high-quality computer-generated holography. *Opt. Express* **31**, 19931–19944 (2023).
- Dong, Z., Xu, C., Ling, Y., Li, Y. & Su, Y. Fourier-inspired neural module for real-time and high-fidelity computer-generated holography. *Opt. Lett.* **48**, 759–762 (2023).
- Liu, Q., Chen, J., Qiu, B., Wang, Y. & Liu, J. Dcpnet: a dual-channel parallel deep neural network for high quality computer-generated holography. *Opt. Express* **31**, 35908–35921 (2023).
- Yu, G., Wang, J., Yang, H., Guo, Z. & Wu, Y. Asymmetrical neural network for real-time and high-quality computer-generated holography. *Opt. Lett.* **48**, 5351–5354 (2023).
- Horisaki, R., Takagi, R. & Tanida, J. Deep-learning-generated holography. *Appl. Opt.* **57**, 3859–3863 (2018).
- Wu, J., Liu, K., Sui, X. & Cao, L. High-speed computer-generated holography using an autoencoder-based deep neural network. *Opt. Lett.* **46**, 2908–2911 (2021).
- Peng, Y., Choi, S., Padmanaban, N. & Wetzstein, G. Neural holography with camera-in-the-loop training. *ACM Trans. Graph.* **39**, 1–14 (2020).

38. Choi, S. et al. Time-multiplexed neural holography: a flexible framework for holographic near-eye displays with fast heavily-quantized spatial light modulators. In *ACM SIGGRAPH Conf. Proc.*, 1–9 (2022).
39. Zhou, M., Zhang, H., Chen, M. K. & Geng, Z. Implicit feature compression for efficient cloud-edge holographic display. *Displays* 103151 (2025).
40. Yuan, G., Zhou, M., Peng, Y., Chen, M. & Geng, Z. Error-compensation network for ringing artifact reduction in holographic displays. *Opt. Lett.* **49**, 3210–3213 (2024).
41. Agustsson, E. & Timofte, R. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog. Workshops* 126–135 (2017).
42. Gerchberg, R. W. A practical algorithm for the determination of the phase from image and diffraction plane pictures. *Optik* **35**, 237–246 (1972).
43. Panayotov, V., Chen, G., Povey, D. & Khudanpur, S. Librispeech: an ASR corpus based on public domain audio books. In *Proc. IEEE Int. Conf. Acoust. Speech Signal Process (ICASSP)*, 5206–5210 (2015).
44. O’Shea, T. J. & West, N. Radio machine learning dataset generation with GNU Radio. In *Proc. GNU Radio Conf.* **1**, 1 (2016).
45. Keydel, E. R., Lee, S. W. & Moore, J. T. Mstar extended operating conditions: A tutorial. *Algorithms Synth. Aperture Radar Imag. III* **2757**, 228–242 (1996).
46. Cai, Z. & Vasconcelos, N. Rethinking differentiable search for mixed-precision neural networks. In *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2349–2358 (2020).
47. Tang, C. et al. Mixed-precision neural network quantization via learned layer-wise importance. In *Eur. Conf. Comput. Vis.*, 259–275 (2022).
48. Tang, C. et al. Seam: Searching transferable mixed-precision quantization policy through large margin regularization. In *Proc. 31st ACM Int. Conf. Multimedia*, 7971–7980 (2023).
49. Johnson, J., Alahi, A. & Fei-Fei, L. Perceptual losses for real-time style transfer and super-resolution. In *Proc. Eur. Conf. Comput. Vis.*, 694–711 (2016).
50. Kim, C., Zimmer, H., Pritch, Y., Sorkine-Hornung, A. & Gross, M. Scene reconstruction from high spatio-angular resolution light fields. *ACM Trans. Graph.* **32**, 73 (2013).

Acknowledgements

We are grateful for financial support from the National Natural Science Foundation of China (62305184) to Z.G., Basic and Applied Basic Research Foundation of Guangdong Province (2023A1515012932) to Z.G., and Science, Technology, and Innovation Commission of Shenzhen Municipality (JCYJ20241202123919027) to Z.G., City University of Hong Kong (Project No. 9610628, and 7005867) to M.K.C., National Natural Science Foundation of China (62501189) to K.J., the POSCO-POSTECH-RIST Convergence Research Center program funded by POSCO to J.R., the National Research Foundation (NRF) grant (RS-2024-00356928) funded by the Ministry of Science and ICT (MSIT) of the Korean government to J.R., the NRF Ph.D. fellowship (RS-2025-25436835) funded by the Ministry of Education (MOE) of the Korean government to C.P., the Science and Technology Innovation 2030-Key Project under Grant 2021ZD0201404 to X.L.

Author contributions

Z.G. conceived the ideas and designed the project. Z.G. and Z.L. developed the algorithm. Z.G., K.J., M.K.C., and J.R. provide guidance and supervision for the research. Z.G., M.Z., and Z.L. conducted hologram experiments and data processing. M.Z., Z.C., G.Y., X.L., K.J., and Z.G. conducted deployment experiments and data processing. M.Z., Z.C., G.Y., X.L., K.J., and Z.G. conducted quantification of multiple complex-valued physical signal experiments and data processing. M.Z., K.J., M.K.C., and Z.G. conducted data processing and visualization. Z.G., Z.L., M.Z., K.J., and M.K.C. wrote the manuscript. H.G., J.N., S.L., C.P., and J.R. provided the working principle optimization and detailed discussion, and they helped with manuscript preparation. All authors read, discussed, and approved the final version of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-026-70319-0>.

Correspondence and requests for materials should be addressed to Zihan Geng, Kui Jiang, Mu Ku Chen or Junsuk Rho.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026