Sensors & Transducers, Vol. 274, Issue 3, October 2026, pp. 42-54

www.sensorsportal.com

Sensors & Transducers

Published by IFSA Publishing, S. L., 2026
http://www.sensorsportal.com

An All-Optical Holographic Video-Graphics Processor for Real-Time Volumetric 3D Display

Ginka IVANOVA, * Branimir IVANOV and Deyan DIMOV
Institute of Optical Materials and Technologies, Laboratory "Holography",
109, Acad. G. Bontchev Str., 1113 Sofia, Bulgaria
Tel.: + 33 601825261
E-mail: branimirivanov@iomt.bas.bg

Received: 31 Aug. 2026 /Revised:26 Oct. 2026 /Accepted: 29 Oct. 2026 /Published:30 Oct. 2026

Abstract: This article presents an extended treatment of an all-optical holographic trigger intended for real-time volumetric three-dimensional display. The architecture combines three holographic triggers: two of them, each generating eight parallel optical channels in three colors through phase and polarization multiplexing, supply 24 virtual color channels each, and the third supplies a single reference plane, yielding forty-nine simultaneously available depth planes from a single input frame. Unlike digital computer-generated holography, where computing depth- and color-resolved content anew for every frame costs more as resolution and depth-plane count grow, the present architecture performs channel multiplication within the optical domain itself, requiring no digital conversion at any stage. The article develops the underlying splitter–encoder–combiner architecture and the physical mechanism generating and assigning discrete polarization states to individual channels, and works through an explicit accounting of physical versus virtual optical channels across the full three-trigger assembly – eight physical channels per full trigger yielding 24 virtual color channels, plus one reference channel – and presents a spectral and polarization characterization methodology for verifying correct channel operation across the visible range, positioned throughout against the digital holography literature addressing the same underlying problem.

Keywords: Holographic optical trigger, Volumetric three-dimensional display, Multichannel volume holography, Splitter–encoder–combiner architecture, Phase and polarization multiplexing.

1. Introduction

Volumetric three-dimensional display has resisted a fully satisfactory solution across several decades of active research, despite its evident value across a broad range of practical applications. A holographic reconstruction reproducing motion parallax, ocular accommodation, and correct occlusion at video frame rates places demands on both the computation and the optics that few systems built to date have met simultaneously. This article is an extended version of two OPAL' 2026 conference papers [1, 2] and develops, in considerably greater depth, a candidate solution built around a holographic optical trigger (HOT): a device performing depth- and color-channel multiplication within the optical domain itself, dispensing with the digital computation step that dominates the cost of most competing approaches. We compare the architecture with digital hologram synthesis and with earlier multichannel optical hardware, and we test its three color-specific units across the visible range. Compared with [1, 2], this work contributes (i) a splitter–encoder–combiner formulation, (ii) an explicit accounting of physical and virtual channels, (iii) wavelength-resolved measurements of transmittance, reflectance, and diffraction efficiency of each color-specific holographic video-graphics processor unit (HVGPU), (iv) polarimetric verification of the input state and of the channel polarization states at the combiner output, and (v) experimental documentation of multi-color synchronization.

1.1. Computational Pathways to Three-Dimensional Hologram Synthesis

Several research efforts have addressed the computational cost of digital hologram synthesis through learned or algorithmically optimized methods, and their results are instructive about where the underlying difficulty actually resides. A deep-learning approach trained on layered-depth-image (LDI) scene representations reaches real-time rates on a desktop graphics card, yet falls to roughly 5 frames per second on a mobile graphics processing unit (GPU) [3]. This is one figure among several illustrating a trade-off that runs through this entire body of work: raising reconstruction quality tends to cost runtime, and the exchange rate is steep. A different class of method, built around overlap-add stereograms, encodes a light field into a single computed two-dimensional hologram per frame, treating depth as a property recovered mathematically through wavefront propagation rather than as a physically distinct output [4]. Reported computation times for this approach span tens of seconds on general-purpose hardware down to under 100 milliseconds with GPU acceleration, for a single hologram at one depth configuration [4]. Both figures describe near-eye display geometries – a virtual accommodation volume of a few millimeters synthesized behind a spatial light modulator (SLM) and viewed through magnifying eyepiece optics – a configuration distinct from the physically projected, free-space multi-plane output pursued here. These numbers are not interchangeable with the depth-plane count developed here.

A related strand extends digital synthesis to layered scene representations propagated toward the hologram plane through a chain of angular-spectrum computations, in which each depth layer's contribution depends recursively on the propagated result of the layer preceding it [5]. That dependency is serial by construction: producing N layers costs N sequential propagation steps, each awaiting the completion of the one before it. The architecture developed in this article generates its depth-resolved channels through physically parallel optical paths carrying no such dependency – no channel waits on another. The handling of specular, non-Lambertian surfaces shows a comparable variety of digital solutions: directionally weighted diffuser modeling in the frequency domain [5], phase-added stereogram techniques incorporating per-ray depth [4], and hybrid schemes switching between plane-wave and spherical-wave primitives according to local surface properties [6]. Three separately engineered answers to a problem that a physically reconstructed wavefront does not need to solve at all, since directional light behavior belongs to the propagating field itself rather than to a computation performed about it.

Lensless imaging offers a useful counterpoint from outside the CGH mainstream. Fourier transform holography (FTH) recovers a full complex object field from a single recorded intensity pattern via one computed Fourier transform, given a reference wave co-located with the object [7]. This demonstrates that full-field reconstruction from a single capture is tractable, yet the reconstruction itself remains computational: software executes this algorithmic solution to the phase problem after acquisition. Thick-medium, Bragg-selective volume holography of the kind employed throughout the present work and its precursors takes the opposite route: wavelength- and angle-selectivity are physical properties of the recording medium, and reconstruction proceeds by simple illumination, with no digital step involved at any point [8, 9]. The Fourier-transform-holography literature itself draws a related distinction between holography and coherent diffraction imaging as two structurally different answers to the same phase-recovery problem. The present architecture belongs to the physical-reconstruction camp, consistent with its broader aim of avoiding analog-to-digital conversion anywhere in the display pipeline.

1.2. Diffraction Accuracy and Off-Axis Wavefront Modeling

The accuracy with which off-axis diffraction can be modeled also bounds reconstructed image quality in holographic systems – a difficulty that arises whenever a diffractive or holographic element must reconstruct content away from its own optical axis. Numerical work on this problem has recently improved substantially through an adaptive sampling strategy built on the angular spectrum framework common to essentially every digital CGH method discussed above, converting an off-axis propagation problem into an equivalent, more efficiently sampled on-axis one [10]. The reported gains are considerable: roughly 35 times faster than a prior state-of-the-art method at 20° of incidence, and a 35° off-axis field computed in about one second on ordinary commercial hardware [10]. One detail in that comparison matters more than the headline speedup. Constrained to the same sampling budget as the proposed method rather than allowed extra samples to match its accuracy, the comparison method produces an incorrect result at an incidence angle as modest as 5° [10].

This consideration informs a design choice developed further in Section 2.5: every channel of the present trigger propagates along its own local optical axis rather than at an angle to it. The physical architecture itself thereby avoids the off-axis correction burden described above, rather than compensating for it after the fact – digitally or otherwise.

1.3. Antecedents of Multichannel Holographic Multiplexing

Multichannel holographic interferometry has a longer history than the display application pursued here. A four-channel holographic optical element built

for phase-stepping digital speckle-pattern interferometry [9] encoded two illumination planes – vertical and horizontal, needed for full three-dimensional displacement-vector recovery – at two near-infrared wavelengths, yielding four independently addressable channels within one compact device. That element originally served precision strain and displacement measurement rather than display; its sixfold scaling into the present 24-channel architecture [1, 2] represents a repurposing of an established multiplexing principle toward a new application area, not a straightforward continuation within the same field.

Demonstrations of coherent beam combination through angle multiplexing of several Bragg gratings recorded sequentially in a thick polymeric substrate separately support the broader feasibility of sustaining many channels within a single volume-holographic element [11], and related work on nanocomposite photoanisotropic materials further reinforces this feasibility in polarization-sensitive holographic media [12]. Such devices combine and split multiple beams with diffraction efficiency well predicted by coupled-wave theory, and the angle-multiplexing literature cited within that work indicates volume-holographic media of this class can in principle sustain well over a thousand independently multiplexed channels [11]. That precedent bears directly on the channel-scaling route outlined in Section 1.4 below and developed architecturally in Section 2, pointing toward substantially more than 24 channels using the same recording medium and without additional laser sources.

1.4. Recording Media: Lineage, Properties, and Scaling Potential

The present prototype records on a silver-halide emulsion (U04) reported to offer 18000 lines/mm resolution and a 4 nm grain, roughly two to three times finer, respectively, than the comparable panchromatic materials used for benchmarking (6000–10000 lines/mm; 8–10 nm grain) [13]. We record on U04, as in [1]. This is a change from [2], which used the LiMo emulsion that work describes as photoanisotropic. We chose U04 because it is initially isotropic, which makes it suitable for later study of anisotropy induced by the recording itself; that study lies outside the present article. The blue source is also changed: the prototype of [1] used 442 nm illumination, whereas the blue hologram of the present HVGPU was recorded and characterized with a 488 nm Coherent OBIS laser. The red and green sources are unchanged, and all blue-channel results reported here correspond to the 488 nm source.

That material sits at a later stage of a materials lineage running through our own earlier work. A four-channel interferometric element built previously [9] used an in-house silver-halide emulsion requiring post-development chemical swelling to shift the reconstruction wavelength from the visible recording wavelength into the near-infrared operating band. A later, systematic study applied the same swelling principle to the considerably smaller, color-balance-critical shifts required for accurate multi-color reflection holography. That study established that controlled post-development swelling can correct shrinkage-induced wavelength shifts of tens of nanometers, and that a residual shift of only a few nanometers is sufficient to visibly disturb color balance [8]. The same study found that recording several wavelengths sequentially, through successive exposures on one plate, yields lower diffraction efficiency than a simultaneous multi-wavelength exposure. This observation bears directly on the present architecture, where a sequential exposure at a distinct half-wave plate setting records each of the eight phase-multiplexed channels within a single-color hologram. That sequential-exposure effect offers a physically grounded reason for the channel-to-channel efficiency variation reported in the fabricated prototype [1, 2], addressed further in Section 3.2. Reported single-wavelength, single-exposure efficiencies of roughly 80 %, 75 %, and 50 % at the red, green, and blue recording wavelengths respectively give a materials-level baseline against which the lower efficiencies measured under repeated [8], multichannel exposure become explicable rather than an unexplained figure.

The present recording setup – three laser sources combined through a dichroic beam-combining stage feeding a shared spatial filter – follows the general optical arrangement of the earlier work discussed above [8], establishing a direct line of design descent between the two setups. Recording the three colors simultaneously within a single shared photosensitive layer, rather than on three separate plates, carries a trade-off. A single shared layer minimizes total optical absorption, since light passes through one holographic layer rather than three; against that, simultaneous multi-color recording demands precise alignment of the fringe systems of all participating wavelengths, which substantially complicates achieving accurate white balance in the reconstructed image. Given that the present stage of this research remains experimental, recording each color on its own holographic plate allows independent characterization and adjustment of amplitude and phase per color. Once this separated-hologram approach establishes amplitude balance and fringe alignment, a single-layer, simultaneous multi-color implementation becomes a reasonable next step, trading the synchronization complexity just described for lower total absorption. Each plate used here consists of a photosensitive layer approximately 4 μm thick over a thin underlayer approximately 0.4 μm thick, both supported on a BK7 glass substrate 3–6 mm thick. Three plates rather than one raise the total absorption of the assembled trigger accordingly, a consideration for the overall optical efficiency budget of the finished system.

Reaching high diffraction efficiency in U04 recording depends on maintaining comparable optical

power between the object and reference beams during exposure, since the two beams achieve maximum interference contrast only when they reach the plate at closely matched intensity. The developers of the emulsion recommend a minimum exposure of 320 μJ/cm² per laser for full-color recording, rising to roughly 600 μJ/cm² for a monochrome exposure [13]. When the available object-beam power falls short relative to the reference beam, extending the exposure time – rather than altering beam geometry – achieves the required energy level.

The same wavelength-shifting swelling technique offers an inexpensive path to a higher channel count without adding more laser sources. Starting from the three recording wavelengths already in use here (632.8 nm, 532 nm, 488 nm), controlled post-development swelling shifts the effective reconstruction wavelength of a given recording in either direction from its own recording wavelength. The available shift is up to roughly 50 nm toward longer wavelengths and, independently, up to roughly 50 nm toward shorter wavelengths, so that each of the three base wavelengths can, in principle, yield three distinct reconstruction wavelengths: for example, the original wavelength variant, a red-shifted variant, and a blue-shifted variant. Applied across all three lasers, this gives up to nine reconstruction wavelengths from three physical sources. With the present recording wavelengths and shifts of ±50 nm, three pairs of these wavelengths (582 and 582.8 nm, 532 and 538 nm, 482 and 488 nm) lie closer together than the spectral width $\Delta\lambda$ of the corresponding diffraction lobe (Section 3.1), so about six of the nine are resolvable. Choosing the shift of each color so that neighboring wavelengths differ by more than $\Delta\lambda$ is an objective for optimization in subsequent work. Paired with the eight-channel phase multiplexing already demonstrated per color, this points toward a trigger of up to 72 channels, nine wavelength states multiplied by eight phase states, as an upper bound reachable with the three lasers already on the bench once the nine wavelengths are spaced by more than $\Delta\lambda$; with the shifts considered here, the resolvable count is about 48 channels. That channel count sits comfortably within the material precedent noted in Section 1.3 [11], and matches the 72-channel scaling direction identified as a target for future development [1]. One caveat belongs here as laboratory practice rather than device design: color-dependent diffraction efficiency, dynamic range, and overall optical response are not fixed properties of an emulsion. They vary between production batches of nominally identical material and drift over a plate's working life, so calibration measurements from one recording session do not guarantee identical results in the next without renewed calibration. Section 5.3 returns to this point in the context of experimental reproducibility.

Beyond this lineage, a dedicated holographic printer has also produced full-color holographic stereograms from AI-generated content on the U04 platform [14], underscoring the material's continued relevance within the broader color-holography literature.

2. Conception and Architecture of the Holographic Video-Graphics Processor Unit

2.1. The Holographic Optical Trigger as a Multichannel Switching Element

A holographic optical trigger, in the sense used throughout this article and its precursor [1], is a holographic image of a real multichannel optical switch. Depending on the direction of use, the same recorded structure functions either as a splitter – one input beam replicated into several outputs – or as a combiner, the reverse optical operation. Its purpose in either direction is the replication or merging of light beams across a defined band of the electromagnetic spectrum; here that band spans roughly 400 nm to 900 nm, the visible range plus a short margin on either side, sufficient for real-space 3D video projection. This element's functions govern the wavelength-dependent behavior of different color beams within the video stream without digitization between input and output; the transit time through the element is negligible compared with the frame period and does not limit the frame rate.

Three such devices combine here into what we term a holographic video-graphics processor unit (HVGPU): a splitter, an encoder, and a combiner, assembled in sequence. Three independent HVGPUs – one per working color – operate in parallel, each producing eight holographic projection planes for its assigned color. We chose the resulting count of 24 virtual color channels per trigger deliberately – eight per color across three colors – to echo the eight-bit byte of conventional digital architectures, easing comparison and future software migration between digital and optical computing platforms. This section develops that architecture well beyond the level of detail the treatment of the HOT concept in [1] allowed. That earlier treatment established the internal structure of the eight-channel base unit – the frequency mode generator, the frequency synchronization and spatial filtering stage, and the trigger proper, comprising the multiplexer, half-wave plate rotators, and depth-focus correction system [1] – which is recast within the splitter-encoder-combiner framework developed in this work.

2.2. Splitter–Combiner Architecture

Each module–the splitter and the combiner–comprises seven non-polarizing beam-splitter cubes, four 90° R100 mirrors, three intermediate path-length equalizers, and seven optical correctors that equalize the total optical path length of each channel. The splitter's task is straightforward to state and

considerably less straightforward to build: one incoming beam, with a defined phase, polarization direction, and wavelength, must emerge as eight parallel outputs in the same output plane, equidistant along one axis, each carrying exactly the input phase, polarization, and wavelength unaltered. The splitter clones one input channel onto eight output channels without modification of any optical parameter. In the present unit, one cube of the cascade divides the S and P amplitudes unequally; the input polarization conditioning described in Section 4.1 compensates for this, so the unit behaves as designed.

The combiner is the same physical unit run in reverse – eight inputs, one output – performing the corresponding merge. Each of the eight input channels' phase, polarization, and wavelength must appear at the single output unchanged relative to its own input. Where two channels of the same color enter the combiner with the same polarization and phase, they superpose coherently at the shared output. Channels of different colors do not interfere in this sense: their relative phase advances at the optical difference frequency, so a red and a green channel with the same polarization co-propagate and superpose in the same geometric channel, and the time-averaged intensity at the output is the sum of the two intensities.

Placed splitter-first, encoder second, combiner third – with the splitter's eight outputs facing the combiner's eight inputs – the assembly defines a net of eight optical channels into which the encoder can insert arbitrary transformations of the wavefront between splitting and combining. The whole assembly preserves phase, polarization, and wavelength, except where a deliberately applied transformation changes them; the assembly itself introduces no error beyond whatever transformation the encoder carries out.

Fig. 1 summarizes this splitter–encoder–combiner path. The splitter's eight output branches each carry, within the encoding stage, a half-wave plate set to produce its own polarization angle (Θ_1 through Θ_8) and the channel's programmed transformation function (01 through 08) – the zoom function of Section 2.3, for instance. The combiner then merges all eight branches back into the single output plane, with every branch remaining on-axis throughout rather than diverging at an angle. This spatial recombination does not erase channel identity: each branch retains the distinct phase and polarization state its half-wave plate assigned it, so the eight channels remain individually distinguishable within the combined output rather than collapsing into a single indistinguishable beam.

2.3. Programmed Wavefunction Transformation and the Zoom Function

The encoder applies a programmed transformation to the wavefunction of the propagating image entirely within the optical domain – a zoom operation, a Fourier transform, a recognition function, or any comparable optical-domain algorithm, depending on what the operator loads into the stage. For converting a flat two-dimensional source into a real-space three-dimensional display, the zoom function is the relevant case. Rescaling the image differently across different channels changes the effective viewing perspective of each channel, so that a viewer perceives an object built from several flat 2D images, each rescaled and repositioned appropriately, as a single three-dimensional form.

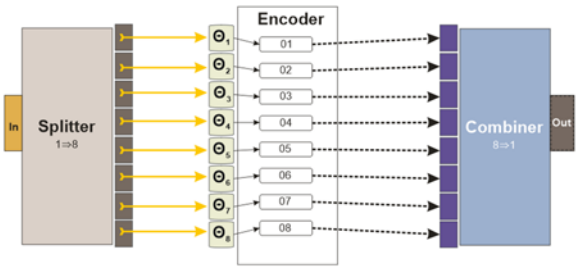


Fig. 1. Splitter–encoder–combiner architecture of the eight-channel HVGPU trigger, showing the polarization angles Θ_1–Θ_8 set by the half-wave plates and per-channel transformation functions 01–08 in the encoding stage. All eight branches propagate on-axis.

Achieving this rescaling optically, without introducing depth distortion, depends on a specific geometric condition. Distortion introduced by angular parallax appears in hologram reconstruction only when the illuminating beam's axis fails to coincide with that of the reference beam used at recording. Where that correspondence is exact, no such distortion occurs, regardless of how strongly the illuminating beam converges or diverges relative to the recording geometry. This is the criterion that makes the zoom function described above tractable directly through the optics. Because axis correspondence, rather than a fixed convergence angle, defines the reconstruction condition, varying the convergence or divergence of the illuminating beam changes the apparent depth of the reconstructed plane without disturbing the geometric correspondence that distortion-free reconstruction requires. This principle underlies the per-channel zoom function implemented in the encoder, and establishes why reconstruction must preserve the recording geometry of the U04 medium introduced in Section 1.4, for each of the eight channels.

2.4. Discrete Polarization States and the Placement of the Half-Wave Plates

The encoder's programmed transformation generates, among its other outputs, a set of distinct polarization states – one per channel. These are eight discrete polarization states, each corresponding to a fixed orientation of the linear polarization of the light in one channel. The number of such states equals the number of channels carrying a distinct polarization

angle. A linear polarization angle spans only the range between 0 and π. The half-wave plate's transversal plane possesses two orthogonal reference directions, the fast axis and the slow axis, perpendicular both to the optical axis of propagation and to the Poynting vector describing the direction of energy flow. Rotating the plate by an angle Θ/2 about the optical axis rotates the linear polarization of the transmitted light by Θ within that same transversal plane, so each channel is specified here by its polarization angle Θ. The eight polarization angles run from 0° to 70° in steps of 10°, which corresponds to plate orientations from 0° to 35° in steps of 5°; the half-wave plate angles quoted in [1, 2] are these polarization angles.

The half-wave plates that generate the eight polarization states sit at the output of the beam-splitter unit, as part of the encoder unit rather than of the splitter unit itself. The splitter unit transforms the single input channel into eight output channels (Section 2.2), and the half-wave plates, positioned immediately at that output, apply their rotation as the first operation the encoder unit performs on each of the eight channels. Because this set of plates belongs functionally to the encoder rather than to the splitter, the combiner remains a true mirror image of the splitter – the same optical design run in reverse – with no half-wave plates of its own at its input.

The rotation these plates apply must remain consistent with whatever other transformations – the zoom function of Section 2.3, for instance – the encoder applies simultaneously. Verification is therefore continuous, confirming that the combined effect of several transformations acting together does not drive any two channels to the same angle, and does not drive any channel past the π limit, beyond which a polarization angle repeats. The encoder already requires this real-time verification for its other programmed transformations regardless of polarization, so folding the rotation function's own consistency checks into that same process costs little in addition. The half-wave plate offers, in this role, the most direct available means of setting the electric field vector to a precise angle.

2.5. On-Axis Propagation within Each Channel

Each of the eight channels defined by the splitter–encoder–combiner assembly propagates along its own local optical axis rather than at an angle to it. The choice is deliberate. An off-axis channel geometry would introduce the distortion and aberration effects associated with off-axis propagation discussed in Section 1.2, and would require correction of the same general kind digital methods apply computationally when modeling off-axis diffraction [10]. Keeping every channel on-axis relative to its own local propagation direction avoids that correction burden at the architectural level, rather than compensating for it afterward.

2.6. Recording the HVGPU Function as Single-Color Holograms

Once assembled and characterized, the entire HVGPU function – splitter, encoder, and combiner together – can be recorded into a single holographic plate for a given color, condensing what is, in its discrete-component form, a fairly large optical bench into one compact element. Because each color is recorded separately (Section 1.4), the result is a set of three holograms: an R-HVGPU hologram, a G-HVGPU hologram, and a B-HVGPU hologram, each recorded with its own laser and each carrying the global wavelength function of the HVGPU encoded into that color's own spectral diffraction characteristic. The red hologram, for instance, records the global HVGPU function as expressed through the spectral diffraction efficiency curve obtained under red-laser recording; the green and blue holograms record the equivalent function at their own wavelengths. Additional holograms, recorded with additional coherent sources, extend color reproduction without altering the geometric parameters of the displayed image – a route to richer color rendering that leaves the zoom function and channel geometry established above untouched.

2.7. Physical and Virtual Channel Accounting Across the Trigger Assembly

The distinction between a physical channel and a color channel needs stating explicitly, since the two figures describing this system – 24 and 49 – arise from counting different things. A physical channel is one of the eight polarization-encoded paths within a single trigger; a color channel, or virtual channel, is one color propagating through one physical channel. Because all three colors travel through each of the eight physical channels of a given trigger, one trigger's eight physical channels yield 8 × 3 = 24 virtual color channels, consistent with the 24-channel figure already established for a single trigger [1].

Three triggers compose the full assembly. The first supplies the 24 positive-depth virtual channels; the second, an identical unit, supplies the 24 negative-depth virtual channels. Between them, this accounts for 48 virtual channels drawn from 16 physical channels (8 + 8). The third trigger is used differently: only one of its eight physical channels is active. Red, green, and blue propagate together through that channel as three spectral components of one beam and remain spectrally distinct. Because the three components perform the same optical function, namely transferring the reference image unmodified from the input to the output for comparison, we group them under one channel number and count them as a single reference plane. Counted by color, they are three virtual color channels. No component appears at an intermediate wavelength: a spectrometer in this channel records three lines at 632.8, 532, and 488 nm, and a detector or the eye receives the sum of the three

intensities, which is perceived as a mixed color by additive mixing.

Total physical channels: 8 (first trigger) + 8 (second trigger) + 1 (third trigger) = 17. Total focal planes: 24 + 24 + 1 = 49, the figure reported in [1], where the three color components of the reference plane share one plane number; counted by color, the assembly comprises 24 + 24 + 3 = 51 virtual color channels.

This reference plane serves as the calibration point for white balance across the other 48 virtual channels, and doubles as the plane-zero reference against which the ±24 depth-displaced planes are measured. It carries no zoom transformation and no color-specific processing of its own, passing the original frame from the source SLM through unmodified, and functions in that respect as an untransformed control against which the transformed planes can be checked. A single optical translation applied only to this reference plane moves the entire 49-plane package as a rigid unit relative to the source, since the spacing between the reference plane and each of the other 48 planes remains fixed. The whole assembly's position along the propagation axis thereby becomes a single adjustable parameter, rather than 48 independently tracked ones.

The finished HOT, in summary, comprises eight physical channels, corresponding to the eight polarization states of the light within one physical device, yielding 24 virtual color channels per device; three such devices, combined as described above, yield the complete 49-plane trigger assembly.

3. Spectral and Polarization Characterization Methodology

3.1. Measurement Configuration

Spectral and polarization characteristics of the assembled HVGPU system are measured across the 400–900 nm range using a three-unit Ocean Optics spectrometer suite (HR4000, HR2000+, FLMS). The measurement examines how the system distributes optical energy among the polarized channels in transmittance, reflectance, and diffraction. Because the encoder applies programmed instructions to the propagating wavefunction, the spectral and polarization response is expected to shift accordingly, tracking whatever function is currently loaded rather than remaining fixed.

The resulting data provide a direct check that each color-specific processor – R-HVGPU, G-HVGPU, and B-HVGPU – executes its programmed function correctly across the visible range. Each unit is, physically, a single hologram: the R-HVGPU, for example, is recorded with a coherent red laser on a U04 emulsion [13], operating through the spectral diffraction efficiency of that recording. Measuring its response across 400–900 nm therefore verifies both the recording and the correctness of the encoded wavefunction transformation.

We characterize the spectral selectivity of each color-specific HVGPU by the full width at half maximum, $\Delta\lambda$, of its diffraction-efficiency lobe in Figs. 2–4. Two functions executed at neighboring wavelengths remain resolvable when their center wavelengths differ by at least $\Delta\lambda$, so each optical "CPU" operates within its own spectral band. The values read from Figs. 2–4 are $\Delta\lambda_R = 15$ nm, $\Delta\lambda_G = 11$ nm, and $\Delta\lambda_B = 8$ nm for the red, green, and blue units.

The diffraction maxima of the three units lie at 627 nm (red), 522 nm (green), and 483 nm (blue), which is 5.8, 10.0, and 5.0 nm below the recording wavelengths of 632.8, 532, and 488 nm. At the recording lines the diffraction efficiency is therefore about 47 %, 18 %, and 29 %, compared with maxima of about 65–66 %. This shift is consistent with residual emulsion shrinkage after processing, which the swelling procedure of Section 1.4 is designed to compensate [8]. Because color balance is sensitive to shifts of a few nanometers [8], it is not optimized in the present prototype; aligning the diffraction maxima with the recording wavelengths is an objective for optimization in subsequent work.

For each of the three color holograms, transmittance, reflectance, and diffraction efficiency were measured as a function of reconstruction beam power, to characterize the operating range over which the recorded gratings respond linearly before saturation or measurable degradation sets in.

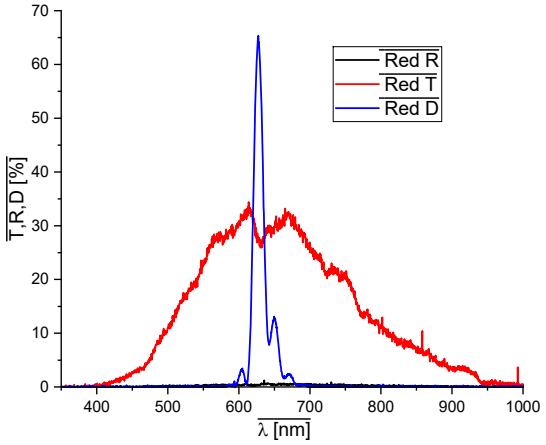


Fig. 2. Spectral dependence of transmittance, reflectance, and diffraction for the Red HVGPU (632.8 nm).

3.2. Anticipated Sources of Channel Non-Uniformity

The measurement campaign described above was designed, in part, to characterize numerically a set of non-uniformities already documented at the fabrication stage of the underlying trigger [1]. Diffraction efficiency across the eight physical channels of a single trigger was previously found to range between roughly 20 % and 50 %, averaging near 35 % – a spread consistent with the sequential-exposure mechanism discussed in Section 1.4. In the

wavelength-multiplexed recordings of [8], successive exposures on one plate lowered the diffraction efficiency relative to a simultaneous exposure; we expect the same mechanism when the eight channels are recorded through successive half-wave plate exposures. Where efficiency falls toward the lower end of the observed range, a proportional increase in source illumination – no more than a factor of 2.5, well within the range of standard diode laser drivers – compensates the hologram's contribution to channel non-uniformity without requiring any change to the recording itself [1].

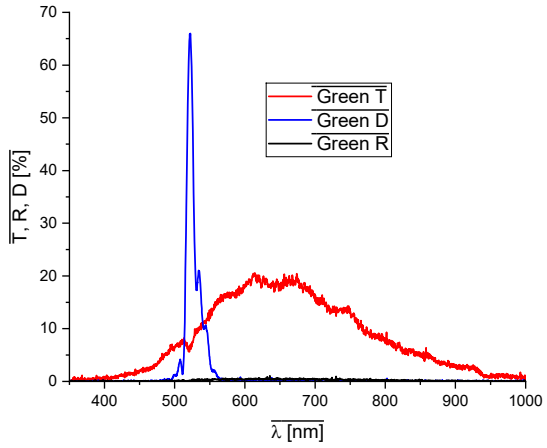


Fig. 3. Spectral dependence of transmittance, reflectance, and diffraction for the Green HVGPU (532 nm).

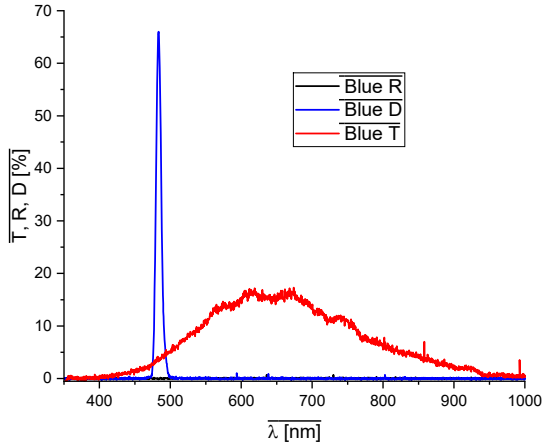


Fig. 4. Spectral dependence of transmittance, reflectance, and diffraction for the Blue HVGPU (488 nm).

We reproduce the measurements reported previously [1] only as the baseline against which the new HVGPU-level spectral and polarization measurements are evaluated. The new measurements extend the characterization to the assembled HVGPU, including the splitter and combiner stages, and to the response of the spectral energy distribution to the wavefunction transformations programmed into the encoder. We leave further functions and manipulations beyond the zoom operation described in Section 2.3 to subsequent work.

For the present experiments, we set a maximum crosstalk of 12.5 % between adjacent channels as the initial technical specification of the optical system. To meet it, we adjusted the polarization azimuth of each channel with a resolution of 0.2° and its ellipticity angle with a resolution of 0.4°, using a Thorlabs PAX1000VIS/M polarimeter, which also gave the degree of polarization. We measured each channel at the output of the beam combiner, where the optically manipulated color images of the eight physical channels of each color merge, with the other seven channels blocked. We define the overlap (crosstalk) as the fraction of the area of the narrower polarization ellipse that the ellipse of the physically adjacent channel of the same color covers, calculated from the azimuths and ellipticity angles read on the polarimeter. This measure describes the separability of the polarization states, not the transferred power. For example, the blue component had an azimuth (referenced to the horizontal axis) of 89.97° and an ellipticity angle of 0.25° in channel 1, and 79.91° and 0.71° in channel 2, a separation of 10.06°. The overlap of the two ellipses is approximately 9 %, below the specification, so these two channels needed no further adjustment. These values are approximate: they show that the overlap can be estimated with simple means, which is the purpose of this demonstration. All adjacent pairs of the three colors met the specification. The crosstalk values of [2] were obtained for different azimuth inclinations and polarization ellipses and therefore do not apply here.

4. Experimental Realization and Verification

This section presents direct experimental verification of the assembled trigger: confirmation that the light entering the HVGPU carries the correct input polarization state (Section 4.1), and photographic documentation of the physical trigger across successive stages of assembly and calibration (Section 4.2).

4.1. Input Polarization Conditioning and Verification

The polarization state at the HVGPU input is a controlled parameter, because the splitter cascade and the encoder both require a defined polarization state. In [2], the illumination was strictly linear, which that work identified as a critical condition. In the present cascade, one non-polarizing cube divides the amplitudes of the S and P components unequally, depending on the polarization direction of the incident light. To keep the division equal in both directions of the cube, we use left-circularly polarized light at the splitter input, which supplies equal S and P components. A quarter-wave plate at each cascade output then restores vertical linear polarization,

identical for all eight channels, as we verified with the polarimeter, so the encoder receives the same linear state as in [2]. These measures compensate for the cube within the splitter unit; the architecture of the trigger remains that of [1, 2]. Fig. 5 shows the input polarization state at the splitter entrance for each wavelength, upstream of the combiner output where we measure the crosstalk with the polarimeter. The polarization chain follows from Jones calculus in the cube-arm s–p basis. The left-circular input is

$$E_{\text{in}} = \left(1/\sqrt{2}\right)(1, i)^{\tau} \quad (1)$$

In arm k the cascade acts as:

$$C_k = diag(a_k, b_k\, e^{i\,\delta_k}), \quad (2)$$

where a_k and b_k are the amplitude factors for s and p and δ_k is their relative phase, so the arm output is:

$$E_k = \left(1/\sqrt{2}\right)\left(a_k, i\, b_k\, e^{i\,\delta_k}\right)^{\tau} \quad (3)$$

For an ideal non-polarizing cube a_k equals b_k and δ_k is equal to 0, and the arm output stays circular. A quarter-wave plate whose fast axis is aligned with a principal axis of the polarization ellipse converts an elliptical state to linear light, so the plate at the output of arm k is set with the polarimeter until the arm delivers linear polarization as close as possible to the vertical reference axis.

The residual azimuth error, determined by the amplitude imbalance, is absorbed in the setting of the half-wave plate. A half-wave plate whose fast axis is turned by $\Theta_k/2$ away from the input polarization then rotates this linear state by Θ_k, giving:

$$\mathrm{E}'_{\mathrm{k}} = (\sin\Theta_{\mathrm{k}}, \cos\Theta_{\mathrm{k}})^{\tau}, \quad (4)$$

up to a common phase factor, with

$$\Theta_{\mathrm{k}} = 0°, 10°, \ldots, 70°, \quad (5)$$

and a polarimeter azimuth of 90° minus Θ_k measured from the horizontal axis. The circular input therefore enters only through the splitting ratio of the cascade and the quarter-wave plate settings, and does not appear in the eight final states, which are linear. The states measured at the combiner output follow this prediction: the blue component had azimuths of 89.97° and 79.91° in channels 1 and 2 (Section 3.2), against 90° and 80° predicted.

4.2. Physical Assembly and Multi-Color Synchronization

Fig. 6 documents the trigger's physical assembly across three stages of adjustment. Fig. 6(a) shows the beam-splitter unit (left) and the beam combiner (right) positioned before adding the encoding stage between them – the configuration's starting point, prior to installing the half-wave-plate rotators and their associated transformation functions. Fig. 6(b) shows the multi-color output before channel synchronization: the red, green, and blue channels each appear as a separate, diffuse cloud of points rather than overlapping into a single spot. The combiner overlaps the arriving channels into a single spot (Section 2.2) only where they share spatial alignment, a condition the unsynchronized setup does not yet meet. Fig. 6(c) shows the trigger in its working, synchronized state, with the optical propagation paths of the blue and green channels visible along their designed trajectories through the assembly.

5. Discussion

5.1. Positioning Relative to Digital Computer-Generated Holography

The present approach and the digital CGH literature address the same problem through different means, and neither is presented here as a simple replacement for the other.

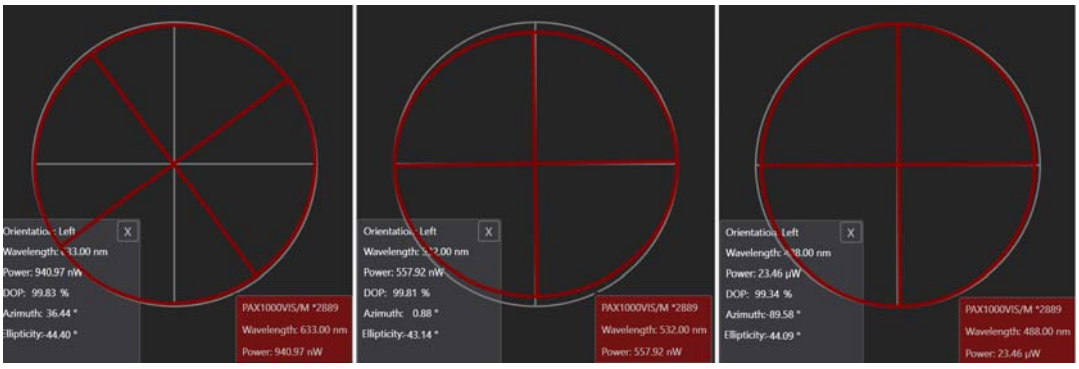


Fig. 5. Polarimeter measurements of the R–G–B input polarization states for the HVGPU at the splitter entrance for left-circularly polarized light.

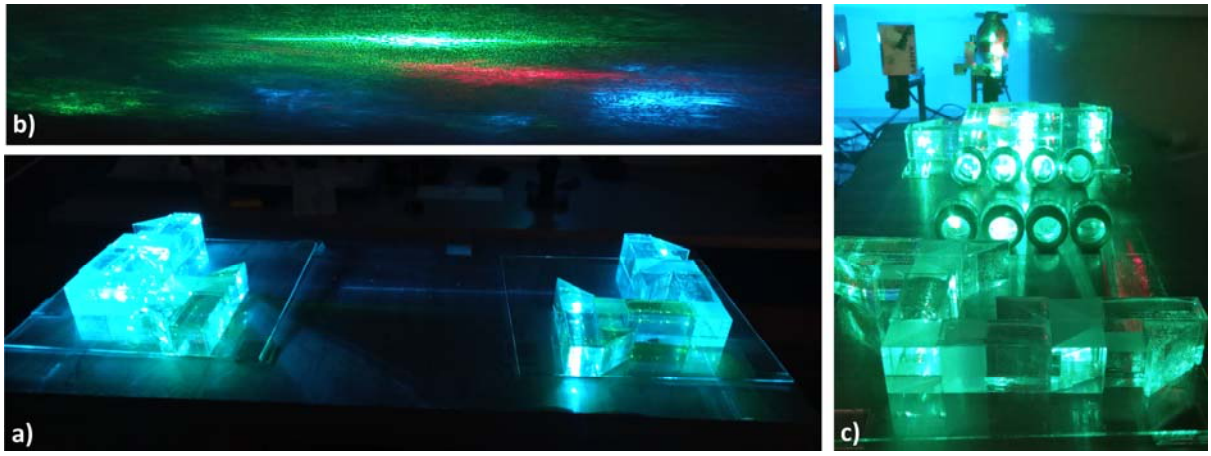


Fig. 6. Trigger assembly across three stages of adjustment: (a) beam-splitter (left) and beam combiner (right) positioned before encoder installation; (b) non-synchronized multi-color output, with red, green, and blue channels visible as separate point clouds; (c) working, synchronized state, showing the propagation paths of the blue and green channels.

A pattern recurs across the digital methods reviewed above: gains in accuracy tend to arrive with a disclosed, often substantial cost in computation time. A hybrid method combining planar and spherical wave primitives to more accurately model both diffuse and specular reflections features an analytically accurate configuration that runs roughly 60 times slower than a widely used overlap-add stereogram baseline [6]. Recovering real-time performance requires training a neural network, over several days of dedicated GPU time, to approximate that slower analytical method. A related line of work on reconstructing a dense light field from a sparse set of captured views – logically upstream of hologram synthesis itself – reports a comparable offline cost, on the order of 40 hours of training on a single high-end GPU before real-time inference becomes available [15]. A fully real-time digital pipeline for high-quality volumetric display, taken as a whole, consists of a chain of separately trained stages – sparse-view capture, learned dense reconstruction, hologram synthesis, phase-only conversion – each carrying its own accuracy–speed trade-off and its own offline training cost. Generating depth- and color-resolved channels within a single physical optical pass, as the present architecture does, requires no equivalent to any of these separately trained stages.

A formal evaluation grounded in wave-optics and sampling theory, rather than empirical benchmarking alone, demonstrates that several widely used image-based CGH methods – holographic stereograms, phase-added stereograms, and the overlap-add approach discussed in Section 1.1 – are different heuristic answers to the same mathematically posed problem. That problem is recovering a complex hologram from a light field by inverting a short-time Fourier transform, whose magnitude-squared formulation discards the phase information exact inversion needs [16]. Under that framework, the authors estimate that artifact-free, full-color reconstruction for a representative near-eye configuration requires on the order of several thousand sampled perspective views [16] – evidence that the accuracy–resolution trade-offs seen throughout this literature follow from the underlying computational formulation itself rather than from any one implementation's shortcomings. A separate method addressing computational cost through an adaptive combination of point-source and wave-field propagation primitives, chosen per depth layer according to local scene content, reports a computation-time reduction of up to 68 % relative to either primitive alone [17]. This reduction lowers the cost of digital synthesis, but it leaves the per-plane computation in place, and that computation remains the dominant cost for continuous volumetric video at high frame rates and high depth-plane counts.

The clearest single illustration of this trade-off, reported for one static two-dimensional hologram at one depth plane, comes from a formal gradient-based optimization framework for phase-only hologram computation. A common fast heuristic encoding technique runs in around 10 milliseconds but at markedly reduced quality, while recovering higher accuracy through iterative gradient-based optimization takes on the order of 30 seconds on comparable hardware [18]. The same authors note that extending their optimization to true multi-depth-plane holograms, by summing multiple computed layers, costs computation proportional to the number of planes, and state plainly that this does not scale well in practice. Rather than pursue simultaneous multi-plane reconstruction, they substitute gaze-contingent single-plane refocusing: eye tracking determines a single depth, and the system computes and displays only that depth at any given moment [18]. This is a considered and effective answer for the near-eye context that work addresses. It illustrates how differently the same multi-plane reconstruction challenge can be approached. A single dynamically refocused plane stands in for true simultaneous multi-plane content in that near-eye work, against

generating a larger number of physically distinct planes directly – the route the present architecture pursues, targeting 49 planes available simultaneously within a single optical pass. Motivating that target is the observation that a modest number of representation layers – five, in one representative neural CGH implementation [3] – may not by itself supply the depth resolution a free-space volumetric display requires. These five layers and the present 49 planes describe physically different quantities, however, a compact internal representation versus physically separated projection planes, and this distinction should be borne in mind when comparing the two figures numerically.

5.2. Structural Comparison to Other Multichannel and Multiplexing Strategies

The efficiency figures reported for the present prototype (Section 3.2) fit a pattern also seen in other reported multichannel volume-holographic devices. The coherent beam combiner discussed in Section 1.3 reports a measured combined-beam output roughly 20 % below its own coupled-wave-theory prediction, attributed to real beams departing from the idealized plane-wave assumptions the theory rests on, together with practically unavoidable imperfections in angular alignment [11]. That offers a documented precedent, from a different device, for the diffraction-efficiency shortfall relative to the recording material's specified theoretical maximum observed here, suggesting the shortfall may be a general feature of angle-, phase-, or wavelength-multiplexed volume-holographic devices rather than something specific to the present fabrication run.

Different lines of work on related multi-viewpoint or multichannel problems have arrived at different multiplexing strategies. One recent approach to eliminating the eyepiece in near-eye holographic displays generates several viewpoints within the eye's pupil through time-division multiplexing, displaying each viewpoint's hologram in sequence, fast enough to exploit visual persistence, to avoid the coherent interference that simultaneous display under shared illumination would otherwise produce [19]. The present architecture multiplexes its channels simultaneously instead – spatially, by polarization, and by wavelength, through physically parallel optical paths rather than a time sequence – a different resolution of the same underlying requirement to generate and separate multiple coherent channels without mutual interference. A smaller-scale precedent for solving a capacity limit through added physical hardware rather than algorithms alone appears in a cascaded dual-spatial-light-modulator configuration used to synthesize a higher-resolution hologram than either device produces alone [18]. On a more modest scale, this reflects the same design instinct – extending optical capability through added physical structure – underlying the present multichannel trigger.

Content-adaptive resource allocation represents a further, different strategy, applied within a different class of display technology. Multilayer, attenuation-based light-field displays, built from stacked LCD panels under incoherent illumination rather than coherent holographic recording media, benefit from weighting reconstruction fidelity according to human visual acuity across the field of view [20]. For displays with a hard physical ceiling on reproducible depth range, they benefit further from automatically positioning that limited budget around the most visually salient scene content rather than a fixed window [21]. These represent a different philosophy applied to a related resource problem: dynamically prioritizing a small, fixed budget by scene content, rather than physically increasing the number of available channels as pursued here. This article notes them for the breadth of strategy the wider display research community has brought to related constraints, given how different the underlying display physics remains between the two device classes.

5.3. On Experimental Reproducibility and Calibration

The efficiency and crosstalk figures reported for the present prototype (Section 3.2) are characteristic of the fabrication run and emulsion batch tested, because the response of the emulsion varies between batches and over the life of a plate (Section 1.4). Any replication or extension of the architecture, including the higher channel counts discussed in Section 1.4, should therefore include routine recalibration.

5.4. Summary

Across the digital CGH methods reviewed in this article, substantial engineering effort has gone into narrowing the trade-off between reconstruction accuracy, achievable frame rate, and the number of depth planes reconstructed simultaneously, and each method shows incremental progress along one or more of these axes. The present architecture addresses the same underlying multi-depth-plane, multi-color reconstruction problem by a fundamentally different route: physically parallel, all-optical multiplexing rather than sequential digital computation. This route removes the per-plane computation from the frame pipeline, at the cost of the fabrication complexity and calibration burden discussed in Sections 2 and 5.3. The throughput remains bounded by the refresh rate of the source of the input frames, a HOLOEYE LETO-3-VIS-009 phase-only spatial light modulator [22], whose manufacturer-specified input frame rate is 60 Hz, or 180 Hz in color-field-sequential mode. The prototype of [1] operates it at 90 monochrome frames per second, which corresponds to 4410 depth-plane sub-frames per second for 49 planes. A complete optical power budget, source modulation,

synchronization, detector response, and scene updating are not characterized in this article. Neither route displaces the other outright: digital CGH retains an advantage in flexibility, since a trained model or algorithm can synthesize arbitrary new scene content without new optical hardware, whereas the HVGPU requires physical fabrication for new channel configurations but delivers its full channel count and depth-plane set from each input frame without per-plane computation, at the frame rate of the source. Both directions read best as complementary parts of a shared effort toward practical, real-time volumetric holographic display.

6. Conclusions

The architecture developed in this work demonstrates that physically parallel, all-optical multiplexing can address the multi-depth-plane, multi-color reconstruction problem without the per-plane computation that constrains digital CGH approaches. By performing depth- and color-channel multiplication entirely within the optical domain, the holographic HVGPU offers a route toward real-time volumetric reconstruction that scales with fabrication rather than computation, with the frame rate bounded by the source of the input frames. Although this approach carries its own fabrication and calibration demands, it complements algorithmic CGH methods by providing a hardware pathway to high-speed, multi-plane holography. Two measured results define the next optimization steps: the diffraction maxima lie 5–10 nm below the recording wavelengths, and swelling shifts of ±50 nm place some of the nine reconstruction wavelengths closer together than the measured $\Delta\lambda$. Aligning the maxima and spacing the shifts are objectives for subsequent work. Together, these directions outline a convergent trajectory toward practical, real-time volumetric display systems capable of supporting dynamic, multi-color holographic scenes.

References

[1]. G. Ivanova, B. Ivanov, D. Dimov, A 24-channel optical trigger for use in 3D holographic display systems, in *Proceedings of the 9th International Conference on Optics, Photonics and Lasers (OPAL)*, 2026, pp. 35-39.

[2]. B. Ivanov, G. Ivanova, D. Dimov, Design and fabrication of a 24-channel holographic optical trigger for high-speed optical computation, in *Proceedings of the 9th International Conference on Optics, Photonics and Lasers (OPAL)*, 2026, pp. 86-91.

[3]. L. Shi, B. Li, C. Kim, P. Kellnhofer, et al., End-to-end learning of 3D phase-only holograms for holographic display, *Light: Science & Applications*, Vol. 11, Issue 1, 2022, 247.

[4]. N. Padmanaban, Y. Peng, G. Wetzstein, Holographic near-eye displays based on overlap-add stereograms, *ACM Transactions on Graphics*, Vol. 38, Issue 6, 2019, 214.

[5]. A. Gilles, P. Gioia, R. Cozot, L. Morin, Computer generated hologram from multiview-plus-depth data considering specular reflections, in *Proceedings of the IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, 2016, pp. 1-6.

[6]. J. Mäkinen, U. Gudelek, L. Varjo, E. Sahin, et al., Computer-generated holography using hybrid planar-spherical wave primitives, *Optics Express*, Vol. 34, Issue 8, 2026, pp. 14780-14800.

[7]. S. Mustafi, T. Latychevskaia, Fourier transform holography: A lensless imaging technique, its principles and applications, *Photonics*, Vol. 10, Issue 2, 2023, 153.

[8]. B. Ivanov, M. Shopova, E. Stoykova, V. Sainov, Reconstruction of silver halide volume reflection holograms at the wavelengths of recording, *Optica Applicata*, Vol. 42, Issue 3, 2012, pp. 643-650.

[9]. B. Ivanov, M. Shopova, A. Baldjiev, E. Stoykova, et al., Four-channel holographic infrared optical element, *Proceedings of SPIE*, Vol. 7747, 2011, 77470K.

[10]. H. Wei, X. Liu, X. Hao, E. Y. Lam, et al., Modeling off-axis diffraction with the least-sampling angular spectrum method, *Optica*, Vol. 10, Issue 7, 2023, pp. 959-962.

[11]. H. N. Yum, P. R. Hemmer, A. Heifetz, J. T. Shen, et al., Demonstration of a multiwave coherent holographic beam combiner in a polymeric substrate, *Optics Letters*, Vol. 30, Issue 22, 2005, pp. 3012-3014.

[12]. D. Nazarova, L. Nedelchev, N. Berberova-Buhova, G. Mateev, Nanocomposite photoanisotropic materials for applications in polarization holography and photonics, *Nanomaterials*, Vol. 13, Issue 22, 2023, 2946.

[13]. P. Gentet, Y. Gentet, S.-H. Lee, Ultimate 04: The new reference for ultra realistic color holography, in *Proceedings of the International Conference on Emerging Trends & Innovation in ICT (ICEI)*, 2017, pp. 162-166.

[14]. P. Gentet, M. Coffin, Y. Gentet, S.-H. Lee, From text to hologram: Creation of high-quality holographic stereograms using artificial intelligence, *Photonics*, Vol. 11, Issue 9, 2024, 787.

[15]. U. Gudelek, E. Sahin, A. Gotchev, Perceptually optimized model for near eye light field reconstruction, in *Proceedings of the IEEE International Workshop on Multimedia Signal Processing (MMSP)*, 2023, pp. 1-6.

[16]. U. Gudelek, E. Sahin, A. Gotchev, A simulation framework for perceptual characterization of holographic near-eye displays, *Proceedings of SPIE*, Vol. 13821, 2026, 1382105.

[17]. A. Gilles, P. Gioia, R. Cozot, L. Morin, Complex modulation computer-generated hologram by a fast hybrid point-source/wave-field approach, in *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, 2015, pp. 4962-4966.

[18]. P. Chakravarthula, Y. Peng, J. Kollin, H. Fuchs, et al., Wirtinger holography for near-eye displays, *ACM Transactions on Graphics*, Vol. 38, Issue 6, 2019, 213.

[19]. J. Zhou, S. Xie, L. Jiang, Y. Luo, et al., High-quality, depth-cue-enhanced lensless holographic near-eye displays via pupil optimization, *Optics Letters*, Vol. 50, Issue 21, 2025, pp. 6811-6814.

[20]. J. Ding, M. Liu, Q. Zhong, H. Li, et al., Optimization algorithm of near-eye light field displays based on

human visual characteristics, *Chinese Optics Letters*, Vol. 14, Issue 4, 2016, 041101.
[21]. S. Wang, W. Liao, P. Surman, Z. Tu, et al., Salience guided depth calibration for perceptually optimized compressive light field 3D display, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 2031-2040.
[22]. LETO-3 phase only spatial light modulator, Datasheet, HOLOEYE Photonics AG, https://holoeye.com/products/spatial-light-modulators/leto-3-phase-only/

Published by International Frequency Sensor Association (IFSA) Publishing, S. L., 2026 (http://www.sensorsportal.com).

