


# A Phasor-based Camera-Space Construction for Making Radiometric Accumulation Assumptions Explicit

Yuxiang Sun<sup>id</sup>, Gladimir V. G. Baranoski<sup>id</sup>, Frank Fan<sup>id</sup>, and Scott Steinfield<sup>id</sup>

University of Waterloo, Canada



**Figure 1: Instantiation of radiometric accumulation assumptions in geometric optics-based rendering.** Conventional geometric optics-based rendering relies on implicit assumptions such as effectively incoherent illumination and sensor accumulation over long exposure times. We make the corresponding accumulation process explicit in a controlled camera-space construction. Starting from a ray-based radiometric input (a), we assign uniformly randomized phases to instantiate a camera-space complex field and propagate it through the aperture and lens. One sampled phase configuration produces an instantaneous sensor intensity (b) with prominent interference structure (c). We then model long-exposure sensor accumulation by averaging independent phase samples over time (d). Under this accumulation, the interference structure is progressively suppressed (e), yielding a stable radiometric result (f) consistent with the path-tracing reference. Although the accumulation returns to the radiometric endpoint, the construction makes the transient camera-space process available for analysis and use.

## Abstract

Geometric optics-based rendering relies on radiometric assumptions that yield linear intensity superposition. These assumptions include effectively incoherent illumination and accumulation over exposure times long compared to the optical period, yet they are usually absorbed into the standard image-formation model rather than explicitly isolated and examined. We present a controlled phasor-based camera-space construction that makes this accumulation process explicit through statistical phasor construction and propagation. This construction replaces only the primary camera segment. Starting from ray-based radiometric inputs whose phase is underdetermined, we assign randomized phases to instantiate statistical camera-space complex fields and propagate them through the aperture and lens to the sensor. Individual phase samples produce interference-dominated patterns on the sensor plane, while temporal or spectral accumulation progressively suppresses this structure. Under matched scene and camera configurations, the accumulated result agrees with the path-tracing reference, providing a fixed radiometric baseline for examining the randomized phase completion. Since the construction starts from path-traced irradiance, the resulting complex field is interpreted as a statistical camera-space completion rather than as a recovered physical source field, a scene-scale wave-transport solution, or a microscopic wave-material model. This restricted interpretation directs attention to the intermediate accumulation process itself while bounding the scope of the construction. Within this scope, the introduced phase degree of freedom gives transient camera-space interference a functional representation and provides a controlled, backward-compatible interface for rendering and analyzing phase-sensitive camera-space effects, including aperture-induced diffraction.

## CCS Concepts

• Computing methodologies → Rendering;

## 1. Introduction

Geometric optics-based rendering has been highly successful in research and industrial applications, even though light propagation can be more generally described in terms of waves. In practice, this success depends on radiometric assumptions under which wave-level fluctuations are not modeled explicitly in the final image-formation process. These assumptions include effectively incoherent illumination [Vea98], accumulation over exposure times long compared to the optical period, and wavelength-wise integration.

In computer graphics, such assumptions are usually adopted as part of the radiometric foundation of rendering, but they are less often isolated and examined in a controlled numerical setting. In this paper, we study this transition through a camera-space construction that makes the accumulation process explicit while preserving the usual geometric treatment elsewhere.

Concretely, we replace only the primary camera segment with statistical wave propagation and keep the remaining light transport in standard ray-based form. Working on top of path tracing, we begin from a radiometric input in camera space (Fig. 1). Since this input contains only radiometric energy, its phase is underdetermined. We therefore instantiate a camera-space complex field by assigning a randomized phase, propagate the resulting field through the aperture and thin lens, and examine both the instantaneous sensor intensity and its accumulated limit.

Within this setting, individual phase samples exhibit structured interference patterns on the sensor, whereas temporal or spectral accumulation progressively suppresses this structure. Under matched camera configurations, the accumulated result remains consistent with the stable irradiance predicted by conventional thin-lens path tracing. This agreement is used as a consistency baseline for the randomized phase completion, rather than as evidence that the missing physical phase has been recovered. The point of the construction is, therefore, not to replace geometric optics-based rendering with a complete wave-optical transport formulation, but to provide a controlled setting in which an otherwise implicit stage of radiometric image formation becomes directly analyzable.

We note that the scope of our work is deliberately restricted. We focus on the primary camera segment and introduce statistical wave propagation only there. Accordingly, the resulting complex field should be understood as a camera-space intermediate representation constructed from radiometric input, rather than as a physical source model or a general-purpose wave-optical rendering framework. This restriction keeps the construction compatible with existing ray-based rendering pipelines while exposing phase-sensitive camera-space propagation effects.

The above positioning guides the remainder of the paper. Sections 2 and 3 review related work and wave-optical background. Section 4 details the proposed construction, with implementation provided in Section 5. We evaluate radiometric consistency in Section 6 and analyze intermediate pattern suppression in Section 7. Finally, Section 8 discusses the construction's scope, limitations, and potential applications, concluding in Section 9.

## 2. Related Work

This work draws on the theoretical foundations of classical radiometry, wave optics, and numerical diffraction methods. Following a brief review of these core fields, we establish the scope and positioning of our approach to clarify how our restricted camera space construction complements existing rendering techniques.

### 2.1. Radiometry and the Geometric Rendering Equation

Radiometry [NRH\*77] provides the physical basis of geometric optics-based rendering. In computer graphics, geometric light transport is commonly expressed through the rendering equation introduced by Kajiya [Kaj86]. Although the original formulation did not explicitly adopt radiometric quantities, the rendering equation is commonly written in radiometric terms in modern physically based rendering [Vea98, PJH23].

### 2.2. Radiometry and Coherence

Radiometric quantities are traditionally defined under assumptions consistent with incoherent light [Wal68] and, therefore, do not explicitly account for coherence effects or phase information. As a result, wave-optical phenomena, such as interference and diffraction, lie outside the descriptive scope of classical radiometry. These limitations have long been recognized in the optics literature and have motivated the development of generalized quantities for coherent and partially coherent fields [Wal68, Wal73, Fri79]. For example, the Wigner Distribution Function [Bas79, CHB\*12] has been explored as a quantity that connects wave fields and radiometric rays.

### 2.3. Wave Optics in Computer Graphics Rendering

Since purely geometric rendering formulations cannot account for interference and diffraction, several works have incorporated wave phenomena into ray-based rendering pipelines. Examples include the embedding of wave-related effects into local appearance or shading models [DTS\*14, WVJH17, YHW\*18, BAGL19, SY22, YXW\*23, XWH\*23, YKPJ25]. Even within such local models, the finite spatial extent of coherent interaction can matter. Dhillon et al. [DTS\*14] explicitly incorporate this aspect through a finite Gaussian coherence window when rendering structural colors from measured biological nanostructures. At a more theoretical level, the rendering equation itself has also been generalized. For instance, Lessig [Les12] established modern mathematical foundations for light transport simulation, formalizing its underlying connections to wave optics. More recently, Steinberg et al. [SY21, SP26] proposed new transport formulations that explicitly extend classical rendering to wave optics.

### 2.4. Numerical Wave Simulation

Rigorous full-wave solvers such as Finite-Difference Time-Domain (FDTD) [TH00], Finite Element Method (FEM) [Jin14], and Boundary Element Method (BEM) [BTW84] are standard tools for wave propagation. However, they remain computationally infeasible for scene-scale rendering. Instead, we adopt layer-based methods, such as the one proposed by Zhao et al. [ZCZ\*15], which have been well developed in computer-generated holography [SCS05, VZT19].

## 2.5. Boundary and Positioning

Positioning this construction at the intersection of these fields requires clearly defining its boundaries. Regarding our theoretical foundation, we do not use radiometric quantities to derive new light transport formulations. They serve merely as inputs to explicitly unbox the implicit accumulation process. Similarly, analyzing mutual coherence is not our objective. We employ the random phasor sum model solely to ensure alignment with the geometric reference. In the context of strict physical wave simulation, our scene scale makes rigorous computation completely infeasible. While we borrow propagation methods from holography, we simulate macroscopically incoherent light rather than coherent wave fields. Within computer graphics, compared to existing wave-optical rendering, we have not yet proposed entirely novel visual effects. Given these boundaries, wave propagation is restricted entirely to the primary camera segment. Under this limitation, the construction retains a practical utility as a backward-compatible foundation for post-processing pipelines, as discussed further in Section 8.

## 3. Background

This section reviews the background from radiometry and wave optics that is needed for the rest of the paper. Rather than presenting a comprehensive physical optics framework, we establish the specific theoretical connections between wave fields and radiometric observables required to formulate our camera-space construction.

### 3.1. Radiometric Quantities

We briefly summarize the radiometric quantities used in this work to establish notation. Further details can be found in related literature [NRH\*77, Vea98, PJH23]. These quantities remain the reference language of the paper since both the precomputed input and the final accumulated result are interpreted radiometrically.

Radiant flux (or power)  $\Phi$  denotes the rate of radiant energy  $Q$  with respect to time, that is  $\Phi = dQ/dt$ .

Irradiance  $E_e(\mathbf{p})$  is defined as the radiant flux per unit area  $A$  incident at position  $\mathbf{p}$ ,  $E_e(\mathbf{p}) = \frac{d\Phi(\mathbf{p})}{dA}$ . In rendering, pixel values are obtained by integrating irradiance over the sensor area and exposure time.

Radiance  $L(\mathbf{p}, \omega)$  is a fundamental quantity of ray-based light transport, and it is defined as the radiant flux per unit projected area  $dA_\perp$  and per unit solid angle  $\omega$ :

$$L(\mathbf{p}, \omega) = \frac{d^2\Phi(\mathbf{p})}{dA_\perp d\omega}, \quad dA_\perp = (\mathbf{n} \cdot \omega) dA, \quad (1)$$

where  $\mathbf{n}$  denotes the surface normal at  $\mathbf{p}$ . Under geometric optics and in the absence of absorption, radiance is invariant along rays. This makes it a convenient quantity for modeling light transport.

Irradiance can equivalently be expressed as an integral of radiance over the incident hemisphere  $\Omega$ :

$$E_e(\mathbf{p}) = \int_{\Omega} L(\mathbf{p}, \omega) (\mathbf{n} \cdot \omega) d\omega, \quad (2)$$

which forms the basis of radiance-based image synthesis. Our construction is designed to remain aligned with this radiometric image-formation target at the accumulated level.

### 3.2. Wave-Optical Quantities

Building on the radiometric foundation outlined above, we introduce the corresponding wave-optical description that explicitly relates optical fields to the aforementioned radiometric observables.

In rendering applications, scene geometry is typically many orders of magnitude larger than the wavelength of light. Under this condition, electromagnetic wave propagation can be approximated using a scalar wave model, that is, by neglecting polarization [Goo17]. In this work, we use this approximation in the camera-space construction described later.

Within this scalar setting, light is described by a time-varying optical field  $E(\mathbf{p}, t)$ , corresponding to a single component of the electric field. For a (quasi-)monochromatic field with temporal frequency  $\nu$ , the physical field at position  $\mathbf{p}$  can be written as:

$$E(\mathbf{p}, t) = A(\mathbf{p}, t) \cos(2\pi\nu t - \varphi(\mathbf{p}, t)), \quad (3)$$

where  $A(\mathbf{p}, t)$  denotes the real-valued field amplitude and  $\varphi(\mathbf{p}, t)$  denotes the optical phase. Here, the optical frequency is given by  $\nu = c/\lambda$ , where  $\lambda$  denotes the wavelength.

For analytical convenience, it is common to factor out the rapidly oscillating carrier term and represent the field using a complex-valued amplitude, or phasor, expressed as:

$$U(\mathbf{p}, t) = A(\mathbf{p}, t) e^{j\varphi(\mathbf{p}, t)}, \quad (4)$$

where  $j$  is the imaginary unit. This representation is useful here because it makes amplitude and phase explicit in the camera-space construction.

This timescale difference is central to the definition of observable quantities. The instantaneous intensity associated with the optical field is defined as [Goo15]:

$$I(\mathbf{p}, t) = |U(\mathbf{p}, t)|^2. \quad (5)$$

However, optical oscillations occur far faster than the response time of typical image sensors. As a result, sensors do not resolve the optical field at this timescale, but instead record a temporal average of the field intensity. Accordingly, irradiance corresponds to the temporal average of the instantaneous intensity:

$$E_e(\mathbf{p}) \propto \langle I(\mathbf{p}, t) \rangle = \langle |U(\mathbf{p}, t)|^2 \rangle, \quad (6)$$

where  $\langle \cdot \rangle$  denotes a temporal average over durations long compared to the optical period. This distinction will later be used to interpret the intermediate sensor intensity and the final accumulated irradiance in the proposed construction.

### 3.3. Mutual Coherence and Linear Intensity Accumulation

In geometric rendering, light sources are typically treated as effectively incoherent [Vea98]. In statistical optics, assuming the optical field to be stationary and ergodic [Wol07], the spatial and temporal coherence of a field is characterized by the mutual coherence function (MCF). This is defined as the correlation of the complex field  $U$  at two spatial positions  $\mathbf{p}_1$  and  $\mathbf{p}_2$  separated by a time delay  $\tau$ :

$$\Gamma(\mathbf{p}_1, \mathbf{p}_2, \tau) = \langle U(\mathbf{p}_1, t) U^*(\mathbf{p}_2, t + \tau) \rangle, \quad (7)$$

where  $U^*$  is the complex conjugate of  $U$ .

Ideally, an optical field is strictly incoherent if  $\Gamma = 0$  for any two distinct spatial positions under an arbitrary time delay. Under this condition, the cross-interference terms between different field contributions vanish, leading to the linear addition of intensities.

### 3.4. Random Phasor Sums as a Statistical Foundation

To interface radiometric inputs with wave-optical propagation, it is necessary to assign complex-valued fields to purely energy-based quantities. To this end, we adopt the random phasor sum model, governed by three classical statistical assumptions [Goo20]:

1. Phasor contributions are statistically independent.
2. For each contribution, amplitude and phase are statistically independent.
3. Phases are uniformly distributed over  $(-\pi, \pi)$ .

Fields satisfying these conditions provide a statistical model for random phasor superposition and can lead to speckle-like intensity statistics after propagation. In this model, the phase-dependent cross-interference terms between distinct contributors vanish after averaging over sufficiently independent phase configurations. This property motivates the accumulated radiometric consistency examined later in the paper. After averaging, the remaining intensity is governed by the camera-space propagation response rather than by a coherent sum of phase-dependent interference terms.

The assumption of phase independence in this statistical model inevitably leads to an approximation of spatial-temporal decoupling. For a rigorous discussion on the coupling of these coherence properties, we refer readers to Wolf [Wol07], as it lies beyond the scope of this work.

### 3.5. Wave Propagation Model

Under scalar wave optics, propagation can be described by the Huygens–Fresnel principle [BWB19]. According to this principle, the optical field at a point in space is given by the superposition of secondary spherical waves emitted from a reference surface [Goo17]. For monochromatic scalar fields, the Rayleigh–Sommerfeld formulation provides a mathematically consistent expression of this principle [Goo17].

We note that the vanishing of coherence mentioned above is only between each contributor. However, propagation itself can still induce spatial coherence effects, as described by the Van Cittert–Zernike theorem [Wol07]. In other words, treating the source as effectively incoherent does not mean that this incoherence is fully preserved during propagation. Accordingly, the accumulated result should not be expected to be exactly identical to the radiometric reference. At the same time, the remaining camera-space propagation response is retained in the construction, which makes aperture-dependent behavior observable in the controlled setting.

In this paper, we use this wave-optical description only for the primary camera segment in a controlled camera-space construction. The remainder of the light path stays in its usual ray-based form. This is the propagation model that underlies the proposed construction described in the next section.

## 4. Controlled Wave-Optical Construction in Camera Space

This section describes the controlled construction used in our experiment. To compare wave-optical and ray-based image formation under matched radiometric conditions, we isolate the primary camera segment while preserving the usual radiometric treatment elsewhere. Fig. 2 provides an overview of the procedure, which consists of three stages: Procedure Input, Layered Scene Representation, and Ensemble Wave Simulation.

In surface rendering, a light path connecting the source to the sensor can be separated into the primary segment from the sensor to the first visible surface and the remaining light transport paths. The proposed construction replaces only the primary segment with a wave-optical model while keeping the rest of the path ray-based. As a result, wave propagation is confined to camera space, while global illumination is precomputed and stored as an irradiance texture.

### 4.1. Procedure Input

Our procedure begins with a scene (Fig. 2a) and two camera configurations (Fig. 2b), namely a pinhole setting and a thin-lens setting. The two share the same camera position and orientation, while the thin-lens setting defines the optical parameters used for both the path tracing reference and the wave-optical simulation.


**Path Tracing Reference Image** The reference image is obtained by rendering the scene with the thin-lens camera configuration using path tracing as shown in Fig. 2k. This image serves as the radiometric reference for the comparisons that follow.

**Camera Space Irradiance Texture** To generate the input for the wave-optical simulation, we use the pinhole camera setting to render an irradiance texture (Fig. 2c) and a depth map (Fig. 2d). These maps are pre-computed using path tracing and are subsequently transformed into the Layered Scene Representation (Fig. 2e) for the wave-optical propagation. Note that the irradiance texture is not itself a wave field. It contains only radiometric energy, while the phase information used in the wave simulation is constructed statistically afterward.

### 4.2. Layered Scene Representation

Directly evaluating the diffraction integral scales as  $O(N^4)$ , where  $N$  is the lateral resolution of the irradiance texture (Fig. 2c). This complexity stems from the pairwise contribution between  $N^2$  source points and  $N^2$  aperture points, rendering the approach impractical for high-resolution rendering. To address this, we adopt a layer-based method [ZCZ\*15] from computer-generated holography, transforming the irradiance texture (Fig. 2c) using the depth map (Fig. 2d) into a discrete layered representation (Fig. 2e). More specifically, we discretize the camera frustum into  $k$  parallel slices according to depth, and each pixel from the irradiance texture is assigned to the nearest slice using its depth value. This layered representation makes it possible to implement the propagation using Fast Fourier Transforms, reducing the complexity to  $O(k \cdot N^2 \log N)$ .





**Figure 2: Overview of our numerical experimental procedure.** The procedure begins with the scene geometry (a) and camera settings (b). We first generate a radiometric irradiance texture (c) and a depth map (d) using a pinhole camera model. These inputs are transformed into a layered scene representation via depth slicing (e), which discretizes the scene into parallel planes. Inside the simulation loop, we perform a single realization (trial) by assigning a uniformly randomized phase  $\phi \sim \mathcal{U}[-\pi, \pi]$  (f) to initialize the wave field on each slice (g). Through slice-wise propagation, these fields are accumulated to form the combined wave field at the aperture plane (h). Finally, this field is modulated and propagated to the sensor plane, yielding the instantaneous sensor intensity (i). Averaging these intensities over multiple realizations produces the final wave-optical irradiance (j), which closely matches the reference thin-lens path tracing result (k).

### 4.3. Ensemble Wave Simulation

In this section, we describe how the ensemble wave simulation produces the final sensor irradiance. We first introduce the averaging procedure and the corresponding ensemble-averaged quantity (Section 4.3.1). We then describe how the instantaneous intensity entering this average is constructed (Section 4.3.2) and propagated (Section 4.3.3).

#### 4.3.1. Incoherent Time Integration

As discussed in Section 3.2, the irradiance (Eq. 6) observed by the sensor (Fig. 3g) is the time average of the instantaneous intensity (Eq. 5). To simulate this process, we employ a Monte Carlo approach [HH64]. We refer to a single simulation pass with one sampled phase configuration as a realization. The final incoherent irradiance  $E_e^{\text{wave}}$  is obtained by averaging the instantaneous intensities over  $M$  independent realizations, which gives a finite-sample Monte Carlo estimate of the temporal average:

$$E_e^{\text{wave}}(\mathbf{p}) \propto \langle |U(\mathbf{p}, t)|^2 \rangle \approx \frac{1}{M} \sum_{m=1}^M |U_{\text{sensor}}^{(m)}(\mathbf{p})|^2. \quad (8)$$

In the following subsections, we describe how the instantaneous intensity  $|U_{\text{sensor}}^{(m)}|^2$  is computed for one realization.

#### 4.3.2. Complex Amplitude Initialization


A representation mismatch arises because numerical wave propagation requires a defined complex field, whereas the irradiance texture (Fig. 3a) provides only a time-averaged statistical quantity. Since the mapping from irradiance to a specific field realization is underdetermined, we adopt the following assignment consistent with the assumptions in Section 3.4 [Goo20]. Accordingly, we write the field at a point  $\mathbf{p}$  for the  $m$ -th realization as:

$$U^{(m)}(\mathbf{p}) = \sqrt{E_e(\mathbf{p})} e^{j\phi_m(\mathbf{p})}, \quad \phi_m \sim \mathcal{U}[-\pi, \pi]. \quad (9)$$

Here, the random phase  $\phi_m$  introduces variation across realizations, while the amplitude  $\sqrt{E_e}$  ensures intensity consistency.

#### 4.3.3. Propagation Sequence

Once constructed, the complex field is propagated through the camera-space setup shown in Fig. 3. Within a single realization, the field is treated as coherent for propagation. As a result, the wave contributions originating from different scene slices (Fig. 3a) can be linearly superposed and accumulated at the aperture plane (Fig. 3d) after propagation from each slice. The accumulated field is then modulated by the optical elements, including the binary aperture mask (Fig. 3e) and a thin-lens phase factor (Fig. 3f), given



**Figure 3: Camera-space propagation setup.** The scene geometry is discretized into parallel depth slices (a). The wave field propagates from these slices to the aperture plane (d), where it is modulated by a binary aperture mask (e) and a thin-lens phase factor (f), before reaching the sensor plane (g). Blue dashed lines (b) indicate the geometric pinhole projection used for input precomputation, while red-blue gradient lines (c) represent the region of wave field propagation.

by  $e^{-j\pi v(x^2+y^2)/(cf)}$ , where  $\mathbf{p} = (x, y, z)$  denotes the spatial position and  $f$  is the focal length of the lens. The modulated field is subsequently propagated to the sensor plane (Fig. 3g), where its squared magnitude yields the instantaneous sensor intensity, denoted by  $|U_{\text{sensor}}^{(m)}|^2$ .

## 5. Implementation Details

This section summarizes the implementation and experimental settings used for reproducibility and to illustrate the computational scale of our simulations (e.g., up to  $8192^2$  resolution), rather than performance evaluation.

### 5.1. Code Implementation

The proposed camera-space construction is implemented using Mitsuba 3 [JSR\*22], NumPy [HMW\*20], CuPy [OUN\*17], and the Diffractio library [SBSGAP\*24]. Wave propagation is evaluated with a Rayleigh–Sommerfeld kernel [Goo17, Och17, SB25] using the Chirp–Z transform [HWW\*20], which is convenient for the sampling and regions of interest required in our setup. Our reference implementation is publicly available [SBFS26].

### 5.2. Scene Configuration


We evaluate the proposed construction on four scenes from the Mitsuba 3 gallery [JSR\*22]. They are the Cornell Box, Veach MIS, Spaceship, and Living Room scenes. Table 1 lists the configuration for each of these scenes used in Section 6 and Figs. 4 and 5, including the spectral range, number of spectral samples ( $N_\lambda$ ), number of axial slices ( $N_z$ ), number of realizations ( $M$ ), field of view (FOV), and texture resolution. These settings define the computational scale of the construction used in the consistency and pattern analyses that follow.

## 6. Consistency Evaluation

Before analyzing the intermediate patterns produced by the construction, we first verify that its accumulated limit remains aligned with conventional thin-lens path tracing under matched optical configurations. This consistency is not the end goal of the paper by itself. Instead, it indicates that the proposed construction provides an appropriate setting for examining the pattern-formation and suppression process discussed in the following sections.

### 6.1. Visual Comparisons

In Figs. 4 and 5, we compare the accumulated result of our construction (left column) with the thin-lens path tracing reference (right column) across several scenes. As shown, the two sets of images demonstrate close visual agreement.




**Figure 4: Cornell Box visual comparisons.** Accumulated results produced by the proposed camera-space wave propagation construction (left column) and the reference thin-lens path tracing results (right column). The top row shows the monochromatic 550 nm setting, and the bottom row shows the visible-spectrum setting.

**Cornell Box** We use the Cornell Box as a base case because it contains only Lambertian reflection and emission. We evaluate both a monochromatic setting at 550 nm (Figs. 4a and 4b) and a visible-spectrum integration from 390 nm to 710 nm at 5 nm intervals (Figs. 4c and 4d). In both cases, our accumulated results show close visual agreement with the thin-lens path tracing reference. A small residual difference appears around the ceiling light in our result (e.g., Figs. 4a and 4c), which is related to the aperture-dependent propagation response discussed in Section 8.

| Scene       | Spectrum        | $N_\lambda$ | $N_z$ | $M$  | FOV   | Resolution  | PSNR  | SSIM   |
|-------------|-----------------|-------------|-------|------|-------|-------------|-------|--------|
| Cornell Box | 550 nm          | 1           | 128   | 4096 | 19.5° | 4096 × 4096 | 40.38 | 0.9965 |
| Cornell Box | 390 nm - 710 nm | 64          | 64    | 64   | 19.5° | 4096 × 4096 | 40.41 | 0.9961 |
| Veach MIS   | 390 nm - 710 nm | 64          | 64    | 64   | 19.5° | 4096 × 4096 | 39.46 | 0.9911 |
| Spaceship   | 390 nm - 710 nm | 64          | 64    | 64   | 19.5° | 4096 × 4096 | 46.44 | 0.9962 |
| Living Room | 390 nm - 710 nm | 64          | 64    | 64   | 50.0° | 8192 × 8192 | 44.61 | 0.9877 |

**Table 1: Numerical configurations and consistency metrics for the evaluated scenes.** Here,  $N_\lambda$  denotes the number of spectral samples,  $N_z$  the number of axial slices, and  $M$  the number of realizations. PSNR and SSIM represent the Peak Signal-to-Noise Ratio and Structural Similarity Index, respectively.



**Figure 5: Additional visual comparisons.** Accumulated results produced by the proposed camera-space wave propagation construction (left column) and the reference thin-lens path tracing results (right column) across scenes with specular transport, transmission, and a larger field of view.

**Additional Scenes** We further evaluate our construction across a diverse set of scenes to test different transport phenomena and camera configurations (Fig. 5). The Veach MIS scene (Figs. 5a and 5b) introduces specular transport. Since its brightest light source region is approximately  $10^4$  times more intense than the background, the residual difference near the source becomes more visible here. The Spaceship scene (Figs. 5c and 5d) tests light transmission. The Living Room scene (Figs. 5e and 5f) provides a broader environmental context rendered with a larger field of view.

## 6.2. Quantitative Comparisons

We use the values calculated for the PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index) metrics [WBSS04] as supporting evidence for accumulated-level agreement, rather than as the primary contribution of the work. Before evaluation, we apply a linear global exposure alignment to account for overall intensity differences. All images are resized to  $1024 \times 1024$  for conformity with the reference resolution. For visible-spectrum results, the comparisons are performed on the Y channel in XYZ space after exposure alignment and log transformation. The values listed in Table 1 show that the accumulated images remain in close agreement with the thin-lens reference across all test scenes.

## 6.3. Observed Low-Pass Behavior

When the accumulated results are examined at a magnified scale, they show a slight low-pass behavior compared with the thin-lens path tracing reference. This observation can likely be attributed to the finite-grid camera-space construction used in our simulation, rather than an additional image-space filtering step. We record it here as an observed limitation and defer its detailed interpretation to Section 8.

## 6.4. Summary

Across all tested scenes, the accumulated results of our construction remain consistent with the thin-lens path tracing reference, up to a small residual difference in some cases near bright sources. This is the main point of the present section. It establishes that the proposed construction remains aligned with the standard radiometric image at the accumulated level. With that accumulated-level consistency established, we can now examine the intermediate pattern-suppression process within the construction.



## 7. Intermediate Interference Patterns


Having established accumulated-level consistency, we now turn to the intermediate sensor patterns produced before that limit is reached. In single realizations, the sensor intensity is not dominated by ordinary Monte Carlo noise, but by realization-dependent interference structure induced by the statistical wave construction. The purpose of this section is to characterize that structure and to examine how it is progressively suppressed under accumulation.

### 7.1. Interference Patterns in a Single Realization


Fig. 6 shows the instantaneous sensor intensity from a single realization. The image contains visible interference structure that should not be interpreted as ordinary path-tracing noise. Instead, it arises from coherent interference within one realization of the statistical field construction in camera space. As Fig. 7 shows, this pattern differs qualitatively from Monte Carlo noise in both spatial scale and connectivity.

Fig. 7a provides a zoomed-in view of Fig. 6 around the upper-left corner of the shorter box. The observed interference pattern differs from conventional Monte Carlo noise produced by path tracing, as shown in Fig. 7b.

More specifically, the interference patterns appear coarser and exhibit more spatial connectivity, whereas Monte Carlo noise from path tracing consists of finer and more isolated grain structures. Note that both images are shown at the same resolution.



**Figure 6: Single realization.** Instantaneous sensor intensity obtained from a single realization. The visible interference patterns differ from Monte Carlo noise in path tracing.




**Figure 7: Interference pattern vs. Monte Carlo noise.** Images shown for comparison, with (a) showing an interference pattern from a single realization as zoomed-in views of the scene depicted in Fig. 6 and (b) showing a typical Monte Carlo path tracing noise pattern cropped from the same image region.

### 7.2. Statistical Signature of the Single Realization Pattern


To further characterize the single realization shown in Fig. 6, we examine the statistics of its instantaneous intensity. Specifically, we select the light-source region in the Cornell Box and plot the normalized histogram shown in Fig. 8. The observed distribution is consistent with a negative-exponential profile. In addition, the contrast of the instantaneous intensity, computed as the standard deviation divided by the mean, is 0.9762. These statistics are similar to those commonly associated with fully developed speckle.

These statistics suggest that a single-realization image (e.g., Fig. 6) is not simply ordinary noise, but contains a realization-dependent interference pattern.



**Figure 8: Statistical signature.** Normalized intensity distribution in a single realization, measured in the light-source region of Fig. 6. The observed distribution is consistent with the negative-exponential distribution  $p(x) = e^{-x}$ .





**Figure 9: Suppression through accumulation.** (a) Temporal suppression of the interference pattern under temporal accumulation. (b) Spectral suppression of the interference pattern under spectral accumulation. (c) Monte Carlo noise reduction in path tracing.

### 7.3. Suppression Through Accumulation

We now examine how the realization-dependent interference pattern identified above is progressively suppressed under accumulation. Figs. 9a, 9b, and 9c show this process in our monochromatic setting, in our spectral setting, and in path tracing, respectively. Although these processes look different, they all become stable as accumulation proceeds. In our construction, the stable image does not come from any single realization. It is obtained only through accumulation, which gradually averages out the interference structure visible in single realizations.

**Temporal Suppression of the Interference Pattern** Fig. 9a shows the temporal accumulation process in our monochromatic setting at 550 nm. Starting from a single realization, the image contains visible interference patterns. As more realizations are accumulated, these patterns are progressively suppressed. The final result approaches the stable irradiance image shown in Fig. 4a.

**Spectral Suppression of the Interference Pattern** Fig. 9b shows a similar suppression process under spectral accumulation from 390 nm to 710 nm. The result is visualized exclusively using the Y channel in XYZ space to prevent color shifts from confounding visual inspection of the pattern suppression. As contributions from different wavelengths are integrated, the interference patterns are also progressively suppressed. The final result approaches the spectral image shown in Fig. 4c. Compared with Fig. 9a, the left-most subfigure here has a smaller pattern size, which is consistent with the shorter wavelength at the beginning of the spectral integration. For a more rigorous and complete treatment of spectral suppression and wavelength-dependent coherence effects beyond the scope of this work, readers are referred to Wolf [Wol07].

**Monte Carlo Noise Reduction** Fig. 9c shows the usual convergence process in path tracing. As the number of samples increases, Monte Carlo noise is gradually reduced, and the result approaches the reference image in Fig. 4b.

### 7.4. Summary

Overall, this section shows that single realizations produce interference structure on the sensor and that this structure is progressively suppressed under temporal or spectral accumulation. The stable irradiance image is obtained only after this suppression. In this way, the proposed construction makes it possible to examine an intermediate stage of image formation that is usually implicit in geometric rendering. The observed patterns are interpreted here at the level of statistical interference behavior within the present construction, not as a complete physical account of speckle formation.

## 8. Discussion

The intermediate patterns discussed above require careful interpretation. The main question is not only what the proposed construction produces, but also what kinds of conclusions it supports. In this section, we clarify the role of the construction, the scope of the resulting patterns, and a limited secondary observation enabled by the same setup.

### 8.1. Relation to Existing Approaches

The proposed construction serves a different role from both rigorous full-wave solvers and full-path wave-optical rendering formulations. Methods such as FDTD [TH00], FEM [Jin14], and BEM

[BTW84] aim at first-principles wave simulation but remain difficult to apply at the scene scale. Existing full-path wave-optical formulations in computer graphics [SY21, SP26] are typically expressed through transport-level statistical or phase-space quantities. By contrast, our method simulates the explicit complex field using a configuration-space formulation, restricted entirely to camera space. For the present purpose, it is more intuitive because it keeps the intermediate interference structure visible and allows us to examine its suppression under accumulation before the stable radiometric image is obtained.

## 8.2. Scope

The scope of the proposed construction is intentionally narrow. It supports statements about how interference structure appears in single realizations and how that structure is suppressed through temporal or spectral accumulation within the present camera-space setting. Accordingly, the results in this paper should be interpreted as results obtained within this construction and under its assumptions. They are not presented as direct statements about full scene-scale wave transport or about wave-optical interactions distributed across the entire light path. The scene-side transport remains in its usual ray-based form, and the wave-optical component is introduced only along the primary camera segment. This restriction is part of the design. It isolates the camera-space accumulation process and makes the relevant assumptions directly analyzable.

## 8.3. Interpretation of the Observed Patterns

The patterns produced by the proposed construction are best interpreted as realization-dependent interference structure arising within statistical camera-space propagation. In this sense, they are meaningful features of the construction rather than arbitrary numerical artifacts. Their intensity statistics are similar to those commonly associated with fully developed speckle, and their appearance and accumulation behavior distinguish them from ordinary image noise. In the present work, this is the level at which these patterns are interpreted.

At the same time, the construction does not model microscopic wave-material interactions. The scene-side transport remains in its usual geometric form, and statistical wave propagation is introduced only in camera space. For this reason, the resulting patterns should not be understood as carrying material-specific wave information. They should instead be understood as statistical consequences of the construction.

## 8.4. Limitations from the Finite-Grid Construction

The finite-grid construction introduces several limitations for the physical interpretation of the observed patterns. In the ideal incoherence condition introduced in Section 3.3, the mutual coherence is zero between any two distinct spatial positions under arbitrary time delay. Our numerical construction does not instantiate this continuous limit directly. Instead, one random phase is assigned to each discrete radiometric sample.

As a result, each sample acts as a finite coherent cell rather than an infinitesimal point of an ideal incoherent source. This explains

why the patterns shown here should not be interpreted as realistic physical speckles. In real physical settings, speckle size typically depends on propagation distance, depth, aperture, sensor integration, and the coherence properties of the source. That behavior is not fully reflected in our images. In particular, the observed pattern scale is tied to the discrete camera-space representation and therefore does not reproduce the expected depth-dependent speckle-size variation.

The same finite-grid setting may also affect the accumulated image. Under the discrete random-phasor construction, accumulation suppresses phase-dependent interference only between distinct discrete samples, rather than between arbitrary spatial positions in a continuous field. Thus, the accumulated image should not be interpreted as the exact continuous-limit result of an ideal incoherent source. Even after this discrete cancellation, the accumulated image is still shaped by the sampled camera-space propagation response. The detailed form of this response depends on the propagation operator, the camera-space discretization, and the optical configuration, including the aperture and thin-lens phase. Evaluating this propagation on a finite grid inherently limits the representable spatial frequencies. This provides one possible finite-grid-related explanation for the mild low-pass appearance observed under magnified inspection.


Another related limitation is the lack of explicit sensor-side spatial integration. A physical sensor would integrate irradiance over a finite pixel area, while the present implementation evaluates the propagated field on a discrete sensor grid. Modeling this integration, together with a finer scene-side representation, would require substantially higher spatial resolution and is beyond the scope of the present setting.

These limitations matter for physical interpretation, but they do not change the main point of the paper. The proposed construction is intended to make the realization-dependent interference structure and its suppression under temporal or spectral accumulation directly analyzable within a controlled camera-space setting. The results should therefore be understood within this construction and at that level of interpretation.

## 8.5. Aperture-Dependent Effects as a Camera-Space Utility

After accumulation, phase-dependent cross-interference terms are suppressed, but the spatial interference structure can remain in the accumulated result. This makes the construction useful as a camera-space interface for aperture-dependent effects, rather than merely as an image-space filtering operation. Since the propagated field interacts directly with the camera aperture before accumulation, changes in aperture geometry can remain visible in the accumulated result. As shown in Fig. 10, changing the pupil to a regular hexagon produces distinct spike patterns.

We include this example to illustrate a possible utility of the construction, rather than as a developed contribution to lens-induced artifact rendering. Camera and lens effects, such as wave-based lens flare, have been studied as dedicated rendering problems, for example by Hullin et al. [HESL11]. A fuller treatment of bloom, diffractive spikes, or other lens-induced effects would require a separate study and lies outside the scope of this work.



**Figure 10: Aperture-related diffraction effects.** After changing the pupil to a regular hexagon, diffractive spikes are naturally induced by the edge of the pupil.

### 8.6. Practical Role of the Camera-Space Interface

Full scene-scale wave-optical transport would be a more complete solution in principle. However, existing rendering pipelines, assets, and workflows are deeply tied to geometric optics-based rendering. A direct transition to a fully wave-optical pipeline is therefore unlikely to be practical in many settings.

The present construction is intended to occupy this gap. It starts from a conventional radiometric rendering result, statistically completes the missing phase, and propagates the resulting complex field only in camera space. The completed field should not be interpreted as the true optical field of a physical light source, nor as a replacement for full wave-optical transport. Instead, it provides a restricted, backward-compatible interface for introducing phase-sensitive camera-space propagation when only radiometric rendering input is available.

This role is similar in spirit to spectral reconstruction from RGB values. The reconstruction is underdetermined and not a fully physical measurement, but it can still be useful when it preserves the original RGB values while exposing additional degrees of freedom for spectral processing. Similarly, the randomized phase completion is not unique or fully physical, but it provides a concrete field representation that can be propagated while remaining consistent with the original radiometric input after accumulation.

This perspective also suggests extensions toward structured phase relationships or partial coherence. Relaxing the independent random-phase assumption to explore structured phase relationships or partial coherence is left as a direction for future work.

## 9. Conclusion

In this paper, we introduced a controlled camera-space construction for making radiometric accumulation assumptions explicit. Working on top of path tracing, the construction replaces only the primary camera segment with statistical wave propagation, while keeping the remaining transport in its usual ray-based form.

Within this setting, we established two main results. First, the accumulated result remains consistent with the reference thin-lens path tracing image under matched camera configurations. This agreement should be interpreted as a consistency baseline for the underdetermined randomized phase completion, rather than as evidence that the missing physical phase has been recovered. Second, individual phase samples produce interference structure on the sensor, and this structure is progressively suppressed through temporal or spectral accumulation before the stable radiometric image is ultimately obtained.

Overall, the construction exposes an intermediate pattern-formation and suppression process that is usually implicit in radiometric image formation. The resulting complex field should be understood as a camera-space intermediate representation constructed from radiometric input, rather than as a recovered physical phase. By preserving the accumulated radiometric objective and modifying only the primary camera segment, the construction remains backward-compatible with existing geometric rendering pipelines. Future work may extend this interface toward richer camera-space wave effects, partial coherence, sensor-side integration, and wave-material interactions.

## Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada under Grant 238337. The authors would like to thank their families. Furthermore, we sincerely thank the EGSR 2026 community and the anonymous reviewers of both the current and earlier versions of this manuscript, whose valuable suggestions significantly improved this paper.

## References

- [BAGL19] BAR C., ALTERMAN M., GKIOULEKAS I., LEVIN A.: A Monte Carlo framework for rendering speckle statistics in scattering media. *ACM Transactions on Graphics* 38, 4 (Aug. 2019), 1–22.
- [Bas79] BASTIAANS M. J.: Wigner distribution function and its application to first-order optics. *Journal of the Optical Society of America* 69, 12 (Dec. 1979), 1710.
- [BTW84] BREBBIA C. A., TELLES J. C. F., WROBEL L. C.: *Boundary Element Techniques*. Springer, Berlin, Heidelberg, 1984.
- [BWB19] BORN M., WOLF E., BHATIA A. B.: *Principles of Optics: Electromagnetic Theory of Propagation, Interference, and Diffraction of Light*, seventh ed. Cambridge University Press, Cambridge, 2019.
- [CHB\*12] CUYPERS T., HABER T., BEKAERT P., OH S. B., RASKAR R.: Reflectance model for diffraction. *ACM Transactions on Graphics* 31, 5 (Aug. 2012), 1–11.
- [DTS\*14] DHILLON D., TEYSSIER J., SINGLE M., GAPONENKO I., MILINKOVITCH M., ZWICKER M.: Interactive diffraction from biological nanostructures. *Computer Graphics Forum* 33, 8 (Dec. 2014), 177–188.



- [Fri79] FRIBERG A. T.: On the existence of a radiance function for finite planar sources of arbitrary states of coherence. *Journal of the Optical Society of America* 69, 1 (Jan. 1979), 192–198.
- [Goo15] GOODMAN J. W.: *Statistical Optics*, second ed. Wiley Series in Pure and Applied Optics. John Wiley & Sons, Hoboken, New Jersey, 2015.
- [Goo17] GOODMAN J. W.: *Introduction to Fourier Optics*, fourth ed. W.H. Freeman, Macmillan Learning, New York, 2017.
- [Goo20] GOODMAN J. W.: *Speckle Phenomena in Optics: Theory and Applications*. SPIE, Bellingham, Washington, USA, Jan. 2020.
- [HESL11] HULLIN M., EISEMANN E., SEIDEL H.-P., LEE S.: Physically-based real-time lens flare rendering. *ACM Transactions on Graphics* 30, 4 (July 2011), 108:1–108:10.
- [HH64] HAMMERSLEY J. M., HANDSCOMB D. C.: *Monte Carlo Methods*. Springer Netherlands, Dordrecht, 1964.
- [HMW\*20] HARRIS C. R., MILLMAN K. J., WALT S. J. V. D., GOMMERS R., VIRTANEN P., COUNAPEAU D., WIESER E., TAYLOR J., BERG S., SMITH N. J., KERN R., PICUS M., HOYER S., KERKWIJK M. H. V., BRETT M., HALDANE A., RÍO J. F. D., WIEBE M., PETERSON P., GÉRARD-MARCHANT P., SHEPPARD K., REDDY T., WECKESSER W., ABBASI H., GOHLKE C., OLIPHANT T. E.: Array programming with NumPy. *Nature* 585, 7825 (Sept. 2020), 357–362.
- [HWW\*20] HU Y., WANG Z., WANG X., JI S., ZHANG C., LI J., ZHU W., WU D., CHU J.: Efficient full-path optical calculation of scalar and vector diffraction using the Bluestein method. *Light: Science & Applications* 9, 1 (July 2020), 119.
- [Jin14] JIN J.-M.: *The Finite Element Method in Electromagnetics*, third ed. John Wiley & Sons, Hoboken, New Jersey, 2014.
- [JSR\*22] JAKOB W., SPEIERER S., ROUSSEL N., NIMIER-DAVID M., VICINI D., ZELTNER T., NICOLET B., CRESPO M., LEROY V., ZHANG Z.: Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>.
- [Kaj86] KAJIYA J. T.: The rendering equation. *SIGGRAPH Comput. Graph.* 20, 4 (Aug. 1986), 143–150.
- [Les12] LESSIG C.: *Modern Foundations of Light Transport Simulation*. PhD thesis, University of Toronto, 2012.
- [NRH\*77] NICODEMUS F. E., RICHMOND J. C., HSIA J. J., GINSBERG I. W., LIMPERIS T.: *Geometrical considerations and nomenclature for reflectance*. Tech. Rep. NBS MONO 160, National Bureau of Standards, Gaithersburg, MD, 1977.
- [Och17] OCHOA N. A.: Alternative approach to evaluate the Rayleigh-Sommerfeld diffraction integrals using tilted spherical waves. *Optics Express* 25, 10 (May 2017), 12008.
- [OUN\*17] OKUTA R., UNNO Y., NISHINO D., HIDO S., LOOMIS C.: CuPy: A NumPy-compatible library for NVIDIA GPU calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)* (Long Beach, California, USA, 2017), Neural Information Processing Systems, pp. 1–7.
- [PJH23] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, fourth ed. The MIT Press, Cambridge London, 2023.
- [SB25] SUN Y., BARANOSKI G. V. G.: Glare pattern depiction: high-fidelity physical computation and physiologically-inspired visual response. *ACM Transactions on Graphics* 44, 6 (Dec. 2025), 1–14.
- [SBFS26] SUN Y., BARANOSKI G., FAN F., STEINFELD S.: A Phasor-based Camera-Space Construction for Making Radiometric Accumulation Assumptions Explicit - Reference Implementation, 2026. <https://github.com/sunyuxiangyx/phasor-construction-egsr-2026>.
- [SBSGAP\*24] SANCHEZ-BREA L. M., SORIA-GARCIA A., ANDRES-PORRAS J., PASTOR-VILLARRUBIA V., ELSHORBAGY M. H., MUÑOZ J. D. H., TORCAL-MILLA F. J., ALDA J.: Diffractio: An open-source library for diffraction and interference calculations. In *Optics and Photonics for Advanced Dimensional Metrology III* (Bellingham, Washington, USA, 2024), Groot P. J. d., Guzman F., Picart P., (Eds.), SPIE, p. 129971B.
- [SCS05] SLINGER C., CAMERON C., STANLEY M.: Computer-generated holography as a generic display technology. *Computer* 38, 8 (Aug. 2005), 46–53.
- [SP26] STEINBERG S., PHARR M.: Wave tracing: Generalizing the path integral to wave optics. *Computer Graphics Forum* (Apr. 2026), e70322.
- [SY21] STEINBERG S., YAN L.-Q.: A generic framework for physical light transport. *ACM Transactions on Graphics* 40, 4 (Aug. 2021), 1–20.
- [SY22] STEINBERG S., YAN L.-Q.: Rendering of subjective speckle formed by rough statistical surfaces. *ACM Transactions on Graphics* 41, 1 (Feb. 2022), 1–23.
- [TH00] TAFLOVE A., HAGNESS S. C.: *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, second ed. Artech House antennas and propagation library. Artech House, Boston, 2000.
- [Vea98] VEACH E.: *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford University, 1998.
- [VZT19] VELEZ-ZEA A., TORROBA R.: Optimized random phase tiles for non-iterative hologram generation. *Appl. Opt.* 58, 32 (Nov. 2019), 9013–9019.
- [Wal68] WALTHER A.: Radiometry and coherence. *Journal of the Optical Society of America* 58, 9 (Sept. 1968), 1256.
- [Wal73] WALTHER A.: Radiometry and coherence. *Journal of the Optical Society of America* 63, 12 (Dec. 1973), 1622–1623.
- [WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (Apr. 2004), 600–612.
- [Wol07] WOLF E.: *Introduction to the Theory of Coherence and Polarization of Light*. Cambridge University Press, Cambridge, 2007.
- [WVJH17] WERNER S., VELINOV Z., JAKOB W., HULLIN M. B.: Scratch iridescence: wave-optical rendering of diffractive surface structure. *ACM Transactions on Graphics* 36, 6 (Nov. 2017), 1–14.
- [XWH\*23] XIA M., WALTER B., HERY C., MAURY O., MICHIELSSEN E., MARSCHNER S.: A practical wave optics reflection model for hair and fur. *ACM Transactions on Graphics* 42, 4 (Aug. 2023), 1–15.
- [YHW\*18] YAN L.-Q., HAŠAN M., WALTER B., MARSCHNER S., RAMAMOORTHY R.: Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics* 37, 4 (Aug. 2018), 1–10.
- [YKPI25] YANG R., KIM J., PEDIREDLA A., JAROSZ W.: A wave-optics BSDF for correlated scatterers. *Computer Graphics Forum* 44, 4 (July 2025), e70167.
- [YXW\*23] YU Y., XIA M., WALTER B., MICHIELSSEN E., MARSCHNER S.: A full-wave reference simulator for computing surface reflectance. *ACM Transactions on Graphics* 42, 4 (Aug. 2023), 1–17.
- [ZCZ\*15] ZHAO Y., CAO L., ZHANG H., KONG D., JIN G.: Accurate calculation of computer-generated holograms using angular-spectrum layer-oriented method. *Optics Express* 23, 20 (Oct. 2015), 25440–25449.