

A Monte Carlo Rendering Framework for Simulating Optical Heterodyne Detection

JUHYEON KIM, Dartmouth College, USA

CRAIG BENKO, MAGNUS WRENNINGE, RYUSUKE VILLEMIN, and ZEB BARBER, Aurora Innovation, USA

WOJCIECH JAROSZ and ADITHYA PEDIREDLA, Dartmouth College, USA

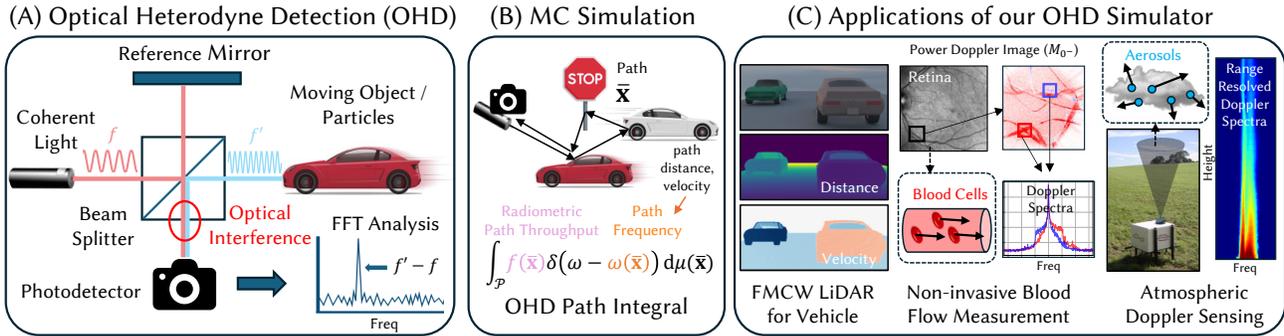


Fig. 1. We propose a ray-tracing based Monte Carlo rendering framework for simulating optical heterodyne detection (OHD). (A) OHD is a technique that measures the frequency modulation of light using optical interferometry. Through spectral analysis of the photon-induced current on the photodetector, we can obtain useful information such as object velocity/distance or Doppler spectra. (B) Our MC simulation is based on the *OHD path integral* that resembles radiometric path integral [Veach 1997], but is resolved by the path frequency which depends on the path length and velocity. (C) Using our simulator, one can simulate various OHD scenarios, such as coherent lidar used in autonomous vehicles, non-invasive blood flow measurement, and atmospheric Doppler sensing.

Optical heterodyne detection (OHD) employs coherent light and optical interference techniques (Fig. 1-(A)) to extract physical parameters, such as velocity or distance, which are encoded in the frequency modulation of the light. With its superior signal-to-noise ratio compared to incoherent detection methods, such as time-of-flight lidar, OHD has become integral to applications requiring high sensitivity, including autonomous navigation, atmospheric sensing, and biomedical velocimetry. However, current simulation tools for OHD focus narrowly on specific applications, relying on domain-specific settings like restricted reflection functions, scene configurations, or single-bounce assumptions, which limit their applicability. In this work, we introduce a flexible and general framework for spectral-domain simulation of OHD. We demonstrate that classical radiometry-based path integral formulation can be adapted and extended to simulate the OHD measurements in the spectral domain. This enables us to leverage the rich modeling and sampling capabilities of existing Monte Carlo path tracing techniques. Our formulation shares structural similarities with transient rendering but operates in the spectral domain and accounts for the Doppler effect (Fig. 1-(B)). While simulators for the Doppler effect in incoherent (intensity) detection methods exist, they are largely not suitable to simulate OHD. We use a microsurface interpretation to show that these two Doppler imaging techniques capture different physical quantities and thus need different simulation frameworks. We validate the correctness and predictive power of our simulation framework by qualitatively comparing the simulations

Authors' addresses: Juhyeon Kim, Dartmouth College, USA, juhyeon.kim.gr@dartmouth.edu; Craig Benko, cbenko@aurora.tech; Magnus Wrenninge, mwrenninge@aurora.tech; Ryusuke Villemin, rvillemin@aurora.tech; Zeb Barber, zbarber@aurora.tech, Aurora Innovation, USA; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu; Adithya Pediredla, adithya.k.pediredla@dartmouth.edu, Dartmouth College, USA.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3731150>.

with real-world captured data for three different OHD applications—FMCW lidar, blood flow velocimetry, and wind Doppler lidar (Fig. 1-(C)).

CCS Concepts: • **Computing methodologies** → **Ray tracing; Computational photography.**

Additional Key Words and Phrases: optical interferometry, Doppler effect, time-of-flight rendering, physically-based rendering, computational imaging

ACM Reference Format:

Juhyeon Kim, Craig Benko, Magnus Wrenninge, Ryusuke Villemin, Zeb Barber, Wojciech Jarosz, and Adithya Pediredla. 2025. A Monte Carlo Rendering Framework for Simulating Optical Heterodyne Detection. *ACM Trans. Graph.* 44, 4 (August 2025), 19 pages. <https://doi.org/10.1145/3731150>

1 INTRODUCTION

Optical heterodyne detection (OHD) is a widely used interferometric technique for measuring frequency and phase modulations in light. OHD mixes a weak signal beam with a strong, coherent reference beam, often referred to as the *local oscillator* [Cummins and Swinney 1970; DeLange 1968]. The interference between these two optical fields generates a *beat* signal at a lower frequency, encoding how the received signal is modulated relative to the reference signal. This often includes useful information such as the velocity or distance of various objects in the scene. This beat frequency, being within the measurable range of electronic devices (few hundred MHz), allows for precise detection of otherwise imperceptible optical signal variations (few hundred THz). While both phase [Crouch and Rupavatharam 2019] and frequency modulations can be employed for OHD, we focus on frequency modulation due to its predominant use in numerous applications [Piggott 2020].

One of the most common causes of optical frequency modulation is the Doppler effect arising from moving objects or particles. OHD is sensitive to these frequency shifts that manifest in a wide range of applications, including velocity sensing in automobiles [Behroozpour et al. 2017; Peng and Shan 2021], atmospheric profiling and turbulence measurements [Chanin et al. 1989; Menzies and Hardesty 1989], and non-invasive biomedical Doppler techniques [Detre et al. 1992; Holloway Jr and Watkins 1977]. Frequency modulation can also be induced by specially designed transmitters, e.g. those used in *frequency-modulated continuous-wave* (FMCW) lidar. In this method, the temporally varying frequency makes *time-of-flight* values correspond to different frequencies, enabling not only velocity but also distance measurement through spectral analysis of the received signal. In general, OHD is pivotal for high-resolution spectral interpretation across various optical systems. While incoherent (intensity) imaging techniques like amplitude-modulated continuous-wave (AMCW) lidar can be used for similar purposes, the coherent detection offered by OHD provides substantial advantages such as high signal-to-noise ratio or immunity to ambient light.

Despite the widespread adoption and growing interest in OHD, a general-purpose simulation framework for OHD does not exist, which is hindering the progress in these imaging techniques. We build a generic and scalable Monte Carlo OHD simulator that is applicable to various scenarios including FMCW lidar measuring multi-bounced light from moving objects, blood flow velocimetry to compute the speed of blood inside vasculature made of diverse tissue layers, and wind Doppler lidar imaging complex atmospheric layers. Our key contribution is to write the OHD measurements in the spectral domain as a *path integral* (Eq. (26)), which allows us to leverage rich modeling and sampling capabilities of existing Monte Carlo path tracers.

This paper is structured into three main parts (Fig. 4): (1) (Sec. 4) We derive the OHD path integral that describes the average behavior (statistical mean) of power spectrum of OHD measurements. (2) (Sec. 5) We develop two approaches to sample a power spectrum from a single micro-surface realization, using speckle statistics. (3) (Sec. 6) We use a micro-surface interpretation following Dolan [2009] to show that OHD Doppler differs fundamentally from incoherent (intensity) Doppler, necessitating OHD simulation to be distinct from incoherent Doppler simulation [Kim et al. 2023]. Finally, We validate the proposed technique on three real-world applications: FMCW lidar, blood flow velocimetry, and wind Doppler lidar (Sec. 7–Sec. 9). We expect that our simulator will assist researchers in various fields who use OHD under different configurations. Our code and data can be found in the following project page¹.

2 RELATED WORKS

2.1 OHD Techniques

Heterodyning is a widely used technique in electrical signal processing that mixes two signals at different frequencies f_1, f_2 , creating a new signal at a lower frequency $f_1 - f_2$, called the *beat* frequency. However, due to the bandwidth limits of electronics, such heterodyning has been limited to radio frequencies (RF). Optical heterodyne

detection (OHD) [DeLange 1968], in principle, is the same as heterodyning in the RF domain, but the heterodyning is now achieved by an optical interferometer (Fig. 1-(A)). The beat frequency arises as an AC component of a photon-induced current on a *square-law* intensity detector. Here, we show three specific use cases of OHD.

2.1.1 FMCW lidar. FMCW techniques were first introduced in the radio-wave spectrum [Barrick 1973], but later applied to the lidar field with the advent of broadband tunable lasers such as distributed feedback laser [Passy et al. 1994] or vertical-cavity surface-emitting laser [Jewell et al. 1991]. Since optical FMCW provides much higher spatial resolution compared to RF, FMCW lidars are widely used for microscale 3D detection [Hariyama et al. 2018; Iiyama et al. 2011; Ula et al. 2019]. FMCW lidars also enable simultaneous distance and velocity estimation, making them suitable for a variety of applications [Karlsson et al. 2000; Piracha et al. 2011]. There are also other works that exploit FMCW lidars for 2D synthetic aperture imaging [Beck et al. 2005], Doppler vibrometry [Zhang et al. 2019], or autonomous driving [Li and Ibanez-Guzman 2020].

2.1.2 Non-invasive biomedical Doppler velocimetry. In the biomedical field, OHD is widely used to measure blood flow, commonly with the name of laser Doppler flowmetry, or velocimetry (LDV) [Riva et al. 1992; Shepherd and Öberg 2013]. LDV operates by directing a coherent laser beam onto a biological tissue, where the light scatters off of moving red blood cells, inducing a Doppler frequency shift. LDV is particularly valuable in assessing microvascular blood flow [Rendell et al. 1989] and detecting abnormalities in circulation [Kruger et al. 2006], offering real-time blood flow measurement [Bonner et al. 1981; Stern et al. 1977] which is critical for diagnosing and monitoring conditions like diabetes [Hu et al. 2017] or peripheral vascular diseases [Cochrane et al. 1986]. Here, we simulate the real-data of laser Doppler holography [Puyo et al. 2021, 2019], a variation of LDV that measures full-field retinal blood flow.

2.1.3 Atmospheric sensing. OHD, or *coherent Doppler wind lidar* (C-DWL) in this literature, plays a crucial role in atmospheric sensing, particularly for measuring wind velocities, turbulence, and aerosol dynamics [Chanin et al. 1989; Frehlich and Cornman 2002; Werner 2005]. These methods rely on the principle of detecting Doppler frequency shifts caused by the motion of air particles or aerosols scattering a coherent laser beam. C-DWL's high sensitivity makes it invaluable for weather forecasting [Yuan et al. 2020], aviation safety [Yuan et al. 2022], and cloud seeding [Yuan et al. 2021], offering a versatile tool for studying complex meteorological phenomena. C-DWL could be either continuous wave [Huffaker et al. 1970; Lawrence et al. 1972] or pulsed [Liu et al. 2019; Menzies and Hardesty 1989]. We simulate the hardware results from Wei et al. [2019], which uses a pulsed C-DWL for wind and rainfall detection.

2.2 OHD Simulation

Though various simulation works in different OHD applications exist, they are highly specialized and typically tailored to specific domains. Separate models have been developed for applications such as blood flow velocimetry [Bonner and Nossal 1981; de Mul et al. 1995; Fredriksson et al. 2007; Jentink et al. 1990; Nilsson et al. 1992; Stern 1985], atmospheric sensing [Banakh and Werner 2005; Szczap

¹https://juhyeonkim95.github.io/project-pages/ohd_rendering

et al. 2021], and coherent lidar for autonomous vehicles [Hofrichter et al. 2024; Rosenberger et al. 2020], yet these models rarely extend their applicability beyond the original contexts. Many of these methods rely on domain-specific assumptions, such as restricting to single-bounced rays or focusing exclusively on volume or surface scattering. Furthermore, due to a lack of modeling capability, they often employ oversimplified geometries (e.g., cylinder or plane to model blood vessels) or limited reflection models (e.g., perfectly diffuse reflection or Mie scattering). In contrast, our proposed algorithm which is based on the path integral formulation [Veach 1997], can take advantage of existing radiometric Monte Carlo path tracers, offering significant scalability across various scenarios.

2.3 Related Rendering Literature

2.3.1 Speckle Rendering. Since OHD uses coherent illumination, we observe *speckle* when light is scattered by a rough surface. Although most rendering models are designed for incoherent light sources and thus ignore speckle effects, several works [Bar et al. 2019, 2020; Steinberg and Yan 2022] render speckles arising from perfectly or partially coherent light. One of the most relevant works to this paper is Bar et al. [2019], who proposed a Monte Carlo rendering framework for fully-developed speckles scattered by participating media. Our work also shares the same philosophy—evaluation of speckle statistics through macroscopic path-space formulation instead of evaluating all of the microscopic configurations. However, we render OHD measurements and not an intensity image like Bar et al. [2019], which requires a different formulation. We also do not explicitly model spatial speckle statistics (covariance), but introduce a global attenuation term—*heterodyne efficiency*, to account for it.

2.3.2 Time-of-Flight and Doppler Rendering. Time-of-Flight (ToF) rendering aims to synthesize physically accurate measurements of ToF cameras, which basically belong to incoherent (intensity) detection. These works have mostly concentrated on transient simulation [Jarabo et al. 2014; Liu et al. 2022; Marco et al. 2017] or time-gated simulation [Liu et al. 2022; Pediredla et al. 2019]. Interestingly, our frequency-resolved OHD path integral exhibits structural similarity with their time-resolved path integral, with the difference that we also consider the Doppler shift. Simulation of dynamic scenes for ToF cameras has rarely been explored, with the exception of Kim et al. [2023], who faithfully reproduced the Doppler effect of heterodyne AMCW ToF cameras [Heide et al. 2015]. However, Kim et al. [2023] lacks on a detailed explanation of underlying physics. We perform a theoretical comparison of the Doppler effect in coherent (OHD) and incoherent (AMCW) detection, extending the discussion of micro-surface models by Dolan [2009], and highlight that we need a different strategy for each simulation.

3 BACKGROUND

In this section, we provide background knowledge necessary to understand the contributions of this paper. Readers familiar with OHD imaging may skip Sections 3.1 and 3.2, while those with a background in wave-optics rendering may skip Sections 3.3 and 3.4.

3.1 Basics of Optical Heterodyne Detection

Let's start with the *incoherent detection* (also called direct or intensity detection) that measures the intensity of the signal. Assume we have received *signal* with an optical field of $E_s(t) = A_s e^{j\phi_s(t)}$, with amplitude A_s and time-varying phase $\phi_s(t)$. Then, the conventional photodetector measures the photocurrent $i(t)$ which is proportional to the electromagnetic power P_{EM} carried by the signal

$$i(t) \propto P_{EM}(t) = |E_s(t)|^2 = A_s^2, \quad (1)$$

which only depends on the *square* of the signal amplitude.

In *coherent detection* or *heterodyne detection* (OHD), we use a beam splitter (or fiber-optic splitter) to split the transmitted light into the object and the reference mirror (Fig. 1-(A)). Therefore, not only the signal from the object, but also the reference local oscillator (LO) field $E_{LO}(t) = A_{LO} e^{j\phi_{LO}(t)}$ (typically $A_{LO} \gg A_s$), falls onto photodetector, generating an *optical interference*. For simplicity, we assume the phase of the LO field is equal to the phase of the transmitted light. The square-law photodetector then measures the $i(t)$ which is from the coherent sum of the signal and the LO fields:

$$i(t) = |E_{LO}(t) + E_s(t)|^2 = A_{LO}^2 + A_s^2 + 2\Re(E_{LO}(t)E_s^*(t)). \quad (2)$$

We ignore the DC terms (A_{LO}^2, A_s^2) and the high-frequency (few hundred THz) term, which is beyond the bandwidth of the electronics. The simplified model for the *heterodyne* measurement is:

$$i_{AC}(t) := A_s A_{LO} \cos(\phi_{LO}(t) - \phi_s(t)). \quad (3)$$

The frequency of this measurement is called the *beat frequency*, and is typically obtained from Fast Fourier Transform (FFT) as illustrated in Fig. 1-(A).

3.2 Velocity and Distance Detection from OHD

There are two types of light sources that are typically used for OHD.

3.2.1 Single-frequency Laser. For a single-frequency laser, the phase terms are given by

$$\phi_{LO}(t) = 2\pi f_0 t, \quad \phi_s(t) = 2\pi (f_0 + \Delta f_0) (t - \tau), \quad (4)$$

where $\Delta f_0 = f_0 2v/c$ is Doppler frequency shift and $\tau = 2d/c$ is time-of-flight, which depends on the object velocity v , distance d and speed of the light c . Then, we have

$$\phi_{LO}(t) - \phi_s(t) = -2\pi \Delta f_0 t + \psi(\Delta f_0, \tau), \quad (5)$$

where $\psi(\Delta f_0, \tau)$ is phase offset. Therefore, performing FFT gives Doppler frequency shift Δf_0 , which is proportional to v .

3.2.2 Swept-Frequency Laser. A swept-frequency laser, which is used for FMCW lidar, emits a linearly increasing frequency signal also known as *linear chirp* (Fig. 2-(A)). It could be modeled as

$$\phi_{LO}(t) = 2\pi f_0 t + \frac{\pi B t^2}{T}, \quad \frac{d}{dt} \phi_{LO}(t) = 2\pi \left(f_0 + \frac{B}{T} t \right), \quad (6)$$

where f_0 is the base frequency, B is the chirp bandwidth and T is the sweep period. We can calculate ϕ_s and the phase difference as

$$\phi_s(t) = 2\pi (f_0 + \Delta f_0) (t - \tau) + \frac{\pi B (t - \tau)^2}{T} \quad (7)$$

$$\phi_{LO}(t) - \phi_s(t) = 2\pi \left(\frac{B\tau}{T} - \Delta f_0 \right) t + \psi(\Delta f_0, \tau), \quad (8)$$

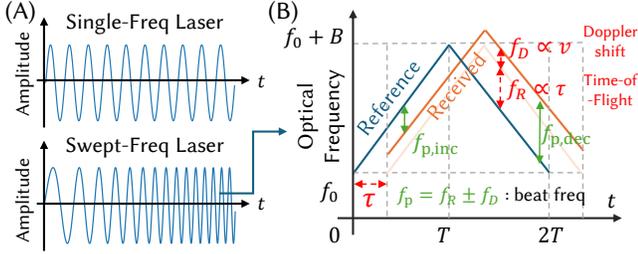


Fig. 2. (A) Two types of laser, single- and swept-frequency, typically used for OHD. We plot the amplitude over time for each laser. (B) Triangular chirp is commonly used for swept-frequency laser in FMCW lidar to decompose Doppler and ToF terms from the beat frequency.

where ψ is again the t independent offset term. The beat frequency f_p now comprises not only the Doppler term $f_D = \Delta f_0 = f_0 \frac{2v}{c}$, but also the distance term $f_R = \frac{2Bd}{cT}$. A commonly employed technique to decompose each term is using a triangular chirp with opposite bandwidths ($\pm B$) as depicted in Fig. 2-(B), resulting in:

$$f_{p,inc} = f_R - f_D, \quad \text{and} \quad f_{p,dec} = f_R + f_D. \quad (9)$$

With these two measurements, we can compute f_R and f_D , which are directly proportional to distance d and velocity v .

Problem of the Simple Model. The primary limitation of the model used to interpret OHD detection here (Eq. (3)) is that it considers only a *single reflector* (Fig. 3-(A)). In reality, the object surface is rough compared to the wavelength of light, which is on the micrometer scale in the visible-infrared range, as illustrated in Fig. 3-(B). As a result, the optical field from the rough surface must be modeled as a composition of multiple reflectors with random phases. This randomness introduces a *speckle* effect, causing the OHD measurement $i_{AC}(t)$ to fluctuate across different micro-surface realizations (e.g., different areas of the same material). It also causes $i_{AC}(t)$ to spatially vary over the photodetector necessitating space-dependent modeling. Moreover, the presence of multiple bounces further complicates the problem (Fig. 3-(C)). The multi-bounce effect is important when simulating the Doppler effect in dense aerosols or blood cells. In such cases, the beat frequency does not peak at a single value but instead exhibits a broad spectrum.

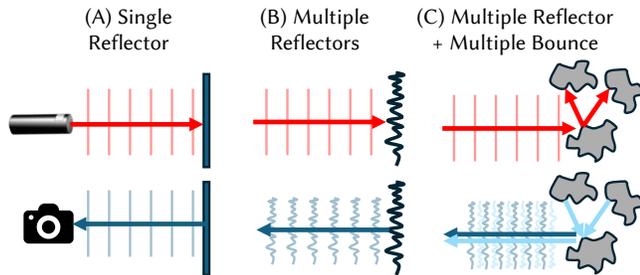


Fig. 3. (A) Single reflector model, (B) rough surface comprise of multiple reflectors and (C) also multiple bounce contribution. Transparent lines represent the wavefront (points with equal phase) of the propagating light.

While the most accurate way to handle these limitations would be to model all microscopic geometries and evaluate $i_{AC}(t)$ using complex light transport, it is not practically feasible due to the associated computational costs. Instead, to address this problem, (1) we treat the optical field as a stochastic process and evaluate its *average* property (or statistics) through (2) radiometric path tracing which incorporates multi-bounce contributions.

3.3 Optical Field as a Stochastic Process and Speckle

The rough surface itself and eventually the resulting optical field are often modeled as a *stochastic process*. A stochastic process is a collection of random variables, typically indexed by time, that describe the evolution of a system. While the exact optical field could be determined if the configurations of all microscopic random processes were known, in practice, the statistical properties of the field are of greater interest than the precise value of any single *realization*. To analyze these statistical properties, we utilize the *ensemble average*, denoted as $\langle \cdot \rangle$, which averages over all possible realizations, or the *ensemble*. Here, we restrict our discussion to the *wide-sense stationary* case, meaning their first- and second-order statistical moments are time-invariant.

Most radiometric reflection models (BRDFs) are based on ensemble-averaged measurements [Nayar 1991; Stam 1999; Ticconi et al. 2011]. Therefore, these models are valid only for incoherent light sources (the sun or an LED lamp), which naturally smooth out the stochastic speckle effect. Nevertheless, explicit ensemble-averaging on reflection from coherent light sources (e.g., repeated measurement on different regions of the same material) should still conform to the BRDF, which can thus provide useful information on the *average behavior* of the measurement.

3.4 Path Integral in Terms of Optical Field Decomposition

Although light inherently exhibits wave-like behavior, most simulations adopt a ray-based approach. Each ray propagates independently, with no interference between rays, and their intensities can be linearly added. A sequence of ray segments forms a *path* \vec{x} , and by adding the contribution of all possible paths in the scene, we arrive at the well-known path integral [Veach 1997]:

$$I = \int_{\mathcal{P}} f(\vec{x}) d\mu(\vec{x}), \quad (10)$$

where \mathcal{P} is *path space*, the set of all possible paths in macroscopic geometry, I is the camera measurement and $f(\vec{x})$ is the path throughput determined by the BRDF, visibility, and geometry terms as well as sensor importance function and emitted radiance.

This ray-based approach can also be interpreted through optical field theory. The key idea is the decomposition of the optical field into *mutually incoherent* components corresponding to different paths, so that the intensity of each component can be incoherently (i.e., linearly) summed. Formally, the decomposition of the optical field can be expressed as: $E_s(t) = \int_{\mathcal{P}} E_s(\vec{x}, t) d\mu(\vec{x})$, where $E_s(\vec{x}, t)$ represents mutually incoherent fields associated with each path, satisfying $\langle E_s(\vec{x}, t) E_s^*(\vec{y}, t + \tau) \rangle = \delta(\vec{x}, \vec{y})$ for any t and τ by definition. Note that while each ray is infinitesimally narrow from a macroscopic perspective, it intersects with numerous random reflectors from a microscopic perspective. Hence, the paths are stochastic in

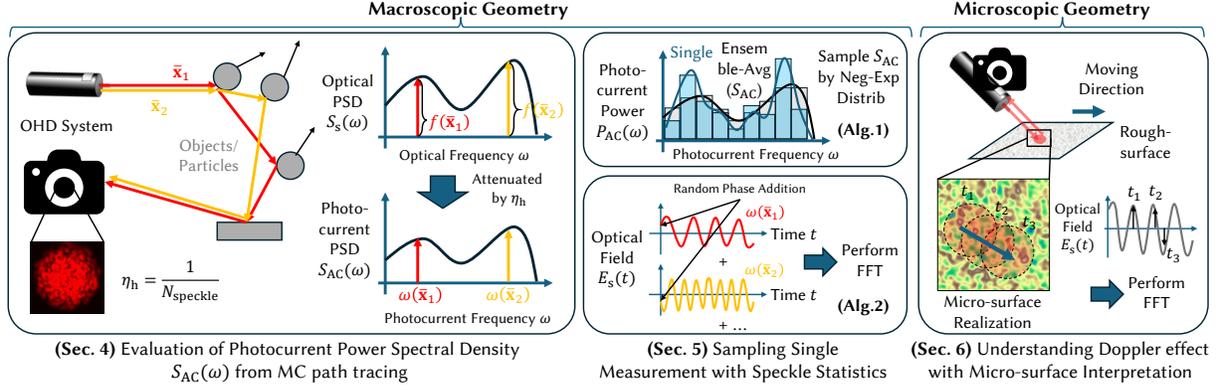


Fig. 4. Overview of our proposed method. In Sec. 4, we show PSD (ensemble-averaged power spectrum) of OHD photocurrent $S_{AC}(\omega)$ is equivalent to an attenuated replica of optical PSD of the signal $S_s(\omega)$. We propose a Monte Carlo path tracing algorithm to evaluate S_{AC} based on OHD path integral (Eq. (26)), which shares structural similarity with and transient rendering [Jarabo et al. 2014]. In Sec. 5, we derive a single measurement with two methods (1) sample $S_{AC}(\omega)$ with negative-exponential distribution or (2) sample optical field with random phase addition and performing FFT. Note that both Sec. 4 and Sec. 5 are based on macroscopic geometry. In Sec. 6, we perform the most thorough Doppler simulation with actual micro-surface realization. We also discuss the difference between the Doppler effect in incoherent (intensity) detection and its simulation strategy.

nature. Also, the micro-reflectors that each path encounters are assumed to be independent, which makes $E_s(\bar{x}, t)$ mutually incoherent.² Under this mutual-incoherence assumption, we have:

$$\langle |E_s(t)|^2 \rangle = \left\langle \left| \int_{\mathcal{P}} E_s(\bar{x}, t) d\mu(\bar{x}) \right|^2 \right\rangle = \int_{\mathcal{P}} \langle |E_s(\bar{x}, t)|^2 \rangle d\mu(\bar{x}) \quad (11)$$

since all cross-terms vanish due to mutual incoherence. Further removing the time-dependence with the stationary assumption gives:

$$\langle P_{EM}(t) \rangle = \langle |E_s(t)|^2 \rangle = \int_{\mathcal{P}} \langle |E_s(\bar{x})|^2 \rangle d\mu(\bar{x}). \quad (12)$$

Equation (12) aligns with the path integral in Eq. (10). As previously discussed in Sec. 3.3, I in Eq. (10) is indeed an ensemble-averaged measurement, where $f(\bar{x})$ is computed using reflection models based on ensemble averaging. Importantly, \mathcal{P} is defined based on macroscopic geometry, while all microscopic variations of the surface are incorporated into the ensemble averaging within the reflection model. The key takeaway here is the *mutual incoherence*, which will also be used to derive the path integral for OHD.

4 SPECTRAL DOMAIN PATH INTEGRAL OF OPTICAL HETERODYNE DETECTION

Here, we derive the spectral domain path integral that characterizes the power spectral density (PSD) of the OHD photocurrent which allows us to use existing ray-based techniques to simulate OHD. The formulation derived here is an ensemble-averaged value and the sampling of a single measurement will be addressed in Sec. 5.

4.1 OHD Modeling - Revisited

We now assume that E_s is a stochastic process. Also we extend Eq. (2) to account for non-zero surface area— $i(t)$ is now integration

²To be precise, a ray should be considered as a beam with an extent of lateral correlation length of the surface to be mutually independent, but we assume it is infinitesimal in macroscopic geometry, following the assumption of conventional path tracing.

of infinitesimal photocurrent at position \vec{r} on the detector \mathcal{A}_{det} :

$$i(t) = \int_{\mathcal{A}_{det}} |E_s(\vec{r}, t) + E_{LO}(\vec{r}, t)|^2 d\vec{r}. \quad (13)$$

Expanding the integrand yields three terms. $|E_{LO}(\vec{r}, t)|^2 = A_{LO}^2$ is still the DC component and hence, ignored. Although $|E_s|^2$ is no longer guaranteed to be a DC term, the LO field is typically much stronger than the signal [Cummins and Swinney 1970; Goodman 2007] and hence, ignored as well. We have the remaining AC term:

$$i_{AC}(t) := \int_{\mathcal{A}_{det}} \Re \{ E_s(\vec{r}, t) E_{LO}^*(\vec{r}, t) \} d\vec{r}. \quad (14)$$

Here, we assume E_s is stationary (further explained in Sec. 5), which makes i_{AC} stationary.

4.2 Spectral Evaluation of Optical Heterodyne Detection

We aim to simulate the spectral distribution of the photocurrent as the downstream algorithms for computing physical parameters (distance, velocity) use spectral information rather than raw i_{AC} . Mathematically, the goal is to evaluate the *power spectral density* (PSD) of $i_{AC}(t)$, denoted as $S_{AC}(\omega)$, which is defined as:

$$S_{AC}(\omega) := \lim_{T \rightarrow \infty} \frac{1}{T} \langle |\mathcal{F}(i_{AC}(t)w_T(t))|^2 \rangle. \quad (15)$$

Here, $w_T(t)$ is a window function of width T , introduced to ensure convergence of the Fourier transform. For stationary processes, the PSD and the *autocorrelation* are related as a Fourier transform pair according to *Wiener-Khinchin theorem*, so we have $S_{AC}(\omega) = \mathcal{F}(C_{AC}(\tau))$, where $C_{AC}(\tau) = \langle i_{AC}(t) i_{AC}^*(t + \tau) \rangle$ is autocorrelation. We begin by evaluating $i_{AC}(t) i_{AC}^*(t + \tau)$ which is given by

$$\int_{\mathcal{A}_{det}^2} E_s(\vec{r}_1, t) E_{LO}^*(\vec{r}_1, t) E_s^*(\vec{r}_2, t + \tau) E_{LO}(\vec{r}_2, t + \tau) d\vec{r}_1 d\vec{r}_2 + c.c. \quad (16)$$

where c.c. represents the complex conjugate and can be computed from the first term. So, we will omit it for notational simplicity.

Using the *mutual coherence function* (MCF), defined as $\Gamma(\vec{r}_1, \vec{r}_2, \tau) := \langle E(\vec{r}_1, t)E^*(\vec{r}_2, t + \tau) \rangle$ and taking the ensemble average leads to

$$C_{AC}(\tau) = \int_{\mathcal{A}_{\text{det}}^2} \Gamma_s(\vec{r}_1, \vec{r}_2, \tau) \Gamma_{LO}^*(\vec{r}_1, \vec{r}_2, \tau) d\vec{r}_1 d\vec{r}_2. \quad (17)$$

Assuming the optical PSD is approximately constant across the detector (this *does not* mean phase is constant), we have

$$C_{AC}(\tau) \approx \eta_h \left(\int_{\mathcal{A}_{\text{det}}} \Gamma_s(\vec{r}, \vec{r}, \tau) d\vec{r} \right) \left(\int_{\mathcal{A}_{\text{det}}} \Gamma_{LO}^*(\vec{r}, \vec{r}, \tau) d\vec{r} \right). \quad (18)$$

Here, η_h is *heterodyne efficiency*, a constant that shows how well the signal and LO match (or *coherent*), which is known to be inverse-proportional to the number of speckles on the photodetector [Cummins and Swinney 1970; Goodman 2007]. As OHD applications use a spectral distribution and the absolute scale is not that important, we do not explicitly compute η_h . Please refer to the supplementary material for details of η_h and also a proof of Eq. (18).

Taking the Fourier transform and using definition of optical PSD $S(\vec{r}, \omega) = \mathcal{F}(\langle E(\vec{r}, t)E^*(\vec{r}, t + \tau) \rangle) = \mathcal{F}(\Gamma(\vec{r}, \vec{r}, \tau))$ yields the important conclusion:

$$S_{AC}(\omega) = \eta_h \mathcal{A}_{\text{det}}^2 S_s(\omega) * S_{LO}^*(-\omega), \quad (19)$$

where $S_{LO}(\omega)$ and $S_s(\omega)$ represent the *optical PSD* of the LO and the signal, which are assumed to be constant over \mathcal{A}_{det} . Note that if the LO field has a delta-peaked spectrum at ω_0 , S_{AC} just becomes a frequency-shifted version of the signal's optical PSD.

4.3 Spectral Domain Path Integral Derivation

Now we derive the path integral formulation for Eq. (19). Before entering into the calculation, we first explicitly define a time-evolving path \bar{x} in *macroscopic* geometry, composed of the vertex $\mathbf{x}_k, k \in [0, K]$ as illustrated in Fig. 5-(A). \mathbf{x}_0 is a vertex at the light source while \mathbf{x}_K is a vertex at the detector. Each path vertex is a function of time $\mathbf{x}(t)$, but we omit time dependence. We denote with $\mathbf{v}_k = \frac{d}{dt} \mathbf{x}_k(t)$ the instantaneous velocity at \mathbf{x}_k . As mentioned in Sec. 3.4, each scattering point includes numerous independent random micro-reflectors. We further assume that these micro-reflectors are *attached* to the vertex so that the composition of micro-reflectors at each vertex does not change over time. The rationale for this assumption will be further explained in Sec. 6.

We now evaluate each term in Eq. (18). For LO term, we have $\int_{\mathcal{A}_{\text{det}}} \Gamma_{LO}(\vec{r}, \vec{r}, \tau) d\vec{r} = P_{LO} e^{j(\phi_{LO}(t) - \phi_{LO}(t + \tau))}$ where $P_{LO} = \mathcal{A}_{\text{det}} \mathcal{A}_{LO}^2$. For the signal term, we decompose the optical field into mutually incoherent components from different paths, similar to Eq. (11),

$$\int_{\mathcal{A}_{\text{det}}} \Gamma_s(\vec{r}, \vec{r}, \tau) d\vec{r} = \int_{\mathcal{A}_{\text{det}}} \langle E_s(\vec{r}, t) E_s^*(\vec{r}, t + \tau) \rangle d\vec{r} \quad (20)$$

$$= \int_{\mathcal{A}_{\text{det}}} \int_{\mathcal{P}(\vec{r})} \langle E_{s, \bar{x}}(\vec{r}, t) E_{s, \bar{x}}^*(\vec{r}, t + \tau) \rangle d\mu(\bar{x}) d\vec{r} \quad (21)$$

$$= \int_{\mathcal{P}} \langle E_s(\bar{x}, t) E_s^*(\bar{x}, t + \tau) \rangle d\mu(\bar{x}) \quad (22)$$

$$= \int_{\mathcal{P}} |E_s(\bar{x})|^2 e^{j(\phi_{\bar{x}}(t) - \phi_{\bar{x}}(t + \tau))} d\mu(\bar{x}), \quad (23)$$

where $\mathcal{P}(\vec{r})$ is the path space that includes \vec{r} as end point and we merged integration domain from Eq. (21) to Eq. (22). $\phi_{\bar{x}}(t)$ is the

phase of the path \bar{x} , which is determined by macroscopic geometry and hence can be factored out of the ensemble average. The remaining phasor term $E_s(\bar{x})$ which is related to the microscopic random perturbation at each scattering point, is time-invariant as we assumed $\mathbf{x}(t)$ and $\mathbf{x}(t + \tau)$ are composed of the same pool of micro-reflectors. One can also understand this as assuming a perfect temporal coherence of $E_s(\bar{x}, t)$.

We can now calculate the final phase difference term. For a single-frequency laser, we have:

$$\begin{aligned} & \{\phi_{\bar{x}}(t) - \phi_{\bar{x}}(t + \tau)\} - \{\phi_{LO}(t) - \phi_{LO}(t + \tau)\} \\ &= 2\pi (f_0 + \Delta f_0(\bar{x})) \{(t - \|\bar{x}\|) - (t - \|\bar{x}\| + \tau)\} - 2\pi f_0 \{t - (t + \tau)\} \\ &= -2\pi \Delta f_0(\bar{x}) \tau. \end{aligned} \quad (24)$$

where $\|\bar{x}\|$ is ToF and $\Delta f_0(\bar{x})$ is Doppler frequency shift from the path \bar{x} . For a swept-frequency laser, E_s is not stationary so we need generalized MCF and repeat the derivation, but eventually, we get $2\pi \left(\frac{\beta}{c} \|\bar{x}\| - \Delta f_0(\bar{x}) \right) \tau$, which makes $i_{AC}(t)$ stationary. Interestingly, the if and only if the condition that makes $i_{AC}(t)$ stationary is $\phi_{LO}(t)$ to be linear (single-frequency laser) or quadratic (linear-chirp laser). We include the proof in the supplementary material. Note that, in all these cases, the phase difference term can be written as $\omega(\bar{x})\tau$, where $\omega(\bar{x})$ is the frequency shift due to both path length and path velocity. Substituting $\omega(\bar{x})\tau$ in Eq. (18), we obtain:

$$C_{AC}(\tau) = \eta_h P_{LO} \int_{\mathcal{P}} |E_s(\bar{x})|^2 e^{j\omega(\bar{x})\tau} d\mu(\bar{x}). \quad (25)$$

Taking the Fourier transform and replacing the ensemble-averaged field intensity with the radiometric path throughput, we arrive at:

$$S_{AC}(\omega) = \eta_h P_{LO} \int_{\mathcal{P}} f(\bar{x}) \delta(\omega - \omega(\bar{x})) d\mu(\bar{x}). \quad (26)$$

which we refer to as the *OHD path integral*. This result demonstrates that conventional radiometric path tracing can be employed to simulate the average spectral distribution of OHD photocurrent!

Similarity with Transient Rendering. OHD path integral exhibits structural equivalence to the transient path integral [Jarabo et al. 2014; Pediredla et al. 2019] which is given by:

$$I(t) = \int_{\mathcal{P}} f(\bar{x}) \delta(t - \|\bar{x}\|) d\mu(\bar{x}). \quad (27)$$

In transient rendering, each path contributes to a corresponding time bin based on its path length. In OHD, each path contributes to a corresponding frequency bin based on the frequency modulation of the path, which depends on the path length and velocity.

Finite Measurement Time. In practice, we can only measure a finite length of i_{AC} , resulting in a convolved spectrum with the window function:

$$S_{AC, T}(\omega) = S_{AC}(\omega) * |W_T(\omega)|^2, \quad (28)$$

where $W_T(\omega)$ is the Fourier transform of $w_T(t)$.

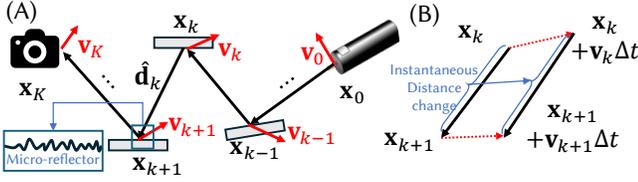


Fig. 5. (A) Time-evolving path consisting of $K + 1$ vertices and their instantaneous velocities. Each vertex contains several micro-reflectors which we assume are not changing significantly with time. (B) By calculating the distance change after infinitesimal time Δt , we derive path velocity in Eq. (32).

Sign Sensitiveness. Reviving the c.c term that we omitted in Eq. (16) modifies Eq. (26) into

$$\eta_{\text{h}} P_{\text{LO}} \int_{\mathcal{P}} f(\bar{\mathbf{x}}) (\delta(\omega - \omega(\bar{\mathbf{x}})) + \delta(\omega + \omega(\bar{\mathbf{x}}))) d\mu(\bar{\mathbf{x}}), \quad (29)$$

which indicates it is not possible to distinguish the sign of the frequency shift. However, such ambiguity can be easily resolved using a quadrature demodulator (90-degree phase shift of LO field), so we will keep using the sign-sensitive formulation of Eq. (26).

4.4 Frequency Contribution from Each Path

In this part, we derive the exact representation of $\omega(\bar{\mathbf{x}})$ from the path. Consider a time-evolving path in Fig. 5-(A) as explained before. We also define the unit ray direction vector $\hat{\mathbf{d}}_k = \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\|\mathbf{x}_{k+1} - \mathbf{x}_k\|}$ and η_k as the refractive index for the path segment $\mathbf{x}_k \mathbf{x}_{k+1}$. We use f_k to represent the observed frequency at \mathbf{x}_k . Then, for the path segment $\mathbf{x}_k \mathbf{x}_{k+1}$, we can calculate the Doppler frequency shift as follows from well-known Doppler formula [Drain 1980]:

$$f_{k+1} = f_k \left(\frac{c/\eta_k - \mathbf{v}_{k+1} \cdot \hat{\mathbf{d}}_k}{c/\eta_k - \mathbf{v}_k \cdot \hat{\mathbf{d}}_k} \right) \approx f_k \left(1 + \frac{\eta_k (\mathbf{v}_k - \mathbf{v}_{k+1}) \cdot \hat{\mathbf{d}}_k}{c} \right), \quad (30)$$

as $\|\mathbf{v}_k\| \ll c$. We can then calculate the final observed frequency as

$$f_K = f_0 \prod_{k=0}^{K-1} \frac{f_{k+1}}{f_k} \approx f_0 \left(1 + \sum_{k=0}^{K-1} \frac{\eta_k (\mathbf{v}_k - \mathbf{v}_{k+1}) \cdot \hat{\mathbf{d}}_k}{c} \right). \quad (31)$$

We define the (*optical*) path velocity $u(\bar{\mathbf{x}})$ as

$$u(\bar{\mathbf{x}}) = \sum_{k=0}^{K-1} \eta_k (\mathbf{v}_k - \mathbf{v}_{k+1}) \cdot \hat{\mathbf{d}}_k, \quad (32)$$

which gives path Doppler shift $\Delta f(\bar{\mathbf{x}}) := f_K - f_0 = f_0 u(\bar{\mathbf{x}})/c$. For correspondence in notation, instead of ToF $\|\bar{\mathbf{x}}\|$, let's also define (*optical*) path length as $l(\bar{\mathbf{x}}) = \sum_{k=0}^{K-1} \eta_k \|\mathbf{x}_{k+1} - \mathbf{x}_k\|$. The path length and velocity are two components that determine frequency in general OHD simulation. Interestingly, they are in a derivative relationship $u(\bar{\mathbf{x}}) = -\frac{d}{dt} l(\bar{\mathbf{x}})$ which is intuitively understandable as the Doppler effect arises due to phase change, which is equivalent to path length change as illustrated in Fig. 5-(B). Using optical path velocity and path length, we can rewrite $\omega(\bar{\mathbf{x}})$ in Eq. (26) as

$$\omega(\bar{\mathbf{x}}) = 2\pi \left(\frac{B}{T} \frac{l(\bar{\mathbf{x}})}{c} - f_0 \frac{u(\bar{\mathbf{x}})}{c} \right). \quad (33)$$

5 SINGLE MEASUREMENT SAMPLING FROM SPECKLE STATISTICS

While the previous section provides an ensemble-averaged power spectrum S_{AC} , it can be different from a single measurement power spectrum P_{AC} . In this section, we first review the speckle statistics from Goodman [2007], and propose two approaches to sample a single measurement for OHD: (1) sampling from the PSD in Eq. (26) using a negative-exponential distribution, and (2) sampling $i_{\text{AC}}(t)$ by incorporating random phase noise, and performing FFT.

5.1 Basics of Speckle Statistics and Extension to Path

We can represent E_s as the sum of $N(t)$ independent microscopic elementary phasor components:

$$E_s(t) = \frac{1}{\sqrt{N(t)}} \sum_{n=1}^{N(t)} A_n(t) e^{j\phi_n(t)}, \quad (34)$$

where $A_n(t)$ is amplitude and $\phi_n(t)$ is phase of n -th element. The normalization factor $1/\sqrt{N(t)}$ is included to ensure energy preservation. The above equation also holds given a macroscopic path $\bar{\mathbf{x}}$. Assuming that all the micro-reflectors the path encounters (Fig. 5-(A)) are not changing significantly over time, we can factor out macroscopic phase $\phi_{\bar{\mathbf{x}}}(t)$ and compute E_s for the path as:

$$E_s(\bar{\mathbf{x}}, t) = e^{j\phi_{\bar{\mathbf{x}}}(t)} \frac{1}{\sqrt{N}} \sum_{n=1}^N A_n e^{j\Delta\phi_n} = e^{j\phi_{\bar{\mathbf{x}}}(t)} E_s(\bar{\mathbf{x}}), \quad (35)$$

where $\Delta\phi_n$ is time-independent phase perturbation with regard to macroscopic properties and $E_s(\bar{\mathbf{x}})$ is *random phasor sum* [Goodman 2007]. The phase $\Delta\phi_n$ is often assumed to be uniformly distributed in the interval $(-\pi, \pi]$. This assumption holds when the surface is sufficiently rough compared to the wavelength of the light. Given the wavelengths of interest ($\sim 1 \mu\text{m}$), it is satisfied for most man-made and natural surfaces [Goodman 2007; Steinberg and Yan 2022]. From the central limit theorem, the distribution of $E_s(\bar{\mathbf{x}})$ converges to the *circular complex Gaussian* $\mathcal{CN}(0, \sigma^2)$ with uniformly distributed phase and *Rayleigh*-distributed amplitudes. Importantly, the intensity $I = |E_s(\bar{\mathbf{x}})|^2$ follows a *negative-exponential* distribution:

$$\text{pdf}(I) = \frac{1}{\langle I \rangle} \exp\left(-\frac{I}{\langle I \rangle}\right). \quad (36)$$

where $\langle I \rangle = 2\sigma^2$ is the mean intensity. Speckles whose intensities follow this PDF are called *fully developed* speckles, and we only consider such speckles here. Note that $\langle E_s(\bar{\mathbf{x}}, t) \rangle = 0$ and as we assume the micro-reflectors do not change over time, its autocorrelation only depends on the time difference, which makes it stationary.

We also discuss the statistics of the sum of M independent fully developed speckles (e.g. different speckle grains belonging to the same pixel), which can arise as a phasor sum or an intensity sum. In phasor sum (coherent sum), the addition of M different speckle patterns is equivalent to a single speckle pattern with a total number of component phasors $N_s = N_1 + N_2 + \dots + N_M$. Therefore, the statistical properties of the speckle do *not* change. On the other hand, for intensity sum (incoherent sum), as M increases, the probability density function transitions from a negative exponential distribution ($M = 1$) to a sharp Gaussian as $M \rightarrow \infty$.

5.2 OHD Speckle from PSD Sampling

For incoherent (intensity) detection, the intensity of the field is integrated over each point on the photodetector region. Therefore, as the photodetector measurement includes more independent components (e.g., \mathcal{A}_{det} covers more speckle grains), we get less noise in the measurement. For coherent heterodyne detection, we have:

$$i_{\text{AC}}(t) = \int_{\mathcal{A}_{\text{det}}} \Re(E_s(\vec{r}, t) E_{\text{LO}}^*(\vec{r}, t)) d\vec{r}. \quad (37)$$

Only considering a specific frequency component ω , we have:

$$i_{\text{AC}}(\omega, t) = A_{\text{LO}} |B| \cos(\omega t + \theta_B), \quad (38)$$

where $B = \int_{\mathcal{A}_{\text{det}}} E_s(\vec{r}) d\vec{r} = |B| e^{j\theta_B}$ or using path-domain formulation $B = \int_{\mathcal{P}} E_s(\vec{x}) d\mu(\vec{x})$. Thus, i_{AC} is a phasor sum, and speckle statistics *do not* change due to integration over \mathcal{A}_{det} or \mathcal{P} —its intensity follows a negative exponential distribution *regardless* of the number of speckles on \mathcal{A}_{det} [Goodman 2007], or more generally, independent path components. Therefore, all we need to do to get P_{AC} is a sampling from $S_{\text{AC}} = \langle P_{\text{AC}} \rangle$ according to the negative exponential distribution at each frequency bin (Algorithm 1). We assume that speckles contributing to each frequency bin are uncorrelated so that the power distribution at each frequency bin is independent.

Real Data Validation. We present real data from an FMCW lidar measurement to support our argument. Fig. 6-(A) shows the FFT result of an FMCW lidar measurement on a specular material, revealing two distinct peaks within 80 FFT bins. We show the average power spectrum for 100 measurements as the red line and two single measurements as transparent lines. Fig. 6-(C) shows the intensity across different measurements for a specific frequency, while Fig. 6-(D) shows the histogram of Fig. 6-(C) data. The orange line in Fig. 6-(D) is the regression fit which qualitatively and quantitatively (chi-square and Kolmogorov–Smirnov tests) aligns well with the histogram data demonstrating that a single measurement frequency follows a negative exponential distribution. Apart from our empirical observations, Baumann et al. [2014]; Dabas et al. [1994]; Flamant et al. [1984]; Hardesty et al. [1981] also reported that OHD measurements follow a negative exponential distribution, though they did not observe that each spectral channel is also statistically independent. Fig. 6-(B) illustrates the covariance map of the power spectrum over frequencies. We calculate the covariance based on the normalized power spectrum (i.e. normalize the power spectrum with a shape of [100, 80] over the first axis). Overall, it showed a diagonal shape especially showing no correlation between the two peaks.³ However, there are “+” shaped artifacts at each peak, which are due to *spectral leakage* inherent to FFT—because of discretization, frequency components not only contribute to a single bin but also to the nearby bins [Lyons 2011]. If a bin has a signal that is strong enough to dominate the spectral leakage from the nearby bin, it remains uncorrelated (dark part of +), otherwise we can see some correlation (bright part of +). Sampling PSD is not capable of reproducing such spectral leakage as each frequency bin is treated independently, but one can use direct evaluation in the next section.

³In fact, the bins other than the peaks are from shot noise of LO field, which have same statistics with the signal. Please see the supplementary for details.

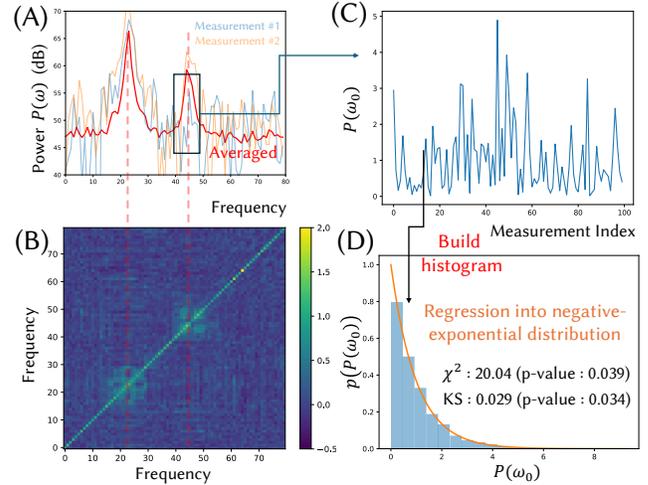


Fig. 6. (A) FFT result for 100 different measurements using FMCW lidar. We plot the averaged spectrum with red lines while transparent lines represent two different single measurements. (B) Covariance map that shows independence between different frequencies. We can observe the “+” shape near the peak due to spectral leakage. (C) The power measurements at the second peak and (D) the histogram of the measurements and curve-fitted negative-exponential plot. We also denote chi-square and KS statistics.

5.3 OHD Speckle from Field Sampling

One way to directly evaluate E_s , i_{AC} , and eventually P_{AC} is by considering all of the microscopic geometries, which would be most accurate, but computationally expensive. Instead, here we mimic the speckle effect keeping the path-tracing formulation with macroscopic geometry. Since $E_s(t) = \int_{\mathcal{P}} E_s(\vec{x}, t) d\mu(\vec{x})$ and $E_s(\vec{x})$ is a circular complex Gaussian, we can sample $E_s(t)$ with N paths:

$$E_s(t) = \frac{1}{\sqrt{N}} \sum_{n=1}^N Z_n e^{j\phi_{\vec{x}_n}(t)}, \quad (39)$$

where Z_n is a random variable that follows $Z_n \sim \mathcal{CN}\left(0, \frac{f(\vec{x}_n)}{2p(\vec{x}_n)}\right)$. The path throughput is $f(\vec{x}_n)$ and $p(\vec{x}_n)$ is the sampling pdf. Also, Z_n s are sampled independently to make all paths mutually incoherent. Compared to expensive evaluation from explicit micro-reflectors in Eq. (34), we now reproduce stochastic optical field $E_s(t)$ using path space in Eq. (39). Note that evaluating the ensemble average of the square of Eq. (39) matches with Monte Carlo evaluation on intensity I which is given by

$$I = \langle |E_s(t)|^2 \rangle = \frac{1}{N} \sum_{n=1}^N \langle |Z_n|^2 \rangle = \frac{1}{N} \sum_{n=1}^N \frac{f(\vec{x}_n)}{p(\vec{x}_n)}. \quad (40)$$

Interestingly, instead of sampling the amplitude term (Z_n), we can use the fact that when we add circular Gaussian variables, from the central limit theorem, their statistics converge to the sum of constant amplitude with uniform random phase random variables. Thus, we can skip sampling Z_n and simplify the sampling as

$$E_s(t) = \frac{1}{\sqrt{N}} \sum_{n=1}^N \sqrt{\frac{f(\vec{x}_n)}{p(\vec{x}_n)}} e^{j(\phi_{\vec{x}_n}(t) + \psi_n^{\text{rand}})}, \quad (41)$$

where $\psi_n^{\text{rand}} \in (-\pi, \pi]$ is uniformly random phase noise term. We can also sample i_{AC} accordingly:

$$i_{AC}(t) = \frac{1}{\sqrt{N}} \sum_{n=1}^N \sqrt{\frac{f(\bar{x}_n)}{p(\bar{x}_n)}} e^{j(\omega(\bar{x}_n)t + \psi_n^{\text{rand}})}, \quad (42)$$

and P_{AC} could be evaluated from FFT of i_{AC} (Algorithm 2). This approach is in fact equivalent to the Fourier tracing with random phase addition proposed by Hofrichter et al. [2024]. However, we provide a more principled explanation grounded in speckle statistics.

PSD Sampling (Algorithm 1) vs Field Sampling (Algorithm 2). The key shortcoming of field sampling is that it is difficult to know the average distribution - we need a lot of samples to get S_{AC} , which PSD sampling already offers. Fig. 7-(A) shows simulation using field sampling under GT PSD (gray line) of two peaks with different widths, and baseline noise. Blue and orange transparent lines represent two different single measurements and the red line represents an averaged power spectrum over 1000 realizations. The green dotted line shows ground truth PSD with convolution as Eq. (28). We could observe that repeated direct evaluation converges to GT PSD, but it took a lot of repeats to converge. We also plot distribution at the peak, which follows negative-exponential distribution as expected. In terms of FFT accuracy, PSD sampling cannot reproduce spectral leakage and the resulting correlation. Field sampling, on the other hand, captures this effect by directly computing i_{AC} and performing the FFT, which is illustrated in Fig. 7-(B), similar to that of Fig. 6-(B). If the spectrum is broad, it is hard to detect spectral leakage because each independent peak cancels out the correlation of others. Users may select between each algorithm based on needs.

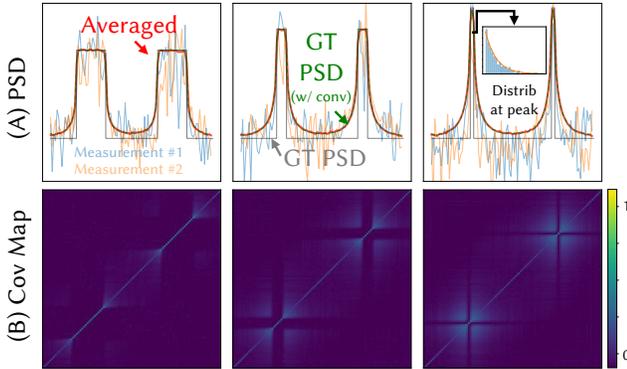


Fig. 7. (A) Simulation result for two-peak scenario as Fig. 6, but with different peak widths. The blue and orange lines are two different direct evaluation results and the red line is averaged over 1000 measurements, which converges to GT PSD with convolution (dotted green line). We also plot distribution at the peak which follows a negative exponential. (B) As width becomes smaller, we can clearly observe + shaped correlation due to spectral leakage as Fig. 6.

6 COMPARISON OF DOPPLER EFFECT OF COHERENT VS INCOHERENT DETECTION

In this section, we compare the Doppler effect in coherent detection (OHD) and incoherent detection methods such as AMCW lidar.

Algorithm 1 OHD Speckle with PSD Sampling

```

result_avg = 0
for iteration = 1 : N do
    Sample path  $\bar{x}$  and evaluate  $f(\bar{x}), p(\bar{x})$ 
     $\omega(\bar{x}) = \text{CalculateFreq}(l(\bar{x}), u(\bar{x}))$ 
    bin_index = FreqToBinIdx( $\omega(\bar{x})$ )
    result_avg[bin_index] +=  $\frac{1}{N} \frac{f(\bar{x})}{p(\bar{x})}$ 
result_single = SampleNegativeExponential(result_avg)
return result_avg, result_single

```

Algorithm 2 OHD Speckle with Field Sampling

```

 $i_{AC} = 0$ 
for iteration = 1 : N do
    Sample path  $\bar{x}$  and evaluate  $f(\bar{x}), p(\bar{x})$ 
     $\omega(\bar{x}) = \text{CalculateFreq}(l(\bar{x}), u(\bar{x}))$ 
    Sample  $\psi^{\text{rand}} \sim \text{Uniform}[0, 2\pi]$ 
    for time_index = 1 :  $N_T$  do
         $t = \text{GetTime}(\text{time\_index})$ 
         $i_{AC}[\text{time\_index}] += \frac{1}{\sqrt{N}} \sqrt{\frac{f(\bar{x})}{p(\bar{x})}} \exp j(\omega(\bar{x})t + \psi^{\text{rand}})$ 
result_single = |FFT( $i_{AC}$ )|2
return result_single

```

While the fundamental nature of the Doppler effect remains the same for both types of measurements—where phase changes across time result in frequency shifts—the significant difference in operating wavelengths leads to distinct Doppler frequency shifts. This disparity necessitates different simulation strategies to accurately model each phenomenon. Our discussion complements the AMCW Doppler simulation detailed in Kim et al. [2023], offering a more comprehensive theoretical understanding of Doppler simulation.

6.1 Nature of the Doppler Effect

The Doppler effect fundamentally arises from phase changes due to geometric displacement. Consider a single reflector with a phase of Φ , measured relatively from the observer. If the object is located at a distance d , the phase is given by $\Phi = \frac{2d}{\lambda}$ where λ is the wavelength of the signal, which could be both optical field (OHD, μm scale) or from intensity modulation (AMCW, m scale). If the object moves away from the camera with velocity v , Φ increases at a rate of $\frac{2v}{\lambda}$, leading to linear growth in the observed frequency. This rate is referred to as the Doppler frequency shift. It is important to note that the measurement is not the direct Φ , but rather $\exp j\Phi$. This sinusoidal waveform can be implemented either via optical fields (OHD) or intensity modulation (AMCW). In both cases, heterodyning is applied through an optical interferometer or electronic mixing. However, we are more interested in typical rough surfaces that are composed of numerous random micro-reflectors. In the following subsections, we analyze the Doppler effect originating from a rough surface. Our discussion is similar to that of Dolan [2009], but we consider a more general case of continuous height map instead of stair-step geometry assumed in their work.

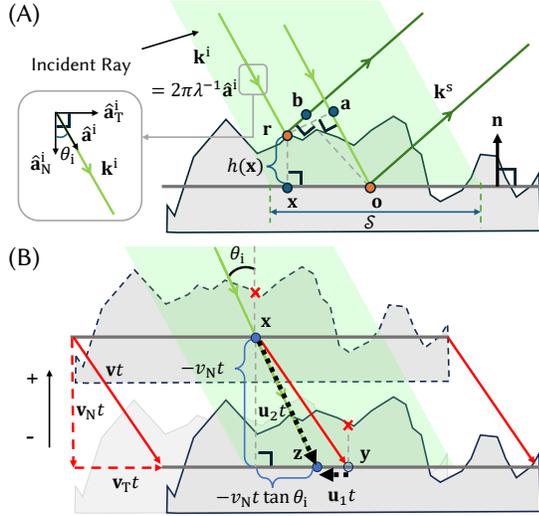


Fig. 8. Light scattering on a rough surface. (A) Assume incident and outgoing wavefront with wavevector of \mathbf{k}^i and \mathbf{k}^s . Then, the wave scattered at $\mathbf{r} = \mathbf{x} + h(\mathbf{x})\mathbf{n}$ has phase variation compared to the origin point \mathbf{o} . This could be geometrically calculated from the distance difference between two different light paths. (B) The surface moves with velocity \mathbf{v} composed of normal component $\mathbf{v}_N = v_N \mathbf{n}$ and transverse component \mathbf{v}_T . At the time t , the ray hit point changes from \mathbf{x} to \mathbf{z} which has a different local coordinate with \mathbf{y} .

6.2 Phase Variation on Rough Surface

First, consider a wavefront incident on a static rough surface (Fig. 8-(A)). The incident wavevector is denoted as $\mathbf{k}^i = 2\pi\lambda^{-1}\hat{\mathbf{a}}^i$ where $\hat{\mathbf{a}}^i$ is the unit vector along the incident direction. Similarly, the scattered wavevector along a specific direction is denoted as \mathbf{k}^s . Each point $\mathbf{x} \in \mathcal{S}$ on the certain region of macro-surface \mathcal{S} generates a phase variation due to differences in the geometric path length traveled by the wavefront. For a specific rough surface realization $h(\mathbf{x}) : \mathcal{S} \rightarrow \mathbb{R}$, the relative optical path length difference $\Delta\Phi(\mathbf{x}, \mathbf{k}^i, \mathbf{k}^s)$ between the scattering point $\mathbf{r} = \mathbf{x} + h(\mathbf{x})\mathbf{n}$ and the origin point \mathbf{o} is given by:

$$\Delta\Phi(\mathbf{x}, \mathbf{k}^i, \mathbf{k}^s) = \overline{\mathbf{o}\mathbf{a}} - \overline{\mathbf{r}\mathbf{b}} = h(\mathbf{x})(\mathbf{k}^m \cdot \mathbf{n}) + (\mathbf{k}^m \cdot \mathbf{x}). \quad (43)$$

For a simpler notation, we use wavevector difference $\mathbf{k}^m = \mathbf{k}^s - \mathbf{k}^i$ and omit \mathbf{o} , assuming that \mathbf{x} is defined relative to \mathbf{o} .

The observed field U at the far-field is then the superposition of the phasors corresponding to each phase variation⁴:

$$U(\mathbf{k}^i, \mathbf{k}^s) = \int_{\mathcal{S}} \exp [j\Delta\Phi(\mathbf{x}, \mathbf{k}^i, \mathbf{k}^s)] d\mathbf{x}, \quad (44)$$

where we assume amplitude at each point is constant over \mathcal{S} . Indeed, the above equation is a continuous version of the discrete phasor sum in the previous section (Eq. (34)). Evaluating the ensemble-averaged field intensity over the statistical surface yields the BRDF, which has been extensively studied using various models [Nayar 1991; Ticconi et al. 2011; Yan et al. 2018]. Instead, we are particularly interested in how the power distribution changes in the spectral domain when the surface undergoes motion over time.

⁴We use U instead of E to cover both coherent (field) and incoherent (intensity) cases.

6.3 Doppler Frequency Shift of Coherent Detection

Now, let's consider the case that the surface moves over time (Fig. 8-(B)) with macroscopic velocity \mathbf{v} , which has a normal component $\mathbf{v}_N = v_N \mathbf{n}$ and a transverse component \mathbf{v}_T . Similar to Eq. (44), we can evaluate U at time t as:

$$U(t) = \int_{\mathcal{S}_t} \exp [j(h(\mathbf{x})(\mathbf{k}^m \cdot \mathbf{n}) + \mathbf{k}^m \cdot (M_t \mathbf{x}))] d\mathbf{x}, \quad (45)$$

where \mathcal{S}_t is a set of local coordinates that is illuminated at t and $M_t \mathbf{x} = \mathbf{x} + \mathbf{v}t$, which is \mathbf{y} in Fig. 8-(B). We omitted $\mathbf{k}^i, \mathbf{k}^s$ for simplicity. If transverse velocity only exists, we have $\mathcal{S}_t = \mathcal{S} - \mathbf{v}_T t$, where $\mathcal{S} = \mathcal{S}_{t=0}$. However, if incident angle θ_i is not zero, normal velocity also affects the local coordinate change. We can calculate this local coordinate change rate $\mathbf{u}_1 = \mathbf{v}_T - \hat{\mathbf{a}}_T^i v_N \tan \theta_i$. Since $\mathcal{S}_t = \mathcal{S} + \mathbf{u}_1 t$ and $M_t \mathbf{x} = \mathbf{x} + \mathbf{v}t$, we can rewrite Eq. (45) on integration domain \mathcal{S} as :

$$U(t) = \int_{\mathcal{S}} \exp [j(h(\mathbf{x} + \mathbf{u}_1 t)(\mathbf{k}^m \cdot \mathbf{n}) + \mathbf{k}^m \cdot (\mathbf{x} + \mathbf{v}t + \mathbf{u}_1 t))] d\mathbf{x}. \quad (46)$$

To make the notation simpler, we also define $\mathbf{u}_2 = \mathbf{v} + \mathbf{u}_1$. In fact, this is the instantaneous velocity of the intersection point on the macroscopic geometry along the ray, or *spot*, $\mathbf{z} = \mathbf{x} + \mathbf{u}_2 t$ which can be calculated as $-v_N / \cos \theta_i \hat{\mathbf{a}}^i = \mathbf{v}_N - \hat{\mathbf{a}}_T^i v_N \tan \theta_i$. Using $\mathbf{u}_1, \mathbf{u}_2$, we can evaluate PSD of $U(t)$ from its ensemble-averaged autocorrelation $\langle U(t)U^*(t + \tau) \rangle$ or just $C_U(\tau)$ with stationary assumption:

$$C_U(\tau) = \int_{\mathcal{S}} \int_{\mathcal{S}} \langle \exp [j(h(\mathbf{x}_1) - h(\mathbf{x}_2 + \mathbf{u}_1 \tau))(\mathbf{k}^m \cdot \mathbf{n})] \exp [j(\mathbf{k}^m \cdot (\mathbf{x}_1 - \mathbf{x}_2 - \mathbf{u}_2 \tau))] d\mathbf{x}_1 d\mathbf{x}_2 \quad (47)$$

If the surface is sufficiently rough (RMS height $\sigma_h \gg \lambda$) and uncorrelated (spatial autocorrelation function is approximated to delta function), which was also assumed in Goodman [1975]; Kliese and Rakić [2012], the ensemble averaged term in last line of Eq. (47) becomes negligible unless $\mathbf{x}_1 = \mathbf{x}_2 + \mathbf{u}_1 \tau$. Substituting this, we get

$$C_U(\tau) \approx \alpha \int_{\mathcal{S}} \exp [j(\mathbf{k}^m \cdot (\mathbf{u}_1 - \mathbf{u}_2) \tau)] d\mathbf{x} = \alpha |\mathcal{S}| e^{-j(\mathbf{k}^m \cdot \mathbf{v}) \tau}, \quad (48)$$

where α is a constant term proportional to BRDF, which is inserted to satisfy the energy conservation law. Taking Fourier transform gives PSD of frequency shift,

$$S(\Delta\omega) = \mathcal{F}(C_U(\tau)) = \alpha |\mathcal{S}| \delta(\Delta\omega - \mathbf{k}^m \cdot \mathbf{v}) \quad (49)$$

This indicates that observed frequency is shifted by the amount of $\Delta\omega = \mathbf{k}^m \cdot \mathbf{v}$. We can use another expression using unit vectors :

$$\Delta\omega = \mathbf{k}^m \cdot \mathbf{v} = (\mathbf{k}^s - \mathbf{k}^i) \cdot \mathbf{v} = \frac{2\pi}{\lambda} (\hat{\mathbf{a}}^s - \hat{\mathbf{a}}^i) \cdot \mathbf{v} = \frac{\omega}{c} (\hat{\mathbf{a}}^s - \hat{\mathbf{a}}^i) \cdot \mathbf{v}, \quad (50)$$

which matches with traditional Doppler shift formula. We can define a scalar value called *target velocity* for this Doppler shift as

$$v_{\text{tg}} = (\hat{\mathbf{a}}^s - \hat{\mathbf{a}}^i) \cdot \mathbf{v}. \quad (51)$$

In fact, this name is from Dolan [2009], and is coined so, as it measures the velocity of microscopic geometry (target) on the surface.

6.4 Doppler Frequency Shift of Incoherent Detection

For incoherent (intensity) detection such as AMCW, the Doppler effect occurs on intensity modulation instead of the optical field itself. Still, we can make use of Eq. (44), considering U as an intensity summation, instead of field summation. Furthermore, in this case, λ is in meter scale (carrier frequency is still in micrometer scale, but amplitude modulation frequency that causes Doppler effect has meter scale), so $\sigma_h \ll \lambda$ and we can ignore height dependent terms in Eq. (45), which gives

$$U(t) = \alpha \int_{\mathcal{S}} \exp [j (\mathbf{k}^m \cdot (\mathbf{x} + \mathbf{u}_2 t))] dx, \quad (52)$$

where BRDF related term α now appears directly at $U(t)$. Assuming λ^2 is much larger than \mathcal{S} , we can further approximate into

$$U(t) \approx \alpha |\mathcal{S}| \exp [j (\mathbf{k}^m \cdot \mathbf{u}_2 t)] \exp [j (\mathbf{k}^m \cdot \mathbf{x}_{\text{mean}})], \quad (53)$$

where $\mathbf{x}_{\text{mean}} = 1/|\mathcal{S}| \int_{\mathcal{S}} \mathbf{x} dx$. Consequently, we have

$$\langle U(t)U^*(t + \tau) \rangle = \alpha^2 |\mathcal{S}|^2 \exp [-j (\mathbf{k}^m \cdot \mathbf{u}_2) \tau]. \quad (54)$$

Similar to target velocity, we can get

$$\mathbf{v}_{\text{sp}} = (\hat{\mathbf{a}}^s - \hat{\mathbf{a}}^i) \mathbf{u}_2 = \frac{v_N}{\cos \theta_1} (1 - \hat{\mathbf{a}}^s \cdot \hat{\mathbf{a}}^i), \quad (55)$$

which is called *spot velocity*. Again, this name is from Dolan [2009], which describes the velocity of *spot* point on a macroscopic surface lit by illumination that is detached from micro-geometry.

6.5 Target vs Spot Velocity

To help with understanding, comparison between target and spot velocity is presented in Fig. 9 for normal and transverse motions illuminated at an oblique angle. For both cases, the target velocity arises from correspondence that conserves local geometry, which we call *attached* correspondence of the microscopic target geometry.⁵ Since the illuminated surface is finite in practice, not all attached correspondences are valid. However, assuming such an invalid region is small enough, the property of the bulk region, where the attached correspondence is valid, governs the whole property. Note that such attached correspondence was already exploited in our discussion in Sec. 4 and Sec. 5. On the other hand, for spot velocity, locality is not preserved as every point is identical (due to long wavelength). Instead, we can find correspondence of macroscopic surface region change, which we call *detached* correspondence of spot point lit by the incident light. Such distinct correspondence leads us to observe different Doppler values for each case.

Monte Carlo Simulation. Here, we also show the MC simulation results to show the difference between target and spot velocity. As shown in Fig. 10-(A), we generated a Gaussian random rough surface $h(\mathbf{x})$ parameterized by RMS height σ_h and correlation length l . The light wavelength is set to $1.55 \mu\text{m}$ and illuminated area is $100 \mu\text{m} \times 100 \mu\text{m}$. The camera and light source are set to be collocated and oriented at 45 degrees with respect to the surface normal and we set the surface to move along its normal direction with a velocity of 10 m/sec. By default, σ_h and l are set to 2λ and 10λ . In

⁵We only considered surface reflection in this section, but one can easily notice that particles in participating media should naturally belong to attached correspondence that follows the well-known Doppler formula [Drain 1980].

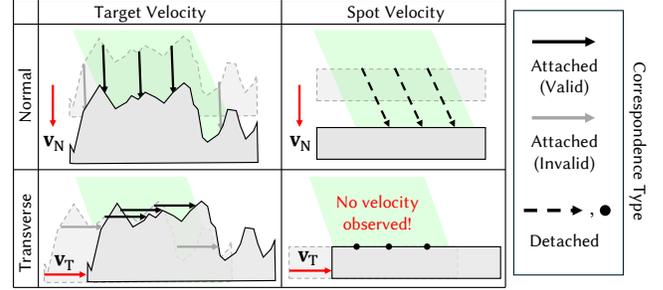


Fig. 9. Spot and target velocity comparison for normal and transverse motions. Target velocity originates from local microgeometry correspondence, which is called attached, while spot velocity comes from detached macroscopic correspondence.

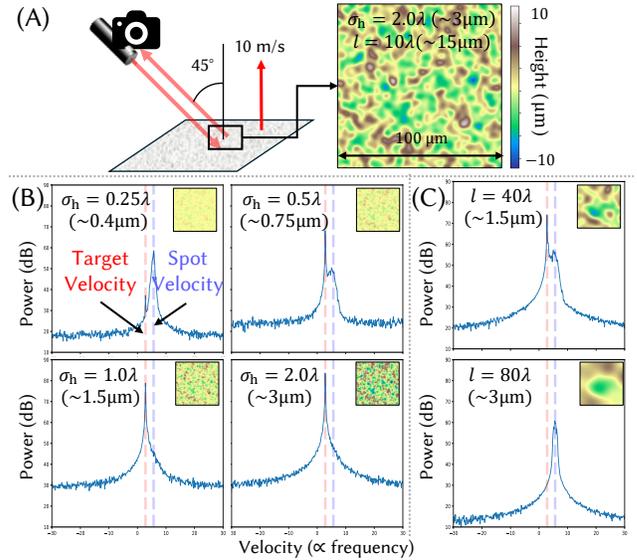


Fig. 10. (A) Experimental setup of MC simulation for coherent detection. We show the height field lit by the laser and collocated camera. (B) PSD with different σ_h s. If σ_h is small compared to λ , the peak appears at the spot velocity, but it gradually translates to the target velocity as σ_h increases. (C) PSD with different l s. If l is too large, we observe spot velocity, as nearby points become highly correlated which invalidates our assumption for target velocity detection.

Fig. 10-(B), we show results for different σ_h averaged over 64 surface realizations. If σ_h is small, we could detect spot velocity, and if σ_h becomes larger, we could detect target velocity. This is because whether each microscopic structure becomes undistinguishable depends on σ_h as explained before. In Fig. 10-(C), we show results for larger l . We found that larger l makes us detect spot velocity. This is because large l makes the neighboring point of the surface highly correlated which invalidates our assumption on target velocity detection. Though we haven't explicitly modeled BRDF, we could also notice that the field intensity decreases for smaller σ_h and larger l , which coincides with the fact that slope σ_h/l determines the BRDF in microfacet theory.

Table 1. Comparison of Doppler effect in coherent and incoherent detection discussed in Sec. 6.

Measurement Type	Incoherent Detection (AMCW)	Coherent Detection (OHD)
Physical Principle	Macro scale depth change rate	Doppler effect on micro geometry
Wavelength for Doppler	Long (intensity modulation: $\lambda = 10$ m)	Short (optical field : $\lambda = 1550$ nm)
What is measured?	Spot velocity v_{sp}	Target velocity v_{tg}
Simulation	Need to explicitly move objects [Kim et al. 2023]	static geometry + calculate Eq. (32)
Treadmill, Rotating wheel	Cannot detect frequency shift	Can detect frequency shift

6.6 Implications on Simulation Strategy

The argument in the previous section implies that we need a different simulation strategy for AMCW and OHD Doppler measurement. First, AMCW cameras do not detect the Doppler effect that occurs due to micro-scale reflectors, and hence, we only have to consider macroscopic object movements over time. As Kim et al. [2023] pointed out, *shift-mapping* in gradient-rendering literature [Kettunen et al. 2015; Lin et al. 2022; Manzi et al. 2016] could be helpful to find macroscopic correspondence under complex path domain. On the other hand, OHD detects the Doppler effect that occurs due to micro-scale reflectors' motion, and hence, we have to either consider their movement at the microscopic scale, which is computationally expensive or use their statistics to simulate the OHD system, which is what we propose in this manuscript. We provide a summary of this section in Tab. 1.

7 APPLICATIONS: (1) FMCW LIDAR

In this section and the following two, we use the proposed technique to simulate three different OHD applications—FMCW lidar, blood flow velocimetry, and wind Doppler lidar. Each part is structured with a brief explanation of the application principle and demonstration using our simulator, except for FMCW lidar, which has been already explained sufficiently before. For FMCW lidar, the real measurements are from our hardware prototype, while for others, they are from previous papers. We implemented our algorithm based on the Mitsuba renderer [Jakob 2013].

7.1 FMCW: Simulation

7.1.1 Demonstration of Proposed Algorithms. We first demonstrate the correctness of our proposed algorithms—PSD evaluation (Algorithm 1) and field sampling (Algorithm 2) for rendering case, similar to Fig. 7. As shown in Fig. 11-(A), we consider a SPECULAR CORNELL-BOX scene, which is the same as the well-known CORNELL-BOX scene, but with a specular floor. The velocity of the moving objects is shown in the scene figure (Fig. 11-(A)). The simulated power spectrum at the target point (red dot) is plotted on Fig. 11-(B), for both up and down chirp. Orange colored lines are from Algorithm 1, the bright solid line represents $S_{AC}(\omega)$, while the dark dashed line represents convolved $S_{AC}(\omega)$ as Eq. (28). Blue colored lines are from Algorithm 2, the lighter one shows a single measurement, while the darker line shows the averaged power spectrum from 1000 realizations. We could find that each algorithm gives the same spectrum, which validates the convergence of the proposed method, even for complex scene rendering. Another important observation is regarding the multi-bounce effect. The green line in Fig. 11-(B) shows a

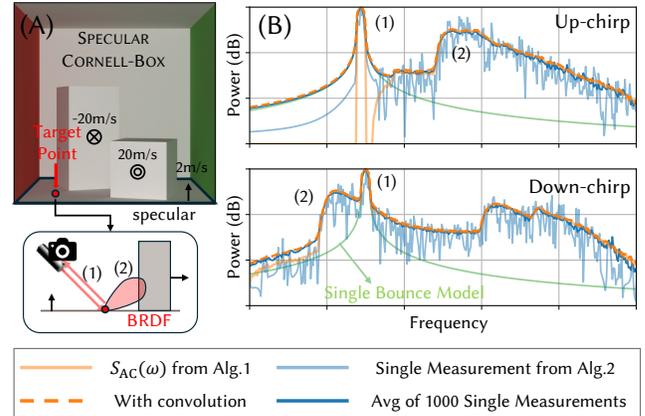


Fig. 11. (A) We show simulated power spectrum evaluated for SPECULAR CORNELL-BOX with proposed method. The camera and laser are outside of the box, looking at the target point. (B) Orange lines represent results from Algorithm 1, where a bright solid line is evaluated PSD and a dark dashed line is convolved with time-window. The light blue line shows a single measurement using Algorithm 2, while dark blue is an averaged result over 1000 measurements, which converged to the result from Algorithm 1.

simple single-bounce model from the floor point, which is denoted as (1). However, it fails to reproduce the spectrum from the multi-bounced light (e.g., from the tall box denoted as (2)), which verifies the importance of considering the multi-bounce effect.

Rendering performance statistics. We also report performance statistics under our hardware (AMD 3955WX). Instead of single-point rendering as Fig. 11, we rendered the whole image for evaluation. The bin number is set to 3072 for both methods with 256 spp. The image size is 128×128 and the maximum bounce is set to 4. For PSD sampling, it took 6 seconds to build the $128 \times 128 \times 3072$ PSD histogram while field sampling took a much longer time, (75 seconds) due to the existence of double for-loops. Considering that it took 1000 iterations to make field sampling converge to GT PSD in the single-point case, Algorithm 2 is not practical for image-wide simulation. Thus, in general, we believe that PSD evaluation and sampling (Algorithm 1) would be a default option for OHD simulation.

7.1.2 Distance and Velocity Reconstruction. We also show more results on the complex scenes in Fig. 12 and perform image-wide distance and velocity reconstruction from the evaluated PSD using Algorithm 1. We used the argmax operator to determine the peak frequency. For the LIVING-ROOM scene, we found that there are lots

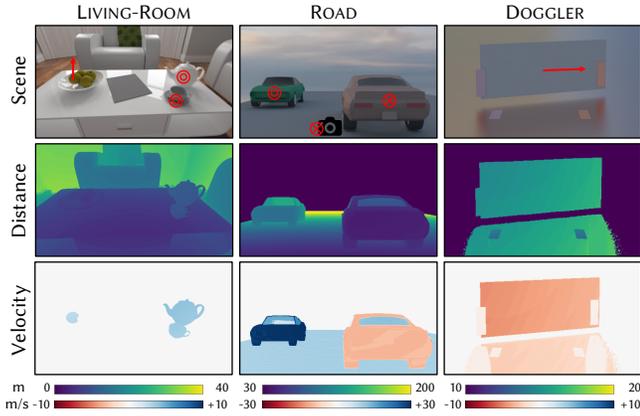


Fig. 12. We show distance and velocity reconstruction results for several different scenes. We denote the object moving direction as red arrows. Overall, due to the multi-path effect, we found that the result deviates from just considering the first hit point.

of inter-reflections at the corners which makes depth reconstruction inaccurate. For ROAD scene, the observer is also moving forward, but the pink car is moving faster. Thus, we see a small Doppler effect for the pink car, but much larger for the green car. The car windows show different behavior from its body as most of the light gets refracted after it hits the window. Finally, for the DOGGLER scene (dog treadmill + Doppler), we could clearly observe the wrong distance and velocity at the shiny floor. We will defer the details for DOGGLER scene to Sec. 7.2.1, where we compare it with real experiments. Overall, these simulations show that when single-bounce assumptions fail, it is important to consider the full power spectrum.

7.1.3 Simulation Comparison with Kim et al. [2023]. We visualize the simulation difference between our FMCW formulation and Kim et al. [2023]’s method for AMCW Doppler simulation for CORNELL-Box scene in Fig. 13. For better visualization, we only considered a direct bounce. For Fig. 13-(A), we used explicitly moving objects over time as Kim et al. [2023] suggested for D-ToF simulation. Note that we still use the FMCW camera but simulate the Doppler effect with explicit movement as their method. For Fig. 13-(B), we used a path velocity calculation from Eq. (32). Both methods give the same result for depth reconstruction but largely differ in velocity reconstruction. (A) gives the spot velocity, which is larger at the grazing angle (side of boxes, floor). On the other hand, (B) gives near-constant results over a single object as v_N is almost the same. Also since the floor moves nearly perpendicularly to the camera, we could observe very small velocity. As discussed in Sec. 6, the correct simulation for FMCW is (B), but for the simulation of the D-ToF camera, we need to use (A).

7.2 FMCW: Real Data Validation

We already showed some speckle statistics from real FMCW data in Sec. 5, here we demonstrate two more experiments. Except for simple cases, it is in general challenging to accurately calibrate

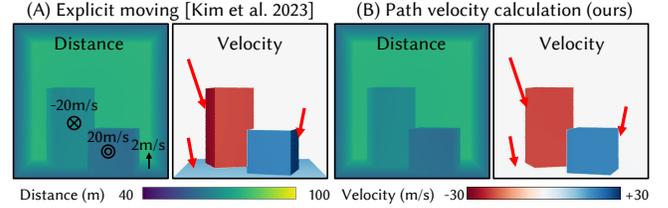


Fig. 13. Comparison between two different simulation strategies-(A) explicitly move objects as Kim et al. [2023] and (B) calculating path velocity from Doppler formula. They give the same depth reconstruction, but different velocities that (A) gives spot velocity, which is larger at the edge of boxes or floor, and (B) gives constant value over a single object.

geometric and photometric aspects of a scene and create its digital twin. So, we resort to qualitative comparisons similar to previous rendering papers [Kim et al. 2023; Pediredla et al. 2020].

7.2.1 FMCW Measurement on Different Materials. We first demonstrate the correctness of our simulator using real FMCW measurement on different materials which may cause a multipath effect as shown in Fig. 14. We have a single-pixel FMCW scanner that is scanning along a line on the floor with varying reflection properties. The ray reflected on the floor material hits the back wall, made of spectralon, a highly diffuse material. We experimented with four different floor materials. Inside the simulator, we tried to match all simulation materials as closely to the real materials as possible using various Mitsuba materials, whose details could be found in the supplementary material.

The average PSD along the scanning line for both real and simulation with various floor materials is shown in Fig. 14. Since the scene is static (ignore scanning speed for now) we can compute the distance from the PSD. If the floor material is specular, most of the

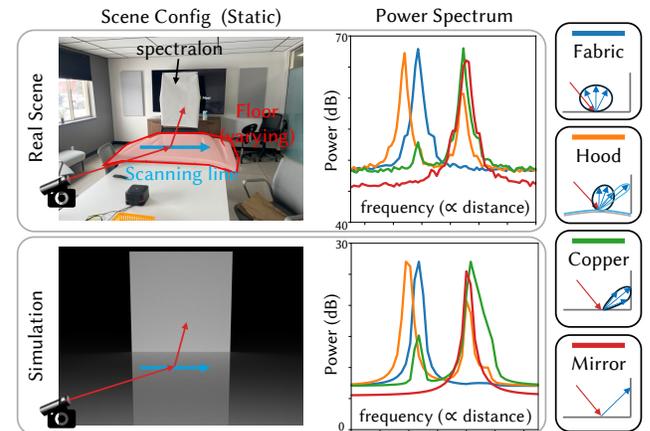


Fig. 14. To show the accuracy of our simulator, we built a multipath scene with different floor materials. The back wall is spectralon, a highly diffuse material. By interpreting the average PSD over the scanning line, we observed that the real-world measurements matched our simulator for a wide range of materials.

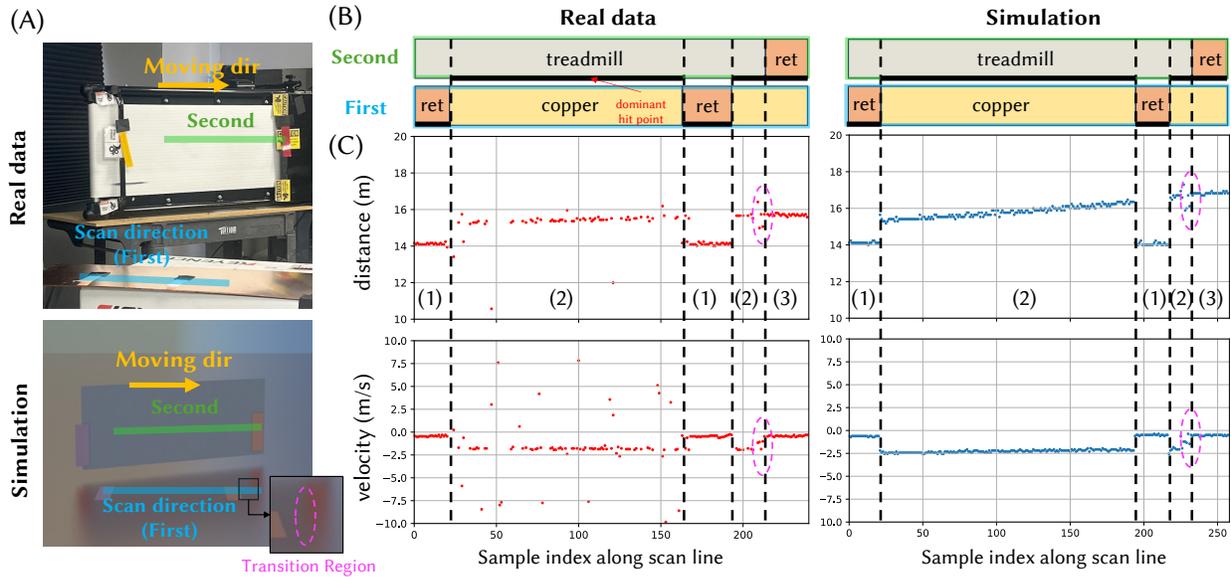


Fig. 15. (A) We demonstrate our simulator for an experiment that measured copper with the treadmill behind. We represent the first hit points which align with the scan direction with a sky-blue color and the second hit points on the treadmill with a light-green color. (B) We represent material on the first and second hit points as a bar plot. We also denote the dominant hit point as a black line. Considering the material property, these dominant hit points could be grouped into three regions with different reflections: (1) retro (2) copper-treadmill, and (3) copper-retro. (C) We show velocity and distance reconstruction from the power spectrum, along the scan line. We found that our simulator faithfully reproduces the result corresponding to each region, even including the smooth transition region (pink dashed circle).

light received by the sensor arrives from the back wall, which is at a larger depth and hence shows up as a peak at a higher frequency. When the floor is made up of a non-reflective material (fabric), the dominant light return is from the first bounce and the peak is at a lower frequency. Overall, we qualitatively found that our simulator matches the real data for varying materials in the spectral domain.

7.2.2 FMCW Measurement on Treadmill after Reflection. We demonstrate our simulator on dynamic scene as illustrated in Fig. 15. The overall setting is similar to the previous static case; the floor material is fixed to be copper and the back wall is replaced with a treadmill at an oblique angle. A total of four retroreflectors are used to indicate the start and end points on both the copper plane and the treadmill. They are all fixed on nearby objects. The only moving object in the scene is the treadmill surface.

Given this scene configuration, we can expect two possible dominant hit points on the copper plane and on the treadmill. We mark the first bounce hit points as sky-blue color in Fig. 15-(A), which is equal to where the scan line is. We also mark the possible second bounce hit points on the treadmill with a light green color. The corresponding material for each region is represented as a bar graph (gray:treadmill, yellow:copper, orange: retroreflector) in Fig. 15-(B). Considering the reflectivity properties, we denote the point that most of the rays are reflected back, as a black line on the bar plot in Fig. 15-(B). Then, we can find three distinguishable regions:

- (1) retro (first) - the laser beam is directly reflected back from the first hit point, which shows a distance around 14 m without velocity

- (2) copper (first), treadmill (second) - the laser beam is reflected on copper and reflects back from the treadmill. It shows distance at secondary hit around 16 m with velocity of 2 m/s.
- (3) copper (first), retro (second) - the laser beam is reflected on copper and reflects back from the retroreflector on the treadmill which is at a distance of 16 m without velocity.

The exact region is denoted in Fig. 15-(B) as (1), (2), (3) for both real data and simulation.

In Fig. 15-(C), we reconstruct the distance and velocity for both real data and simulation from the evaluated PSD. We set the bin with the maximum power to be the representative hit point. Overall, we find that our simulator reliably reproduces the experiment for each region that distinguishes one another, even though there are some disparities due to imperfect calibration. There are three things to point out here. First, there is a slight increment of distance over the scan line. This is due to the treadmill being at an oblique angle and thus not at a fixed distance from the lidar. Second is the presence of background velocity - we could observe a small amount of velocity even when detecting the fixed retroreflector on the copper plane. This is because the imaging system is actively scanned at a high rate. To reproduce this, we also added a small velocity component for the sensor for simulation. Finally, between (2) and (3), there is a smooth transition region, which is not strictly divided by a retro-reflector (marked as a pink dashed circle in Fig. 15). For this region, distance and velocity give unexpected values that deviate from both (2) and (3). We found that our simulator even managed to reproduce these artifacts, which could be useful for future studies to resolve them.

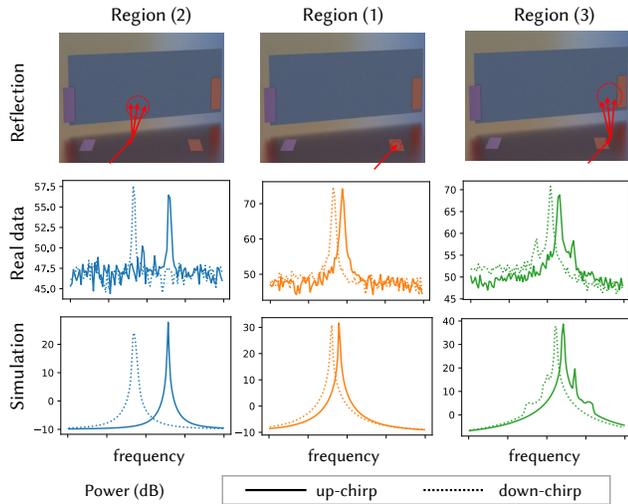


Fig. 16. We show frequency domain distribution for three different points, using FFT on real data and simulation. We also plot GT simulation frequency in the last row. For region (2), we could observe a split between up and down chirps as the ray hits the treadmill after the reflection. For region (1), the ray hits the retroreflector, but we still could observe a small peak difference which is due to the laser scan. For region (3), strong frequency distribution broadening was observed as reflected light can hit both treadmill (moving) or retroreflector (fixed).

In Fig. 16, we show the spectral analysis of the data and simulation. We plot the aforementioned three different regions for both up and down chirps. For the copper-treadmill case (Region 2), we could not observe that much reflection from copper as most of the rays are reflected due to the grazing angle (Fresnel reflection). For the retroreflector case (Region 1), we could expect and observe a single strong peak at the retroreflector. Also since the camera is moving, we could clearly observe that the down and up chirp peaks do not match which causes small bias velocity in Fig. 15-(B). Lastly, for the copper-retroreflector case (Region 3), we could clearly observe the effect of frequency distribution spreading. To elaborate, there is one dominant peak from the retroreflector on the treadmill, but we could also observe a secondary peak from the treadmill. Overall, we could verify that our FMCW simulator faithfully reproduces real-world measurements in both distance-velocity reconstructions and power spectrum.

8 APPLICATIONS: (2) BLOOD FLOW VELOCIMETRY

In this section, we show a biomedical OHD application called laser Doppler holography [Puyo et al. 2021, 2018, 2019], a type of laser Doppler velocimetry, but provides full-field Doppler image with high temporal resolution.

8.1 Laser Doppler Holography: Principle

Laser Doppler holography (LDH) uses the OHD principle to non-invasively measure the retinal blood flow dynamics. Traditional biomedical laser Doppler velocimetry (LDV) is point-based technique that measures Doppler frequency shift at a specific location.

Thus, we need scanning systems to monitor blood flow over a full field (similar to FMCW). In contrast, LDH is inherently a full-field imaging technique based on digital holography, where the entire interference pattern is recorded as a sequence of images from a digital camera, which is called *hologram stream*. The space-resolved PSD can be obtained by performing a short-time Fourier transform (STFT) along the temporal axis. LDH offers the advantage of high temporal resolution (1.6 ms), en-face imaging, and the ability to analyze dynamic flow changes, such as during a cardiac cycle, which traditional LDV cannot achieve.

The overall optical setup is similar to that of FMCW lidar. The difference is that since blood flow ($1\text{--}100\text{ mm s}^{-1}$) is much slower than target velocity in FMCW applications ($1\text{--}100\text{ m s}^{-1}$), beat frequency also is much smaller in $10\text{--}20\text{ kHz}$ [Riva et al. 1985]. Thus, unlike FMCW, which requires a sampling rate of GHz, a fast CMOS camera that operates in tens of the range kHz can be used for LDH.

8.2 Laser Doppler Holography: Simulation

We replicate and demonstrate experimental results from Puyo et al. [2018, 2019] using our simulator. We use the same settings from these papers in our simulation. A fast CMOS camera has a frame rate of 39 kHz, 512×512 size. Laser has a wavelength of $\lambda = 785\text{ nm}$. Single STFT time window σ_w is set to 1.6, 13.1ms which results in number of STFT samples $N_w = 64, 512$ respectively.

We show the overall simulation configuration in Fig. 17-(A-C). We modeled a retina as Fig. 17-(A), whose radiometric image at the camera is shown in Fig. 17-(B). The circular shape in Fig. 17-(B) represents the pupil of the eye, and we can see the retinal blood vessels behind it. Fig. 17-(C) shows examples of three different possible paths in the simulation, the path from other retinal tissues which are static (p_1 , solid line), from moving red blood cell (p_2 , dashed line) and other multi-bounced paths (p_3 , dotted line). We modeled light transport at RBC and eye tissue/matrix as volumetric path tracing with the Henyey-Greenstein model following Hammer et al. [1998, 2001]. We adjusted several coefficients to avoid path tracing too dense particles, but one can apply advanced rendering algorithms such as subsurface scattering or translucent rendering [Jimenez et al. 2010; Wann Jensen et al. 2023] in the future.

The simulation results are presented in Fig. 18-(A-D). First, we evaluate zeroth-order moment $M_{0\pm}$ which is often used for both

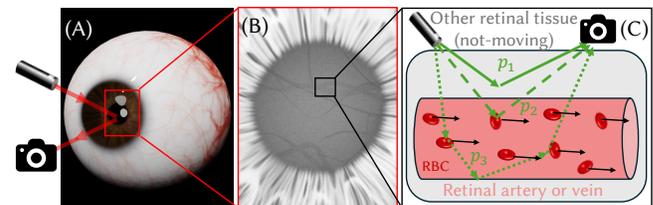


Fig. 17. (A) A coherent light illuminates the eye pupil and we detect it with a high frame rate CMOS camera. (B) An image from standard rendering. We can see retinal tissues over the pupil. (C) We visualize simplified light transport inside the eye. The red blood cells (RBC) in a retinal artery or vein cause a Doppler frequency shift. We also mark three different paths p_1 (static object), p_2 (direct bounce from RBC), and p_3 (multi bounced path).

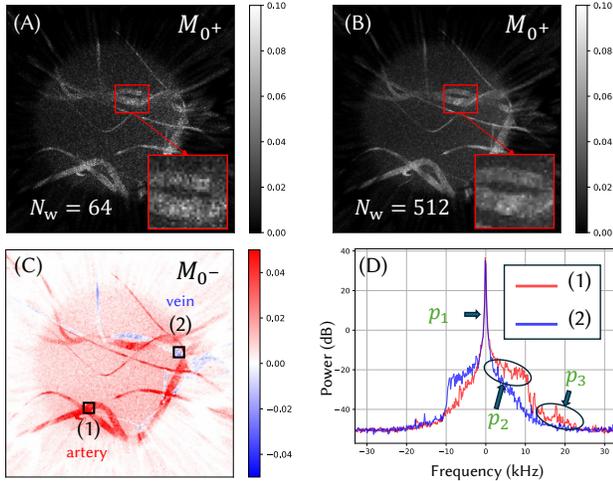


Fig. 18. (A, B) M_{0+} for different STFT ($N_w = 64, 512$). Similar to Puyo et al. [2018], longer N_w gives less noisy result. (C) M_{0+} with $N_w = 512$ shows the direction of the flow. We can see different colors for the artery and vein as the blood flow direction is opposite in these two vessels. (D) PSD of two ROIs at an artery and a vein. We can see a high peak at zero as most of the tissue is not moving (from p_1 in Fig. 17), and different frequency shifts for artery and vein. We could also see contribution at a high frequency due to multi-bounce (p_3 in Fig. 17).

analysis and visualization, that is defined as

$$M_{0\pm} = \int_{\omega_1}^{\omega_2} S(\omega) \pm S(-\omega) d\omega. \quad (56)$$

$M_{0\pm}$ is also known as *power Doppler images* (PDIs) in this literature. M_{0+} represents the total power of the Doppler PSD, while M_{0-} represents the asymmetry (direction) of the Doppler PSD so that we can obtain full Doppler information. Fig. 18-(A,B) shows M_{0+} for two different STFT windows ($N_w = 64, 512$). We could find less noisy results for $N_w = 512$ which matches [Puyo et al. 2018]’s empirical observations. Note that since the sampling frequency remains the same, we have the same frequency range from FFT for both cases. Though each frequency bin has the same distribution statistics, larger N_w has more frequency bins so when we calculate M_{0+} , we effectively have 8 times more samples, which gives the less noisy image. However, for full LDH, we need to time-average the sequence of M_{0+} , and this gives the same noise for both cases.

Fig. 18-(C) shows M_{0-} for $N_w = 512$. We can see two different colors, red and blue which show the direction. Typically artery and vein have different flow directions, so red corresponds to the artery while blue corresponds to the vein in this case. Compared to blue, red has smaller values as the vein typically has slower velocity compared to the artery. Also due to multiple scattering, we can see Doppler shift even for static tissues. In fact, the diminished Doppler shift in choroidal vessels is caused by the presence of pigment layers and denser tissue. Overall, the results resemble real results from Puyo et al. [2018] (Fig. 5-(c,d)). Fig. 18-(D) shows PSD of two ROIs in Fig. 18-(C). We can see a high peak at zero frequency as most tissues are static (p_1 in Fig. 17). Also, we can see two different distributions

at vein and artery that resemble results from Puyo et al. [2021] (Fig. 1-(C)) and Puyo et al. [2018] (Fig. 7-(d,e)). Compared to single scattering (p_2 in Fig. 17), we could also see the contributions of multiple scattering (p_3 in Fig. 17) at higher frequency region which is over the maximum velocity range.

9 APPLICATIONS: (3) ATMOSPHERIC SENSING

For atmospheric applications, we show pulsed coherent wind Doppler lidar (C-DWL), which has been widely used in applications like wind profiling, weather forecasting, and atmospheric research [Dabas et al. 1994; Shangguan et al. 2022].

9.1 Pulsed Coherent Wind Doppler Lidar : Principle

Pulsed C-DWL is a laser-based (usually $\lambda = 1.5 \mu\text{m}$) remote sensing system used to measure wind speed and direction at various distances (range-resolved) by exploiting the Doppler shift of backscattered light. The operating principle of pulsed C-DWL is plotted in Fig. 19. It operates by emitting a short, coherent laser pulse into the atmosphere, where the light interacts with aerosols or particles in the air and is reflected back to the lidar system. Same as other OHD applications, the backscattered signal is mixed with the LO beam and we perform STFT to measure the frequency shift caused by the motion of the scatterers within a certain time window. By analyzing this Doppler shift and time-of-flight, the system detects range-resolved 2D Doppler spectra as illustrated in Fig. 19. Pulsed C-DWL is widely used for line-of-sight wind speed detection, but its spectral information also allows us to observe precipitation [Wei et al. 2019], or identify clouds [Yuan et al. 2020] and raindrops [Wei et al. 2021].

9.2 Pulsed Coherent Wind Doppler Lidar : Simulation

Here, we aim to reproduce the results from Wei et al. [2019], who demonstrated C-DWL PSD under different weather conditions. We modeled the atmosphere from [0, 1000]m, with 100 different layers with different properties. We use the Henyey-Greenstein model for aerosol, with varying $\sigma_t \in [0.001, 0.01]$ and $g = 0.7$. Each air layer

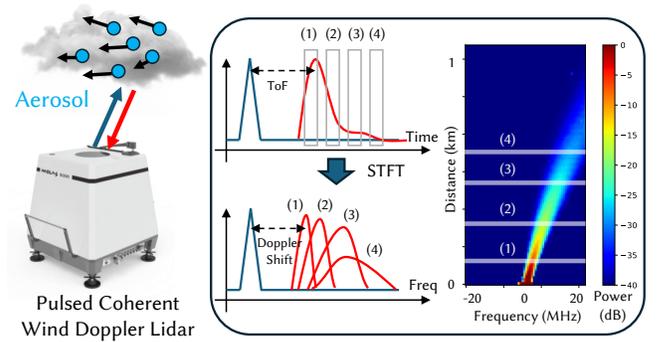


Fig. 19. The principle of pulsed C-DWL. We shoot a pulsed laser and measure the transient photocurrent signal. Then, we perform STFT to get the power spectrum at each ToF. Finally, we get ToF (distance), and frequency (velocity) resolved 2D spectrum as shown. We also denoted 4 regions and corresponding lines on the 2D spectrum.

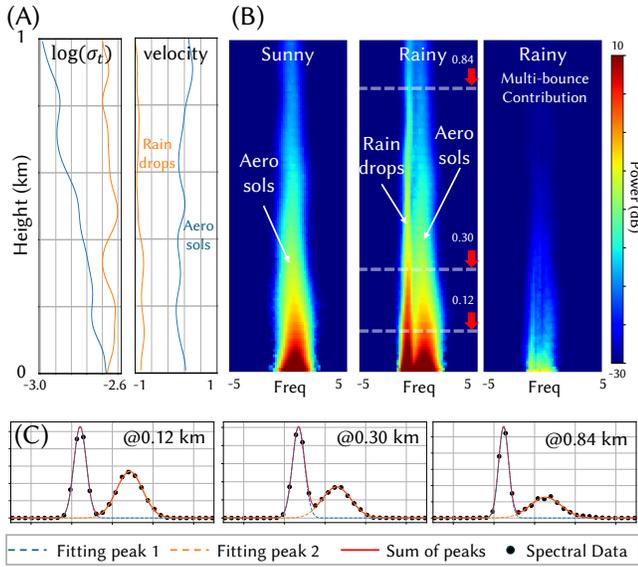


Fig. 20. (A) Aerosol density and velocity profile used for simulation. (B) Distance-resolved 2D PSD result for sunny and rainy days. Compared to a sunny day, a rainy day shows another sharp peak from the raindrops. We also plot multi-bounce components for the rainy day. (C) We show Gaussian curve fitting at three different height points (0.12, 0.30, 0.84 km) for rainy days, which faithfully replicates [Wei et al. 2019]’s observation.

is set to have a different velocity $v \in [0.0, 0.5]$ with some variation. The aerosol profile used for simulation is plotted in Fig. 20-(A). We simulated two scenarios: one for a sunny day and another for a rainy day, similar to Wei et al. [2019]’s measurements. For a rainy day, we additionally simulated raindrops, which have different profiles compared to aerosols, as plotted in Fig. 20-(A). Fig. 20-(B) shows ToF resolved PSD for each weather condition. We could observe a single peak in sunny day PSD due to aerosols and a double peak in rainy day PSD due to both aerosols and raindrops. We also plot PSD at three different height points (0.12, 0.30, 0.84 km) as a black dot in Fig. 20-(C). We perform curve-fitting with two Gaussians (blue and orange dashed lines) and find that the sum of two Gaussians (red line) well fits with the original data, which coincides with Wei et al. [2019].

We also plot the contribution of multi-bounced paths in the third row of Fig. 20-(B). In fact, decomposing light into such multi-bounce components is typically not possible with real data, but in simulation, we could easily compute it by controlling the total number of bounces. Such new information from simulators could be useful in improving the imaging systems, as shown for virtual optical waveguides (acousto-optics) by Pediredla et al. [2023].

10 DISCUSSION AND CONCLUSION

In this paper, we presented physically accurate Monte Carlo simulator for OHD techniques and demonstrated its utility on three OHD imaging systems. Our key contribution is the path-integral formulation [Veach 1997] for spectral-domain OHD simulation, which brings Monte Carlo path sampling techniques to accurately model

and simulate various OHD systems. We also investigated sampling using speckle statistics and conducted a theoretical comparison of the Doppler effect in coherent (OHD) and incoherent (AMCW) detection. Our findings suggest that different rendering techniques are needed for each Doppler effect—[Kim et al. 2023] for AMCW Doppler and the proposed method for OHD Doppler.

Here, we discuss several points that we have not considered. First, we do not explicitly model the micro-surface itself and instead model its statistics (Sec. 6). While this reduces the computational complexity, our simulator cannot directly reproduce some resulting phenomena. One of them is the Doppler broadening due to time-varying micro-geometry. We ignored the effect of invalid correspondence in target velocity, but if the illuminated region moves quite fast over time compared to its area, we cannot ignore this effect—new micro-geometry contributes each time. This makes the Doppler frequency shift not single-peaked, but it will have a broadened spectrum. One can also understand this as time-varying perturbation ψ in phasor calculation which effectively decays temporal coherence and hence causes frequency broadening. In fact, such Doppler broadening has been reported for fast-moving objects or small beam-waist by Kliese and Rakić [2012]. Though our current model cannot handle this case directly, one can easily incorporate this in post-processing by replacing the δ function in our path integral Eq. (26) with the broadening kernel. Second, we assumed a constant heterodyne efficiency, coming up with a better model for more accurately evaluating heterodyne efficiency is an interesting future direction. To do so, we need to consider accurate spatial speckle statistics, which may be done similarly to Bar et al. [2019]. Third, we stick to the path-tracing framework in this manuscript and did not consider other wave-optics effects such as diffraction. Incorporating the same, similar to Steinberg et al. [2023]; Steinberg and Yan [2021], but for OHD would be an interesting future direction. Finally, we did not address non-idealities of lasers, such as non-linear frequency modulation [Amann et al. 2001], phase noise [Vasilyev 2013] or wide-band spectrum (white light interferometry), which could be further studied in the future works.

In spite of these limitations, we noticed that our rendered data matches the real hardware data and our model performs much better than the simple formation models and provides a scalable OHD simulation framework under various scenarios. We anticipate that our OHD simulator can act as a digital twin and empower imaging researchers to investigate OHD systems more effectively, both to improve them and to find new avenues for these systems.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their feedback. We also thank the authors of Hu et al. [2022] for helpful discussions regarding Doppler effects in AMCW cameras. This work is supported by NSF awards 2403122, 2326904, and 1844538.

REFERENCES

- Markus-Christian Amann, Thierry M Bosch, Marc Lescure, Risto A Myllylae, and Marc Rioux. 2001. Laser ranging: a critical review of unusual techniques for distance measurement. *Optical engineering* 40 (2001), 10–19.
- Viktor A Banakh and Christian Werner. 2005. Computer simulation of coherent Doppler lidar measurement of wind velocity and retrieval of turbulent wind statistics. *Optical Engineering* 44, 7 (2005), 071205–071205.

- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38, 4 (July 2019), 39:1–39:22. <https://doi.org/10/gf5jkb>
- Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Rendering Near-Field Speckle Statistics in Scattering Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 39, 6 (Nov. 2020). <https://doi.org/10/kn55>
- Donald E Barrick. 1973. *FM/CW radar signals and digital processing*. US Department of Commerce, National Oceanic and Atmospheric Administration
- Esther Baumann, Jean-Daniel Deschênes, Fabrizio R Giorgetta, William C Swann, Ian Coddington, and Nathan R Newbury. 2014. Speckle phase noise in coherent laser ranging: fundamental precision limitations. *Optics letters* 39, 16 (2014), 4776–4779.
- Steven M Beck, Joseph R Buck, Walter F Buell, Richard P Dickinson, David A Kozlowski, Nicholas J Marechal, and Timothy J Wright. 2005. Synthetic-aperture imaging laser radar: laboratory demonstration and signal processing. *Applied optics* 44, 35 (2005), 7621–7629.
- Behnam Behroozpour, Phillip AM Sandborn, Ming C Wu, and Bernhard E Boser. 2017. Lidar system architectures and circuits. *IEEE Communications Magazine* 55, 10 (2017), 135–142.
- RF Bonner, TR Clem, PD Bowen, and RL Bowman. 1981. Laser-Doppler continuous real-time monitor of pulsatile and mean blood flow in tissue microcirculation. *Scattering techniques applied to supramolecular and nonequilibrium systems* (1981), 685–701.
- R Bonner and R Nossal. 1981. Model for laser Doppler measurements of blood flow in tissue. *Applied optics* 20, 12 (1981), 2097–2107.
- Marie-Lise Chanin, Anne Garnier, Alain Hauchecorne, and Jacques Porteneuve. 1989. A Doppler lidar for measuring winds in the middle atmosphere. *Geophysical research letters* 16, 11 (1989), 1273–1276.
- T Cochran, SB Sheriff, AJM Boulton, JD Ward, and RM Atkins. 1986. Laser Doppler flowmetry: in the assessment of peripheral vascular disorders? A preliminary evaluation. *Clinical Physics and Physiological Measurement* 7, 1 (1986), 31.
- Stephen C Crouch and Krishna Rupavatharam. 2019. Method and system for time separated quadrature detection of doppler effects in optical range measurements. US Patent 10,401,495.
- Herman Z Cummins and Harry L Swinney. 1970. III light beating spectroscopy. In *Progress in optics*. Vol. 8. Elsevier, 133–200.
- Alain Dabas, Pierre H Flamant, and Philippe Salamitou. 1994. Characterization of pulsed coherent Doppler lidar with the speckle effect. *Applied optics* 33, 27 (1994), 6524–6532.
- FFM de de Mul, MH Koelink, ML Kok, PJ Harmsma, Jan Greve, R Graaff, and JG Aarnoudse. 1995. Laser Doppler velocimetry and Monte Carlo simulations on models for blood perfusion in tissue. *Applied optics* 34, 28 (1995), 6595–6611.
- OE DeLange. 1968. Optical heterodyne detection. *IEEE spectrum* 5, 10 (1968), 77–85.
- John A Detre, John S Leigh, Donald S Williams, and Alan P Koretsky. 1992. Perfusion imaging. *Magnetic resonance in medicine* 23, 1 (1992), 37–45.
- DH Dolan. 2009. What does “velocity” interferometry really measure?. In *AIP Conference Proceedings*, Vol. 1195. American Institute of Physics, 589–594.
- Leslie E Drain. 1980. The laser Doppler techniques. *Chichester* (1980).
- Pierre H Flamant, Robert T Menzies, and Michael J Kavaya. 1984. Evidence for speckle effects on pulsed CO 2 lidar signal returns from remote targets. *Applied optics* 23, 9 (1984), 1412–1417.
- Ingemar Fredriksson, Carina Fors, and Johannes Johansson. 2007. Laser doppler flowmetry—a theoretical framework. *Department of Biomedical Engineering, Linköping University* (2007), 6–7.
- Rod Frehlich and Larry Cornman. 2002. Estimating spatial velocity statistics with coherent Doppler lidar. *Journal of Atmospheric and Oceanic Technology* 19, 3 (2002), 355–366.
- Joseph W Goodman. 1975. Statistical properties of laser speckle patterns. In *Laser speckle and related phenomena*. Springer, 9–75.
- Joseph W Goodman. 2007. *Speckle phenomena in optics: theory and applications*. Roberts and Company Publishers.
- Martin Hammer, Dietrich Schweitzer, Bernhard Michel, Eike Thamm, and Achim Kolb. 1998. Single scattering by red blood cells. *Applied Optics* 37, 31 (1998), 7410–7418.
- Martin Hammer, Anna N Yaroslavsky, and Dietrich Schweitzer. 2001. A scattering phase function for blood with physiological haematocrit. *Physics in Medicine & Biology* 46, 3 (2001), N65.
- RM Hardesty, RJ Keeler, MJ Post, and RA Richter. 1981. Characteristics of coherent lidar returns from calibration targets and aerosols. *Applied Optics* 20, 21 (1981), 3763–3769.
- Tatsuo Hariyama, Phillip AM Sandborn, Masahiro Watanabe, and Ming C Wu. 2018. High-accuracy range-sensing system based on FMCW using low-cost VCSEL. *Optics express* 26, 7 (2018), 9285–9297.
- Felix Heide, Wolfgang Heidrich, Matthias Hullin, and Gordon Wetzstein. 2015. Doppler Time-of-Flight Imaging. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (2015). <https://doi.org/10/gfz5k9>
- Kristof Hofrichter, Clemens Linnhoff, Lukas Elster, and Steven Peters. 2024. *FMCW Lidar Simulation with Ray Tracing and Standardized Interfaces*. Technical Report. SAE Technical Paper.
- G Allen Holloway Jr and Dennis W Watkins. 1977. Laser Doppler measurement of cutaneous blood flow. *Journal of Investigative Dermatology* 69, 3 (1977), 306–309.
- Hsiao-Feng Hu, Hsin Hsiu, Ciao-Jyuan Sung, and Chien-Hsing Lee. 2017. Combining laser-Doppler flowmetry measurements with spectral analysis to study different microcirculatory effects in human prediabetic and diabetic subjects. *Lasers in medical science* 32, 2 (2017), 327–334.
- Yunpu Hu, Leo Miyashita, and Masatoshi Ishikawa. 2022. Differential Frequency Heterodyne Time-of-Flight Imaging for Instantaneous Depth and Velocity Estimation. *ACM Transactions on Graphics* 42, 1 (Sept. 2022), 9:1–9:13. <https://doi.org/10/gspq2z>
- RM Huffaker, AV Jelalian, and JAL Thomson. 1970. Laser-Doppler system for detection of aircraft trailing vortices. *Proc. IEEE* 58, 3 (1970), 322–326.
- Koichi Iiyama, Shin-ichiro Matsui, Takao Kobayashi, and Takeo Maruyama. 2011. High-resolution FMCW reflectometry using a single-mode vertical-cavity surface-emitting laser. *IEEE Photonics Technology Letters* 23, 11 (2011), 703–705.
- Wenzel Jakob. 2013. Mitsuba Renderer. <http://www.mitsuba-renderer.org>
- Adrian Jarabo, Julio Marco, Adolfo Munoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A Framework for Transient Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 33, 6 (Nov. 2014), 177:1–177:10. <https://doi.org/10/gfzn8>
- HW Jentink, FFM De Mul, RGAM Hermsen, R Graaff, and Jan Greve. 1990. Monte Carlo simulations of laser Doppler blood flow measurements in tissue. *Applied optics* 29, 16 (1990), 2371–2381.
- Jack L Jewell, JP Harbison, Axel Scherer, Yong-Hee Lee, and LT Florez. 1991. Vertical-cavity surface-emitting lasers: design, growth, fabrication, characterization. *IEEE Journal of Quantum Electronics* 27, 6 (1991), 1332–1346.
- Jorge Jimenez, David Whelan, Veronica Sundstedt, and Diego Gutierrez. 2010. Real-time realistic skin translucency. *IEEE Computer Graphics and Applications* 30, 4 (2010), 32–41.
- Christer J Karlsson, Fredrik ÅA Olsson, Dietmar Letalick, and Michael Harris. 2000. All-fiber multifunction continuous-wave coherent laser radar at 1.55 μm for range, speed, vibration, and wind measurements. *Applied optics* 39, 21 (2000), 3716–3726.
- Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-Domain Path Tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015), 123. <https://doi.org/10/gfzrhv>
- Juhyeon Kim, Wojciech Jarosz, Ioannis Gkioulekas, and Adithya Pediredla. 2023. Doppler Time-of-Flight Rendering. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–18.
- Russell Kliese and AD Rakić. 2012. Spectral broadening caused by dynamic speckle in self-mixing velocimetry sensors. *Optics express* 20, 17 (2012), 18757–18771.
- A Kruger, J Stewart, R Sahityani, E O’riordan, C Thompson, S Adler, R Garrick, P Vallance, and MS Goligorsky. 2006. Laser Doppler flowmetry detection of endothelial dysfunction in end-stage renal disease patients: correlation with cardiovascular risk. *Kidney international* 70, 1 (2006), 157–164.
- TR Lawrence, DJ Wilson, CE Craven, IP Jones, RM Huffaker, and JAL Thomson. 1972. A laser velocimeter for remote wind sensing. *Review of Scientific Instruments* 43, 3 (1972), 512–518.
- You Li and Javier Ibanez-Guzman. 2020. Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems. *IEEE Signal Processing Magazine* 37, 4 (2020), 50–61.
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized Resampled Importance Sampling: Foundations of ReSTIR. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022), 75:1–75:23. <https://doi.org/10/gqjn7b>
- Yang Liu, Shaojie Jiao, and Wojciech Jarosz. 2022. Temporally Sliced Photon Primitives for Time-of-Flight Rendering. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 41, 4 (2022). <https://doi.org/10.1111/cgf.14584>
- Zhengliang Liu, Janet F Barlow, Pak-Wai Chan, Jimmy Chi Hung Fung, Yuguo Li, Chao Ren, Hugo Wai Leung Mak, and Edward Ng. 2019. A review of progress and applications of pulsed Doppler wind LiDARs. *remote sensing* 11, 21 (2019), 2522.
- Richard G. Lyons. 2011. *Understanding Digital Signal Processing* (3 ed.). Prentice Hall, Upper Saddle River, NJ.
- Marco Manzi, Markus Kettunen, Frédo Durand, Matthias Zwicker, and Jaakko Lehtinen. 2016. Temporal Gradient-Domain Path Tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (Dec. 2016), 246:1–246:9. <https://doi.org/10/fc9psw>
- Julio Marco, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2017. Transient Photon Beams. In *Congreso Espanol de Informatica Grafica*. Eurographics Association. <https://doi.org/10/gfzncc>
- Robert T Menzies and R Michael Hardesty. 1989. Coherent Doppler lidar for measurements of wind fields. *Proc. IEEE* 77, 3 (1989), 449–462.
- Shree K. Nayar. 1991. Surface Reflection: Physical and Geometrical Perspectives. *PAMI* 13, 7 (July 1991), 611–634. <https://doi.org/10/b3hrqz>
- Gert E Nilsson, Anneli Jakobsson, and Karin Wardell. 1992. Tissue perfusion monitoring and imaging by coherent light scattering. In *Bioptics: optics in biomedicine and environmental sciences*, Vol. 1524. SPIE, 90–109.

- Rogério Passy, Nicholas Gisin, Jean-Pierre von der Weid, and HH Gilgen. 1994. Experimental and theoretical investigations of coherent OFDR with semiconductor laser sources. *Journal of lightwave technology* 12, 9 (1994), 1622–1630.
- Adithya Pediredla, Yasin Karimi Chalmiani, Matteo Giuseppe Scopelliti, Maysamreza Chamanzar, Srinivasa Narasimhan, and Ioannis Gkioulekas. 2020. Path tracing estimators for refractive radiative transfer. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.
- Adithya Pediredla, Matteo Giuseppe Scopelliti, Srinivasa Narasimhan, Maysamreza Chamanzar, and Ioannis Gkioulekas. 2023. Optimized virtual optical waveguides enhance light throughput in scattering media. *Nature Communications* 14, 1 (2023), 5681.
- Adithya Pediredla, Ashok Veeraraghavan, and Ioannis Gkioulekas. 2019. Ellipsoidal Path Connections for Time-Gated Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38, 4 (July 2019), 38:1–38:12. <https://doi.org/10/gf5jbm>
- Xiaoyi Peng and Jie Shan. 2021. Detection and tracking of pedestrians using doppler lidar. *Remote Sensing* 13, 15 (2021), 2952.
- Alexander Y Piggott. 2020. Understanding the physics of coherent LiDAR. *arXiv preprint arXiv:2011.05313* (2020).
- Mohammad U Piracha, Dat Nguyen, Ibrahim Ozdur, and Peter J Delfyett. 2011. Simultaneous ranging and velocimetry of fast moving targets using oppositely chirped pulses from a mode-locked laser. *Optics express* 19, 12 (2011), 11213–11219.
- Léo Puyo, Michel Paques, and Michael Atlan. 2021. Retinal blood flow reversal quantitatively monitored in out-of-plane vessels with laser Doppler holography. *Scientific Reports* 11, 1 (2021), 17828.
- L Puyo, M Paques, M Fink, J-A Sahel, and Michael Atlan. 2018. In vivo laser Doppler holography of the human retina. *Biomedical optics express* 9, 9 (2018), 4113–4129.
- Léo Puyo, Michel Paques, Mathias Fink, José-Alain Sahel, and Michael Atlan. 2019. Choroidal vasculature imaging with laser Doppler holography. *Biomedical optics express* 10, 2 (2019), 995–1012.
- Marc Rendell, Tom Bergman, Greg O'Donnell, ED Drobný, John Borgos, and Robert F Bonner. 1989. Microvascular blood flow, volume, and velocity measured by laser Doppler techniques in IDDM. *Diabetes* 38, 7 (1989), 819–824.
- CE Riva, S Harino, BL Petrig, and RD Shonat. 1992. Laser Doppler flowmetry in the optic nerve. *Experimental eye research* 55, 3 (1992), 499–506.
- Charles E Riva, Juan E Grunwald, Stephen H Sinclair, and BL Petrig. 1985. Blood velocity and volumetric flow rate in human retinal vessels. *Investigative ophthalmology & visual science* 26, 8 (1985), 1124–1132.
- Philipp Rosenberger, Martin Friedrich Holder, Nicodemo Cianciaruso, Philip Aust, Jonas Franz Tamm-Morschel, Clemens Linnhoff, and Hermann Winner. 2020. Sequential lidar sensor system simulation: a modular approach for simulation-based safety validation of automated driving. *Automotive and Engine Technology* 5, 3 (2020), 187–197.
- Mingjia Shangguan, Jiawei Qiu, Jinlong Yuan, Zhifeng Shu, Lingfeng Zhou, and Haiyun Xia. 2022. Doppler wind lidar from UV to NIR: A review with case study examples. *Frontiers in Remote Sensing* 2 (2022), 787111.
- Albert P Shepherd and P Åke Öberg. 2013. *Laser-Doppler blood flowmetry*. Vol. 107. Springer Science & Business Media.
- Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. 101–110.
- Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Eugene d'Eon, Ling-Qi Yan, and Matt Pharr. 2023. A Generalized Ray Formulation for Wave-Optics Rendering. <https://doi.org/10/10.1145/3472293> arXiv:2303.15762 [cs]
- Shlomi Steinberg and Ling-Qi Yan. 2021. Physical Light-Matter Interaction in Hermite-Gauss Space. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 40, 6 (Dec. 2021). <https://doi.org/10/gsmm8w>
- Shlomi Steinberg and Ling-Qi Yan. 2022. Rendering of Subjective Speckle Formed by Rough Statistical Surfaces. *ACM Trans. Graph.* 41, 1, Article 2 (feb 2022), 23 pages. <https://doi.org/10.1145/3472293>
- Michael D Stern. 1985. Laser Doppler velocimetry in blood and multiply scattering fluids: theory. *Applied optics* 24, 13 (1985), 1968–1986.
- MICHAEL D Stern, DONALD L Lappe, PATRICK D Bowen, JOHN E Chimosky, GA Holloway Jr, HR Keiser, and RL Bowman. 1977. Continuous measurement of tissue blood flow by laser-Doppler spectroscopy. *American Journal of Physiology-Heart and Circulatory Physiology* 232, 4 (1977), H441–H448.
- Frédéric Szczip, Alaa Alkasem, Guillaume Mioche, Valery Shcherbakov, Céline Cornet, Julien Delanoë, Yahya Gour, Olivier Jourdan, Sandra Banson, and Edouard Bray. 2021. McRALI: a Monte Carlo high-spectral-resolution lidar and Doppler radar simulator for three-dimensional cloudy atmosphere remote sensing. *Atmospheric Measurement Techniques* 14, 1 (2021), 199–221.
- F Ticcioni, L Pulvirenti, and N Pierdicca. 2011. Models for scattering from rough surfaces. *Electromagnetic waves* 10 (2011), 203–226.
- Rini Khamimatul Ula, Yusuke Noguchi, and Koichi Iiyama. 2019. Three-dimensional object profiling using highly accurate FMCW optical ranging system. *Journal of Lightwave Technology* 37, 15 (2019), 3826–3833.
- Arseny Vasilyev. 2013. *The optoelectronic swept-frequency laser and its applications in ranging, three-dimensional imaging, and coherent beam combining of chirped-seed amplifiers*. Ph. D. Dissertation. California Institute of Technology.
- Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph. D. Dissertation. Stanford University.
- Henrik Wann Jensen, Stephen R Marschner, Marc Levoy, and Pat Hanrahan. 2023. A practical model for subsurface light transport. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 319–326.
- Tianwen Wei, Haiyun Xia, Jianjun Hu, Chong Wang, Mingjia Shangguan, Lu Wang, Mingjiao Jia, and Xiankang Dou. 2019. Simultaneous wind and rainfall detection by power spectrum analysis using a VAD scanning coherent Doppler lidar. *Optics express* 27, 22 (2019), 31235–31245.
- Tianwen Wei, Haiyun Xia, Bin Yue, Yunbin Wu, and Qi Liu. 2021. Remote sensing of raindrop size distribution using the coherent Doppler lidar. *Optics Express* 29, 11 (2021), 17246–17257.
- Christian Werner. 2005. Doppler wind lidar. In *Lidar: range-resolved optical remote sensing of the atmosphere*. Springer, 325–354.
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–10.
- Jinlong Yuan, Lian Su, Haiyun Xia, Yi Li, Ming Zhang, Guangju Zhen, and Jianyu Li. 2022. Microburst, windshear, gust front, and vortex detection in mega airport using a single coherent Doppler wind lidar. *Remote Sensing* 14, 7 (2022), 1626.
- Jinlong Yuan, Kenan Wu, Tianwen Wei, Lu Wang, Zhifeng Shu, Yuanjian Yang, and Haiyun Xia. 2021. Cloud seeding evidenced by coherent Doppler wind lidar. *Remote Sensing* 13, 19 (2021), 3815.
- Jinlong Yuan, Haiyun Xia, Tianwen Wei, Lu Wang, Bin Yue, and Yunbin Wu. 2020. Identifying cloud, precipitation, windshear, and turbulence by deep analysis of the power spectrum of coherent Doppler wind lidar. *Optics Express* 28, 25 (2020), 37406–37418.
- Fu-Min Zhang, Ya-Ting Li, Hao Pan, Chun-Zhao Shi, and Xing-Hua Qu. 2019. Vibration compensation of the frequency-scanning-interferometry-based absolute ranging system. *Applied Sciences* 9, 1 (2019), 147.