

Rendering of Subjective Speckle Formed by Rough Statistical Surfaces

SHLOMI STEINBERG, University of California, Santa Barbara

LING-QI YAN, University of California, Santa Barbara




Fig. 1. A Stanford Dragon made of chromium rendered under a D65 illuminant. The light source is moderately coherent with a coherence radius of roughly $\approx 30 \mu\text{m}$ on average when incident upon the Dragon's surface. The surface was modelled statistically only, therefore the scattered intensity, I , can be considered as a stochastic process. In order to render this scattered intensity we decompose it, in a physically and mathematically consistent manner, into its ensemble average, $\langle I \rangle$, and a *fluctuating intensity*, \mathfrak{I} : (left) The ensemble average of the process, $\langle I \rangle$, dominates the scattered energy and is the averaged scattered intensity over all possible realizations of the surface. (middle) The fluctuating intensity is a zero-mean process (only positive values were visualised) that gives rise to diffraction patterns—known as *subjective optical speckle*—that depend on the statistical properties of the light, surface and the imaging device. (right) The final intensity is then the superposition of the ensemble averaged lobe and fluctuating field.

Tremendous effort has been extended by the computer graphics community to advance the level of realism of material appearance reproduction by incorporating increasingly more advanced techniques. We are now able to re-enact the complicated interplay between light and microscopic surface features—scratches, bumps and other imperfections—in a visually convincing fashion. However, diffractive patterns arise even when no explicitly defined features are present: Any random surface will act as a diffracting aperture and its statistics heavily influence the statistics of the diffracted wave fields. Nonetheless, the problem of rendering diffraction effects induced by surfaces

Authors' addresses: Shlomi Steinberg, p@shlomisteinberg.com, University of California, Santa Barbara, 2119 Harold Frank Hall, Santa Barbara, California, 93106; Ling-Qi Yan, lingqi@cs.ucsb.edu, University of California, Santa Barbara, 2119 Harold Frank Hall, Santa Barbara, California, 93106.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

XXXX-XXXX/2021/6-ART \$15.00

<https://doi.org/10.1145/1122445.1122456>

that are defined purely statistically remains wholly unexplored. We present a thorough derivation, from core optical principles, of the intensity of the scattered fields that arise when a natural, partially-coherent light source illuminates a random surface. We follow with a probability theory analysis of the statistics of those fields and present our rendering algorithm. All of our derivations are formally proven and verified numerically as well. Our method is the first to render diffraction effects produced by a surface described statistically only, and bridges the theoretical gap between contemporary surface modelling and rendering.

CCS Concepts: • **Computing methodologies** → **Rendering; Computer graphics**; • **Applied computing** → *Physics*.

Additional Key Words and Phrases: diffractions, wave optics, speckle, scatter, iridescence, statistical surfaces, appearance reproduction, coherence, Monte Carlo

ACM Reference Format:

Shlomi Steinberg and Ling-Qi Yan. 2021. Rendering of Subjective Speckle Formed by Rough Statistical Surfaces. 1, 1 (June 2021), 23 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Appearance reproduction remains a challenging, important task in computer graphics. Over the years a significant body of work has been dedicated to modelling the interaction between light and a material’s surface. In particular, rendering of diffraction phenomena—the rainbow patterns that appear on oily surfaces, colour shifts on oxidised metals, and iridescent glints from scratches or rough surfaces—has received significant attention from the computer graphics community, as well as the applied optics community.

However, a disconnect arises between how the scientific and professional communities model surfaces and the way current work addresses rendering surface diffraction. Representing surfaces via statistically-defined, instead of explicitly-defined, models carries significant practical appeal. Such models are used extensively and have been proven to be exceptionally useful [Dorsey et al. 2007]: A cornerstone of modern surface modelling in computer graphics is the *microfacet theory*, which postulates that surfaces are a collection of microfacets and the facet orientation is described statistically. It has been shown that microfacet models effectively reproduce the appearance of some classes of materials [Ngan 2006], and a myriad of extensions to the microfacet theory have been developed, e.g., thin-film interference [Belcour and Barla 2017], multi-layered microfacet materials [Guo et al. 2019; Weidlich and Wilkie 2007] and multiple scattering [Lee et al. 2018]. Additionally, more physically rigorous modern scatter theories have been developed by the applied optics community, and have seen use in rendering (see Section 2). Amongst the most capable of those are the *modified Beckmann-Kirchhoff* model [Harvey et al. 2007] and *generalised Harvey-Shack* model [Krywonos 2006]. These models consider a surface’s power spectral density (that is, the autocorrelation function in Fourier space) to formulate a BRDF. In sharp contrast, most current computer graphics techniques that aim to render surface diffraction effects require knowledge of the explicit geometry: for example, analytic 1D scratch profiles [Velinov et al. 2018; Werner et al. 2017] or high-resolution height fields [Yan et al. 2018]. This data needs to be provided or procedurally generated at the sub-wavelength resolution required to model wave interactions, and can be difficult to measure, compute and store.

We present the first work that considers only the statistical properties of the surface, specifically, the surface power spectral density, to render diffraction patterns. To do so, we derive an optically accurate formalism that relates the statistics of the scattered intensity to the statistics of the scattering surface, taking into account the statistical properties of the incident radiation, i.e. its optical coherence, and the parameters of the employed imaging system (e.g., eye or camera). We then proceed with rendering the scattered intensity by drawing a pattern that conforms to the required statistics. The rendering is done via a Monte Carlo process inspired by Bar et al. [2019] and the rendered pattern corresponds to scatter produced by some realization of the surface statistics. Nonetheless, throughout the entire process we do not draw or consider any explicit geometry. Our motivation for this work is accurate reproduction of surfaces under natural lighting, targeting applications that strive for photo-realistic rendering. Scientific applications are secondary, as those

would typically employ fully coherent (e.g., laser) radiation, which is out of scope of this paper and is left for future work.

Speckle. When rendered using current methods, statistical surfaces appear smooth and unblemished. However, numerical experiments performed by Lanari et al. [2017] show noise-like variations in the intensity of light scattered off rough surfaces that neither the microfacet model nor the modified Beckmann-Kirchhoff model can reproduce. Those intensity variations are *optical speckle* [Goodman 2007], noise-like patterns that arise in practically any coherent imaging modality. When the incident radiation arises from natural light sources, this speckle becomes superposed with the incoherently scattered energy, i.e. the energy that dominates the scatter and shapes the scattered lobe (that is, the BRDF). Following this insight, we mathematically decompose the intensity of the scattered field into its ensemble averaged mean field and a fluctuating field. We show that for many classes of rough surfaces, the fluctuating field is indeed fully developed speckle. While speckle may appear chaotic, it possesses strong statistical properties, and the bulk of our theoretical contributions are the result of a statistical analysis of the scattered fluctuating field.

2 RELATED WORK

Explicit models. There has been a large body of work focusing on the rendering of glints, scratches and fine surface structures, however most of this work has been confined to a scenario where the geometry is known a priori. We summarize this work here.

Werner et al. [2017] present a framework for the rendering of scratches: Under a Fourier optics formalism, the optical response of a surface that contains a multitude of scratches is formulated as a superposition of the optical transfer functions of those scratches. Then, the (spatially varying) BRDF is formulated via distinct base and scratch response terms. This separation conceptually bears some likeness to our decomposition of the scattered intensity into its mean and fluctuating speckle field, however our decomposition is physically and mathematically exact. Additional related work has also focused on rendering surface diffraction effects that result from arbitrary micro-scale structures modelled by a heightmap [Falster et al. 2020; Yan et al. 2018] and real-time rendering of diffractive scratches [Velinov et al. 2018]. Along another line of research, the rendering of non-iridescent glints and scratches, with explicitly defined geometries, has received extensive focus in computer graphics [Chermain et al. 2019; Gamboa et al. 2018; Yan et al. 2014, 2016].

Statistical models. In contrast to work on modelling diffraction effects from explicitly formulated models, work that considers only the statistics of the scattering medium or surface to render interference phenomena is scarce. Stam [1999] presents one of the first works to consider diffraction from a random surface. Jakob et al. [2014] describe a discrete variant to the microfacet model for the purpose of multiscale rendering of non-diffractive glints. To maintain temporal correlation between rendered frames, they generate a deterministic seed. Their approach shares some similarity with our technique: The integration phases (discussed in Subsection 6.1) that we draw serve to ensure our rendered speckle fields maintain correct and accurate (spatial and angular) correlation. Multiscale methods

for rendering granular materials that are composed of large homogeneous collections of randomly oriented grains is discussed by Meng et al. [2015], and generalised to dynamic heterogeneous grain assemblies with spatially varying concentration and size [Müller et al. 2016]. Holzschuch and Pacanowski [2017]; Löw et al. [2012] employ the Harvey-Shack scatter theory to render surface diffractions, however the Harvey-Shack model only considers the ensemble averaged scattered field, while we draw a field that corresponds to some realization of the statistical surface. Raymond et al. [2016] present a framework for multi-scale rendering of scratches defined statistically. The microfacet model is extended with the aim of accounting for statistical flakes by Guo et al. [2018] and a method for the rendering of metal flakes suspended in car paint is presented by Golla and Klein [2018]. A framework for the rendering of iridescence in pearlescent materials is presented by Guillén et al. [2020].

Optical speckle. Of particular interest is the work by Bar et al. [2019, 2020], who introduce a Monte Carlo framework for drawing fully developed speckle patterns that arise on scatter by participating media under idealised conditions. Those conditions can be summarised as a perfectly coherent radiation source and a scattering medium. In contrast, we deal with scattering off a random surface illuminated by a natural light source, a scenario where none of those assumptions apply. Therefore, while our rendering algorithm is based on the Monte Carlo framework that was presented by them, our presented theoretical formulations are entirely different (Section 4). Furthermore, we extend the algorithm to handle the different requirements that arise in our framework, and we formally prove its correctness under general conditions (Section 6).

Other work that aims to render speckle considers the simplified case where we are bereft of any (explicit or statistical) knowledge of the scattering medium and assume perfect optical coherence. Synthesis of arbitrary, realistic-looking speckle patterns is discussed by Bergmann et al. [2016], and simple general algorithms are introduced by Duncan and Kirkpatrick [2008].

Optical speckle has also been used for non-light-of-sight imaging [Smith et al. 2018] and motion tracking [Smith et al. 2017].

Outside computer graphics, speckle phenomena have given rise to a massive body of research and applications, e.g., methods in speckle reduction [Dainty 2013], extraction of surface roughness information via speckle interferometry [Dhanasekar and Ramamoorthy 2008]. Speckle has also been used to carry or store information [He et al. 2003], perform sensing, e.g., to detect malaria-infected blood cells [Cojoc et al. 2012], and to deduce information in very short-exposure photography as used in astronomy [Scott et al. 2018].

Partial optical coherence. Also relevant are formalisms that aim to handle partially coherent light sources in a physically consistent manner. A common approach in computer graphics is to model the spatial coherence as a Gaussian footprint centred at the sampled point [Dhillon et al. 2014; Toisoul et al. 2018; Werner et al. 2017]. Similarly, a surface patch of interest can be partitioned into “coherence kernels”, modelled as Gaussian windows [Yan et al. 2018], or simply as rect functions [Levin et al. 2013]. Partitioning the integration area into kernels only remains physically sound when the coherence area is very small compared to the surface patch in question, however it simplifies the formulation as the mutual coherence between pairs

of arbitrary points does not need to be considered. Nonetheless, a two-point characterization of the mutual coherence provides much richer information about the diffracted field’s properties [Mandel and Wolf 1995]. Our derivations consider mutual intensities with respect to their mutual spatial and temporal coherence. Furthermore, we avoid making any assumptions regarding the mutual coherence function during our theoretical formulations.

Other related work. We list here additional related work in the realm of computer graphics that model different wave-interference related phenomena. Belcour and Barla [2017] present an analytic formulation for the spectral integration of phase shifts produced by reflections off a base layer covered by a dielectric thin-film, and an extension capable of handling rough surfaces is developed by Kneiphof et al. [2019]. Steinberg [2019] present another spectral integration approach under the context of optical anisotropy, and a framework for real-time rendering of surface diffraction from measured data was presented by Toisoul and Ghosh [2017].

As part of our statistical analysis, we also present an extension to Isserlis’ theorem [Isserlis 1918] for circularly-symmetric complex Gaussian random variables. Extensions for other distributions have been developed: For mixed-Gaussians [Michalowicz et al. 2009] and Gaussian matrix mixtures [Grigelionis 2009].

Table 1. List of symbols and notation (location of definition on the right)

SYMBOLS	
$\langle \mathbf{a}, \mathbf{b} \rangle$	Inner product between \mathbf{a} and \mathbf{b}
$\delta(x)$	Dirac delta
δ_{xy}	Kronecker delta
\mathcal{F}	Fourier transform operator
\mathbf{k}	Wavevector
λ	Wavelength
$\text{Re}\{\cdot\}$	Real value operator
$\text{Im}\{\cdot\}$	Imaginary value operator
$\arg(\cdot)$	Complex argument
$\langle \cdot \rangle$	Ensemble averaging operator
$E[\cdot]$	Expected value of a random variable
$\sigma[\cdot]$	Standard deviation of a random variable
$\text{cov}[\cdot, \cdot]$	Covariance
$\text{Pr}\{X = a\}$	Probability of random variable X taking value a
\star	Complex conjugate
\dagger	Conjugate transpose
$*$	Convolution operator
Γ	Gamma function
K_ν	Modified Bessel function of second kind
P_2	Surface power spectral density (PSD)
C_S	Surface autocorrelation function
σ_{rel}	Surface relative roughness
l_{cor}	Surface correlation length
Γ_c	Mutual coherence function in the object (surface) plane

3 BACKGROUND: SURFACE STATISTICS AND THE POWER SPECTRUM

Surface roughness is a key surface attribute that affects the optical properties and wave scattering characteristics of a surface. Nonetheless, despite being an intuitive concept, optical roughness of a surface is difficult to capture analytically, and different disciplines and applications employ different formalisms. The measure of roughness commonly used in computer graphics—the variance of the slope or height profile—is inadequate for appropriately describing the optical behaviour, especially under the context of “wave optics”. Instead we describe the surface micro-scale features statistically using a more flexible construct that is employed in contemporary scatter theories: the power spectrum of the surface’s spatial frequencies.

Let $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ denote a height field function describing a deterministic surface. The Fourier transform of the height field decomposes the signal into the contributing frequencies and is defined as

$$\hat{h}(\mathbf{f}) = \mathcal{F}\{h\}(\mathbf{f}) = \int_{\mathbb{R}^2} h(\mathbf{x}) e^{-i2\pi\langle \mathbf{f}, \mathbf{x} \rangle} d\mathbf{x}, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product. When h and its transform pair \hat{h} are absolutely integrable, the surface profile can be recovered via the inverse transform:

$$h(\mathbf{x}) = \mathcal{F}^{-1}\{\hat{h}\}(\mathbf{x}) = \int_{\mathbb{R}^2} \hat{h}(\mathbf{f}) e^{i2\pi\langle \mathbf{x}, \mathbf{f} \rangle} d\mathbf{f}, \quad (2)$$

with $\mathbf{x} \in \mathbb{R}^2$ being a point and $\mathbf{f} \in \mathbb{R}^2$ representing the spatial frequency.

The key quantity in analysing surface characteristics is the surface *power spectral density* (PSD), defined as follows:

$$P_2(\mathbf{f}) = \lim_{d \rightarrow +\infty} \frac{1}{d} \left| \int_{\|\mathbf{x}\| < d} h(\mathbf{x}) e^{-i2\pi\langle \mathbf{f}, \mathbf{x} \rangle} d\mathbf{x} \right|^2, \quad (3)$$

The contributions from the spatial frequencies composing a surface are responsible for the surface’s scattering behaviour [Bass et al. 2010], and this information is fully described by the PSD. For this reason, the surface PSD plays a crucial role in predicting average scattered energy, i.e. the BRDF, from surface characteristics [Church et al. 1990; Harvey and Pfisterer 2016]. Furthermore, using PSDs to formulate surface statistics is practical as a PSD can be directly measured from a physical surface [Siewert et al. 2008]. However, the phase information is lost in the process of computing (or measuring) the PSD, therefore the PSD does not define an explicit surface, but instead describes the surface statistically.

Useful statistical properties can be inferred from the PSD. The *surface autocorrelation function* C_S is the spectral decomposition of the PSD (a consequence of the *Wiener–Khinchin theorem*). Then, assuming its Fourier transform exists, the following holds:

$$C_S(\mathbf{x}) = \mathcal{F}^{-1}\{P_2\}(\mathbf{x}). \quad (4)$$

In this paper we use the *K-correlation model* to describe the surface and its autocorrelation function. See Appendix A for details.

The surface *relative roughness*—the standard deviation of the surface height function h —is the 0th-moment of the PSD:

$$\sigma_{\text{rel}}^2 = \int_{\|\mathbf{f}\| < \frac{1}{\lambda}} P_2(\mathbf{f}) d\mathbf{f}, \quad (5)$$

where the integration is bandwidth-limited as frequencies greater than λ^{-1} produce evanescent waves and do not contribute to surface scatter [Harvey et al. 2012]. Finally, another related statistic is the *correlation length*—the spatial distance required to decrease the surface autocorrelation by e^{-1} —which can be extracted from the PSD as follows:

$$l_{\text{cor}} = \sigma_{\text{rel}}^{-1} \int_{\|\mathbf{f}\| < \frac{1}{\lambda}} P_2(\mathbf{f})^2 d\mathbf{f}. \quad (6)$$

Drawing surfaces from the PSD. While the source profile cannot be recovered from the PSD, drawing explicit, statistically indistinguishable surfaces from the PSD is nonetheless useful. The PSD is a real function with all the phase information discarded, therefore in order to synthesise conforming surfaces we need to introduce a phase for each frequency. Being Fourier transform pairs, the PSD and the correlation function, C_S , convey essentially the same information and contain no insight into the distribution of the phases or the height, h , of the surface. Devoid of any such explicit formulation, it is reasonable to assume that the phases are independently, identically and uniformly distributed:

$$h(\mathbf{x}) = \mathcal{F}^{-1}\left\{e^{i2\pi\langle \mathbf{x}, \mathbf{f} \rangle} \sqrt{P_2(\mathbf{f})}\right\}(\mathbf{x}), \quad (7)$$

where $\mathbf{r}(\mathbf{f}) \sim \mathcal{U}[0,1)$ is an anti-symmetric function of random phases.

Surface scatter theory: the generalised Harvey-Shack model. The key quantity of interest in formulating scatter characteristics is the scattered intensity observed at the far-field region [Stover 2012], expressed by a bidirectional reflectance distribution function (BRDF), i.e. the power ratio between the outgoing radiance and incident irradiance. As discussed briefly in the previous subsection, modern surface scatter theories attempt to relate the scattering characteristics of a surface—the BRDF—to the surface’s statistics described by the PSD. While scattering off rough surfaces generally gives rise to high-frequency details and fluctuations in the scattered intensity (as will be discussed in Section 4), these fine-scale details are lost in the BRDF as the scattered energy is averaged over all possible statistically identical realizations of the surface [Bass et al. 2010]. We will take advantage of this fact later in our discussion.

We now briefly outline the generalised Harvey-Shack scatter BRDF. The Harvey-Shack scatter theory is a linear system formulation of the scatter behaviour, which came into being by the empirical observation that scatter is shift-invariant in direction cosine space [Krywonos 2006]. Some noteworthy simplifying assumptions are made: First, the optical principles in effect are assumed to conform to the Rayleigh hypothesis and the Kirchhoff approximation; likewise, a few assumptions are made with regards to the surface statistics, namely the Gaussianity of the height function, stationarity and ergodicity [Harvey and Pfisterer 2016]. Multiple and subsurface scattering effects are also ignored. Furthermore, note that this is a

scalar diffraction theory that fails to properly account for the direction of polarization of the incident field. Nonetheless, experimental validations of the theory show good results even at far-from-normal incident angle and with rough surfaces that violate the classical Rayleigh-Rice smooth surface approximation [Choi and Harvey 2013], and the Harvey-Shack surface scatter theory has also been used effectively in rendering [Holzschuch and Pacanowski 2017; Löw et al. 2012; Yan et al. 2018]. See Krywonos [2006] for a comprehensive overview of some of the current scatter theories and additional information.

The generalised Harvey-Shack BRDF is defined as follows: Given the surface PSD P_2 , a wavelength λ , the normal \mathbf{n} of the mean of the scattering surface, direction of incident irradiance \mathbf{i} and outgoing direction of scatter \mathbf{o} , the BRDF becomes

$$\text{BRDF}_{\text{gHS}}(\mathbf{o}, \mathbf{i}) = sF \cdot \delta(\text{refl}(\mathbf{i}) - \mathbf{o}) + (1 - s)Q \cdot P_2(\mathbf{f}), \quad (8)$$

where $s = e^{-(2\pi\sigma_{\text{rel}}\lambda^{-1}(\mathbf{i}+\mathbf{o}, \mathbf{n}))^2}$ is the energy fraction that remains in the specularly reflected beam [Harvey et al. 2012]. As we will discuss later, for the rough surfaces that we are interested in and non-grazing angles we can assume that $s \approx 0$, i.e. the specular term is negligible. F is the well-known specular Fresnel power term, δ is the Dirac delta and $\text{refl}(\mathbf{v})$ is the reflection operator. While the first term is the specular reflectance, the second term of the BRDF is the scattered energy, with \mathbf{f} being the spatial frequency of the diffraction lobe. \mathbf{f} can be expressed succinctly in vector form as the projection [Holzschuch and Pacanowski 2017]:

$$\mathbf{f} = \frac{1}{\lambda} (\mathbf{I} - \mathbf{nn}^T) (\text{refl}(\mathbf{i}) - \mathbf{o}) = -\frac{1}{\lambda} (\mathbf{I} - \mathbf{nn}^T) (\mathbf{i} + \mathbf{o}), \quad (9)$$

where \mathbf{I} denotes the identity matrix. Finally, the term Q is the Fresnel power term that arises in the Rayleigh-Rice theory [Krywonos 2006]. For completeness, we provide the expressions for Q in Appendix D.

Assumptions. We make a few assumptions about the surface model and we summarize them here: In similar fashion to the generalised Harvey-Shack theory, we assume that the surface statistics describe a stationary and weakly ergodic (different regions of the surface are weakly correlated) stochastic process. Further, we assume that the height profile, h , of the random scattering surface is point-wise Gaussian as well as any two points, $h(\mathbf{x}_1)$, $h(\mathbf{x}_2)$, are jointly Gaussian. Note that this also follows directly from our assumption that the phases of $\mathcal{F}\{h\}$ are independent (Equation (7)), due to the functional extension of the central limit theorem [Billingsley 1995]. We also explicitly formalise the roughness of the surfaces that we limit our discussion to, via the order-of-magnitude relation

$$\sigma_{\text{rel}} \gg \frac{1}{2}\lambda. \quad (10)$$

Given a mean wavelength of $\sim 0.50 \mu\text{m}$ this implies that the standard deviation of the surface height profile is greater than $0.25 \mu\text{m}$. This implies that the random scattering surface is rough with respect to the light's wavelength. That is, very smooth surfaces and artificial surfaces with a highly periodic structure, e.g., surfaces used for diffraction gratings, are ignored.

Note that those are very reasonable assumptions: Consider the BRDF listed in Equation (8), for smooth surfaces or grazing angles

the term s becomes significant, meaning a greater fraction of energy is lost to the specular beam causing the speckle pattern to diminish.

4 FUNDAMENTALS: SPECKLE ON SCATTERING FROM RANDOM ROUGH SURFACE

In this section we formulate our theoretical optical foundation. The study of light scatter off rough surfaces is concerned with understanding the physical relationship between the surface detail and the amplitudes or intensities of the scattered waves. In the case of a random surface described statistically, the interest is in understanding the statistics of the scattered light. Those statistics model the spatial distribution of the intensities that form due to scattering off a rough surface and give rise to a granular pattern known as *speckle*, which is a widely-studied optical phenomenon [Goodman 2007].

4.1 Far-Field Speckle Statistics

We first briefly review the basic statistical properties of speckle patterns, formed due to scattering from a random rough surface and observed at the Fraunhofer region (far field). In the far-field region, we are discussing only fully formed propagating electromagnetic radiation fields (fields that decay in amplitude as $\frac{1}{r}$, with r being distance from the source). See Dainty [2013]; Goodman [2007] for a more complete discussion of optical speckle.

For this subsection we assume coherent light, that is monochromatic light that consists only of in-phase contributions (waves), e.g., light generated by a (single-mode) laser. Let u be the phasor of a scattered electric field, observed far from the scattering surface. u is then a superposition of a multitude of reflected contributions—called *elementary phasors*—each arising from reflection off a distinct microscopic element of the scattering surface:

$$u = \sum_k a_k e^{-i\phi_k}, \quad (11)$$

with $a_k, \phi_k \in \mathbb{R}$. A few practical assumptions are typically made with regards to the phasor u when discussing speckle [Goodman 2007]: First, it is assumed that the phases ϕ_k of the elementary phasors contributing to a scattered field are uniformly distributed over $[-\pi, \pi)$; and second, the contributions contain many uncorrelated or weakly correlated phasors. Both these assumptions follows directly from our roughness and weakly ergodic surface statistics assumptions (discussed in Section 3). Speckle fields conforming to those assumptions are described by Goodman [2007] as “fully developed”. Under those assumptions, the amplitudes and phases in the sum in Equation (11) can be considered as being mutually independently distributed and thus model a recurrent random-walk in the complex plane [Dainty 2013]. Therefore u averaged over statistically-identical realizations of the random surface has negligible value:

$$\langle u \rangle \simeq 0, \quad (12)$$

where the brackets $\langle \cdot \rangle$ denote *ensemble averaging* with respect to the stochastic process of scatter off a random surface. The ensemble average is the average over all surfaces that conform to the given surface statistics, i.e. have the same PSD, P_2 (Equation (3)). For a weakly ergodic stochastic process, this is equivalent to averaging




Fig. 2. Speckle angular correlation: The Dragon at the top left is made of iron and is illuminated by a D50 illuminant that slowly rotates around the Dragon at an angular velocity of $0.10^\circ \text{ s}^{-1}$. Close-ups on the area marked by an orange square are shown on the right. The close-ups (displayed with artificially increased brightness and contrast for visualization purposes) were rendered at very high resolution of $50 \mu\text{m} \times 50 \mu\text{m}$ per pixel. As the observed intensity at each pixel is composed of many unresolved speckle, the effect is subtle. Nonetheless, as the light rotates, the speckle pattern shifts and changes until it dissolves into a different pattern and visible examples of this correlation are circled. The correlation of the centre pixel in each marked circle as a function of time are plotted on the bottom left. We expect the (long range) correlation of speckle to follow a negative exponential correlation [Akkermans 2007], therefore we have fitted a negative exponential to the red and blue plots (dashed black line). The fitting suggests that small errors cause overestimation by our method (especially visible in the green plot), however the exponential drop in correlation is well captured.

over the entirety of one such surface:

$$\langle u \rangle = \int d^2\mathbf{x} u(h(\mathbf{x})), \quad (13)$$

where h is any drawn surface height field that conforms to the given surface PSD, and $u(h(\mathbf{x}))$ is the scattered electric field that is produced due to scattering from the surface at height $h(\mathbf{x})$.

The *average intensity* of the field described by u can be written as the ensemble average

$$\langle I \rangle = \langle uu^\star \rangle, \quad (14)$$

where the superscript \star denotes the complex conjugate.

We turn our attention to the second-order statistics of speckle observed at two different directions and far from the surface. Let u_1 and u_2 be a couple of phasors. Then, the second-order moment of the phasors that form the speckle pattern is known as the *mutual intensity* [Mandel and Wolf 1995] and can be expressed as:

$$C_2(u_1, u_2) = \langle u_1 u_2^\star \rangle. \quad (15)$$

The mutual intensity describes not only the spatial correlation of speckle viewed by a single observer, but also the correlation between speckle observed from different directions. That is, if the observer's position, light source's position or surface characteristics were to change slowly over time the speckle pattern would shift and morph as well up until it bears no relation to the original pattern. This correlation of speckle is known as the *memory effect* [Feng et al. 1988] (see Figure 2). For our context, of most interest is the correlation of the formed speckle pattern between multiple successive

rendered frames with a slowly moving light source. We refer to this phenomenon as the *angular correlation* of speckle.

Our assumptions imply that the spatial correlation of the rough surface that contributes to the formation of speckle is low and therefore the set of contributing elementary phasors contains many weakly-correlated phasors. Thus, assuming favourable mixing conditions of the weakly-dependent random phasors [Billingsley 1995], by the central limit theorem, as the number of weakly-correlated phasors increases, u asymptotically tends towards a circularly-symmetric complex Gaussian distribution. Therefore, speckle field obey multivariate Gaussian statistics and we do not consider higher-order moments. Treating u as a circularly-symmetric complex Gaussian random variable, the intensity uu^\star becomes a negative exponential random variable [Goodman 2007].

4.2 Subjective Speckle from Partially Coherent Sources

The theory reviewed in Subsection 4.1 considers idealised conditions where the light source is a monochromatic point light source, ergo the incident radiation is assumed to be perfectly spatially and temporally coherent. Furthermore the scattered fields are observed at the far-field region. However, both those assumptions do not hold in practical conditions: Physical light sources have positive extent and spectral width and therefore a limited coherence size and length. In addition, an imaging system will be used to measure the intensities of the speckle pattern at discrete positions and times. This results in a speckle pattern known as a *subjective speckle pattern*. The term “subjective” refers to the fact that the imaging system's

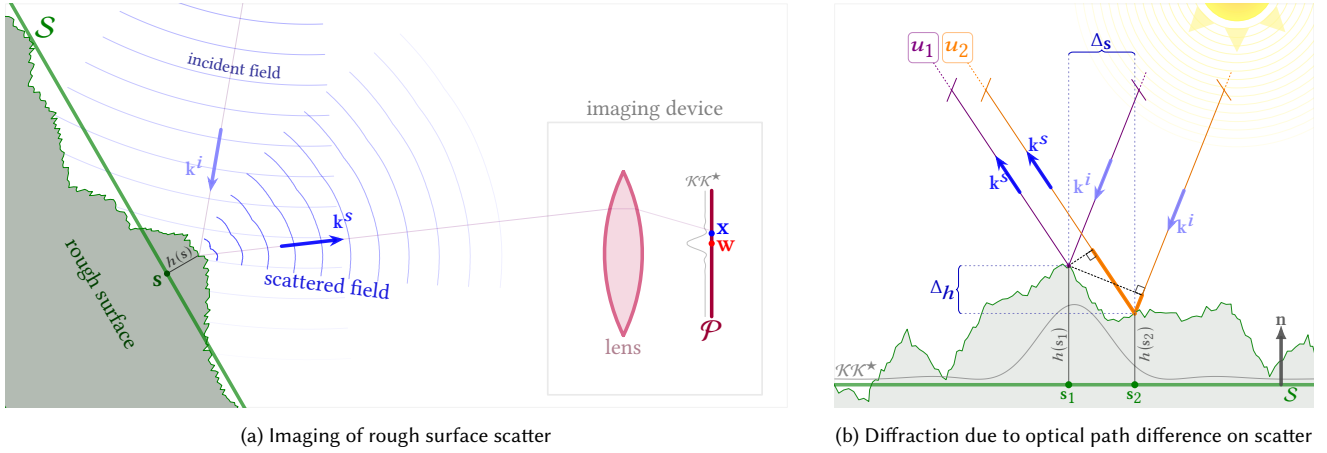


Fig. 3. (a) Depiction of a focused optical system used to image scatter off a rough surface. The surface plane \mathcal{S} (green thick line) is defined as the mean of the random surface and an explicit realization (green surface) is described by $h(s) : \mathcal{S} \rightarrow \mathbb{R}$, the height deviations from the mean. A light source (not shown) illuminates the rough surface giving rise to an incident field which scatters off the surface. The light fields are depicted in blue and some incident and scattered wavevectors—directions of electromagnetic flux—are illustrated as \mathbf{k}^i (light blue arrow), and \mathbf{k}^s (blue arrow), respectively. Imaging is then performed via sensing of the incident flux by a discrete set of image sensors located on the image plane, \mathcal{P} (purple thick line). Because the impulse response of some imaging sensor $\mathbf{w} \in \mathcal{P}$ (red point), described by the *impulse response function* \mathcal{K} , has positive extend, all image points within the impulse response power (depicted in grey, left side of the image plane) contribute to the intensity sensed by that sensor. (b) A far-field light source (top right) gives rise to a couple of incident and scattered wavefronts, illustrated in violet and orange, and the scattered fields result in the phasors u_1, u_2 as they are observed by a distant imaging device (not shown). Both scatters fall within the impulse response of a single image sensor (illustrated in grey above the surface plane), thus their superposition results in wave interference effects when imaged by the imaging system. The optical path difference between the phasors (depicted in thick orange highlight) can be computed geometrically by taking into account the surface and height differences, Δ_s and Δ_h respectively, between the scatter points.

properties, e.g., aperture size and resolution, influence the imaged speckle pattern.

In this subsection we discuss speckle generated by a light source with limited optical coherence and observed through an imaging system. For simplicity, we begin by assuming a quasi-monochromatic light source. Quasi-monochromatic radiation is defined as light with a low temporal frequency spread, and therefore can be assumed to have very good temporal coherence. Polychromatic radiation will be discussed at the end of this subsection.

Imaging and the scale-normalised coordinate system. Before we proceed with our discussion of subjective (imaged) speckle, we briefly describe a general optical imaging system and the geometry we consider (see Figure 3). Let \mathcal{S} be the *surface plane*, which is defined to be the mean of the random rough surface. Furthermore we denote by \mathcal{P} the image plane and let $\{\mathbf{w}_1, \mathbf{w}_2, \dots\} \subset \mathcal{P}$ be a finite set of imaging sensors, referred to as the *image elements* or just pixels henceforth, located on the image plane. The optical properties of the (possibly focused) imaging system are defined via the *impulse response function*, denoted as \mathcal{K} , and we refer to $\mathcal{K}\mathcal{K}^*$ as the *impulse response power*. The impulse response function is a complex-valued function that describes the aperture size and shape as well as imperfections and aberrations induced by the optical system.

For convenience, we will at times abuse notation and refer to points on the image plane interchangeably with points on the surface (in similar fashion to Born and Wolf [1999]). For example, the

notations $u(s)$ and $u(\mathbf{x})$ refer to the phasor scattered at point $s \in \mathcal{S}$ and observed at point $\mathbf{x} \in \mathcal{P}$, and are equivalent.

Transmission of mutual intensity. Let $u(\mathbf{x})$ be the phasor giving rise to a scattered electric field observed at a point $\mathbf{x} \in \mathcal{P}$ on the image plane, and denote by $I(\mathbf{w})$ the intensity measured at an image element centred at coordinate $\mathbf{w} \in \mathcal{P}$ of the formed image. Assuming a stationary process, the quantity $I(\mathbf{w})$ can be regarded as a partially coherent sum, weighted by the impulse response function, of all speckle patterns produced by each point of the scattering surface. That is, the intensity I can be written as an operation that consists of convolving u with the complex impulse response function of the imaging system with respect to the mutual coherence between the phasors [Mandel and Wolf 1995]:

$$I(\mathbf{w}) = \iint_{\mathcal{P}} \iint_{\mathcal{P}} \Gamma_c(\xi, \zeta) u(\xi) u(\zeta)^* \cdot \mathcal{K}(\mathbf{w} - \xi) \mathcal{K}(\mathbf{w} - \zeta)^* d\xi d\zeta, \quad (16)$$

where $\Gamma_c \in \mathbb{C}$ is the *mutual coherence function* [Born and Wolf 1999], which describes the cross-correlation between wave ensembles on the surface plane, and it holds that $\Gamma_c(\mathbf{x}_1, \mathbf{x}_2) = \Gamma_c(\mathbf{x}_2, \mathbf{x}_1)^*$. The mutual coherence function will be discussed in more detail in Subsection 4.3. Equation (16) describes the transmission of the mutual intensity through an optical system under partial coherence [Born and Wolf 1999]. Note that while the integrands are complex-valued quantities, due to the integration symmetry the intensity $I(\mathbf{w})$ is real-valued, as expected. Looking at the limit case where the scattered radiation is perfectly incoherent, viz. $\Gamma_c(\xi, \zeta) \equiv \delta^2(\xi - \zeta)$,

with δ^2 being the two-dimensional Dirac delta function, we can observe that the intensity reduces to

$$I_{\text{incoherent}}(\mathbf{w}) = [(uu^*) * (\mathcal{K}\mathcal{K}^*)](\mathbf{w}),$$

with $*$ denoting convolution. Therefore in the incoherent case the intensity is simply the convolution of the intensity of the contributions with the impulse response power. At the other extreme, when $\Gamma_c = 1$ the intensity becomes

$$I_{\text{coherent}}(\mathbf{w}) = [(u * \mathcal{K})(\mathbf{w})][(u * \mathcal{K})(\mathbf{w})]^*,$$

that is, the intensity is the intensity of the coherent convolution of the contributions with the impulse response.

Polychromatic light. We consider the correlation of speckle intensity formed by a polychromatic light source with limited spatial and temporal coherence. Under polychromatic radiation a scattered electric field u becomes wavelength-dependent and the observed intensity becomes

$$I(\mathbf{w}) = \int_{\Lambda} \int_{\mathcal{P}} \int_{\mathcal{P}} S(\lambda) \Gamma_c(\xi, \zeta, \tau) u(\xi, \lambda) u(\zeta, \lambda)^* \cdot \mathcal{K}(\mathbf{w} - \xi) \mathcal{K}(\mathbf{w} - \zeta)^* d\xi d\zeta d\lambda, \quad (17)$$

where Λ is the spectrum and $S(\lambda)$ is the intensity of each spectral line. The integration over the spectral contributions is justified when the polychromatic light is observed over a period long compared to the temporal coherence of the light. The mutual coherence functions is now replaced with its temporal cross-correlation equivalent, $\Gamma_c(\xi, \zeta, \tau)$, i.e. is the time-averaged mutual coherence function [Mandel and Wolf 1995] and with τ being the time difference between the observations of the two phasors, $u(\xi, \lambda)$ and $u(\zeta, \lambda)$. Spatial and temporal coherence will be discussed further in Subsection 4.3.

Statistics of partially-coherent speckle. While for coherent contributions it is the phasors that are linearly summed up, incoherent contributions are linear with respect to their intensities. Our discussion is centred on surfaces lit by a weakly-coherent incident radiation (natural light sources) and imaged by an imaging device. Under most typical scenarios the resolution of the imaging device will be lower than the light's spatial coherence on the surface. In such a scenario, the intensity of a speckle pattern, imaged by a single image element, is actually a superposition of multiple independent speckle patterns. Such independent speckle patterns arise in a couple of ways: first, incoherent contributions from different surface regions, due to the light's limited spatial coherence; and, second, contributions from uncorrelated spectral components (which, for a high-bandwidth natural light source, can be many). Because each such contribution is a fully-coherent, monochromatic speckle pattern, each contributing intensity follows negative exponential statistics, and thus the superposition of those is gamma distributed. The order parameter of such a gamma distribution will typically be large, and therefore the distribution is very-well approximated by a Gaussian. See Appendix B where we formalise this argument.

4.3 The Mutual Coherence Function Γ_c

In this paper we are satisfied with approximating the mutual coherence based on the spatial and temporal distance differences between the observed electric fields u and do not take a more rigorous approach. Therefore, for simplicity, we neglect here an accurate explicit expression for $\Gamma_c(\xi, \zeta, \lambda)$ and the reader is referred to Mandel and Wolf [1995] for a far more in-depth discussion about optical coherence.

Let $u(\xi)$ and $u(\zeta)$ be the observed phasors at the points ξ and ζ , respectively. The mutual coherence function Γ_c characterizes the correlation between electromagnetic vibrations at two space-time points, thus both the spatial distance $\|\xi - \zeta\|$ as well as the time difference between the observations play a role. However, instead of time difference we consider the (physically equivalent) difference in the spatial distance travelled by the light from the source to the two points of observations, denoted Δ_l . The optical coherence properties of a light source are then related via the light's *coherence radius* and *coherence length* [Mandel and Wolf 1995], given by the order of magnitude relations:

$$\epsilon_r \approx \frac{\bar{\lambda}}{\sqrt{\pi\sqrt{\Delta\alpha}}} \quad \epsilon_l \approx \frac{\bar{\lambda}^2}{\Delta\lambda}, \quad (18)$$

where $\Delta\alpha$ is the solid angle subtended by the light source, $\bar{\lambda}$ is the mean wavelength and $\Delta\lambda$ is the effective wavelength range. Clearly, $\epsilon_r \xrightarrow{\Delta\alpha \rightarrow 0} \infty$, that is radiation from a point light is perfectly spatially coherent, and likewise a monochromatic light source is perfectly temporally coherent ($\epsilon_l = \infty$). While not a very physically rigorous approach, those quantities are directly related to experimental data, are easy to understand and simplify the importance sampling of the mutual coherence Γ_c (will be discussed in Subsection 6.1), as it can now be approximated simply by a product of Gaussians:

$$\Gamma_c(\xi, \zeta, \lambda) \approx \frac{1}{2\pi\epsilon_r^2} e^{-\frac{\|\xi-\zeta\|^2}{2\epsilon_r^2}} e^{-\frac{\Delta_l^2}{2\epsilon_l^2}}, \quad (19)$$

with \mathbf{l} being the direction to the light source. At non-normal incidence anisotropy arises in the coherence area, however for simplicity this was ignored. In our implementation we use Equation (19) with a preset source-dependent coherence length, e.g., $\epsilon_l \sim 0.50 \mu\text{m} - 10 \mu\text{m}$ for a natural unfiltered white source, and we assume that ϵ_r remains constant over a single image element.

5 THE FLUCTUATING INTENSITY

Further analysis can be performed in order to obtain more tractable approximate expressions for Equation (16) and its covariance under various simplifying assumptions [Dainty 2013; Goodman 2007]. We take a different approach, and our primary technical contributions are developed in this section. We decompose the intensity I into a pair of intensities as follows:

$$I = \langle I \rangle + \mathfrak{I}, \quad (20)$$

where the ensemble average is over all possible realizations of the random scattering surface. By taking the ensemble averages on both sides of Equation (20) we immediately observe that $\langle \mathfrak{I} \rangle \equiv 0$, therefore \mathfrak{I} , referred to as the *fluctuating intensity*, represents the deviation from the mean of the imaged speckle pattern. $\langle I \rangle$ is then

simply the average (with respect to the random surface) reflected energy and disregards effects that arise from a distinct realization of the surface. We proceed by examining the statistics of \mathfrak{I} .

Consider the second-order moment of \mathfrak{I} when speckle is observed at two image elements centred at $\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{P}$. As $\langle \mathfrak{I} \rangle \equiv 0$ the autocovariance becomes

$$\begin{aligned} \tilde{C}(\mathbf{w}_1, \mathbf{w}_2) &= \langle \mathfrak{I}(\mathbf{w}_1) \mathfrak{I}(\mathbf{w}_2) \rangle - \langle \mathfrak{I}(\mathbf{w}_1) \rangle \langle \mathfrak{I}(\mathbf{w}_2) \rangle = \\ &= \langle I(\mathbf{w}_1) \rangle \langle I(\mathbf{w}_2) \rangle + \langle I(\mathbf{w}_1) I(\mathbf{w}_2)^* \rangle, \end{aligned} \quad (21)$$

and by plugging in Equation (16) we get the following expression:

$$\begin{aligned} \tilde{C}(\mathbf{w}_1, \mathbf{w}_2) &= \langle I(\mathbf{w}_1) \rangle \langle I(\mathbf{w}_2) \rangle + \\ &+ \int_{\mathcal{P}} \int_{\mathcal{P}} \int_{\mathcal{P}} \int_{\mathcal{P}} \Gamma_c(\xi_1, \zeta_1) \Gamma_c(\xi_2, \zeta_2)^* C_4(\xi_1, \zeta_1, \xi_2, \zeta_2) \cdot \\ &\cdot \mathcal{K}_2(\xi_1, \zeta_1, \mathbf{w}_1) \mathcal{K}_2(\xi_2, \zeta_2, \mathbf{w}_2)^* d\xi_1 d\zeta_1 d\xi_2 d\zeta_2, \end{aligned} \quad (22)$$

where the integration is over tuples of four points on the image plane and C_4 is the fourth-order moment of the observed fields

$$C_4(\xi_1, \zeta_1, \xi_2, \zeta_2) = \langle u(\xi_1) u(\zeta_1)^* (u(\xi_2) u(\zeta_2)^*)^* \rangle. \quad (23)$$

The fourth-order moment will be discussed in greater detail in Subsection 5.1. \mathcal{K}_2 is defined for convenience as the shorthand

$$\mathcal{K}_2(\xi, \zeta, \mathbf{w}) = \mathcal{K}(\mathbf{w} - \xi) \mathcal{K}(\mathbf{w} - \zeta)^*. \quad (24)$$

In an analogous manner to Equation (22), the autocovariance of the fluctuating intensity in the case of polychromatic light becomes:

$$\begin{aligned} \tilde{C}(\mathbf{w}_1, \mathbf{w}_2) &= \langle I(\mathbf{w}_1) \rangle \langle I(\mathbf{w}_2) \rangle + \\ &+ \int_{\Lambda} \int_{\mathcal{P}} \int_{\mathcal{P}} \int_{\Lambda} \int_{\mathcal{P}} \int_{\mathcal{P}} S(\lambda_1) S(\lambda_2) \Gamma_c(\xi_1, \zeta_1, \tau_1) \Gamma_c(\xi_2, \zeta_2, \tau_2)^* \cdot \\ &\cdot C_4(\xi_1, \zeta_1, \lambda_1, \xi_2, \zeta_2, \lambda_2) \mathcal{K}_2(\xi_1, \zeta_1, \mathbf{w}_1) \mathcal{K}_2(\xi_2, \zeta_2, \mathbf{w}_2)^* \cdot \\ &\cdot d\xi_1 d\zeta_1 d\lambda_1 d\xi_2 d\zeta_2 d\lambda_2, \end{aligned} \quad (25)$$

where integration is now over tuples of positions and wavelengths and the fourth-order moments C_4 now become wavelength-dependent as well. The derived autocovariance of an imaged speckle field (Equation (25)) is the fundamental theoretical foundation in our discussion and will be at the core of our rendering framework.

Why decompose into mean and fluctuating intensities? Decomposing a field into its mean and fluctuating parts is a common technique in scatter theory [Tsang et al. 2002] and our reasons for taking this approach are as follows: We are, essentially, overlaying a speckle field—corresponding to a specific realization of the surface—on top of the scattered mean field in an optically and mathematically accurate manner. The overlaid speckle field is described by the zero-mean fluctuating intensity \mathfrak{I} , and the scattered mean field, $\langle I \rangle$, is then the “smooth” BRDF lobe—the lobe that arises under fully incoherent lighting. This makes the mean field highly sensitive to approximation errors, as small inaccuracies would adversely effect the smooth appearance of the specular lobe. On the other hand, small errors in the standard deviation of the fluctuating speckle field do not produce significant artefacts (see Figure 8). Therefore, while integrating the total intensity I (Equation (20)) directly appears to be, at worst, no more challenging than integrating the autocovariances

(Equations (22) and (25)), generating visually-pleasing results this way is difficult. Our decomposition allows us to overcome this difficulty by applying contemporary surface scatter theories to produce decently-accurate, computationally-tractable approximations to $\langle I \rangle$ (as discussed in Section 3). This immediately solves the problem of computing the mean of the distribution, and we are left with the task of approximating the overlaid (zero-mean) speckle field \mathfrak{I} .

5.1 The Scattered Phasors and Their Moments

In this subsection we discuss the phasors scattered off a random rough surface and their moments in greater detail. The moments are of interest as the fourth-order moment, C_4 , appears in the expressions for the autocovariance of the fluctuating intensity (Equations (22) and (25)). Our primary theoretical contributions are developed here and will be used later in Section 6 to derive our rendering algorithm.

Consider a wavefront incident upon the surface. Denote by \mathbf{k}^i the incident wavefront’s wavevector—the vector describing the direction of propagation of a wave and its spatial frequency. Assume the wave’s wavelength is λ and, therefore, its wavenumber is $k = |\mathbf{k}^i| = 2\pi\lambda^{-1}$. Also assume that the wavefront scatters off the surface at some point $\mathbf{x} \in \mathcal{S}$ and let a scattered wave’s wavevector be \mathbf{k}^s . The scattered wave gives rise to phasors that are observed at the image plane \mathcal{P} . Different phasors will admit a different phase, and a phase difference between observed phasors manifests due to the geometric distance difference travelled by the wavefronts (see illustration in Figure 3b), as well as a phase shift potentially induced on scatter. Let a specific surface realization be defined via its point-wise height deviation from the mean, $h(\mathbf{x}) : \mathcal{S} \rightarrow \mathbb{R}$. Then, by employing simple geometry, the relative phase variation of the wave scattered off the surface can be expressed as [Harvey and Pfisterer 2016; Krywonos 2006]:

$$\Phi(\mathbf{x}, \mathbf{k}^i, \mathbf{k}^s) = h(\mathbf{x}) \langle \mathbf{n}, \mathbf{k}^s - \mathbf{k}^i \rangle + \langle \mathbf{x}, \mathbf{k}^s - \mathbf{k}^i \rangle, \quad (26)$$

(in the rightmost inner product the projection of the term $\mathbf{k}^s - \mathbf{k}^i$ onto the plane \mathcal{S} was omitted as the projection operator is self-adjoint). In general, only a fraction of the energy carried by the incident wave is present in the scattered wave and the observed phasor. Neglecting any energy loss due to the imaging system and volumetric attenuation, the peak amplitude ratio between the incident wave, denoted a_0 , and the scattered wave, a , is the complex Fresnel coefficient \mathfrak{f} :

$$a = \mathfrak{f}(\mathbf{x}, \mathbf{k}^i, \mathbf{k}^s) a_0. \quad (27)$$

Note that the Fresnel coefficient is generally a complex-valued function, representing a phase shift induced by the surface scatter (see Appendix D for explicit expressions for the polarization-dependent coefficients \mathfrak{f} and their relation to the factor Q in Equation (8)).

Moments. Denote $u_j = a_j e^{-i\phi_j - i\Phi(\mathbf{x}_j, \mathbf{k}^i, \mathbf{k}^s)}$ as some observed phasors and we proceed by considering the phasors’ moments. Under the context discussed above we observe the following well-known corollary [Goodman 2007].

COROLLARY 5.1. *Treating $h(\mathbf{x})$ as a random variable, a phasor's phase can be regarded as uniformly distributed on $(-\pi, +\pi]$, that is*

$$\arg(u_j) \sim \mathcal{U}(-\pi, +\pi].$$

Proof is given in Appendix C. By Corollary 5.1 a phasor u_j can be assumed to follow circularly-symmetric complex Gaussian statistics. As a side note, we can observe that the first-order moment vanishes, viz. $\langle u_j \rangle \equiv 0$, when averaged over all possible realizations of the surface, as expected of fully developed speckle.

We turn our attention to the fourth-order moment, C_4 , which plays a crucial role in the speckle autocovariance function (Equation (25)). In general, evaluating the fourth-order moment can be difficult, however we use the fact that under the rough surface assumption the phasors follow circularly-symmetric complex Gaussian statistics in order to derive a simpler analytic expression to C_4 . To that end, by remembering a beautiful result in computer graphics—that uniform values on a sphere can be drawn by drawing Gaussian vectors [Muller 1959]—we present a generalization to the well known Isserlis' Theorem in probability theory:

THEOREM 5.2. *Denote $X = [a_1 e^{i\theta_1}, \dots, a_{2m} e^{i\theta_{2m}}]^T$ to be a $2m$ -dimensional ($m > 1$) multivariate complex random vector, such that the magnitudes a_j are real-valued constant values and the phases are uniformly distributed, viz. $\theta_j \sim \mathcal{U}(-\pi, +\pi]$. Then, the $2m$ -order moment can be decomposed in an identical fashion to Isserlis' Theorem [Isserlis 1918], as follows*

$$\mathbb{E} \left[\prod_{j=1}^{2m} a_j e^{i\theta_j} \right] = \left(\prod_{j=1}^{2m} a_j \right) \cdot \sum_{p \in P_{2m}^2} \prod_{\{j,l\} \in p} \mathbb{E} \left[e^{i(\theta_j + \theta_l)} \right],$$

where the sum is over all possible pairings of indices and the product is over each pair in the pairings.

PROOF. Proof given in Appendix E □

COROLLARY 5.3. *The fourth-order moment can be decomposed as follows:*

$$\begin{aligned} C_4(u_1, u_2, u_3, u_4) = & C_2(u_1, u_2) \cdot C_2(u_3, u_4)^* \\ & + C_2(u_1, u_3) \cdot C_2(u_2, u_4)^* \\ & + C_2(u_1, u_4) \cdot C_2(u_2, u_3)^*. \end{aligned}$$

PROOF. By Corollary 5.1, the phasor u_j follows circularly-symmetric complex Gaussian statistics with fixed magnitude, therefore the desired decomposition is obtained directly by applying Theorem 5.2. □

Corollary 5.3 allows us to write the fourth-order moment C_4 as an expression of second-order moments C_2 only, and we proceed by examining C_2 .

As discussed, we consider the height function h at any two points to be jointly Gaussian, thus a closed-form expression for second-order moment, C_2 , can be easily derived. The derivation of C_2 is outlined in Appendix F and the final expression is listed here:

$$C_2(u_1, u_2) = \langle u_1 u_2^* \rangle = W_2(u_1, u_2) Z_1(u_1) Z_1(u_2)^*, \quad (28)$$

where

$$\begin{aligned} Z_1(u) &= a e^{-i\phi} e^{-i\langle \mathbf{x}, \mathbf{k}^s - \mathbf{k}^i \rangle} \\ W_2(u_1, u_2) &= e^{\sigma_{\text{rel}}^2 C_S(\mathbf{x}_1, \mathbf{x}_2) \langle \mathbf{n}, \mathbf{k}^{s_1} - \mathbf{k}^{i_1} \rangle \langle \mathbf{n}, \mathbf{k}^{s_2} - \mathbf{k}^{i_2} \rangle} \\ &\quad \cdot e^{-\frac{\sigma_{\text{rel}}^2}{2} \langle \mathbf{n}, \mathbf{k}^{s_1} - \mathbf{k}^{i_1} \rangle^2} e^{-\frac{\sigma_{\text{rel}}^2}{2} \langle \mathbf{n}, \mathbf{k}^{s_2} - \mathbf{k}^{i_2} \rangle^2}, \end{aligned} \quad (29)$$

with $W_2 \in \mathbb{R}$ and $Z_1 \in \mathbb{C}$. Equation (28) is applicable for any surface correlation function C_S . It is noteworthy that while the assumption of Gaussianity of the surface height is used to obtain a closed-form expression for C_2 , Gaussianity is not strictly required for our discussion and any other distribution will do with minor modifications as long as C_2 can be written in closed-form.

5.2 Spatial and Angular Covariances

We now decompose the autocovariance \tilde{C} into a *spatial autocovariance* and *angular autocovariance*, viz. $\tilde{C} = \tilde{C}_{\text{spatial}} + \tilde{C}_{\text{ang}}$, and the motivation will become apparent once we present our rendering algorithm in Section 6.

Consider a couple of observations (\mathbf{w}_j, t_j) and (\mathbf{w}_l, t_l) such that $\mathbf{w}_j \neq \mathbf{w}_l$. Inline with the discussion in Subsection 4.1, we assume $l_{\text{cor}} \ll \|\mathbf{w}_j - \mathbf{w}_l\|$, that is the surface correlation length (Equation (6)) is small compared to the spatial distance between the image elements. Then, given a sharp enough impulse response function, observe that for most points \mathbf{x}, \mathbf{y} where $|\mathcal{K}(\mathbf{w}_j - \mathbf{x})|$ and $|\mathcal{K}(\mathbf{w}_l - \mathbf{y})|$ are not negligible it holds that $l_{\text{cor}} \ll \|\mathbf{x} - \mathbf{y}\|$, ergo $C_S(\mathbf{x}, \mathbf{y}) \approx 0$, and thus $W_2 \approx 0$ (directly from Equation (29), for rough surfaces at non-grazing angles). Therefore, the fourth-order moment for some observed phasors $u_{1,2,3,4}$ where $u_{1,2}$ are far, with respect to the correlation length, from $u_{3,4}$ simplifies to

$$C_4(u_1, u_2, u_3, u_4) \approx C_2(u_1, u_2) \cdot C_2(u_3, u_4)^*. \quad (30)$$

Apply now Equation (30) to the covariance between the fluctuating imaged speckle intensities formed at a couple of observations $\tilde{C}(\mathbf{w}_j, t_j, \mathbf{w}_l, t_l)$ (where we abuse notation and introduce the time-dependence of each observation into \tilde{C}). Then, the autocovariance decomposes into

$$\begin{aligned} \tilde{C}_{\text{spatial}}(\mathbf{w}_j, t_j, \mathbf{w}_l, t_l) &= \langle I(\mathbf{w}_j) \rangle \langle I(\mathbf{w}_l) \rangle + \\ &\quad + \iiint \mathcal{X}(\mathbf{w}_j, \xi, \zeta, \lambda) \iiint \mathcal{X}(\mathbf{w}_l, \xi, \zeta, \lambda)^*, \end{aligned} \quad (31)$$

with

$$\begin{aligned} \mathcal{X}(\mathbf{w}, \xi, \zeta, \lambda) &= S(\lambda) \Gamma_c(\xi, \zeta, \lambda) \mathcal{K}_2(\xi, \zeta, \mathbf{w}) \cdot \\ &\quad \cdot C_2(u(\xi, \lambda), u(\zeta, \lambda)). \end{aligned} \quad (32)$$

The “spatial” subscript expresses that $\tilde{C}_{\text{spatial}}$ approximates \tilde{C} under the assumption of $\mathbf{w}_j \neq \mathbf{w}_l$.

Consider now the case when $\mathbf{w}_j = \mathbf{w}_l$: We would like to add a term, denoted \tilde{C}_{ang} , which would account for the second-order moments that were ignored in $\tilde{C}_{\text{spatial}}$. Observe from Equation (25) that when $\Gamma_c = \delta^2$ we expect that $\tilde{C} \approx 0$ and thus $\tilde{C}_{\text{ang}} \approx 0$ (with the physical interpretation being that under incoherent radiation no fluctuating field arises), while when the coherence area is on the order of magnitude of the image element, Γ_c is approximately constant, and by relabelling the integration variables in Equation (25)

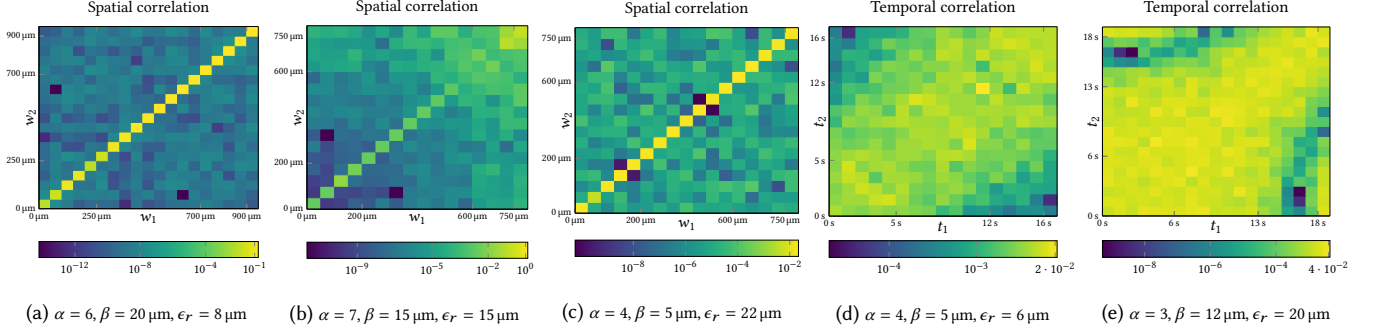


Fig. 4. Plots of the relative errors $e = |\tilde{C} - \tilde{C}'| / \tilde{C}'$ of the covariance matrices \tilde{C} , computed using the decomposition into spatial and angular covariances (see Subsection 5.2), compared against the ground-truth \tilde{C}' , which is the numerically integrated Equation (22). The simulations were done in flatland for different coherence radii ϵ_r and a variety of statistical rough surfaces, with the surface autocorrelation analytically described via the K-correlation model and parametrized by α and β (see Section 3). The comparison is performed in flatland as numerical integration of the 8-dimensional integral (Equation (22)) is infeasible. Plots (a-c) plot the behaviour of the spatial covariance of an imaging device consisting of multiple image elements where each image element covers a surface area of $50 \mu\text{m}$ and the light source is located 1 mm above the surface at position $w = 0 \text{ mm}$. In plot (b) the light is moved to position $w = -1 \text{ mm}$ and some loss of accuracy can be seen at grazing angles. In (d-e) the angular covariance is plotted for a single image element with a moving light source. The time difference between two subsequent time points is 1 s and the light rotates at an angular velocity of 1° s^{-1} around the surface. Plots (c) and (e) show that some errors, up to 0.10 relative to the ground truth, arise with larger coherence areas.

we get that $\tilde{C} \approx 3\tilde{C}_{\text{spatial}}$. Thus, we approximate \tilde{C}_{ang} via a simple first-order approximation:

$$\tilde{C}_{\text{ang}}(\mathbf{w}_j, t_j, \mathbf{w}_l, t_l) = 2g(\mathbf{w}_j) \cdot \tilde{C}_{\text{spatial}}(\mathbf{w}_j, t_j, \mathbf{w}_l, t_l), \quad (33)$$

with

$$g(\mathbf{w}_j) = \frac{\pi\epsilon_r^2}{\int_S \mathcal{K}\mathcal{K}^*},$$

i.e. the ratio between the coherence area and the area of a surface patch that falls under the impulse response power of an image element. From the above analysis we can also deduce that under partially coherent radiation it holds that $|\tilde{C}_{\text{ang}}| \ll |\tilde{C}_{\text{spatial}}|$ as Γ_c constricts the integration (in Equation (25)) to a region small compared to the impulse response \mathcal{K} . We are, therefore, satisfied with the rough approximation for \tilde{C}_{ang} listed in Equation (33). Despite its simplicity numerical analysis suggests that $\tilde{C}_{\text{ang}} + \tilde{C}_{\text{spatial}}$ approximates \tilde{C} very well, especially when dealing with natural light source with small coherence areas (see Figure 4).

6 RENDERING SPECKLE FORMED ON PARTIALLY COHERENT SCATTER FROM ROUGH SURFACES

In this section we present our rendering framework to render optically accurate speckle patterns produced by some distinct realization of the random surface. As before, let \mathcal{P} be the image plane and let the plane \mathcal{S} , with normal vector \mathbf{n} , be the mean of the random scattering surface. Given a sequence of N observations, $O = [(\mathbf{w}_1, t_1), \dots, (\mathbf{w}_N, t_N)]^T$, each observation being a spacetime point, where $\mathbf{w}_j \in \mathcal{P}$ is the spatial position of an image element and t_j is the time point (e.g., pixels in different frames of a rendered animation), our problem statement is as follows: We would like to draw a sequence of fluctuating intensities \mathfrak{I} , one for each observation, from the distribution of the fluctuating speckle field described by the autocovariance function \tilde{C} . The superposition of the drawn fluctuating speckle field with the mean intensity field

$\langle I \rangle$ —interpreted as the convolution of the imaging impulse response power with the BRDF—would then be a physically correct scattered intensity field for some realization of the random scattering surface.

6.1 Monte Carlo Rendering of a Speckle Pattern

The purpose of the decomposition outlined in Equation (31) is to enable us to Monte Carlo approximate the integral $\iiint \mathcal{X}$ and draw values from the fluctuating field distribution, as discussed. However, observe that our decomposed covariance is a superposition of two covariances, the spatial covariance, $\tilde{C}_{\text{spatial}}$, and its angular counterpart, \tilde{C}_{ang} , and there is a crucial conceptual difference between the two: The former applies to all observation pairs, while the second applies only to pairs where $\mathbf{w}_j = \mathbf{w}_l$. A major difficulty in rendering the speckle pattern is devising a feasible approach for estimating those covariances, and drawing values from that distribution for a large count of observations, while correctly accounting both for the spatial and angular correlations. In this subsection we present our general-purpose extension to the Monte Carlo sampling algorithm by Bar et al. [2019] that is capable of drawing values under such constraints as well as handling temporal adaptation induced by scene change.

Our rendering algorithm (see Algorithm 1) comprises three steps: We first compute the intensity of the mean field $\langle I \rangle$. We continue with drawing a value, denoted \mathfrak{I} , from the distribution of the fluctuating intensity field \mathfrak{I} . Finally, those quantities give rise to the final intensity of the speckle field \tilde{I} .

Mean field intensity. We start with computing the mean intensity $\langle I(\mathbf{w}) \rangle$, which is required to compute both the final intensity $I(\mathbf{w})$ (Equation (20)) and the fluctuating field covariance. As discussed in Subsection 4.2, we employ the tractable expressions for the surface BRDF, provided by existing scatter theories, to compute a good approximation for the mean intensity field. Then, for each observation (\mathbf{w}_j, t_j) , we compute the mean intensity, $\langle I(\mathbf{w}_j) \rangle$, by convolving

the BRDF (e.g., Equation (8)) with the impulse response power:

$$\begin{aligned} \langle I(\mathbf{w}_j) \rangle &= \int_{\mathcal{P}} \text{BRDF}(\mathbf{x}) \mathcal{K}_2(\mathbf{x}, \mathbf{x}, \mathbf{w}_j) d\mathbf{x} = \\ &= [\text{BRDF} * (\mathcal{K}\mathcal{K}^*)](\mathbf{w}_j). \end{aligned} \quad (34)$$

Drawing fluctuating intensity values. Let $M > 0$ be the count of Monte Carlo iterations, which we consider to be known a priori and constant throughout the rendering process, and let $H > 0$ be the count of distinct image elements \mathbf{w}_j , which we can also reasonably assume to remain constant. We begin by drawing a couple of M -dimensional discrete random vectors, referred to as the *mean integration* and *spatial integration phases* and denoted $p^{(\text{mean})}$ and $p^{(\text{spatial})}$, respectively. We further draw a matrix of $M \times H$ discrete random values, denoted $p^{(\text{ang})}$, i.e. the *angular integration phases*. All integration phases are i.i.d. from the two-point uniform distribution in $\{-1, +1\}$ (known as the *Rademacher distribution*), i.e.

$$p_m^{(\text{mean})}, p_m^{(\text{spatial})}, p_{m,q}^{(\text{ang})} \sim \mathcal{U}\{-1, +1\}. \quad (35)$$

($1 \leq m \leq M, 1 \leq q \leq H$). Observe that all integration phases are zero-mean and pair-wise orthonormal, that is

$$\mathbb{E}[p_j] = 0 \quad \mathbb{E}[p_j p_l] = \delta_{jl}, \quad (36)$$

where $p_{j,l}$ are any (mean, spatial or angular) integration phases, and δ_{jl} is the Kronecker delta. The way we draw the phases is the key to our algorithm and will be discussed further in Section 7.

We then proceed via the following Monte Carlo process: For each observation (\mathbf{w}_j, t_j) , we draw the data needed to evaluate the integrand, \mathcal{X} , in Equation (31); that is an M -tuple consisting of wavelengths $\lambda_{j,m}$, as well as pairs of points $\xi_{j,m}, \zeta_{j,m}$ (with m indexing the Monte Carlo iteration), and finally, we also need to draw incident wavevectors $\mathbf{k}_{\xi_{j,m}}^i, \mathbf{k}_{\zeta_{j,m}}^i$ for each point by selecting a source of radiation. Note that the image element \mathbf{w}_j and the drawn surface points uniquely determine the scattered wavevectors, $\mathbf{k}_{\xi_{j,m}}^s$ and $\mathbf{k}_{\zeta_{j,m}}^s$. Denote

$$\mathfrak{X}_j = \begin{bmatrix} \mathcal{X}(\mathbf{w}_j, \xi_{j,1}, \zeta_{j,1}, \lambda_{j,1}) \\ \mathcal{X}(\mathbf{w}_j, \xi_{j,2}, \zeta_{j,2}, \lambda_{j,2}) \\ \vdots \\ \mathcal{X}(\mathbf{w}_j, \xi_{j,M}, \zeta_{j,M}, \lambda_{j,M}) \end{bmatrix}, \quad (37)$$

and we draw values from the speckle fluctuating intensity distribution via the following procedure

$$\begin{aligned} \check{\mathfrak{I}}_j &= \frac{1}{\sqrt{M}} \sum_{m=1}^M \left[\sqrt{2g(\mathbf{w}_j)} \cdot p_{m,\mathbf{w}_j}^{(\text{ang})} \left(p_m^{(\text{mean})} \langle I(\mathbf{w}_j) \rangle + \mathfrak{X}_{j,m} \right) \right. \\ &\quad \left. + p_m^{(\text{spatial})} \left(p_m^{(\text{mean})} \langle I(\mathbf{w}_j) \rangle + \mathfrak{X}_{j,m} \right) \right], \end{aligned} \quad (38)$$

where we slightly abuse notation and we index the angular integration phases, $p^{(\text{ang})}$, directly via the image element. Equation (38) outlines our rendering procedure and the correctness of it is ensured formally by the corollary that follows.

COROLLARY 6.1. *The drawn fluctuating intensity $\check{\mathfrak{I}}$ (using the procedure outlined in Equation (38)) converges to the true fluctuating*

Algorithm 1: Speckle rendering algorithm

```

1 begin
  // Draw random integration phases as discussed in
  // Subsection 6.1
2 draw  $p^{(\text{mean})}, p^{(\text{spatial})}$  and  $p^{(\text{ang})}$ ;
  // Render observation at image element  $\mathbf{w}$  at time  $t$ 
3 for  $(\mathbf{w}, t) \in \mathcal{O}$  do
4    $\langle I \rangle \leftarrow \text{BRDF}_{\text{gHS}}(\mathbf{w})$ ; // BRDF
5    $\check{\mathfrak{I}} \leftarrow 0$ ; // Fluctuating intensity
  // Monte Carlo iterations
6   for  $m = 1$  to  $M$  do
  // Draw a couple of positions, a wavelength and
  // light source
7   Draw  $\mathbf{x}_1, \mathbf{x}_2, \lambda, \mathbf{k}_1^i, \mathbf{k}_2^i$ ;
  // Compute scattered wavevectors
8    $\mathbf{k}_1^s \leftarrow \frac{2\pi}{\lambda} \frac{\mathbf{w} - \mathbf{x}_1}{\|\mathbf{w} - \mathbf{x}_1\|}$ ;
9    $\mathbf{k}_2^s \leftarrow \frac{2\pi}{\lambda} \frac{\mathbf{w} - \mathbf{x}_2}{\|\mathbf{w} - \mathbf{x}_2\|}$ ;
  // Evaluate sample and accumulate
10   $t \leftarrow p_m^{(\text{mean})} \langle I \rangle + \mathcal{X}(\mathbf{w}, \mathbf{x}_1, \mathbf{x}_2, \lambda)$ ;
11   $\check{\mathfrak{I}} \leftarrow \check{\mathfrak{I}} + \sqrt{2g(\mathbf{w})} p_{m,\mathbf{w}}^{(\text{ang})} t + p_m^{(\text{spatial})} t$ ;
  // Final intensity for the observation
12   $\check{I} \leftarrow \langle I \rangle + \frac{1}{\sqrt{M}} \text{Re} \{ \check{\mathfrak{I}} \}$ ;

```

intensity $\check{\mathfrak{I}}$ in L^1 asymptotically as M —the count of Monte Carlo samples—increases, viz. $\check{\mathfrak{I}} \xrightarrow{M \rightarrow \infty} \check{\mathfrak{I}}$.

PROOF. Clearly $\mathbb{E}[\check{\mathfrak{I}}] = 0$ and we proceed by examining the covariance. In Appendix G we prove Theorem G.1, which shows that under any probability space and for any zero-mean orthonormal integration phases the covariance of two values drawn via the procedure in Equation (38) is as follows:

$$\begin{aligned} \text{cov}[\check{\mathfrak{I}}_j, \check{\mathfrak{I}}_l] &= \frac{1}{M} \sum_{m=1}^M \left[\langle I(\mathbf{w}_j) \rangle \langle I(\mathbf{w}_l) \rangle + \mathfrak{X}_{j,m} \mathfrak{X}_{l,m}^* \right. \\ &\quad \left. + 2\sqrt{g(\mathbf{w}_j)g(\mathbf{w}_l)} \mathbb{E} \left[p_{m,\mathbf{w}_j}^{(\text{ang})} p_{m,\mathbf{w}_l}^{(\text{ang})} \right] \left(\langle I(\mathbf{w}_j) \rangle \langle I(\mathbf{w}_l) \rangle + \mathfrak{X}_{j,m} \mathfrak{X}_{l,m}^* \right) \right]. \end{aligned} \quad (39)$$

Using the fact that the integration phases are orthonormal as well as Equations (31) and (33) we can deduce that

$$\begin{aligned} \text{cov}[\check{\mathfrak{I}}_j, \check{\mathfrak{I}}_l] &\xrightarrow{M \rightarrow \infty} \left(1 + 2\delta_{\mathbf{w}_j, \mathbf{w}_l} \cdot g(\mathbf{w}_j) \right) \tilde{\mathcal{C}}_{\text{spatial}}(\mathbf{w}_j, t_j, \mathbf{w}_l, t_l) = \\ &= \tilde{\mathcal{C}}_{\text{spatial}} + \delta_{\mathbf{w}_j, \mathbf{w}_l} \tilde{\mathcal{C}}_{\text{ang}} \approx \tilde{\mathcal{C}}, \end{aligned} \quad (40)$$

as required. \square

It is worth noting that while \mathfrak{I} represents a physical stochastic process, $\check{\mathfrak{I}}$ is simply a random variable. This distinction is purely semantic with no practical consequences, however that change of semantics serves to draw attention to the fact that $\check{\mathfrak{I}}$, in contrast to \mathfrak{I} , carries no physical meaning and is drawn via a strictly mathematical process (Equation (38)). This is also the reason behind us switching to the use of the expected value operator $\mathbb{E}[\cdot]$ in place of ensemble averaging.




Fig. 5. Simple spheres made of (top) silver and (bottom) titanium rendered using our method. The spheres are illuminated with a (top) CIE D65 and (bottom) CIE D50 standard white daylight sources, and the sources admit increasing spatial coherence areas (right-to-left). All figures are rendered with $M = 2048$ Monte-Carlo samples and the right part of each figure (on the right of the grey line) is rendered without speckle (only mean field) for comparison.

Computing the final intensity. Note that even though the fluctuating intensity \mathfrak{I} computed directly from Equation (16) or Equation (17) is real as expected—due to the complex parts cancelling out by the integration symmetry—the integrand is nonetheless complex and values drawn using the discrete Monte Carlo process in Equation (38) are complex as well. Nonetheless, as $\check{\mathfrak{I}} \xrightarrow{M \rightarrow \infty} \mathfrak{I}$ we deduce that $\text{Im} \left\{ \check{\mathfrak{I}} \right\} \xrightarrow{M \rightarrow \infty} 0$ and then the final intensity trivially is

$$\check{I} = \langle I \rangle + \text{Re} \left\{ \check{\mathfrak{I}} \right\}. \quad (41)$$

Clearly $\check{I} \xrightarrow{M \rightarrow \infty} I$ in L^1 and it is easy to see that $\text{E} [\check{I}] = \langle I \rangle$. Furthermore, the covariance between the drawn intensities of a couple of observations, \check{I}_j and \check{I}_l , becomes

$$\begin{aligned} \text{cov} [\check{I}_j, \check{I}_l] &= \text{E} [\check{I}_j \check{I}_l] - \text{E} [\check{I}_j] \text{E} [\check{I}_l] = \\ &= \text{E} \left[\text{Re} \left\{ \check{\mathfrak{I}}_j \right\} \text{Re} \left\{ \check{\mathfrak{I}}_l \right\} \right] \xrightarrow{M \rightarrow \infty} \text{cov} [\check{\mathfrak{I}}_j, \check{\mathfrak{I}}_l] \end{aligned} \quad (42)$$

and thus the final intensity \check{I} admits the desired statistics.

Temporal adaptation. The angular integration phases (Equation (35)) give rise to the $\text{E} \left[\hat{p}_{m, \mathbf{w}_j}^{(\text{ang})} \hat{p}_{n, \mathbf{w}_l}^{(\text{ang})} \right] = \delta_{\mathbf{w}_j, \mathbf{w}_l}$ term in Equation (39). As described in Subsection 5.2, this term is designed to selectively factor in the angular covariance approximation, \mathcal{C}_{ang} (Equation (33)), into Equation (40) when $\mathbf{w}_j = \mathbf{w}_l$. The implication of this approach is that we effectively assume that the surface patch giving rise to scatter observed at a couple of space-time points (\mathbf{w}_j, t_j) and (\mathbf{w}_l, t_l) is identical when $\mathbf{w}_j = \mathbf{w}_l$, i.e. the scene is static. We would like then to adapt to angular changes in the covariance between phasors observed by the same image element but at different time points.

Let (\mathbf{w}, t_j) and (\mathbf{w}, t_l) be a couple of observations, which share an image element. Assume that the surface patch that contributes to \mathbf{w} was displaced by $\mathbf{d} = \delta \mathbf{x}$ between time points t_j and t_l . By studying Equation (29) and using the same arguments as in Subsection 5.2 we claim that W_2 , and in turn the angular covariance, decays, when

the underlying surface shifts by \mathbf{d} , approximately by a factor of

$$\Delta_{\mathbf{x}}(\mathbf{d}) = e^{2\pi\sigma_{\text{rel}}^2 \bar{\lambda}^{-1} [\mathcal{C}_s(\|\mathbf{d}\|) - 1]}, \quad (43)$$

where $\bar{\lambda}$ is the mean wavelength and we simplified the expression via the order-of-magnitude relation $\langle \mathbf{n}, \mathbf{k}^s_1 - \mathbf{k}^s_2 \rangle \sim \langle \mathbf{n}, \mathbf{k}^s_2 - \mathbf{k}^s_2 \rangle \sim \frac{2\pi}{\bar{\lambda}}$. \mathcal{C}_s is the surface autocorrelation function at surface point \mathbf{x} . Our intention is then to scale the Kronecker delta in Equation (40) by $\Delta_{\mathbf{x}}$. To that end, instead of drawing a single set of angular integration phases $p^{(\text{ang})}$, we draw a set for each time point, denoted $\hat{p}^{(t_j)}$, recursively: Given $\hat{p}^{(t_j)}$ the angular phases for the next time point t_{j+1} are drawn as follows

$$\hat{p}_{m, \mathbf{w}}^{(t_{j+1})} = \hat{p}_{m, \mathbf{w}}^{(t_j)} \cdot (-1)^{\text{Bernoulli} \left(\frac{1 - \Delta_{\mathbf{x}}(\mathbf{d})}{2} \right)}, \quad (44)$$

where \mathbf{x} is the surface point observed by image element \mathbf{w} at time $t = t_{j+1}$, \mathbf{d} is the surface displacement from $t = t_j$ and $\text{Bernoulli}(k)$ ($0 \leq k \leq 1$) denotes a draw from the Bernoulli distribution, i.e.

$$\Pr\{\text{Bernoulli}(k) = 1\} = k \quad \text{and} \quad \Pr\{\text{Bernoulli}(k) = 0\} = 1 - k. \quad (45)$$

Then, the following property of the angular integration phases holds:

$$\begin{aligned} \text{E} \left[\hat{p}_{m, \mathbf{w}_l}^{(t_j)} \hat{p}_{n, \mathbf{w}_q}^{(t_{j+1})} \right] &= \text{E} \left[\hat{p}_{m, \mathbf{w}_l}^{(t_j)} \hat{p}_{n, \mathbf{w}_q}^{(t_j)} \right] \text{E} \left[(-1)^{\text{Bernoulli} \left(\frac{1 - \Delta_{\mathbf{x}}(\mathbf{d})}{2} \right)} \right] = \\ &= \delta_{mn} \delta_{\mathbf{w}_l, \mathbf{w}_q} \Delta_{\mathbf{x}}(\mathbf{d}) \end{aligned} \quad (46)$$

as desired.

6.2 Practical Considerations

Importance sampling. To evaluate \mathcal{X} (Equation (32)) we require drawing data from a few sources: First, we need to select a source of radiation; and second, we need to draw a couple of scattering points on the surface. The former is relatively straightforward to handle, a point on a light source is chosen and a wavelength is importance sampled with respect to the light's spectral density. For the latter we start by drawing a position, ξ , with respect to the impulse response

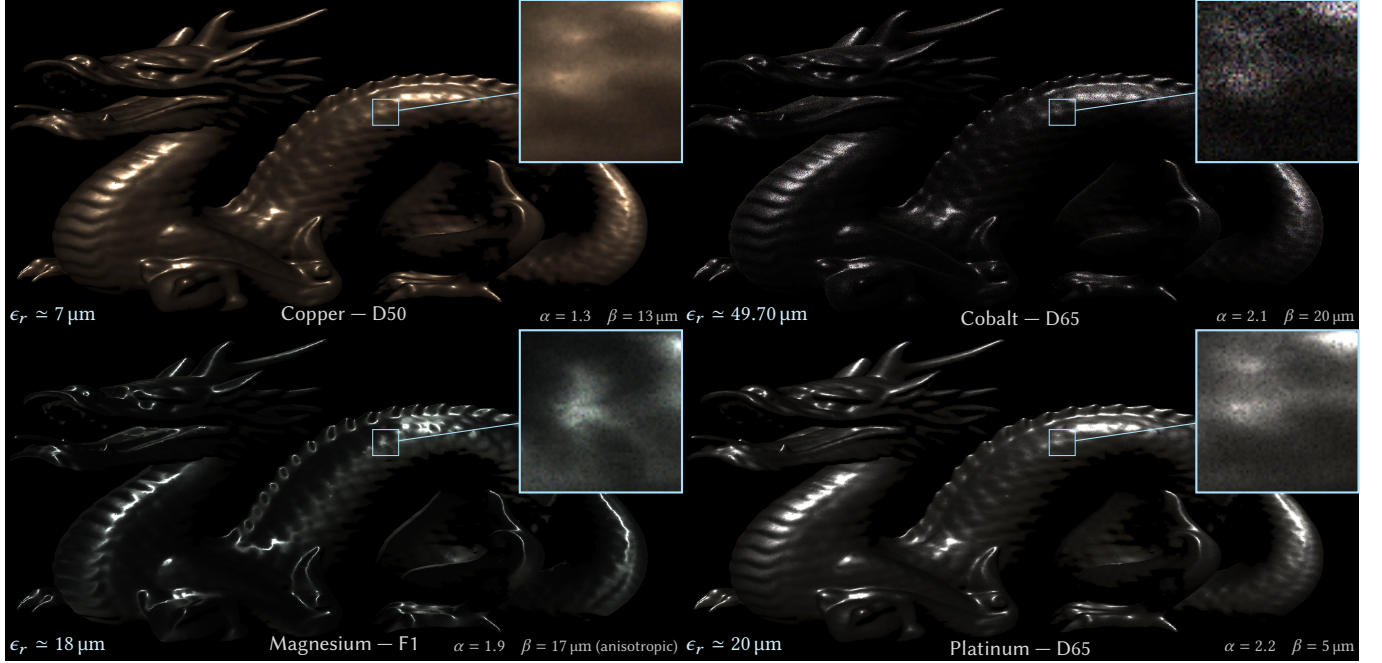


Fig. 6. The Stanford Dragon made of different metals and surface statistics (bottom right of each Dragon), rendered under incident illumination consisting of different spectra and coherence properties (bottom left of each Dragon). The CIE D50 and D65 are standard white daylight light sources with wide spectrum, while the F1 illuminant represents a white Fluorescent light source with sharp peaks in its spectrum. The insets at the top right corner of each Dragon display a close-up of the produced patterns. The Dragon’s dimensions are 50 cm \times 29.50 cm (L \times H) and each rendered pixel covers roughly $\sim 250 \times 250 \mu\text{m}^2$ surface area. The Magnesium Dragon (bottom left) exhibits a high degree of anisotropy in its surface statistical properties. All figures are rendered with $M = 2048$ Monte-Carlo samples.

function $|\mathcal{K}(\mathbf{w} - \xi)|$. The importance sampling strategy for the second position, ζ , depends on the light coherence properties as well as the surface statistics:

- When the rough order-of-magnitude relation $\epsilon_r \ll l_{\text{cor}}$ holds, that is when the light coherence radius ϵ_r is less than the surface correlation length l_{cor} , then the contributions to the integral in Equation (31) are dominated by the coherence radius. Then, we importance sample ζ with respect to the coherence area Gaussian (which is formulated in Subsection 4.3).
- More commonly, however, the surface correlation length is small and the opposite relation $\epsilon_r \gg l_{\text{cor}}$ holds. In this case our implementation samples ζ with respect to the distribution of the surface autocorrelation function, approximated by an exponential distribution (as described in Section 3). This however, is a rough approximation as the actual distribution is dictated by the moment C_2 (Equations (28) and (29)). Better importance sampling of the second-order moment C_2 could potentially greatly speed up convergence and is left for future work.

When the coherence radius and the surface correlation length are roughly similar, any strategy can be used. See our supplemental material for a sample implementation.

Random-number generators. Our integration phases are sensitive to non-optimal random-number generators and small biases can

induce significant errors in the drawn fluctuating intensity $\check{\mathfrak{I}}$. As an experiment, assume that the random-number generator used to draw the integration phases has a constant bias b , such that it draws a value of 1 with probability $\frac{1+b}{2}$ and a value of -1 with probability of $\frac{1-b}{2}$. Observe the sum in Equation (38) and note that when the integration phases are drawn using the biased random-number generator the following holds:

$$\Pr\{p_m^{(\text{spatial})} p_m^{(\text{mean})} = 1\} = \frac{1+b^2}{2}, \quad \Pr\{p_m^{(\text{spatial})} p_m^{(\text{mean})} = -1\} = \frac{1-b^2}{2}.$$

Therefore, we deduce that

$$\mathbb{E} \left[\frac{1}{\sqrt{M}} \sum_m p_m^{(\text{spatial})} p_m^{(\text{mean})} \langle I \rangle \right] = b^2 \sqrt{M} \langle I \rangle, \quad (47)$$

that is, a bias of b introduces an error of magnitude $b^2 \sqrt{M}$ in the final intensity \check{I} . For a large count of Monte Carlo samples, M , this error can become significant and the magnitude of the error grows non-linearly with b .

To alleviate the problem, we simply keep redrawing the integration phases until each set of phases is balanced (sums up to zero), and we leave a more sophisticated solution for future work.

Rendering coloured speckle. For a typical tristimulus renderer, we redefine the function $S(\lambda)$ (the intensity of a spectral line, see Equation (17)) to be a function that maps wavelengths to colour tuples, usually of the CIE XYZ colour-space, scaled by the intensity of the




Fig. 7. A Buddha statue, $12.20\text{ cm} \times 37.10\text{ cm}$ ($W \times H$), made of iron and partially coated in gold (coating assumed thick enough to neglect layered effects). The gold coating admits a smoother surface with greater correlation length compared with the base iron. The Buddha is illuminated by a couple of light sources: A large 1 m^2 D65 illuminant from the right, and a small F2 (warm fluorescent) located behind and above the camera. The large light source is placed close to the statue initially and then is quickly retracted away. The distance from the light source, d , is displayed at the bottom left of each figure. Being close to the surface it admits poor spatial coherence, however its spatial coherence increases as it is moved farther away. This is appreciable by examining the produced speckle pattern, which is visible in the insets to the left of each image (artificially increased contrast and brightness for visualisation purposes). We can see additional speckle appear both on the iron head (purple and red insets) and the gold coating (blue inset). Also notice that the speckle transitions smoothly between the metals, and different surfaces as well as different light sources induce starkly different speckle statistics: The smoother surface gives rise to more vibrant and tightly clustered speckle compared to the rougher surface. See our supplemental material for a rendered video of the Buddha.

spectral line λ . This in-turn also redefines \mathcal{X} , \tilde{I} and $\tilde{\mathfrak{S}}$ to also be colour-space tuples. The rest remains unchanged.

Saving the integration phases. A requirement that arises at times in practice is to be able to re-render a portion of a rendered sequence, or append additional frames. However, if we were to do so with a newly drawn set of integration phases, the new speckle would be inconsistent with the previously rendered one as the speckle would correspond to a different realization of the random surface. To re-render speckle correctly we would need to store the integration phases and reuse the same phases when rendering the new speckle.

7 RESULTS AND DISCUSSION

We have presented a method for rendering speckle formed on scatter by rough statistical surfaces. Some assumptions with regards to the surface statistics were made along the way the most notable of which is the surface roughness assumption formalised as the order-of-magnitude relation $\sigma_{\text{rel}} \gg \frac{1}{2}\lambda$ (Equation (10)). This assumption defines what surfaces fall under our purview, and is also required for our extension to Isserlis' theorem (Theorem 5.2) to apply. Nonetheless, the moment decomposition that follows (Corollary 5.3) remains valid also at the rough limit and for smoother surfaces and the rest of our framework still applies. A decomposition for smooth surfaces would require a slightly different mathematical treatment of the moments, however that would be of limited interest as smooth surfaces do not exhibit any significant speckle.

The only assumption with regards to the properties of the incident radiation is that it is a natural (i.e. weakly coherent) light source.

Indeed, as the rendering is done via Monte Carlo integration, our method is flexible and can work with light sources of arbitrary shape, spectrum and coherence properties (see Figure 7). Some of our results are demonstrated in Figure 6 and rendered videos are available in our supplemental material. For simplicity, we used a simple rect function for the impulse response function \mathcal{K} , though any other function can be used (as long as the imaging resolution remains low enough to give rise to speckle that is Gaussian in intensity). To render the metal surfaces accurately, we use databases of measured wavelength-dependent complex refractive indices, which are easy to incorporate into our rendering framework as we draw singular wavelengths during integration. The rendered Cobalt Dragon (top right in Figure 6) is illuminated by a highly-coherent light source with a coherence radius of $\sim 50\ \mu\text{m}$ which implies that most of the surface subtended by a single pixel gives rise to non-negligible mutual intensities. This stretches our assumption of weakly coherent light, nonetheless the rendered Dragon exhibits colourful speckle, as expected, with stable statistics.

As discussed, an important part of our theoretical formulation is the decomposition of the intensity field into its mean and fluctuating parts in an optically and mathematically consistent manner. This formulation leverages substantial existing work done in applied optics in order to compute the mean field in closed-form, and the rendering algorithm focuses on the fluctuating field. This has multiple advantages: As discussed in Section 3, existing scatter theories are well studied, extensively validated and are being employed both in computer graphics and applied optics. In addition, errors induced

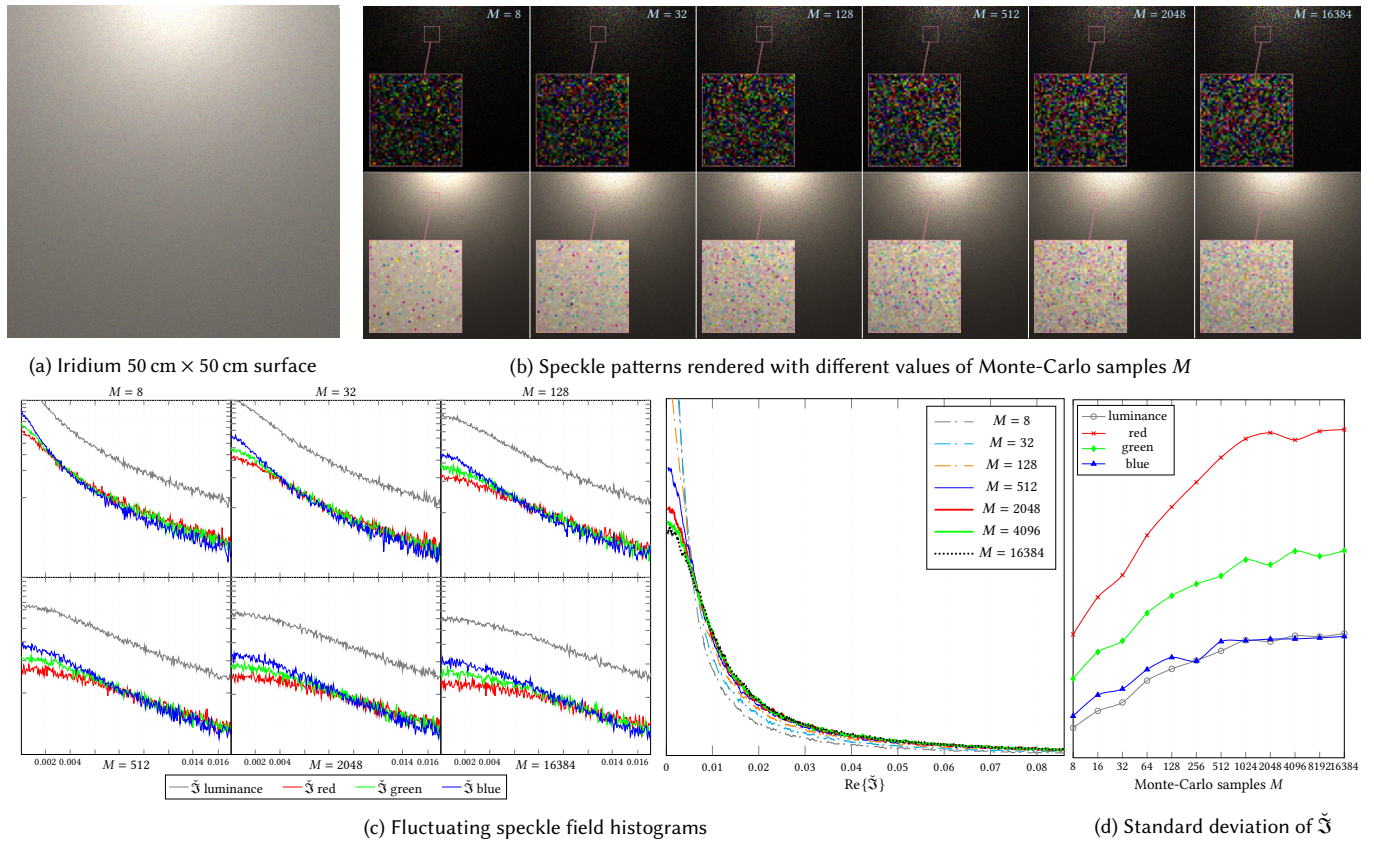


Fig. 8. Analysis of the statistics and convergence of our rendering algorithm. (a) A flat 50 cm x 50 cm surface made of iridium and illuminated by a moderately-coherent light source (illuminant D65) rendered using our method with $M = 16384$ samples, as a reference. (b) The fluctuating field $\tilde{\mathfrak{S}}$ (top) as well as the final intensity \tilde{I} (bottom) were captured after various counts of Monte-Carlo iterations. Note that the fluctuating speckle patterns only show the positive part of the field, and the displayed images have artificially increased brightness and contrast for visualization purposes. Visual inspection reveals that colourful speckle outliers are polluting the image for $M \leq 512$. We also studied the statistics of the rendered fluctuating speckle fields: (c) Histograms and (d) standard deviations of the $\tilde{\mathfrak{S}}$ field were computed for various values of M . The six plots (plots (c) left) display logarithmic histograms of the low intensity range of the fluctuating intensity distribution. The luminance histogram (plot (c) right) shows that the speckle field reproduces a Gaussian curve, as theoretically expected, beginning at roughly $M = 512$ samples. The shape of the histogram as well as the standard deviation stabilises at around $M = 2048$ samples, and the pattern effectively becomes fully developed and converges in statistics: Higher count of samples do not have an effect on the statistics of the pattern.

by the Monte Carlo integration process only affect the fluctuating field, which is generally of lower intensity than the mean field. Finally, the decomposition ensures that our method is agnostic to the way we compute the mean field, and essentially any method or scatter theory would work.

Some of our results (Figures 1, 5 and 7) were rendered with image-based lighting, in the form of an environment map, in addition to a natural light source. We assume that the incident light that arises due to the environment map is incoherent, and thus we do not draw a speckle field for the environment map reflections. However, our method can be applied to image-based lighting methods if the coherence properties (i.e. the coherence radius ϵ_r) of the incident light are quantified. This can be done, for example, by specifying the distance to the source for every pixel in the environment map.

Integration into a rendering framework. At its core our rendering method is simply a spectral Monte Carlo sampler, with one important caveat: We sample mutual intensities, meaning that we draw a pair of points for each evaluation. Nonetheless, the incident radiation can be sourced directly form a light source, or potentially be arriving from multiple directions and secondary sources, therefore the method is compatible with essentially any spectral ray-tracing rendering framework.

Our implementation draws wavelengths by importance sampling the spectrum of the light source. Though we have not tried other sampling strategies, a fixed spectral sampling should work as well as long as the sampling is dense enough to avoid spectral aliasing. Another potential difficulty arises due to the integration phases. As discussed, the integration phases are highly sensitive to imperfect random-number generators and it is advisable to store the drawn integration phases for future reuse.

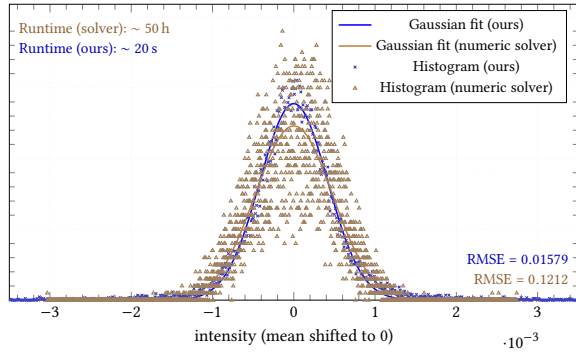


Fig. 9. Explicit methods are inefficient in reproducing the high-frequency speckle noise that arises when partially-coherent light scatters off rough surfaces: In flatland, we modelled a surface patch of length $200\ \mu\text{m}$ explicitly using a $150\,000$ pixel heightmap, drawn from a K-correlation model with $\alpha = 2$, $\beta = 8\ \mu\text{m}$. A light source is positioned directly above the surface and the incident radiation is assumed to admit a coherence radius of $\epsilon_r = 15\ \mu\text{m}$. We rendered the radiation scattered into an arc of 0.10° (centered about 6° off the normal) as a BRDF slice consisting of $20\,000$ pixels using our method (blue plot) and established methods in Fourier optics (brown plot), and the resulting histograms are plotted above. The rendering was done with 40 spectral samples. Due to discrepancies between the means produced by the generalised Harvey-Shack and the explicit numeric solver, a highpass filter is used to extract the relevant data from each of the rendered BRDFs, effectively normalizing the data to the same mean. While the explicit method produces a histogram of roughly Gaussian shape, the poor fit suggests that non-Gaussian noise (due to numeric inaccuracies and resolved surface features) is still present, despite the massive heightmap resolution.

Table 2. Rendering runtimes of our method for the different scenes that appear throughout the paper as well as the videos in the supplemental material. Rendering was done on an Intel® Core™ i9-9900K CPU. The rendering time of the Buddha is dominated by integrating the area lights.

Scene	Rendering time (per frame)
Dragon (2560×1440 , $M = 2048$, 1 light)	3 min 59 s
Buddha (1600×2560 , $M = 2048$, 2 area lights)	25 min 56 s
Teapot (1920×1080 , $M = 128$, 1 light)	6.50 s
Surface (Figure 8) (2048×2048 , $M = 16384$, 1 light)	6 min 37 s

Complexity and performance. Time-wise the rendering algorithm requires $\mathcal{O}(NM)$ operations, where N is the count of observations and M is the count of Monte Carlo samples. As M is chosen a priori, the algorithm is linear with respect to resolution and frame count of the rendered sequence. Space-wise we only store the drawn integration phases, as described in Subsection 6.1. This requires M -bits of space for the mean and spatial phases, as well as $(M \times H)$ -bits for the angular phases, where H is the count of image elements. To perform temporal adaptation we could also be required to redraw some of the angular integration phases. However, once a frame

has been rendered, we can draw the next frame’s phases (using the procedure outlined in Equation (44)) and discard the rendered frame’s phases. Therefore, with M and H assumed constant, the space requirements can also be considered to be constant.

Nonetheless, evaluating \mathcal{X} (Equation (32)) can be expensive in practice, and a large number of Monte Carlo samples might be needed. All of our results have been rendered on a desktop CPU and the rendering times are summarised in Table 2.

7.1 Validation and convergence

The speckle covariance decomposition into $\tilde{C}_{\text{spatial}}$ and \tilde{C}_{ang} , as described in Subsection 5.2, is the driving formalism behind our rendering algorithm. We have numerically tested the correctness of this decomposition on a variety of surfaces, see Figure 4. In addition, the Monte Carlo rendering framework described in Subsection 6.1 has been formally proven to draw fully developed speckle patterns with the desired statistics (see Corollary 6.1). We empirically test the convergence of the Monte Carlo integration in Figure 8. Convergence is slow and the speckle does not converge to a singular pattern, however it does converge in statistics: That is, after a certain number of samples, M , the statistics of the rendered field remain constant. This, however, can give rise to a small amount of residual temporal noise that can be observed in the rendered sequences (see our supplemental material). Nonetheless, our results (see Figure 6) are founded on optically rigorous theory and produce plausible speckle with consistent statistics.

We also compare our method against the explicit surface diffraction rendering method of Yan et al. [2018], see Figure 10. The rendering method by Yan et al. [2018] operates on a heightmap, and in order to perform the comparison we draw an explicit surface height profile. However, to capture the statistics of the surface, a high-resolution heightmap is needed, which results in a large integration area (roughly $250\,000$ texels) as well as very high rendering times. Despite that, visible surface features are still resolved and the statistics strongly suggest that those features are not fully developed speckle.

This is not surprising: It well-known that the (far-field) angular distribution of the scattered fields is related to the field distribution on a diffracting aperture via a Fourier transform [Goodman 2017]. A scattering surface can be readily considered as such an aperture under the Born approximation, i.e. when the interaction of the scattered radiation with the surface is negligible (away from steep angles). Thus, it is evident that the angular frequency resolution of the scattered radiation is proportional to the spatial extend and sampling frequency of the surface heightmap, and simple calculations show that heightmaps with extremely fine resolution are needed to reproduce speckle. In an attempt to reproduce optical speckle with an explicit method, we work in flatland and draw a high-resolution $150\,000$ pixel heightmap. We compute the scattered field by taking the discrete Fourier transform of the incident field distribution on the surface. To simulate partially-coherent radiation, the light source is assumed to consist of many uncorrelated radiators. Assigning random initial phase to each such radiator can be thought of as applying a random phase screen to a coherent source—an established technique for simulating partially-coherent radiation [Xiao




Fig. 10. (a) We compare against Yan et al. [2018] by drawing an explicit surface heightmap that conforms to the K-correlation model statistics with $\alpha = 2$, $\beta = 12 \mu\text{m}$, and we use their method to render the drawn surface. The drawn surface is visualised inside the cyan frame in Figure (a), and the histogram inset confirms the point-wise Gaussianity of the heightmap. The dimensions of the rendered surface patch are $25 \text{ mm} \times 25 \text{ mm}$, with the viewer located about 250 mm above the centre and the surface illuminated by a directional light source that admits a coherence radius of $15 \mu\text{m}$. To capture the surface statistics, the density of the heightmap is set to 30 texels per $1 \mu\text{m}$, nonetheless surface features are clearly resolved when rendering at a resolution of 512×512 . Due to the high-density, rendering of that image takes 40 h, and we did not render with higher density heightmaps. Note, as that the surface subtends a solid angle of 0.01 sr , it is unlikely that those features are speckle. Under an identical configuration, we also rendered the surface using our method. The difference in colour is attributed to different shading models. As expected, speckle is produced when the resolution of the imaging device increases, and the contributions to each pixel are more coherent. In contrast to Yan et al. [2018], our method does not suffer from aliasing at lower resolutions. (b) To study the statistics of the produced patterns, we artificially decompose the 512×512 explicit surface rendering into mean (by averaging over 128×128 moving window) and speckle (top). We then plot the histograms of the total intensities as well as the speckle intensities (bottom). Note that while our rendering clearly follows Gaussian statistics, as expected, the explicit surface renderings do not. This, in addition to the visible features, suggests that despite the high-density heightmap the rendered explicit surface images do not exhibit fully developed speckle.

and Voelz 2006]—and is a very good model for a natural spontaneous emission light source. However, in practice such sources consist of very many independent radiators, and due to the resolution both of the heightmap and the rendered BRDF slice, it is computationally feasible to simulate only hundreds of such sources. We used 128 independent radiators, uniformly spatially distributed inside a light source of radius $1 \mu\text{m}$. The average scattered intensity computed by the numerical method described above is not expected to agree with the average intensity computed using our method (i.e., as we use the generalised Harvey-Shack model). However, extracting the high-frequency details from the computed fields via a highpass filter yields some agreement in the noise variance between both methods, see Figure 9. The noise that arises in the numerically computed field is not fully Gaussian (fit RMSE of 0.1212). It is reasonable to expect that with additional radiators or with a greater heightmap resolution the noise histogram will tend closer to a Gaussian shape, but we did not investigate further. Rendering using this explicit method takes over 50 hours, while our method requires about 20 seconds.

We conclude that such explicit methods are not suitable for the reproduction of high-frequency details that arise on scatter by rough surfaces. Optical speckle is developed due to (essentially random) reflections by tiny details, such as the conducting electrons at a

metal’s surface, but the exact topology of those nanoscale structures and particles is of little importance. It is merely their existence that gives rise to the random phasor sums that are speckle. The summation of these speckle patterns, reflected from different areas of the surface under partially-coherent light, is what ensures the Gaussianity of the final speckle pattern. The above analysis also makes our decomposition into a mean field—dominated by the first-order diffraction from low-frequency surface details—and fluctuating field—dominated by scatter by high-frequency details—pleasing from a physical perspective: they arise due to rather different processes. On the other hand, explicitly modelling those high-frequency surface details is neither feasible nor methodologically sound (for our purposes).

7.2 Limitations

We reiterate the assumptions that were made throughout the paper:

- (1) **Moderate surface roughness:** As formalised by Equation (10). This also implies that the specular lobe is ignored.
- (2) **Weakly-coherent light:** Under strong spatial coherence, i.e. $g \gg 0$ (see Equation (33)), the decomposition of the autocovariance into spatial and angular autocovariance becomes less accurate. For coherent light sources, the approximation for g will have to be reworked. However, as our focus is on

material appearance reproduction under natural lighting, we leave this for future work.

It is of theoretical interest to note that while our method draws a speckle pattern that corresponds to some realization of the surface, we have no means of deducing which specific surface realization induces the drawn pattern. Put differently, it would be of interest to be able to perform inverse rendering by relating the PSD phases, required to draw an explicit surface (see Equation (7)), to the drawn speckle pattern. This important theoretical problem is an active area of research [Dong et al. 2015] and is left for future work.

It is also important to note that we ignore multiple scattering and subsurface scattering when integrating the fluctuating intensity. Contributions from multiply scattered radiation will typically exhibit weaker optical coherence, as well as intensity, than the directly scattered fields. We, therefore, expect multiple scattering effects to mainly reduce the speckle contrast—the ratio between the standard deviation of the fluctuating intensity and the mean intensity. Nonetheless, the question of how does multiple scattering influence the statistics of the speckle pattern is an interesting open problem, and is also left for future work.

8 ACKNOWLEDGEMENTS

The authors would like to thank Simon Sherwin for the permission to use the equirectangular environment maps used in the renderings.

REFERENCES

- Milton Abramowitz. 1974. *Handbook of Mathematical Functions, With Formulas, Graphs, and Mathematical Tables*. Dover Publications, Inc., USA.
- E Akkermans. 2007. *Mesoscopic physics of electrons and photons*. Cambridge University Press, Cambridge.
- Chen Bar, Marina Alterman, Ioannis Gkioulekas, and Anat Levin. 2019. A Monte Carlo Framework for Rendering Speckle Statistics in Scattering Media. *ACM Trans. Graph.* 38, 4, Article 39 (July 2019), 22 pages. <https://doi.org/10.1145/3306346.3322950>
- Chen Bar, Ioannis Gkioulekas, and Anat Levin. 2020. Rendering near-field speckle statistics in scattering media. *ACM Transactions on Graphics* 39, 6 (Nov 2020), 1–18. <https://doi.org/10.1145/3414685.3417813>
- Michael Bass, Casimer DeCusatis, Jay Enoch, Vasudevan Lakshminarayanan, Guifang Li, Carolyn Macdonald, Virendra Mahajan, and Eric Van Stryland. 2010. *Handbook of Optics, Third Edition Volume I: Geometrical and Physical Optics, Polarized Light, Components and Instruments(Set)* (3 ed.). McGraw-Hill, Inc., New York, NY, USA.
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph.* 36, 4, Article 65 (July 2017), 14 pages. <https://doi.org/10.1145/3072959.3073620>
- Stephan Bergmann, Mahsa Mohammadikaji, Stephan Irgefried, Heinz Wörn, Jürgen Beyerer, and Carsten Dachsbacher. 2016. A Phenomenological Approach to Integrating Gaussian Beam Properties and Speckle into a Physically-Based Renderer. In *VMV*.
- P. Billingsley. 1995. *Probability and Measure*. Wiley.
- Max Born and Emil Wolf. 1999. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, Cambridge New York.
- Xavier Chermain, Frédéric Claux, and Stéphane Mérillou. 2019. Glint Rendering based on a Multiple-Scattering Patch BRDF. *Computer Graphics Forum* 38, 4 (2019), 27–37. <https://doi.org/10.1111/cgf.13767> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13767>
- Narak Choi and James E. Harvey. 2013. Numerical validation of the generalized Harvey–Shack surface scatter theory. *Optical Engineering* 52, 11 (2013), 1 – 15. <https://doi.org/10.1117/1.OE.52.11.115103>
- E. L. Church, P. Z. Takacs, and T. A. Leonard. 1990. The Prediction Of BRDFs From Surface Profile Measurements. In *Scatter from Optical Components*, John C. Stover (Ed.), Vol. 1165. International Society for Optics and Photonics, SPIE, 136 – 150. <https://doi.org/10.1117/12.962842>
- Dan Cjoc, Sara Finaurini, Pavel Livshits, Eran Gur, Alon Shapira, Vicente Mico, and Zeev Zalevsky. 2012. Toward fast malaria detection by secondary speckle sensing microscopy. *Biomed. Opt. Express* 3, 5 (May 2012), 991–1005. <https://doi.org/10.1364/BOE.3.000991>
- J Christopher Dainty. 2013. *Laser speckle and related phenomena*. Vol. 9. Springer science & business Media.
- B. Dhanasekar and B. Ramamoorthy. 2008. Digital speckle interferometry for assessment of surface roughness. *Optics and Lasers in Engineering* 46, 3 (2008), 272 – 280. <https://doi.org/10.1016/j.optlaseng.2007.09.003>
- D. S. Dhillon, J. Teyssier, M. Single, I. Gaponenko, M. C. Milinkovitch, and M. Zwicker. 2014. Interactive Diffraction from Biological Nanostructures. *Computer Graphics Forum* 33, 8 (2014), 177–188. <https://doi.org/10.1111/cgf.12425> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12425>
- Michael G. Dittman. 2006. K-correlation power spectral density and surface scatterer model. In *Optical Systems Degradation, Contamination, and Stray Light: Effects, Measurements, and Control II*, O. Manuel Uy, John C. Fleming, and Michael G. Dittman (Eds.), Vol. 6291. International Society for Optics and Photonics, SPIE, 226 – 237. <https://doi.org/10.1117/12.678320>
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2015. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Transactions on Graphics* 35, 1 (Dec 2015), 1–13. <https://doi.org/10.1145/2815618>
- Julie Dorsey, Holly Rushmeier, and Francois Sillion. 2007. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Donald D. Duncan and Sean J. Kirkpatrick. 2008. Algorithms for simulation of speckle (laser and otherwise). In *Complex Dynamics and Fluctuations in Biomedical Photonics V*, Valery V. Tuchin and Lihong V. Wang (Eds.), Vol. 6855. International Society for Optics and Photonics, SPIE, 23 – 30. <https://doi.org/10.1117/12.760518>
- V. Falster, A. Jarabo, and J. R. Frisvad. 2020. Computing the Bidirectional Scattering of a Microstructure Using Scalar Diffraction Theory and Path Tracing. *Computer Graphics Forum* 39, 7 (Oct 2020), 231–242. <https://doi.org/10.1111/cgf.14140>
- Shechao Feng, Charles Kane, Patrick A. Lee, and A. Douglas Stone. 1988. Correlations and Fluctuations of Coherent Wave Transmission through Disordered Media. *Phys. Rev. Lett.* 61 (Aug 1988), 834–837. Issue 7. <https://doi.org/10.1103/PhysRevLett.61.834>
- Luis E. Gamboa, Jean-Philippe Guertin, and Derek Nowrouzezahrai. 2018. Scalable Appearance Filtering for Complex Lighting Effects. *ACM Trans. Graph.* 37, 6, Article 277 (Dec. 2018), 13 pages. <https://doi.org/10.1145/3272127.3275058>
- Tim Golla and Reinhard Klein. 2018. Interactive Interpolation of Metallic Effect Car Paints. In *Vision, Modeling and Visualization*, Fabian Beck, Carsten Dachsbacher, and Filip Sadlo (Eds.). The Eurographics Association. <https://doi.org/10.2312/vm.20181248>
- Joseph Goodman. 2017. *Introduction to Fourier Optics*.
- Joseph W. Goodman. 2007. *Speckle Phenomena in Optics: Theory and Applications*.
- Bronius Grigelionis. 2009. On the Wick theorem for mixtures of centered Gaussian distributions. *Lithuanian mathematical journal* 49, 4 (2009), 372.
- Ibón Guillén, Julio Marco, Diego Gutierrez, Wenzel Jakob, and Adrian Jarabo. 2020. A General Framework for Pearlescent Materials. *ACM Transactions on Graphics* 39, 6 (2020). <https://doi.org/10.1145/3414685.3417782>
- Jie Guo, Yanjun Chen, Yanwen Guo, and Jingui Pan. 2018. A Physically-based Appearance Model for Special Effect Pigments. *Computer Graphics Forum* 37, 4 (2018), 67–76. <https://doi.org/10.1111/cgf.13476> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13476>
- Yu Guo, Miloš Hašan, and Shuang Zhao. 2019. Position-free monte carlo simulation for arbitrary layered BSDFs. *ACM Transactions on Graphics* 37, 6 (Jan 2019), 1–14. <https://doi.org/10.1145/3272127.3275053>
- James Harvey, Andrey Krywonos, and Cynthia I. Vernold. 2007. Modified Beckmann-Kirchhoff scattering model for rough surfaces with large incident and scattering angles. *optical engineering* 46, 7 (2007), 078002.
- James E. Harvey, Narak Choi, Sven Schroeder, and Angela Duparré. 2012. Total integrated scatter from surfaces with arbitrary roughness, correlation widths, and incident angles. *Optical Engineering* 51, 1 (2012), 1 – 12. <https://doi.org/10.1117/1.OE.51.1.013402>
- James E. Harvey and Richard N. Pfisterer. 2016. Evolution of the transfer function characterization of surface scatter phenomena. In *Reflection, Scattering, and Diffraction from Surfaces V*, Leonard M. Hanssen (Ed.), Vol. 9961. International Society for Optics and Photonics, SPIE, 102 – 118. <https://doi.org/10.1117/12.2237083>
- Qingsheng He, Jinnan Wang, Peikun Zhang, Jiangang Wang, Minxian Wu, and Guofan Jin. 2003. Dynamic speckle multiplexing scheme in volume holographic data storage and its realization. *Optics Express* 11, 4 (2003), 366–370.
- Nicolas Holzschuch and Romain Pacanowski. 2017. A Two-scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Trans. Graph.* 36, 4, Article 66 (July 2017), 12 pages. <https://doi.org/10.1145/3072959.3073621>
- L. Isserlis. 1918. On a Formula for the Product-Moment Coefficient of any Order of a Normal Frequency Distribution in any Number of Variables. *Biometrika* 12, 1-2 (11 1918), 134–139. <https://doi.org/10.1093/biomet/12.1-2.134>
- Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete stochastic microfacet models. *ACM Transactions on Graphics* 33, 4 (Jul 2014), 1–10. <https://doi.org/10.1145/2601097.2601186>
- Tom Kneiphof, Tim Golla, and Reinhard Klein. 2019. Real-time Image-based Lighting of Microfacet BRDFs with Varying Iridescence. *Computer Graphics Forum* 38, 4 (2019), 77–85. <https://doi.org/10.1111/cgf.13772>

- arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13772
- Andrey Krywonos. 2006. *Predicting surface scatter using a linear systems formulation of non-paraxial scalar diffraction*. Ph.D. Dissertation. University of Central Florida.
- Ann M. Lanari, Samuel D. Butler, Michael Marciniak, and Mark F. Spencer. 2017. Wave optics simulation of statistically rough surface scatter. In *Earth Observing Systems XXII*, James J. Butler, Xiaoxiong (Jack) Xiong, and Xingfa Gu (Eds.), Vol. 10402. International Society for Optics and Photonics, SPIE, 329 – 337. <https://doi.org/10.1117/12.2274234>
- Joo Ho Lee, Adrian Jarabo, Daniel S. Jeon, Diego Gutierrez, and Min H. Kim. 2018. Practical Multiple Scattering for Rough Surfaces. *ACM Trans. Graph.* 37, 6, Article 275 (Dec. 2018), 12 pages. <https://doi.org/10.1145/3272127.3275016>
- Anat Levin, Daniel Glasner, Ying Xiong, Fredo Durand, Bill Freeman, Wojciech Matusik, and Todd Zickler. 2013. Fabricating BRDFs at High Spatial Resolution Using Wave Optics. *ACM Transaction of Graphics* (2013).
- Joakim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger. 2012. BRDF Models for Accurate and Efficient Rendering of Glossy Surfaces. *ACM Trans. Graph.* 31, 1, Article 9 (Feb. 2012), 14 pages. <https://doi.org/10.1145/2077341.2077350>
- L. Mandel and E. Wolf. 1995. *Optical Coherence and Quantum Optics*. Cambridge University Press, Cambridge [England]; New York.
- A. A. Maradudin and D. L. Mills. 1975. Scattering and absorption of electromagnetic radiation by a semi-infinite medium in the presence of surface roughness. *Phys. Rev. B* 11 (Feb 1975), 1392–1415. Issue 4. <https://doi.org/10.1103/PhysRevB.11.1392>
- Johannes Meng, Marios Pappas, Ralf Habel, Carsten Dachsbacher, Steve Marschner, Markus Gross, and Wojciech Jarosz. 2015. Multi-scale modeling and rendering of granular materials. *ACM Transactions on Graphics* 34, 4 (Jul 2015), 1–13. <https://doi.org/10.1145/2766949>
- J. V. Michalowicz, J. M. Nichols, F. Bucholtz, and C. C. Olson. 2009. An Isserlis' Theorem for Mixed Gaussian Variables: Application to the Auto-Bispectral Density. *Journal of Statistical Physics* 136, 1 (01 Jul 2009), 89–102. <https://doi.org/10.1007/s10955-009-9768-3>
- Mervin E. Muller. 1959. A Note on a Method for Generating Points Uniformly on N-dimensional Spheres. *Commun. ACM* 2, 4 (April 1959), 19–20. <https://doi.org/10.1145/377939.377946>
- Thomas Müller, Marios Pappas, Markus Gross, Wojciech Jarosz, and Jan Novák. 2016. Efficient rendering of heterogeneous polydisperse granular media. *ACM Transactions on Graphics* 35, 6 (Nov 2016), 1–14. <https://doi.org/10.1145/2980179.2982429>
- Wai Kit Addy Ngan. 2006. *Acquisition and modeling of material appearance*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- Boris Raymond, Gaël Guennebaud, and Pascal Barla. 2016. Multi-Scale Rendering of Scratched Materials Using a Structured SV-BRDF Model. *ACM Trans. Graph.* 35, 4, Article 57 (July 2016), 11 pages. <https://doi.org/10.1145/2897824.2925945>
- Nicholas J Scott, Steve B Howell, Elliott P Horch, and Mark E Everett. 2018. The NN-Explore Exoplanet Stellar Speckle Imager: Instrument Description and Preliminary Results. *Publications of the Astronomical Society of the Pacific* 130, 987 (2018), 054502.
- A.A. Scribot. 1974. First-order probability density functions of speckle measured with a finite aperture. *Optics Communications* 11, 3 (Jul 1974), 238–241. [https://doi.org/10.1016/0030-4018\(74\)90171-0](https://doi.org/10.1016/0030-4018(74)90171-0)
- Frank Siewert, Heiner Lammert, and Thomas Zeschke. 2008. *The Nanometer Optical Component Measuring Machine*. Springer Berlin Heidelberg, Berlin, Heidelberg, 193–200. https://doi.org/10.1007/978-3-540-74561-7_11
- Brandon M. Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. 2017. CoLux: Multi-Object 3D Micro-Motion Analysis Using Speckle Imaging. *ACM Trans. Graph.* 36, 4, Article 34 (July 2017), 12 pages. <https://doi.org/10.1145/3072959.3073607>
- Brandon M. Smith, Matthew O'Toole, and Mohit Gupta. 2018. Tracking Multiple Objects Outside the Line of Sight using Speckle Imaging. In *2018 CVF and IEEE Conference on Computer Vision and Pattern Recognition*.
- Jos Stam. 1999. Diffraction shaders. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99*. ACM Press. <https://doi.org/10.1145/311535.311546>
- S. Steinberg. 2019. Analytic Spectral Integration of Birefringence-Induced Iridescence. *Computer Graphics Forum* 38, 4 (Jul 2019), 97–110. <https://doi.org/10.1111/cgf.13774>
- John C. Stover. 2012. *Optical Scattering: Measurements and Analysis, Third Edition*. Society of Photo-Optical Instrumentation Engineers. <https://doi.org/10.1117/3.975276>
- Antoine Toisoul, Daljit Singh Dhillon, and Abhijeet Ghosh. 2018. Acquiring Spatially Varying Appearance of Printed Holographic Surfaces. *ACM Trans. Graph.* 37, 6, Article 272 (Dec. 2018), 16 pages. <https://doi.org/10.1145/3272127.3275077>
- Antoine Toisoul and Abhijeet Ghosh. 2017. Practical Acquisition and Rendering of Diffraction Effects in Surface Reflectance. *ACM Trans. Graph.* 36, 5, Article 166 (July 2017), 16 pages. <https://doi.org/10.1145/3012001>
- Leung Tsang, Jin Au Kong, and Kung-Hau Ding. 2002. *Scattering of Electromagnetic Waves: Theories and Applications*. John Wiley & Sons, Ltd. <https://doi.org/10.1002/0471224286.ch3>
- Z. Velinov, S. Werner, and M. B. Hullin. 2018. Real-Time Rendering of Wave-Optical Effects on Scratched Surfaces. *Computer Graphics Forum* 37, 2 (2018), 123–134. <https://doi.org/10.1111/cgf.13347>
- arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13347
- Andrea Weidlich and Alexander Wilkie. 2007. Arbitrarily Layered Micro-Facet Surfaces. In *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia (Perth, Australia) (GRAPHITE '07)*. Association for Computing Machinery, New York, NY, USA, 171–178. <https://doi.org/10.1145/1321261.1321292>
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin. 2017. Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure. *ACM Trans. Graph.* 36, 6, Article 207 (Nov. 2017), 14 pages. <https://doi.org/10.1145/3130800.3130840>
- Xifeng Xiao and David Voelz. 2006. Wave optics simulation approach for partial spatially coherent beams. *Optics Express* 14, 16 (2006), 6986. <https://doi.org/10.1364/oe.14.006986>
- Ling-Qi Yan, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi. 2014. Rendering Glints on High-Resolution Normal-Mapped Specular Surfaces. *ACM Trans. Graph.* 33, 4, Article 116 (July 2014), 9 pages. <https://doi.org/10.1145/2601097.2601155>
- Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-Normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Trans. Graph.* 35, 4, Article 56 (July 2016), 9 pages. <https://doi.org/10.1145/2897824.2925915>
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2018)* 37, 4 (2018).

A K-CORRELATION MODEL

While Gaussianity is a common assumption, most physical surfaces do not admit a Gaussian autocorrelation function and a better model is required. The *K-Correlation* model of the surface PSD is a common analytic expression to a physically reasonable surface power spectrum [Church et al. 1990; Stover 2012]. The model admits multiple physically intuitive control parameters to describe different surface finishes. This allows to model different conventional and fractal surfaces with a single analytic expression [Dittman 2006]:

$$P_2(f) = \frac{\sigma_{\text{rel}}^2 \beta^2}{1 - (4\beta^2 + 1)^{\frac{1-\alpha}{2}}} \frac{\alpha - 1}{2\pi \left(1 + \beta^2 \|f\|^2\right)^{\frac{\alpha+1}{2}}}, \quad (48)$$

which is a remapping from the more common form given by Church et al. [1990], derived by plugging-in the surface relative roughness σ_{rel} (Equation (5)). The remapping is in similar fashion to Holzschuch and Pacanowski [2017], except that we use the bandwidth-limited relative roughness. Note that this introduces cumbersome wavelength-dependence into the PSD (as σ_{rel} is wavelength dependent), and to avoid this dependence in P_2 we set $\lambda = 0.50 \mu\text{m}$ in Equation (48). The control parameters are then α , the slope of the PSD in log-log space, and β , which is related to the correlation length l_{cor} (Equation (6)). The autocorrelation function of the K-Correlation PSD is then [Stover 2012]:

$$C_S(\mathbf{x}) = \frac{2}{\Gamma\left(\frac{\alpha}{2}\right)} \left(\pi \frac{\|\mathbf{x}\|}{\beta}\right)^{\frac{\alpha}{2}} \cdot K_{\frac{\alpha}{2}}\left(2\pi \frac{\|\mathbf{x}\|}{\beta}\right), \quad (49)$$

where Γ is the Gamma function and K_ν is the modified Bessel function of the second kind. Due to its popularity and flexibility, the K-Correlation PSD has been used in all of our results.

Importance sampling the K-Correlation C_S function. A difficulty that will arise later is the need to importance sample the function C_S . In the case of the K-Correlation model, the behaviour of the surface autocorrelation is dictated by the Bessel function K_ν , which admits an asymptotic expansion [Abramowitz 1974]:

$$K_\nu(x) \sim \sqrt{\frac{\pi}{2x}} e^{-x} \left[1 + \frac{1 - 4\nu^2}{8x} + \mathcal{O}(x^{-2})\right].$$

Using the above the autocorrelation function can be written as the following expression

$$C_S(\mathbf{x}) \sim \frac{\pi^{\frac{\alpha}{2}}}{\Gamma(\frac{\alpha}{2})} \left(\frac{\|\mathbf{x}\|}{\beta} \right)^{\frac{\alpha-1}{2}} e^{-2\pi \frac{\|\mathbf{x}\|}{\beta}}, \quad (50)$$

which is dominated by the exponent. Thus for $\|\mathbf{x}\| \gg 0$ we can deduce that the K-correlation surface autocorrelation function is distributed similarly to an exponential distribution, i.e. $C_S \sim e^{-2\pi \frac{\|\mathbf{x}\|}{\beta}}$.

B STATISTICS OF PARTIALLY-COHERENT SPECKLE

As stated in Subsection 4.1, when light is monochromatic and fully-coherent (spatially and temporally), the intensity of the speckle follows negative exponential distribution, i.e. the probability density function is

$$p_{\text{coherent}}(I) = \frac{1}{2\sigma_I^2} e^{-\frac{I}{2\sigma_I^2}}. \quad (51)$$

with $I \geq 0$. For monochromatic light that exhibits partial spatial coherence, the speckle pattern can be regarded as a sum of multiple independent speckle patterns, each produced by coherent light contributions and therefore each follows negative exponential statistics. The count of such independent patterns is readily approximated as $N \simeq \lceil \frac{a_{\text{surface}}}{a_{\text{coherence}}} \rceil$, where $a_{\text{coherence}}$ is the coherence area of the light (roughly the spatial region where light remains mutually coherent and contributions are summed up on an amplitude basis) and a_{surface} is the area of the imaged surface that falls under a single image element. Then, $I_{\text{monochromatic}} = \sum_{j=1}^N I_{(j)}$ where the $I_{(j)}$ -s are the intensities of the contributing coherent speckle patterns. The sum of N negative exponential random variables is known as the *gamma* distribution of order N (or more specifically, the *Erlang* distribution with shape parameter N) and its probability density function is

$$p_{\text{monochromatic}}(I) = \frac{N^N}{\Gamma(N)I_0^N} I^{N-1} e^{-N \frac{I}{I_0}}, \quad (52)$$

where $I \geq 0$, Γ is the gamma function and for simplicity we assume that the mean intensities of the coherent patterns $I_{(j)}$ are identical and denote the total mean intensity $I_0 = N \langle I_{(j)} \rangle$. In speckle optics, this is known as a “sum of independent speckle patterns” [Goodman 2007].

When speckle is formed by polychromatic light (of partial temporal and spatial coherence), the total intensity can still be considered as a sum of independent (spatially) coherent speckle patterns but each produced by polychromatic light. Polychromatic speckle produced by spatially coherent light (e.g. multi-mode laser) is discussed in detail by Dainty [2013]. When the spectrum is continuous and can be assumed to consist of many “cells”—correlated spectral regions—the following familiar solution to the probability density function was obtained by Scribot [1974]:

$$p_{\text{polychromatic}}(I) = \frac{M^M}{\Gamma(M)I_0^M} I^{M-1} e^{-M \frac{I}{I_0}}, \quad (53)$$

where $I \geq 0$, M is the count of such pair-wise uncorrelated spectral cells and I_0 is defined as the total mean intensity over all cells, analogously to the case of monochromatic partially coherent speckle. The mean intensities of the speckles produced by the cells are assumed




Fig. 11. Plots of the probability density function of the Erlang distribution for different shape parameters (and scaled by a constant for visualization). The dashed plots are Gaussian fits. The RMSE of the fits are 6×10^{-3} , 3×10^{-3} , 2×10^{-3} , 5×10^{-4} , 2×10^{-4} and 1×10^{-4} for values 1,2,4,8,16 and 32 of Θ , respectively.

to be identical. We then deduce that the intensity of a polychromatic (spatially coherent) speckle pattern is also distributed as an Erlang distribution, with the physical interpretation being that such speckle is the sum of M quasi-monochromatic fully-coherent speckle patterns, pair-wise independent due to spectral decorrelation and each with a negative exponentially distributed intensity.

We can conclude now that the statistics of the intensity of polychromatic partially-coherent speckle follow an Erlang distribution with shape parameter $\Theta = N \times M$. Indeed, when the light is fully spatially-coherent, i.e. $N = 1$, and quasi-monochromatic, i.e. $M = 1$, the shape parameter is $\Theta = 1$ and the distribution reduces to the negative exponential distribution of coherent speckle intensity. It is reasonable to assume that for high-bandwidth spontaneous emission light sources M would be “not small” (though difficult to estimate). Therefore, even when N is small, i.e. the light is strongly spatially correlated, Θ would still be significant. For numerical examples of N , consider Figure 6: The Cobalt dragon is lit by a strongly-coherent light that admits a coherence radius of roughly $\epsilon_r \simeq 49.70 \mu\text{m}$ and thus coherence area of $a_{\text{coherence}} \simeq 8000 \mu\text{m}^2$. Each rendered pixel subtends $\simeq 250 \times 250 \mu\text{m}^2$ surface area, thus $N \simeq 8$. The Copper dragon, on the other hand, is lit by a weakly-coherent light admitting $N \simeq 400$.

It is well-known that the Erlang distribution tends towards a Gaussian distribution as the shape parameter increases. We plot the Erlang distribution for different values of the shape parameter Θ , and fit Gaussians to each (see Figure 11) showing very good agreement even for shape parameters as low as $\Theta \geq 16$. A more quantitative argument, as well as an analytic error upper limit, could be deduced (e.g., as an immediate application of the Berry–Esseen inequality). However, as exact values of N and especially M are hard to predict, this is of little value. We satisfy ourselves then by reiterating that the intensity distribution of a weakly-coherent speckle pattern is the sum of tens pair-wise independent (spectrally

or spatially) speckle patterns, and is statistically approximated very well by a Gaussian.

C A PHASOR'S PHASES IS UNIFORMLY DISTRIBUTED

In this appendix we briefly prove the Corollary 5.1.

COROLLARY C.1. *Treating $h(\mathbf{x})$ as a random variable, a phasor's phase can be regarded as uniformly distributed on $(-\pi, +\pi]$, that is*

$$\arg(u_j) \sim \mathcal{U}(-\pi, +\pi].$$

PROOF. Let $\sigma[\cdot]$ denote the standard deviation, then

$$\sigma[\arg(u_j)] = \sigma[\Phi(\mathbf{x}_j, \mathbf{k}^s, \mathbf{k}^i)] \geq \sigma[h(\mathbf{x})] \frac{2\pi}{\lambda} \gg \pi,$$

where we assume that $\langle \mathbf{n}, \mathbf{k}^s - \mathbf{k}^i \rangle \geq 2\pi\lambda^{-1}$, that is, we ignore grazing angles. Therefore the standard deviation of the principal phase value of the phasor is large, and we can assume that the phase takes any value in $(-\pi, +\pi]$ with equal probability. \square

D THE RAYLEIGH-RICE POLARIZATION FACTOR

The Fresnel power term Q is known as the Rayleigh-Rice polarization factor in literature [Maradudin and Mills 1975; Stover 2012]. It can be thought of as an approximative analogue to the Fresnel term when the reflecting (or more appropriately–scattering) surface cannot be considered as a long (with respect to wavelength) perfectly smooth slab. It reduces to the common Fresnel term F (in Equation (8)) when evaluated along the specular direction (i.e. $\mathbf{o} = \text{refl}(\mathbf{i})$). Because Q is polarization-dependent, Krywonos [2006] suggested using it in-place of the reflectance term in the Harvey-Shack BRDF expression to “quasi-vectorize” the theory, that is admit awareness of polarization effects, resulting in Equation (8).

Because Q depends on polarization, it consists of four complex terms \mathbf{f} that relate the incident and scattered field amplitudes for the polarized components, that is

$$Q = |\mathbf{f}_{ss}|^2 + |\mathbf{f}_{pp}|^2 + |\mathbf{f}_{sp}|^2 + |\mathbf{f}_{ps}|^2. \quad (54)$$

Those coefficients are given by Stover [2012] and are listed here in full:

$$\begin{aligned} \mathbf{f}_{ss} &= \frac{(\eta - 1) \cos \phi_s}{\left(\cos \theta_i + \sqrt{\eta - \sin^2 \theta_i} \right) \left(\cos \theta_s + \sqrt{\eta - \sin^2 \theta_s} \right)} \\ \mathbf{f}_{pp} &= \frac{(\eta - 1) \left(\sqrt{\eta - \sin^2 \theta_i} \sqrt{\eta - \sin^2 \theta_s} \cos \phi_s - \eta \sin \theta_i \sin \theta_s \right)}{\left(\eta \cos \theta_i + \sqrt{\eta - \sin^2 \theta_i} \right) \left(\eta \cos \theta_s + \sqrt{\eta - \sin^2 \theta_s} \right)} \\ \mathbf{f}_{sp} &= \frac{(\eta - 1) \sqrt{\eta - \sin^2 \theta_s} \sin \phi_s}{\left(\cos \theta_i + \sqrt{\eta - \sin^2 \theta_i} \right) \left(\eta \cos \theta_s + \sqrt{\eta - \sin^2 \theta_s} \right)} \\ \mathbf{f}_{ps} &= \frac{(\eta - 1) \sqrt{\eta - \sin^2 \theta_i} \sin \phi_s}{\left(\eta \cos \theta_i + \sqrt{\eta - \sin^2 \theta_i} \right) \left(\cos \theta_s + \sqrt{\eta - \sin^2 \theta_s} \right)} \end{aligned} \quad (55)$$

where θ_i, θ_s are the incident and scatter angles, respectively, ϕ_s is the azimuth between the plane of incidence and scatter direction and η is the (possibly complex) refractive-index ratio between the destination medium and the source medium. The subscripts denote first the source polarization component and, second, the destination

component, e.g., \mathbf{f}_{sp} is the scatter term for an s-polarized field to a p-polarized field.

Clearly, when $\phi_s = 0$, that is the scatter direction lies in the incidence plane, we observe that the cross-terms vanish, i.e. $\mathbf{f}_{sp} = \mathbf{f}_{ps} = 0$.

E EXTENSION TO ISSERLIS' THEOREM

In this appendix we prove Theorem 5.2. As before, we denote $X = [X_1, \dots, X_{2m}]^T = [a_1 e^{i\theta_1}, \dots, a_{2m} e^{i\theta_{2m}}]^T$ to be a multivariate complex random vector of dimension $2m$ ($m > 1$) such that the magnitudes a_j are fixed real values and the phases are uniformly distributed random variables, i.e. $\theta_j \sim \mathcal{U}(-\pi, +\pi]$.

Denote the d -dimensional unit sphere as \mathcal{S}^{d-1} . It is a well-known result in computer graphics and topology, first presented by Muller [1959], that uniformly distributed values on \mathcal{S}^{d-1} can be drawn by drawing and normalizing a d -dimensional Gaussian zero-mean vector. That is, each X_j can be written as

$$X_j = a_j \frac{R_j + iI_j}{\sqrt{R_j^2 + I_j^2}}, \quad (56)$$

where $\forall j \rightarrow R_j, I_j \sim \mathcal{N}(0, 1)$ and $\text{cov}[R_j, I_j] = 0$, i.e. R_j, I_j are real standard Gaussian random variables such that each pair is independent. The independence arises due to the circular symmetry of each X_j .

Then, it can be trivially shown by induction that the following holds

$$\begin{aligned} \prod_{j=1}^{2m} X_j &= \left(\prod_{j=1}^{2m} \frac{a_j}{\sqrt{R_j^2 + I_j^2}} \right) \\ &\cdot \sum_{j=0}^{2m} \prod_{\sigma \in A_{2m,j}} i^j I_{\sigma(1)} I_{\sigma(2)} \dots I_{\sigma(j)} R_{\sigma(j+1)} \dots R_{\sigma(2m)}, \end{aligned} \quad (57)$$

where $A_{2m,j}$ is the set of all permutations that partition a $2m$ element set into two sets of j and $2m - j$ elements, and clearly there are $|A_{2m,j}| = \binom{2m}{j}$ ways to do so.

By polar representation of a complex random variable observe that the magnitude $\sqrt{R_j^2 + I_j^2}$ is independent of R_j and I_j , that is

$$\text{cov} \left[R_j, \frac{1}{\sqrt{R_j^2 + I_j^2}} \right] = \text{cov} \left[I_j, \frac{1}{\sqrt{R_j^2 + I_j^2}} \right] = 0$$

(readily verified by direct computation). Furthermore, the magnitudes of X_j are independent of each other, by assumption. Thus, by taking the expected value on each side of Equation (57), applying Isserlis' theorem and recombining the terms the desired result follows immediately:

$$\mathbb{E} \left[\prod_{j=1}^{2m} a_j e^{i\theta_j} \right] = \left(\prod_{j=1}^{2m} a_j \right) \cdot \sum_{p \in P_{2m}^2} \prod_{\{j,l\} \in p} \mathbb{E} \left[e^{i(\theta_j + \theta_l)} \right]. \quad (58)$$

Likewise, we also deduce that the odd moments vanish.

Note that this theorem also holds when the magnitudes are Gaussian random variables, such that the phases and the magnitudes are independent.

F DERIVATION OF SECOND-ORDER MOMENT C_2

Under the same notation as in Subsection 5.1 and the assumption of Gaussianity of the height profile h we derive in this appendix an expression for the second-order moment C_2 of a couple of observed phasors u_1 and u_2 .

Let $v_{\mathbf{x}_1, \mathbf{x}_2}$ be the Gaussian joint probability density, $\sigma = \sigma_{\text{rel}}$ the surface height standard deviation and $\rho = C_S(\mathbf{x}_1, \mathbf{x}_2)$ the surface autocorrelation between the two points, then

$$\begin{aligned} \langle u_1 u_2^* \rangle &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} u_1 u_2^* v_{\mathbf{x}_1, \mathbf{x}_2}(h_1, h_2) dh_1 dh_2 = \\ &= \frac{A}{2\pi\sigma^2\sqrt{1-\rho^2}} \int_{-\infty}^{\infty} e^{-ih_1 \langle \mathbf{n}, \mathbf{k}^s_1 - \mathbf{k}^i_1 \rangle} e^{-\frac{h_1^2}{2\sigma^2(1-\rho^2)}} dh_1 \cdot \\ &\quad \cdot \int_{-\infty}^{\infty} e^{ih_2 \langle \mathbf{n}, \mathbf{k}^s_2 - \mathbf{k}^i_2 \rangle} e^{-\frac{h_2^2 - 2\rho h_1 h_2}{2\sigma^2(1-\rho^2)}} dh_2 = \\ &= \frac{A}{\sqrt{2\pi}\sigma^2} \int_{-\infty}^{\infty} e^{-ih_1 (\langle \mathbf{n}, \mathbf{k}^s_2 - \mathbf{k}^i_2 \rangle - \rho \langle \mathbf{n}, \mathbf{k}^s_1 - \mathbf{k}^i_1 \rangle)} e^{-\frac{h_1^2}{2\sigma^2}} dh_1 = \\ &= A e^{-\frac{\sigma^2}{2} (\langle \mathbf{n}, \mathbf{k}^s_1 - \mathbf{k}^i_1 \rangle^2 + \langle \mathbf{n}, \mathbf{k}^s_2 - \mathbf{k}^i_2 \rangle^2)} e^{\rho\sigma^2 \langle \mathbf{n}, \mathbf{k}^s_1 - \mathbf{k}^i_1 \rangle \langle \mathbf{n}, \mathbf{k}^s_2 - \mathbf{k}^i_2 \rangle} \quad (59) \end{aligned}$$

where we factored out the terms of $u_1 u_2^*$ that do not depend on the integration variables:

$$A = a_1 a_2 e^{-i(\phi_1 - \phi_2)} e^{-i(\langle \mathbf{x}_1, \mathbf{k}^s_1 - \mathbf{k}^i_1 \rangle - \langle \mathbf{x}_2, \mathbf{k}^s_2 - \mathbf{k}^i_2 \rangle)}$$

The integral identity $\int_{-\infty}^{\infty} e^{cx-bx^2} dx = \sqrt{\frac{\pi}{b}} e^{\frac{c^2}{4b}}$ (for $b > 0$) was also used.

G DRAWING VALUES FROM A DISTRIBUTION VIA A MONTE-CARLO PROCESS

To show the validity of Corollary 6.1 we prove a simple but more general theorem:

THEOREM G.1. *Let (Ω, Π, μ) be a probability space consisting of a sample space Ω , set of events Π and the probability measure μ . Let $X = [x_1, x_2, \dots]^T$ be a (possibly infinite) vector of random variables, $x_j : \Omega \rightarrow \mathbb{F}$, with $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$, s.t. $\mathbf{E}[X] = 0$. The autocovariance $\text{cov}[X, X] = \int_{\Omega} X(\omega) X(\omega)^\dagger d\mu(\omega)$ can be Monte-Carlo integrated by drawing M samples $\omega_1, \dots, \omega_M$ with respect to the measure μ . Then, the autocovariance of the random vector drawn as follows*

$$Y = \frac{1}{\sqrt{M}} \sum_{m=1}^M \mathbf{p}_m X(\omega_m)$$

converges to C in L^1 , i.e. $\text{cov}[Y, Y] \xrightarrow{M \rightarrow \infty} C$, where \mathbf{p}_m are some i.i.d (of any distribution) random variables that are zero-mean and orthonormal, i.e. $\forall j, m \rightarrow \mathbf{E}[\mathbf{p}_j] = 0$ and $\mathbf{E}[\mathbf{p}_j \mathbf{p}_m] = \delta_{jm}$.

PROOF. It is easy to see that $\mathbf{E}[Y] = 0$. We examine the second-order statistics of Y :

$$\begin{aligned} \text{cov}[Y, Y] &= \frac{1}{M} \mathbf{E} \left[\sum_{j=1}^M \sum_{m=1}^M \mathbf{p}_j \mathbf{p}_m X(\omega_j) X(\omega_m)^\dagger \right] = \\ &= \frac{1}{M} \sum_{j=1}^M \sum_{m=1}^M \mathbf{E}[\mathbf{p}_j \mathbf{p}_m] \mathbf{E}[X(\omega_j) X(\omega_m)^\dagger] = \\ &= \frac{1}{M} \sum_{m=1}^M \mathbf{E}[X(\omega_m) X(\omega_m)^\dagger] \xrightarrow{M \rightarrow \infty} \text{cov}[X, X] \quad (60) \end{aligned}$$

in L^1 . \square