Validity of Wigner Distribution Function for Ray-based Imaging

Tom Cuypers[1] Roarke Horstmeyer[2] Se Baek Oh[3] Philippe Bekaert[1] Ramesh Raskar[2]
[1] Hasselt University - tUL - IBBT Expertise Centre for Digital Media
[2] Camera Culture Group MIT MediaLab MIT
[3] Mechanical Engineering MIT

Abstract

In this work we provide an introduction to the Wigner Distribution Function (WDF) using geometric optics principles. The WDF provides a useful model of wave–fields, allowing simulation of diffraction and interference effects. We attempt to explain these Fourier optics concepts to computational photography researchers by clarifying the relationship between the WDF and position–angle representations. We demonstrate how the WDF can be used to simulate diffraction effects using a light field representation and discuss its validity in the near–field, far–field, and under the paraxial approximation. Finally, we demonstrate that although the WDF representation contains negative values, any projection always yields a non-negative intensity value.

1. Introduction

In computer science, light propagation is often simulated using ray optics, or geometric optics, for simplicity and efficiency purposes. Using this representation of light traveling in a straight line, we are able to model various phenomena including light reflection, transmission and refraction. For example, raytracers are commonly used to simulate these effects in computer graphics and computational photography. Geometric optics provides convincing and realistic results as long as the size of the geometry is much larger than the wavelength of the emitted light, which is between 400nm and 700nm for visible light. Thus, this representation is well suited to solve common computational photography and computer vision problems, where real-world structures of interest typically fulfill this size criterion.

Occasionally, phenomena beyond those predicted with simple geometric optics can be observed due to the wave-like behavior of light. Most of the wave–effects of light are not typically noticeable in natural scenes due to the small size of the wavelength of visible light. But when light interacts with micro structures at scales close to its wavelength, some of these more interesting effects become apparent:

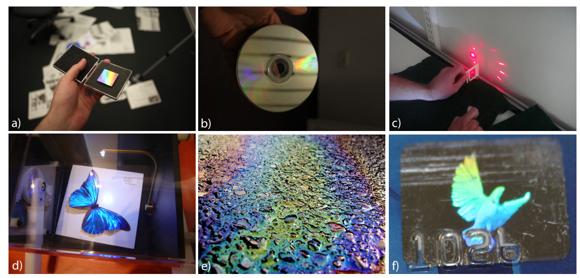

Figure 1. Examples of materials producing wave phenomena: (a) Blazed phase grating, (b) Compact disc, (c) Sinusoidal phase grating, (d) wings of a morpho butterfly, (e) oil slick on water and (f) rainbow hologram.

- Diffraction: bending of waves as they interact with obstacles in their path.
- Interference: when waves combine, they can constructively or destructively interfere.
- Dispersion: Scattering of different wavelengths into different directions, white light into a rainbow.
- Doppler effect: Frequency shifting based upon the relative traveling speed of the observer and the receiver.
- Polarization effects: Passing of only a specific oscillation component.

In this work we address diffraction, interference and dispersion effects, although doppler and polarization effects could be integrated into the same framework. They can be observed with naturally occurring surfaces such as feathers and butterfly wings, and man–made objects like holograms. Figure 1 presents a few examples of these effects.

1.1. Contribution

The goal of this work is to narrow the gap between modeling wave optics and geometric optics, by analyzing the potential and limitations of using the Wigner Distribution Function (WDF) for ray-based imaging (Section 4). Geometric optics allows fast and efficient simulation, but is not complete. Wave optics is more accurate and allows for the simulation of additional phenomena, but is more complex and less efficient. Building upon the previous research of Zhang and Levoy [32] and Oh et al. [21] who demonstrated a connection between the WDF and light fields, we address the validity and limitations of this technique for computer graphics and computational photography. We show how our interpretation of the WDF is valid in the paraxial region in both the near and far–field. In addition, even though the WDF contains negative values, we demonstrate that images rendered using this technique always end up non–negative. We further explain an effective way to visualize the WDF when it is extended to model 2D scenes, in which case it becomes a 4D function. Finally, we offer a simple code base for modeling wave–effects in a graphics setting, which takes a diffracting screen as an input and yields a light field (i.e., a set of weighted rays) as output.

2. Related Work

Light Propagation in Graphics: Ray–based rendering systems such as ray tracing are popular for rendering photorealistic images in computer graphics due to their simplicity and efficiency. They are particularly convenient for simulating reflection and refraction. The simulation is based on the well–known rendering equation [12]. In geometric optics, the function containing all light rays at every position, for every angle, for each wavelength and at each moment in time is called the plenoptic function [2]. A related representation, the light field [13], has an increasing interest in the computer science research community, and is a parametrized plenoptic function describing radiance of rays in position–angle space. It is a 4D subset of the plenoptic function, containing a ray's 2D position and it's direction in terms of two angles. A number of applications of the light field are well-studied, including their capture [13, 31], post processing to achieve image refocussing [20], 3D shape acquisition [26] and 3D display [16].

Wave Phenomena in Computational Photography: Multiple techniques have been proposed to render wave phenomena in computer graphics. Moravec proposed a wave model to render complex light transport efficiently [19]. This technique is based on phase tracking, otherwise called optical path differencing (OPD). This technique keeps track of the travel distance of a ray to calculates its phase. Ziegler et al. [34] developed a wave–based framework where complex values can be assigned for occluders to account for phase effects. They also implemented hologram rendering based on wave propagation utilizing the concept of spatial frequency [33]. Stam [24] implemented a diffraction shader based on the Kirchhoff integral for random or periodic patterns. Diffraction shaders assume both observer and lightsource to be at infinity, simplifying the related equations. Cuypers et al. [9] implemented a ray based diffraction model based on the Wigner Distribution Function.

Light Propagation in Optics: Wave optics describes light as an electromagnetic field with amplitude and phase. The Huygens–Fresnel principle is often used to represent wave propagation, which is a convolution of point scatterers [11]. Among the extensive efforts to connect wave and ray optics [29], two notable ones are the generalized radiance proposed by Walther [27] and the Wigner Distribution Function [6], which describes light in terms of its position and local spatial frequency, leading to a simple relationship with the angular domain [11]. Several alternative representations also describe an electromagnetic field in phase space, which is a joint space-spatial frequency representation [15]. The Short-Time Fourier Transform (STFT) or spectrogram [22] and the Rihaczek Distribution [23] are two other alternative space–spatial frequency representations. A good summary of a variety of phase space functions is offered by Cohen [8], where their use in optics has been considered by Accardi and Wornell [1]. Because it's generality and simplicity, we will focus our analysis on the WDF, whose relationship to geometric optics has been considered in depth by Alonso [4, 3].

Wigner Distribution Function in Optics: The Wigner Distribution Function was first introduced by E. Wigner in 1932 as a description of quantum mechanical systems [28]. The WDF has been exploited in a variety of analysis and design problems in optics: 3D display [10], digital holography [17], generalized sampling problems [25], and super resolution [30]. Recently, an important connection between the light field and WDF was made by Zhang and Levoy [32], which introduced the concept of an observable light field to the graphics community. Indirectly, they describe the space of effects spanned by today's photon and ray-based rendering methods. Oh et al. [21] extended this concept and described augmented light fields containing both positive as well as negative rays, allowing simulation of wave phenomena in a ray framework. In this paper, we explore this idea in more detail, by directly addressing the properties, benefits and limitations of this representation in graphics.

3. Wigner Distribution Function Basics

The Wigner Distribution as considered in this paper relates the space (x,y) and spatial frequency (u,v) content of a given function that defines an optical wavefront. For simplicity, we will begin by considering a quasi-monochromatic, completely coherent optical wavefront at a 2D plane. As in most imaging applications, we will assume that the optical signal we are interested is time-invariant over the finite time window. The Wigner Distribution of this 2D complex optical function, $t(x, y)$, can be defined as

$$W(x, y, u, v) = \iint J(x, y, x', y') e^{-i2\pi(x'u+y'v)} \mathrm{d}x'\mathrm{d}y', \tag{1}$$

where the function in the integrand,

$$J(x, y, x', y') = t(x + \frac{x'}{2}, y + \frac{y'}{2}) t^*(x - \frac{x'}{2}, y - \frac{y'}{2}) \tag{2}$$

is often called the mutual intensity function. Since we've assumed a completely coherent optical wave, our mutual intensity can be represented as a multiplication of two functions t and t^*. The case of partial coherence will be discussed in Section 4. The * operation represents complex conjugation. Note that after the Fourier transform of the mutual intensity function, the WDF contains only real values, positive as well as negative.

With our assumptions explicitly stated, let's now take a close look at Eq. 1. The WDF of a 2D function is 4D, and as we will see is directly related to the geometrical light field well understood in computational photography. First, let's consider the two spatial dimensions (x, y). The spatial frequency variables (u, v) will be examined in Section 3.1. Two simple interpretations for $t(x, y)$ exist: it can be considered a function that describes the amplitude and phase of an optical wave at some plane in space, or it can be considered a function that describes a surface or aperture [21]. The latter interpretation is of more interest from a computational photography viewpoint. Under this assumption, $t(x, y)$ can describe a surface like the grating atop a butterfly wing, or the fine mesh of a fabric. Using the same spatial coordinates, it is clear that the WDF light field $W(x, y, u, v)$ given by Eq. (1) will describe rays with coordinates that start at (x, y), immediately after reflection from or transmission through these thin surfaces. This light field will be consistent with physical optics theory up to certain approximations, discussed in more detail in Section 4.

Specifically, $t(x, y)$ incorporates the surface's ability to absorb or transmit light into its absolute value, and the surface's ability to impart a phase delay to the light into its complex angle:

$$A_s = |t(x, y)| \tag{3}$$

$$\phi_s = \arctan \frac{Im[t(x, y)]}{Re[t(x, y)]}. \tag{4}$$

Here, (A_s, ϕ_s) is the amplitude transmittance and phase delay, respectively, of the surface, and Re and Im represent the real and complex projection operators. The complex portion of t indicates a phase delay due to either a change in the refraction index or the thickness of the surface's material at position (x, y). For example, an amplitude grating (i.e., a series of black and transparent stripes) can be represented by a real-valued $t(x, y)$ that periodically varies between 0 and 1, while a phase grating (i.e., a series of raised glass ridges) can be represented by a $t(x, y)$ with $|t(x, y)| = 1$ for all (x, y) and a complex angle ϕ_s that varies between $-\pi$ and π.

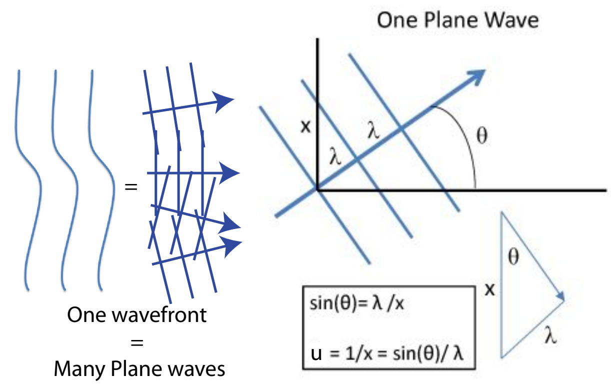

Figure 2. (left) A wavefront can be split up into many plane waves as part of a Fourier decomposition (See Section 3.10 of [11]) This is similar to Huygens's principle, but uses plane waves instead of spherical waves as a basis. (right) Each plane wave in the decomposition can be related to a ray traveling at a certain angle, shown as an arrow. Spatial frequency is given as the ratio of the sine of this angle and the wavelength of light. Its units are 1/meters.

3.1. One Wavefront as Many Plane Waves

The WDF relates the spatial variables of the surface function $t(x, y)$ to the spatial frequency variables (u, v), which can be understood by considering the wave-like nature of light. If a wavefront of light has a single wavelength, as we are assuming, then it is considered monochromatic, or temporally coherent. A great property of a coherent wavefront of monochromatic light is that it can be represented by a sum of plane waves, each traveling at a slightly different angle (Figure 2). Decomposing a wave of light into a sum of plane waves at different angles is very similar to decomposing an arbitrary signal into a sum of sine waves at different frequencies with a Fourier transform. However, since we are dealing with a wave over space, each plane wave traveling at a different angle provides a unique *spatial* frequency to the wavefront. This basic representation of a wavefront is known as an angular spectrum [11], and can be visualized in Figure 2. Again, as our Fourier decomposition of a coherent wave happens across space, the definition of the spatial frequency term u is in units of m^{-1}.

Besides it's elegance in Fourier optics, this plane wave decomposition also offers a simple tie to the ray–space picture of geometric optics. From Figure 2, it is clear that each plane wave set can be described by a single bisecting ray at a certain angle θ. For example, if the wavefront is propagating directly to the right (with $\theta = 0$), the ray is also at $\theta = 0$, and the spatial frequency u is zero. As the angle with

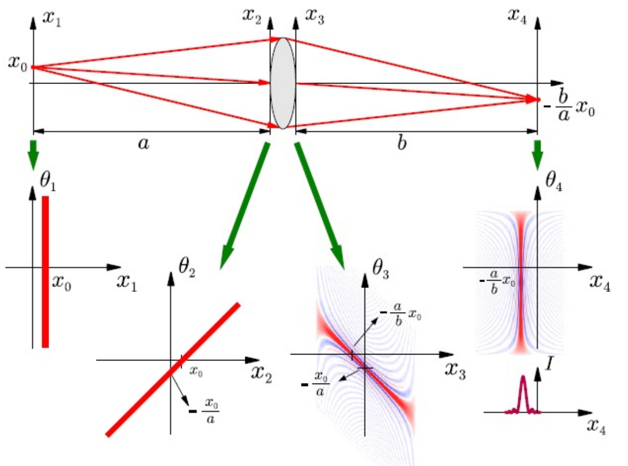

Figure 3. Overview of light propagation through a quadratic lens. (a) The WDF of a point light source (b) gets sheared by propagating through free space (c) get transformed due to the geometry of the lens (d) and gets projected on to the camera lens by integrating over all angles. Red indicates positive values and blue indicates negative.

respect to horizontal changes, u grows. The simple formula connecting θ and u is,

$$u = sin(\theta)/\lambda \approx \theta/\lambda, \tag{5}$$

where the approximation holds in the paraxial region, discussed in Section 4. This relationship allows for the transfer of wave phenomena to a ray treatment. In general, the Fourier decomposition of any coherent wavefront into spatial frequency components (i.e., a sum of plane waves) is indirectly a decomposition of any wavefront into a group of rays at different angles.

Putting it all together, the WDF function $W(x, y, u, v)$ describes a bundle of rays, which start at a surface $t(x, y)$, and leave with a 2D angle, ($\theta_x = sin^{-1}(\lambda u)$, $\theta_y = sin^{-1}(\lambda v)$). Each ray is given a weight from the computation of Eq. 1, which effectively provides a Fourier decomposition of $t(x, y)$ into a bundle of rays at each (x, y).

3.2. A Simple Light Propagation Model

As with ray–based light fields, the WDF also follows linear transformations that can be represented in the well-known ABCD matrix formalism [6, 15]. Following is a WDF light propagation model built using 3 different transformations: propagation through free-space, propagation through a grating, and projection. This model can trace light from an initial source, through a diffracting element, to a screen or image sensor where all rays are projected into an intensity measurement.

Free Space Propagation: The WDF $W_z(x, y, u, v)$ of a complex wavefront will shear due to traveling a distance z, similar to the rays of a light field, with,

$$W_z(x, y, u, v) = W(x - \lambda uz, y - \lambda vz, u, v) \tag{6}$$

Propagation Through a Thin Grating: In Section 3, we saw that the propagation of a plane wave through a grating $t(x, y)$ is given by its WDF, $W_t(x, y, u, v)$, from Eq. 1. If something besides a plane wave hits $t(x, y)$, we can still find the resulting output WDF, W_o. It is defined by a convolution along spatial frequency of the incoming WDF (W_i), and the grating WDF (W_t), with

$$W_o(x, y, u, v) = \iint W_i(x, y, a-u, b-v)W(x, y, a, b)\mathrm{d}a\mathrm{d}b \tag{7}$$

Projection onto a Surface: The intensity of light described by the WDF is found with it's projection along spatial frequencies (u, v):

$$I(x, y) = \int W(x, y, u, v)\mathrm{d}u\mathrm{d}v \tag{8}$$

Even though the Wigner Distribution Function contains negative values, the observed intensity $I(x, y)$ on a surface is always non–negative [5]. This is demonstrated in Figure 3 and also explained in section 4.2. The outgoing WDF however does contain negative values, which are marked in blue, positive values in red.

4. WDF Applicability for Ray-based Imaging

4.1. Valid in Paraxial Region

The above three transformations of the WDF are valid in what is termed the "paraxial region". This is the region where light propagates close to the normal of the diffracting surface (i.e., at smaller diffracting angles). As the region of interest moves away from the paraxial region, direct connections between the WDF and geometric optics become more difficult to apply, as errors are introduced. Degredation is slow, since there is no definite boundary between the paraxial and non-paraxial region, but typically remains low in most practical rendering applications. If the small error in the WDF is not tolerable, then a more rigorous treatment with the angle-impact WDF can be used [4]. Figure 5 demonstrates the impact of off-axis rendering compared to ground truth calculated with phase tracking. Error arises in a geometric treatment of the WDF outside of the paraxial region since it approximates any spherical wavefront (from a point source) as a quadratic surface. In the paraxial zone, where the observation is near the optical axis (Figure 4), we can approximate the distance between the surface and the

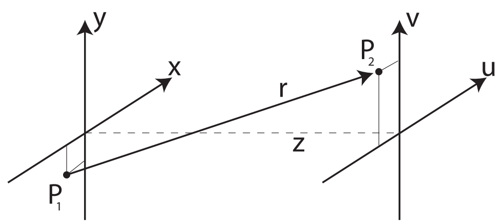

Figure 4. The paraxial region is the region where light is propagating close to the normal of the diffracting surface.

observation point r with,

$$r = z\sqrt{1+\left(\frac{x-u}{z}\right)^2+\left(\frac{y-v}{z}\right)^2} \quad (9)$$

$$\approx z\left[1+\frac{1}{2}\left(\frac{x-u}{z}\right)^2+\frac{1}{2}\left(\frac{y-v}{z}\right)^2\right], \quad (10)$$

This approximation is used in the Fresnel formula:

$$U(x,y)=\frac{\mathrm{e}^{ikz}}{i\lambda z}\int\int U(u,v)e^{\left\{i\frac{k}{2z}\left[(x-u)^2+(y-v)^2\right]\right\}}\mathrm{d}u\mathrm{d}v. \quad (11)$$

More detailed discussion on the accuracy of the Fresnel approximation can be found in Goodman [11].

4.2. Non-negative Projected Intensity

Although the WDF contains negative values, the projection of the WDF along any arbitrary direction within the 4D phase space (x,y,u,v) is always non-negative. Since the light field's intensity is determined at any plane by projecting the WDF along the (u,v) axes, demonstrating all projections are non-negative is similar to demonstrating the light field will have non-negative intensity anywhere along it's propagation axis. Following are 3 projections of the WDF $W(x,y,u,v)$: at 0°, 90° and at an arbitrary angle, which is related to a Fractional Fourier Transform [18] [14]. First, projection at 0° along (u,v) yields:

$$\begin{aligned} & \iint W(x,y,u,v)\mathrm{d}u\mathrm{d}v \\ = & \iint J(x,y,x',y')\delta(x')\delta(y')\mathrm{d}x'\mathrm{d}y' = t(x,y)t^*(x,y) \\ = & |t(x,y)|^2 \end{aligned}$$

Projection at 90°, along $(x,y), gives$:

$$\begin{aligned} & \iint W(x,y,u,v)\mathrm{d}x\mathrm{d}y \\ = & \iint \widetilde{J}(u,v,u',v')\delta(u')\delta(v')\mathrm{d}u'\mathrm{d}v' = \widetilde{t}(u,v)\widetilde{t}^*(u,v) \\ = & |\widetilde{t}(u,v)|^2 \end{aligned}$$

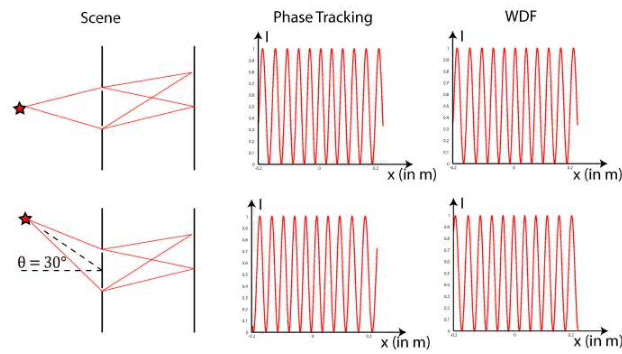

Figure 5. Simulation of Young's two pinhole experiment within and outside of the paraxial region. (Top) For light diffracting in paraxial zone (i.e., at smaller angles), simulation with phase tracking and the WDF yield identical results. (Bottom) Outside the paraxial zone, the pattern created by the WDF is slightly translated compared to phase tracking.

Projection along any arbitrary angle is also always non-negative. This is shown with the help of the Fractional Fourier Transform, where has very close ties to a rotating WDF [14]. Note that this transformation is primarily observed with light propagating through a GRIN lens with a quadratic refractive index profile, as in Figure 6.

$$\begin{aligned} & \textstyle\iint W(x\cos(\alpha)-u\sin(\alpha), y\cos(\alpha)-v\sin(\alpha), \\ & \quad x\sin(\alpha)-u\cos(\alpha), y\sin(\alpha)-v\cos(\alpha))\mathrm{d}u\mathrm{d}v \\ = & \textstyle\iint W_{fr}(x,y,u,v)\mathrm{d}u\mathrm{d}v = |\mathcal{F}_\alpha(t(x,y))|^2 \end{aligned}$$

where $\mathcal{F}_\alpha$ represents the fractional Fourier transform of angle α and W_{fr} is the WDF after a fractional Fourier transform. In conclusion, the projection of the WDF to determine intensity always yields a non-negative function, although the WDF itself can be negative. Thus, while rendering with the WDF requires the use of negative rays, any

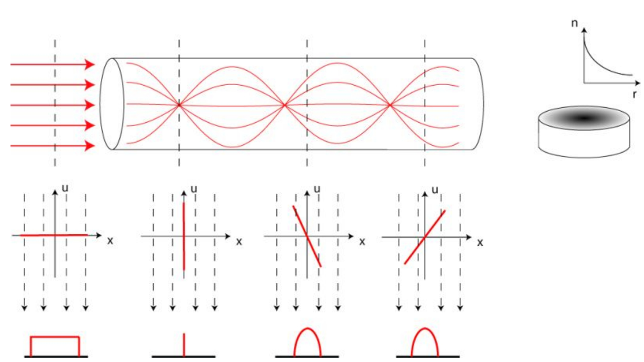

Figure 6. Within a parabolic GRIN lens, the WDF of the wavefront at any location is a rotated version of the input WDF with respect to the origin. The projection along any angle of rotation yields a non-negative intensity.

accurate local measurement of the intensity will require an integration of many rays, which will yield a positive value.

4.3. Partially Coherent Light

So far, our treatment of the WDF has focused on spatially and temporally coherent light, as explained at the beginning of Section 3. The WDF can also be extended to model partially coherent light. Coherence is a statistical relationship, and refers to the cross-correlation of a wave, or how correlated all points on a wave are with other points on the wave in both space and time. A mutual coherence function ($\mathcal{T}$) of a field (E) can describe this with,

$$\mathcal{T}(x_1, x_2, \tau) = < E(x_1, t)E^*(x_2, t + \tau) > \tag{12}$$

While the two are closely related, often temporal coherence deals with the chromatic nature of a wavefront (i.e., how much of a color spread it has), and is related to τ. Spatial coherence, on the other hand, is related to the statistical correlation between two points on a given wave, x_1 and x_2.

Temporal coherence When one is only interested in the correlation with respect to time delay, then temporal coherence can be found from Eq. 12 setting $x_1 = x_2$. Since Eq. 1 assumes quasi-monochromatic light, then it is temporally coherent. Polychromatic light can simply be modeled with a sum of WDF light fields, each using a different wavelength in Eq. 5.

Spatial coherence When we are interested in the correlation with respect to two points in space, then $\tau = 0$. A simple way to characterize spatial coherence is with the relative size of the original light source - the smaller the source, the more coherent it is. For example, the WDF in Eq. 1 is modeled from an ideal point source. If the original source is not a point, but is instead N independent point sources, then a partially coherent W_{pc} can be modeled from a sum of the coherent WDFs, W_c, from each point source:

$$W_{pc}(x, y, u, v) = \sum_{i=0}^{N} W_c^{(i)}(x, y, u, v) \tag{13}$$

Furthermore, light from incoherent sources (i.e. many point sources) becomes more coherent with an increased propagation distance. This is called the van Cittert-Zernike Theorem. More details on partial spatial coherence and the WDF is described by Bastiaans [7].

4.4. The WDF in 4D

While the WDF we've offered in equations is a 4D function, most of the included plots are 2D, along one space and one angle variable, for ease of visualization. To extend visualization to 4D, a display method common with the geometric light field is borrowed [20]. A macroscopic array can define the spatial content (x, y), while the angular content (u, v) can be plotted locally at each spatial location, essentially yielding a tiled grid of (u, v) plots. Figure 7 offers such visualization for a rotating beam, which rotates in (x, y) as it propagates, and would otherwise not be visualizable in one spatial dimension. Note that the angular content appears to rotate around the center, providing a nice display of the angular momentum associated with the beam.

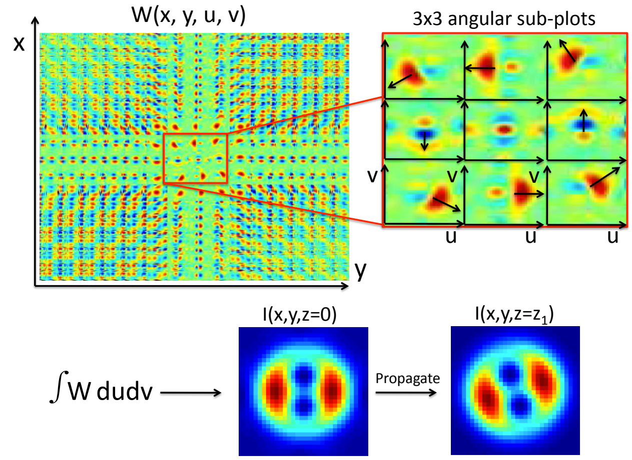

Figure 7. Displaying a 4D WDF in the tiled arrangement common to light field photographs can help with visualization. This example shows the 4D WDF of rotating beam at 32^4 resolution, where each (u, v) subplot is normalized to the same maximum value. It is clear the angular momentum rotates around the center in the zoomed inset (right), where arrows have been added to clarify the direction of rotation. (bottom) The two-point intensity pattern of the beam rotates with propagation.

4.5. Validity, Near–field and Far–field

Before we discuss the validity of the WDF in the near and far fields, let us clarify the different notions of near–field, far–field and paraxial region. Near–field refers to regions very close to a diffracting feature, where evanescent waves are still significant. After evanescent waves fully decay, only propagating waves contribute to diffraction. Hence, beyond the near–field (typically an order of few or tens of wavelengths depending on media and geometry), the Fresnel formula in Eq. 11 is typically used to model diffraction, where it explicitly assumes a quadratic approximation. More precisely, the propagation distance is assumed to be longer than the difference of lateral extent between the diffracting object and the observation plane. This is similar to, but not identical with, the small angle approximation associated with the paraxial region. More details on the accuracy on the Fresnel diffraction formula is described by Goodman [11]. The Fraunhofer diffraction formula is used under a far–field assumption, where the propagation

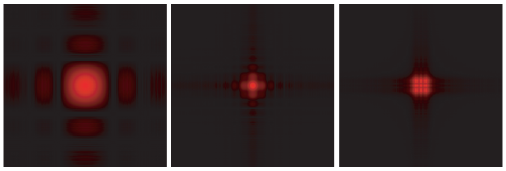

Figure 8. The Wigner Distribution Function is valid under the Fresnel approximation, which is applicable up to very short distances from a diffracting aperture. We show the PSF of a rectangular aperture with a size of $0.5mm$ measured at three different depths.

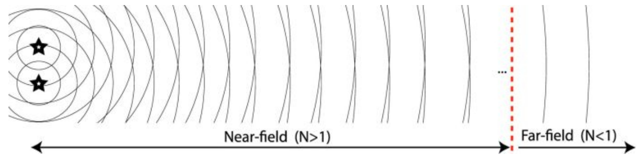

Figure 9. At close distances to a light source, waves appear spherical, and do not match the parabolic approximation that the WDF assumes. The WDFs connection between ray and wave optics (laid out in Section 3) becomes increasingly accurate at larger distances, where wavefronts begin to resemble parabolas

distance is significantly longer than the lateral extent of the diffracting object. Note that the lateral offset of the diffracting object does not change the amplitude of the far–field diffraction pattern; it changes only the phase.

Within the computational photography community, diffractive effects are often ignored. For example, under a ray-based model, it is often assumed that the point-spread function of a in-focus lens is a delta function. This is a valid assumption as long as a large enough sensor pixel size is also assumed. When diffraction is modeled, it is typical to see the Fraunhofer formula, which is based upon a Fourier-transform relationship, and assumes operation in the far–field. A more rigorous approach will apply the Fresnel diffraction formula, which is a very accurate model of a wave at distances that are very close to the diffracting screen [11] (i.e., much closer than the far-field).

In general, it is a safe bet to assume the WDF can provide an accurate connection between wave and ray optics whenever the paraxial approximation is valid, which is often the case under the Fresnel approximation, and almost always the case under the Fraunhofer equation. It certainly has been shown that the WDF's validity can be extended into regions much closer to the diffracting aperture and away from the paraxial region [4]. However, certain properties of the WDF, like it's useful shear–propagation relationship in Eq. 6, may not remain accurate. Thus, the WDF can be safely applied to geometric optic modeling anywhere the paraxial approximation is valid, which remains accurate up to short distances from the diffracting surface, and certainly falls within the physical–optics rigor of previous work in computational photography.

Fresnel Number: The border between where the Fraunhofer equation remains valid and where the Fresnel equation remains valid is characterized well by the Fresnel number (N). This frequently used number is given by [11],

$$N = \frac{D^2}{4\lambda L} \tag{14}$$

where D is the diameter of the diffraction grating, λ the wavelength of the propagating wave and L the propagation distance. If the Fresnel number N is much smaller than 1 (i.e., the propagation distance is much larger than the extent of the diffraction grating), then a far–field, Fraunhofer model is accurate. If the N is similar or larger than 1, then the observation plane is somewhat close to the diffraction grating. This implies that the Fresnel diffraction formula is the more accurate choice. Finally, if N is much larger than 1, then even the Fresnel diffraction formula may be inaccurate. Thus, the WDF is a valid connection between waves and rays as long as N is not much larger than 1.

5. Implementation

We demonstrate rendering with the WDF technique using the concept of an augmented light field [21], which is a framework to convert WDF values to ray values for raytracing. The WDF of a grating is calculated from the 2D microstructure geometry $t(x, y)$, as discussed in Section 3. For simplicity and efficiency our demonstrations are applied to gratings that are separable in the x and y directions. Figure 10 shows several precomputed WDFs for a variety of amplitude and phase gratings. We transform the spatial frequency of the WDF to the angle of the light field using paraxial approximation in Eq. 5. These values are input for a traditional raytracer, which traces rays from the grating towards a second plane at a distance z. We perform this for 30 wavelengths and integrate them along a camera response curve to convert them into an RGB value. For rendering the example in Figure 10, we perform raytracing from the source through the grating until it reaches a sensor. The intensity of each ray is calculated by the incoming and outgoing angle θ_i and θ_o at the grating using Eq. 7. The calculated ray intensities, in 2D, is $W(x, \theta_o - \theta_i)$.

6. Conclusion

In summary, we've attempted to explain the WDF using geometric optics principles, with the aim of introducing physical optics effects to a ray-based graphics and rendering pipeline. We underline the connection between the spatial frequency of a wave and the angle of a ray, thus joining the WDF's variables with the light field's. We have demonstrated how the Wigner Distribution can simulate diffraction and interference under a light field representation. We have also discussed how accurate this representation is outside the paraxial region, and also in the near and far–field. Furthermore, we demonstrated how the WDF's negative values cannot lead to negative intensity, offered a novel method of visualizing a WDF in 4D, and provide grating-to-light field generation code, which we hope others will take advantage of as they explore the connection between geometrical and physical optics in computational photography.

Acknowledgements

Tom Cuypers and Philippe Bekaert acknowledge financial support by the European Commision (FP7 IP 2020 3D Media), the European Regional Development Fund (ERDF), the Flemish Government and the Flemish Interdisciplinary Institute for Broadband Technology (IBBT). Ramesh Raskar is supported by an Alfred P. Sloan Research Fellowship. We also like to thank George Barbastathis and Markus Testorf and all our colleagues at the Camera Culture Group for their time and input.

References

[1] A. Accardi and G. Wornell. Quasi light fields: extending the light field to coherent radiation. 2

[2] E. Adelson and J. Bergen. The plenoptic function and the elements of early vision. *CMVP*, pages 3–20, 1991. 2

[3] M. Alonso. Diffraction of paraxial partially coherent fields by planar obstacles in the wigner representation. In *Journal of Optical Society of America A*, pages 1588–1597, 2009. 2

[4] M. A. Alonso. Wigner functions for nonparaxial, arbitrarily polarized electromagnetic wave fields in free space. *Journal of Optical Society of America A*, 21(11), 2004. 2, 4, 7

[5] M. Bastiaans. Application of the wigner distribution function in optics. In *The Wigner Distribution - Theory and Applications in Signal Processing*. Elsevier Science, 1997. 4

[6] M. J. Bastiaans. Wigner Distribution Function and Its Application to 1st-order Optics. *J. Opt. Soc. Am.*, 1979. 2, 4

[7] M. J. Bastiaans. Application of the Wigner Distribution Function to Partially Coherent-Light. *J. Opt. Soc. Am. A*, pages 1227–1238, 1986. 6

[8] L. Cohen. Time-frequency distributions - a review. In *Proc. IEEE*, pages 941–981, 1989. 2

[9] T. Cuypers, S. B. Oh, T. Haber, P. Bekaert, and R. Raskar. Ray-based reflectance model for diffraction. Technical Report arXiv:1101.5490v1, 2011. 2

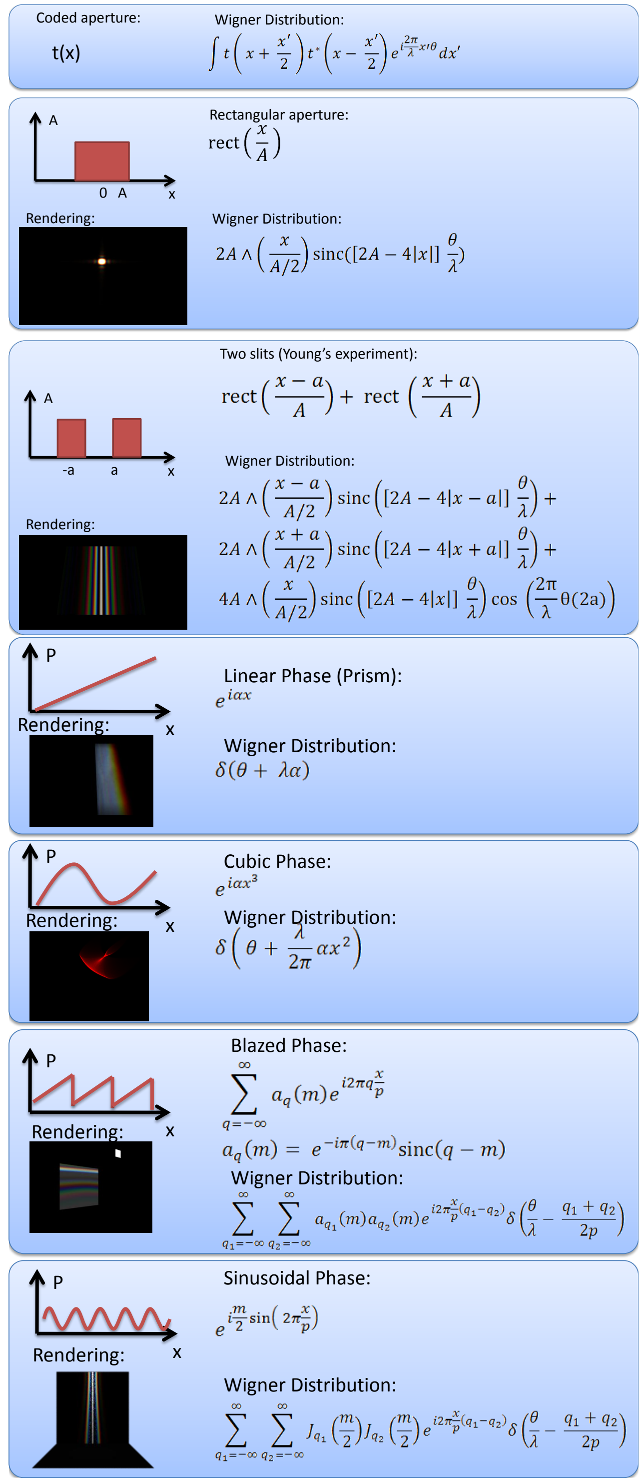

Figure 10. Renderings for light propagation through several coded apertures and phase masks.

[10] W. D. Furlan, M. Martinez-Corral, B. Javidi, and G. Saavedra. Analysis of 3-d integral imaging dis-

plays using the wigner distribution. *J. Display Technol.*, 2006. 2

[11] J. W. Goodman. *Introduction to Fourier optics*. Roberts & Co., Englewood, Colo., 3rd edition, 2005. 2, 3, 5, 6, 7

[12] J. T. Kajiya. The rendering equations. In *SIGGRAPH*, volume 20, pages 143–150. ACM, 1986. 2

[13] M. Levoy and P. Hanrahan. Light field rendering. In *SIGGRAPH*, 1996. 2

[14] A. W. Lohmann. Image Rotation, Wigner Rotation, and the Fractional Fourier-Transform. *J. Opt. Soc. Am. A*, 10(10):2181–2186, 1993. 5

[15] J.-C. M. E. Testorf, B.M. Hennelly. Phase space optics: fundamentals and applications. In *McGraw Hill*, 2009. 2, 4

[16] W. Matusik and H. Pfister. 3d tv: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Trans. Graph.*, 2004. 2

[17] J. Maycock, C. P. McElhinney, B. M. Hennelly, T. J. Naughton, J. B. McDonald, and B. Javidi. Reconstruction of partially occluded objects encoded in three-dimensional scenes by using digital holograms. *Appl. Opt.*, (13), 2006. 2

[18] D. Mendlovic and H. M. Ozaktas. Fractional Fourier-Transforms and Their Optical Implementation .1. *J. Opt. Soc. Am. A*, 10(9):1875–1881, 1993. 5

[19] H. P. Moravec. 3d graphics and the wave theory. In *SIGGRAPH*. ACM, 1981. 2

[20] R. Ng. Fourier slice photography. 24:735–744, 2005. 2, 6

[21] S. B. Oh, S. Kashyap, R. Garg, S. Chandran, and R. Raskar. Rendering wave effects with augmented light fields. *EuroGraphics*, 2010. 2, 3, 7

[22] L. Rabiner and R. Shafer. Digital processing of speech signals. *Prentice-Hall*, 1978. 2

[23] A. Rihaczek. Signal energy distribution in time and frequency. *IEEE Trans. Inform. The Theory*, 1968. 2

[24] J. Stam. Diffraction shaders. In *SIGGRAPH*, 1999. 2

[25] A. Stern and B. Javidi. Sampling in the light of wigner distribution. *J. Opt. Soc. Am. A*, 21(3):360–366, 2004. 2

[26] V. Vaish, M. Levoy, R. Szeliski, L. Zitnick, and S. Kang. Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures. 2006. 2

[27] A. Walther. Radiometry and Coherence. *Journal of Optical Society of America*, 63(12):1622–1623, 1973. 2

[28] E. Wigner. On the quantum correction for thermodynamic equilibrium. *Physical Review*, 40:749–759, 1932. 2

[29] E. Wolf. Coherence and Radiometry. *Journal of Optical Society of America*, 68(1):6–17, 1978. 2

[30] K. B. Wolf, D. Mendlovic, and Z. Zalevsky. Generalized wigner function for the analysis of superresolution systems. *Appl. Opt.*, 37(20):4374–4379, 1998. 2

[31] C. Zhang and T. Chen. Light field capturing with lensless cameras, 2005. 2

[32] Z. Zhang and M. Levoy. Wigner distributions and how they relate to the light field. In *IEEE ICCP*, 2009. 2

[33] R. Ziegler, S. Bucheli, L. Ahrenberg, M. Magnor, and M. Gross. A bidirectional light field - hologram transform. *CG Forum*, 2007. 2

[34] R. Ziegler, S. Croci, and M. H. Gross. Lighting and occlusion in a wave-based framework. *CG Forum*, pages 211–220, 2008. 2