

Wigner Distributions and How They Relate to the Light Field

Zhengyun Zhang
Electrical Engineering
Stanford University
Stanford, CA 94305

zaltor@graphics.stanford.edu

Marc Levoy
Computer Science
Stanford University
Stanford, CA 94305

<http://graphics.stanford.edu/~levoy/>

Abstract

In wave optics, the Wigner distribution and its Fourier dual, the ambiguity function, are important tools in optical system simulation and analysis. The light field fulfills a similar role in the computer graphics community. In this paper, we establish that the light field as it is used in computer graphics is equivalent to a smoothed Wigner distribution and that these are equivalent to the raw Wigner distribution under a geometric optics approximation. Using this insight, we then explore two recent contributions: Fourier slice photography in computer graphics and wavefront coding in optics, and we examine the similarity between explanations of them using Wigner distributions and explanations of them using light fields. Understanding this long-suspected equivalence may lead to additional insights and the productive exchange of ideas between the two fields.

1. Introduction

Simulation and analysis of optical systems have long been an intrinsic part of computer graphics. The light field[15], a mapping from rays to radiance, provides a powerful tool in this regard. For example, pinhole images can be generated through slicing of the light field[24], finite aperture images with depth-of-field effects can be generated through integration[19], and light propagation and lensing can be described using linear coordinate transformations[17] of the light field.

Similarly, in the optics literature, the Wigner distribution[33] of the scalar field and its Fourier dual, the ambiguity function[35], have been used extensively to simulate optical systems because they remove the need to compute complicated integrals of the scalar field directly[30, 5]. In particular, light propagation and lensing can be described using linear coordinate transformations of the Wigner distribution[5].

These similarities are supported by a body of literature in the optics community that relates the concept of radiance to wave optics[32, 34]. However, it turns out that radiance as a concept in wave optics cannot be well defined[13] and cannot be measured directly. Therefore, we describe an alternate path from wave optics to the light field by using radiant intensity. Based on this path, we show that the light field is equivalent to a smoothed Wigner distribution, which approaches the Wigner distribution under a geometric optics approximation.

Using this conceptual equivalence, we can reinterpret previous work in the optics literature and show how it corresponds to recent work in the computer graphics literature. For example, different slices in the Fourier transform of the light field have been shown to yield images focused at different planes[28]. Different slices of the ambiguity function yield the same results[30, 9]. Furthermore, while the original analysis of cubic phase plate wavefront coding was performed using the ambiguity function[11], we can do the same using light fields. Understanding these similarities gives us a framework for analyzing other works in the two fields for which the parallels are not yet recognized.

We are not the first to apply wave optics concepts to light fields. There is previous work exploring and extending the depth of field of integral imaging systems through the use of Wigner distributions[14]. However, that paper does not investigate whether the four-dimensional function represented by the light field is related to the four-dimensional Wigner distribution. There is also previous work in deriving a light field from a scalar field and vice versa[36], but it does not touch upon Wigner distributions.

2. Wigner Distributions in Optics

Before we show how the Wigner distribution is similar to a light field, let us review the Wigner distribution function as it is used in Fourier optics. Light propagates through three-dimensional space as a wave, and by measuring the

scalar field of the wave at some plane, we can obtain information about how the light will propagate past this plane. For this two-dimensional scalar field, the Wigner distribution is a four-dimensional function that describes the field's positional information along two of the axes and its frequency information along the other two.

Let us assume we have narrowband polychromatic light. We can represent its wave function, which is a description of oscillations in the electric and magnetic fields, as $U(x, y, z, \tau)$, a time varying phasor of a scalar field[16]. Suppose we have measurements $U_1(x, y, \tau)$ of the scalar field at some plane $z=z_1$, and we wish to derive the intensity $I(x, y)$ at some other plane $z=z_2$ with optical elements in between. The standard method is to apply a series of Fresnel diffraction integrals and phase mask multiplications[16]. To simplify this situation, we can instead derive the output image through the use of the input scalar field's four-dimensional *Wigner distribution*, defined as:

$$W_{U_1}(x, y, f_\xi, f_\eta) = \iint J_{U_1}(x, y, \xi, \eta) e^{-j2\pi(f_\xi \xi + f_\eta \eta)} d\xi d\eta \quad (1)$$

where

$$J_{U_1}(x, y, \xi, \eta) = \left\langle U_1\left(x + \frac{\xi}{2}, y + \frac{\eta}{2}, \tau\right) U_1^*\left(x - \frac{\xi}{2}, y - \frac{\eta}{2}, \tau\right) \right\rangle_\tau \quad (2)$$

is the *mutual intensity*¹ of the input scalar field. The mutual intensity describes how coherent and energetic any two points in the scalar field are; two points which are less coherent will yield smaller values in the mutual intensity. We obtain this function by multiplying the scalar field $U_1(x, y)$ by a shifted complex conjugate copy of itself for all possible shifts (ξ, η) and then averaging each result over time.

In the case of fully coherent light, the product of the scalar field and its shifted complex conjugate is time-invariant, and thus we can remove the extraneous time average from the mutual intensity:

$$J_{U_1}(x, y, \xi, \eta) = U_1\left(x + \frac{\xi}{2}, y + \frac{\eta}{2}\right) U_1^*\left(x - \frac{\xi}{2}, y - \frac{\eta}{2}\right) \quad (3)$$

The above expression is sometimes referred to as the *instantaneous autocorrelation*, since integrating along the shift axes produces the autocorrelation of the field.

The effect of coherence on the Wigner distribution and the mutual intensity can be seen in Fig. 1. Since they are quadratic representations of the scalar field, we cannot, for example, expect the sum of Wigner distributions to be the Wigner distribution of a sum. This nonlinearity explains why there are cross terms (interference) in a system with two coherent slits, as explained in the figure caption. If we were to simply add the mutual intensity of the two slits

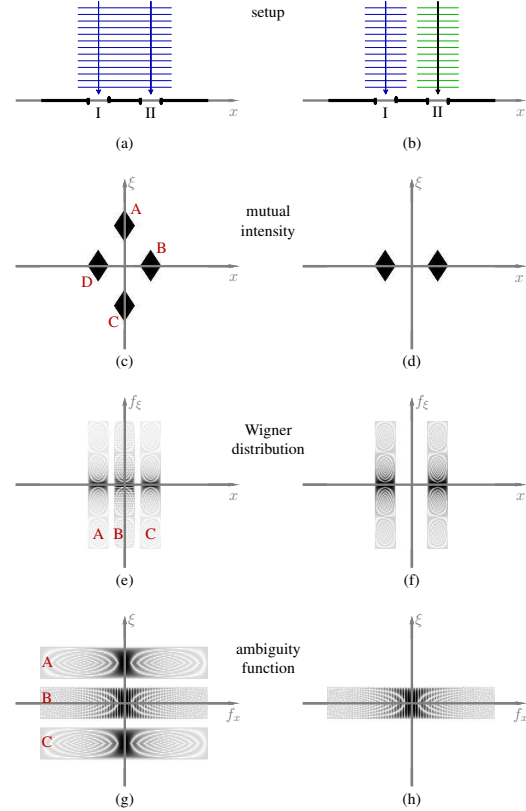


Figure 1. A single coherent plane wave approaches two slits I and II in (a) while plane waves blue (left) and green (right), incoherent relative to each other, hit slit I and II, respectively, in (b). The mutual intensities for the two cases can be seen in (c) and (d), the Wigner distributions in (e) and (f), and the ambiguity functions in (g) and (h). In all cases, darker equals greater magnitude. The presence of cross terms caused by interference between the two slits in the left column give rise to diamonds A and C in the mutual intensity, band B in the Wigner distribution and the sidebands A and C in the ambiguity function. The periodic features in the Wigner distribution and ambiguity function are aliasing artifacts due to discrete sampling in the numerical simulation.

separately, we would obtain only two diamond-shaped regions, exactly like the incoherent case in Fig. 1(d). In fact, adding Wigner distributions is equivalent to treating their scalar fields as being incoherent with respect to each other.

Another useful function related to the Wigner distribution is its Fourier dual, the *ambiguity function*, which can be written as:

$$\begin{aligned} A_{U_1}(f_x, f_y, \xi, \eta) &= \iint J_{U_1}(x, y, \xi, \eta) e^{-j2\pi(f_x x + f_y y)} dx dy \\ &= \mathcal{F}_{\xi, \eta}^{-1} \{ \mathcal{F}_{x, y} \{ W_{U_1}(x, y, f_\xi, f_\eta) \} \} \end{aligned} \quad (4)$$

We will cover one particular use of the Wigner distribution in the following section. Other properties, uses and examples of the Wigner distribution in optics can be found

¹ We use average and difference variables here instead of two separate spatial coordinates, which is more common.

in [6].

2.1. Radiance and the Wigner Distribution

Since the light field is a mapping from rays to radiance, an attractive approach to our overall goal is to explore works in the optics literature that attempt to define radiance in terms of wave optics and see how these definitions relate to the Wigner distribution. However, we will see that radiance is not well-defined in wave optics and is not directly measurable. This motivates us to investigate an alternate approach in Section 3.

Most of the work tying radiance and wave optics concentrate on finding a radiance function that satisfies a set of ideal properties[34]. The first such work, by Walther[32], attempts to find a radiance function that when integrated along angle and space, becomes the flux exiting a surface. In his work, Walther derives the following expression for radiance of some scalar field U :

$$B = N \iint \frac{1}{\lambda^2} J_U(x, y, \xi, \eta) e^{-j \frac{2\pi}{\lambda} (L\xi + M\eta)} d\xi d\eta \quad (5)$$

where B is the radiance, λ is the wavelength of light, J_U is the mutual intensity and the “ray” is parametrized by a point $(x \ y)^T$ and a unit vector $(L \ M \ N)^T$ indicating the direction. Using the definition of the Wigner distribution in (1), we obtain that radiance can be written in terms of the Wigner distribution of the scalar field:

$$B = N \frac{1}{\lambda^2} W_U \left(x, y, \frac{L}{\lambda}, \frac{M}{\lambda} \right) \quad (6)$$

In the small angle approximation, where $N \approx 1$, we obtain that the radiance is of the following form:

$$B = \frac{1}{\lambda^2} \hat{W}_U^{(\lambda)}(x, y, u, v) \quad (7)$$

where $u = \frac{L}{\lambda}$ and $v = \frac{M}{\lambda}$ are the slopes in x and y respectively of the ray direction and

$$\hat{W}_U^{(\lambda)}(x, y, u, v) = W_U \left(x, y, \frac{u}{\lambda}, \frac{v}{\lambda} \right) \quad (8)$$

is what we will call the *slope-form Wigner distribution*.

Strictly speaking, the concept of a ray does not exist in wave optics – a ray is defined by a point and a direction, but one of the fundamental results of Fourier optics is that an infinitesimal point must emit isotropically in all directions. We can see a similar effect when a pebble dropped into a pond produces only circular waves. Therefore, there cannot be a strict mapping of radiance onto wave optics.

In fact, the expression for radiance given by (5), being a function of the Wigner distribution of some arbitrary scalar field, can at times become negative, which violates the positivity of radiance and thus cannot be a physically measurable quantity[27]. Furthermore, there cannot be a radiance function satisfying all the ideal properties of radiance[13].

A summary of the idea of radiance in wave optics can be found in [34].

That said, in the short-wavelength limit (geometric optics), Walther’s equation for radiance (5) has been shown to satisfy all the ideal properties of radiance[12, 20], thereby proving the Wigner distribution equivalent to radiance (and thus light fields) at this limit. However, since radiance cannot be directly measured and a radiance function cannot be well-defined for all cases, we propose an alternate approach using radiant intensity to tie together light fields and Wigner distributions.

3. The Light Field as a Wigner Distribution

Rays in the light field can be parametrized by a point and a direction, which we’ve just seen to be problematic in wave optics. However, while fundamental wave optics principles state that a infinitesimal pinhole must transmit light isotropically, a finite aperture *can* transmit light with angular variation. Therefore, instead of trying to find the angular variation of light transmitted through a point, as we would in the case of radiance, let us find the angular variation of light transmitted through an arbitrary finite aperture centered on this point by using radiant intensity.² We will call this the *observable light field*. Using this approach will help us easily analyze spatioangular tradeoffs inherent in capturing light fields (Section 3.3), as well as the relationship between a discrete, physically captured light field and the Wigner distribution (Section 4.1).

3.1. The Observable Light Field

We shall start with a three-dimensional system in which light propagates in the positive z direction. Let us use the two-plane parametrization of the light field[24], placing the (s, t) , or reference plane, on the (x, y) plane and the (u, v) plane at $z=\infty$. This makes (u, v) coordinates represent angles or slopes. To define the opening through which rays will pass, we will use a two-dimensional aperture transmission function $T(x, y)$, which we will translate along the reference plane for particular values of (s, t) . Note that this aperture needs to have finite area in order to capture any directional information, as discussed earlier. We then define the *observable light field from aperture T* as the observed radiant intensity of light emanating in a particular direction (u, v) from a translated aperture $T(x-s, y-t)$, as shown in Fig. 2.

Given a scalar field $U(x, y)$ at the $z = 0$ plane and an aperture function $T(x, y)$, we can derive an expression for the radiant intensity, which is the amount of power emitted

²Radiant intensity is defined in computer graphics to be some angular variation of light coming from an imaginary point in the scene[18]. However, in optics, it is simply the angular variation of light from an area emitter, with this imaginary point being the origin of the coordinate system[7].




Figure 2. The light field is defined as the radiance along a ray (blue) for all the possible rays, as shown in (a). Rays can be parameterized by an intersection with a reference plane at (s, t) and a slope (u, v) . Since point sources can only be isotropic emitters in wave optics, we must define a new concept called the observable light field, which is a measure of the radiant intensity along the angle given by the slope (u, v) emanating from a small region on the reference plane centered on (s, t) , as shown in (b). The vertical graph along the reference plane in (b) is of a translated aperture transmission function, $T(x-s, y-t)$.

from a finite area in a given direction. In order to obtain this angular distribution, we can invoke the idea of the angular spectrum from wave optics, which is a decomposition of an arbitrary scalar field propagating past a plane into a set of plane waves propagating in various directions. This is in turn equivalent to the Fourier transform of the scalar field at that plane[16]. Therefore, we can obtain a measure of the radiant intensity as seen through the aperture by performing a Fourier transform of the scalar field and then taking the magnitude squared of the amplitude of each plane wave to obtain the power along that direction.

Let us now derive this mathematically. We will omit constant scale factors to simplify notation. The time-varying scalar field just after the aperture is $U(x, y, \tau)T(x-s, y-t)$. We can derive its angular spectrum by applying the Fourier transform to our scalar field[16]:

$$\tilde{U}(u, v, \tau) = \iint U(x, y, \tau)T(x-s, y-t)e^{-j2\pi(\frac{u}{\lambda}x + \frac{v}{\lambda}y)}dxdy \quad (9)$$

where u and v are slopes in x and y respectively, as in Fig. 2, and we've applied a small angle approximation for paraxial optics. Since we defined the observable light field as the radiant intensity from an aperture and the radiant intensity is the intensity of each of the plane waves, we can thus write the observable light field from aperture T as:

$$I_{obs}^{(T)}(s, t, u, v) = \left\langle \left| \tilde{U}(u, v, \tau) \right|^2 \right\rangle_{\tau} \quad (10)$$

3.2. From the Observable Light Field to the Wigner Distribution

So far, we have derived an expression for the observable light field in terms of the scalar field and an aperture function. We now show that this expression is actually a convolution of two Wigner distributions through the following set of transformations. First, we expand the magnitude squared of $\tilde{U}(u, v, \tau)$ in (10) into a product of complex conjugates:

$$\begin{aligned} & \iiint \langle U(x_1, y_1, \tau)U^*(x_2, y_2, \tau) \rangle_{\tau} \\ & \times T(x_1-s, y_1-t)T^*(x_2-s, y_2-t) \\ & \times e^{-j2\pi(\frac{u}{\lambda}(x_1-x_2) + \frac{v}{\lambda}(y_1-y_2))}dxdydy_1dy_2 \quad (11) \end{aligned}$$

Rewriting (11) into average and difference variables yields:

$$\begin{aligned} & \iiint \langle U(x+\frac{\xi}{2}, y+\frac{\eta}{2}, \tau)U^*(x-\frac{\xi}{2}, y-\frac{\eta}{2}, \tau) \rangle_{\tau} \\ & \times T(x+\frac{\xi}{2}-s, y+\frac{\eta}{2}-t)T^*(x-\frac{\xi}{2}-s, y-\frac{\eta}{2}-t) \\ & \times e^{-j2\pi(u\xi/\lambda + v\eta/\lambda)}dxdy d\xi d\eta \\ & = \iiint J_U(x, y, \xi, \eta)J_T(x-s, y-t, \xi, \eta) \\ & \times e^{-j2\pi(u\xi/\lambda + v\eta/\lambda)}dxdy d\xi d\eta \quad (12) \end{aligned}$$

This is the Fourier transform of a product of two mutual intensities. By invoking the convolution theorem and the definition of the Wigner distribution in (1), we can rewrite (12) as the convolution of two Wigner distributions:

$$\begin{aligned} & \iint W_U\left(x, y, \frac{\xi}{\lambda}, \frac{\eta}{\lambda}\right) \otimes_{\xi, \eta} W_T\left(x-s, y-t, \frac{\xi}{\lambda}, \frac{\eta}{\lambda}\right) dxdy \\ & = W_U\left(s, t, \frac{\xi}{\lambda}, \frac{\eta}{\lambda}\right) \otimes W_T\left(-s, -t, \frac{\xi}{\lambda}, \frac{\eta}{\lambda}\right) \quad (13) \end{aligned}$$

Using the definition of the slope-form Wigner distribution (8), we arrive at our final result:

$$I_{obs}^{(T)}(s, t, u, v) = \hat{W}_U^{(\lambda)}(s, t, u, v) \otimes \hat{W}_T^{(\lambda)}(-s, -t, u, v) \quad (14)$$

Therefore, the observable light field is equivalent to the slope-form Wigner distribution of the scalar field blurred by the spatially inverted slope-form Wigner distribution of the aperture transmission function. To complete our proof of equivalence between the light field and the Wigner distribution, we must study the effect of this blur caused by the aperture function.

3.3. The Uncertainty Principle and the Geometric Optics Limit

The presence of a spatioangular blur in the observable light field, as shown by the convolution in (14), is a consequence of the uncertainty principle – we cannot observe precisely both the location and direction of a particle, such

as a photon. Let us look at the extent of this blur by looking at the spread of the slope-form Wigner distribution.

This spread of the Wigner distribution can be obtained by treating it as a statistical distribution and looking at its variance along each axis. The Wigner distribution, while real, can take on negative values, so it is not entirely accurate to treat it directly as a statistical distribution. However, we'll see that this does not negatively impact our results, as the variance induced by a Wigner distribution is equivalent to the variance induced by the energy of the original signal.

For this derivation, we will use a one-dimensional unit-energy signal $h(x)$ and its two-dimensional Wigner distribution $W_h(x, f_\xi)$ to simplify notation, as the same ideas can be applied to the two-dimensional signal case and variance is invariant when a signal is scaled by a constant factor. For such a signal $h(x)$, the variance σ_x^2 in x of its Wigner distribution can be shown to be equal to the variance in x of the signal's energy, $|h(x)|^2$, using only the definition of variance and properties of the Wigner distribution[6]:

$$\begin{aligned}\sigma_x^2 &= \frac{\int \int x^2 W_h(x, f_\xi) dx df_\xi}{\int \int W_h(x, f_\xi) dx df_\xi} \\ &= \int x^2 \int W_h(x, f_\xi) df_\xi dx = \int x^2 |h(x)|^2 dx\end{aligned}\quad (15)$$

where the last line is the definition of variance of $|h(x)|^2$ along x . Through a similar derivation, the variance $\sigma_{f_\xi}^2$ in f_ξ of the Wigner distribution can be shown to be equal to the variance in f_ξ of the power spectrum of the original signal:

$$\sigma_{f_\xi}^2 = \int f_\xi^2 |H(f_\xi)|^2 df_\xi \quad (16)$$

One form of the Fourier uncertainty relation[8] states that the product of the variances along axes related by a Fourier transform has a lower bound[3]:

$$\left(\int x_i^2 |h(\vec{x})|^2 d\vec{x} \right) \left(\int f_{x_i}^2 |H(\vec{f}_x)|^2 d\vec{f}_x \right) \geq \frac{1}{16\pi^2} \quad (17)$$

where $h(\vec{x})$ is a function from \mathbb{R}^n to \mathbb{C} and $H(\vec{f}_x)$ is its Fourier transform and x_i and f_{x_i} are the i^{th} entries of \vec{x} and \vec{f}_x , respectively. Applying (17) to (15) and (16), we obtain:

$$\sigma_x^2 \sigma_{f_\xi}^2 \geq \frac{1}{16\pi^2} \quad (18)$$

Therefore, it is possible for the Wigner distribution to have a small extent in either frequency or position, but not both. For the slope-form Wigner distribution $\hat{W}_h^{(\lambda)}(s, u)$, we can apply the coordinate transform inherent in its definition (8) to obtain a similar uncertainty bound:

$$\sigma_s^2 \sigma_u^2 \geq \frac{\lambda^2}{16\pi^2} \quad (19)$$

These bounds illustrate that the slope-form Wigner distribution cannot be narrow in both s and u . Since the convolution kernel in the expression for the observable light

field (14) is a slope-form Wigner distribution, we can think of the variance along each axis to be the amount of blur in each axis. Since we have a lower bound for the product of the variance in s and u , we must trade off blur in s and u of the observable light field through the selection of an appropriate aperture transmission function. This tradeoff was reported in [25] for the case of light field microscopy.

However, if we step away from microscopic imaging and concentrate on macroscopic photographic applications, then we can make a geometric optics approximation where the wavelength of light is much smaller than the features in a scene we are imaging. Specifically, let's consider the case:

$$\epsilon^2 \Delta_s \Delta_u \gg \lambda \quad (20)$$

where Δ_s, Δ_u are the sizes of features along s, u respectively in the slope-form Wigner distribution of the scalar field at some plane after the scene, and ϵ is a small number. In this case, we can choose an aperture transmission function for the observable light field where:

$$\Delta_s \gg \sigma_s \quad , \quad \Delta_u \gg \sigma_u \quad (21)$$

and still be able to satisfy (19). Since the features in the slope-form Wigner distribution of the scalar field are much larger than the size of the convolution kernel, we can approximate the convolution kernel with a Dirac delta function. Applying this approximation to (14), we obtain:

$$l_{obs}^{(T)}(s, u) \approx \hat{W}_U^{(\lambda)}(s, u) \otimes \delta(-s, u) = \hat{W}_U^{(\lambda)}(s, u) \quad (22)$$

Therefore, at the geometric optics limit, the observable light field is equal to the slope-form Wigner distribution.

4. Applications

Now that we have derived an expression for the observable light field in terms of Wigner distributions and used it to show the equivalence between the light field and the Wigner distribution at the geometric optics limit, let us now apply these results to real applications. We will first investigate the relationship between the observable light field and a physical light field capture system, the plenoptic camera. Then, we will investigate the equivalence between the light field and the Wigner distribution by looking at similar papers from both the optics and computer graphics literatures.

4.1. The Aperture and Plenoptic Imaging

The first application we will consider is a light field capture system that uses an array of microlenses conjugate with the original imaging plane and an imager conjugate with the microlenses' back focal plane[1]. The output of such a system is a discrete light field: $l_{discrete}[m, n, p, q]$. The integer coordinates m, n enumerate which particular microlens the

light passed through to reach the sensor, and the p, q coordinates enumerate which pixel behind microlens is capturing this light.

The m, n coordinates correspond to the spatial coordinates s, t in the light field and the p, q coordinates correspond to the angular coordinates u, v in the light field. Let Δ_m, Δ_n be the pitch of the microlenses along the x and y axes and Δ_p, Δ_q be the pitch of the imaging pixels along the same two axes. To simplify notation, we will omit all constant scale factors in the equations in this section.

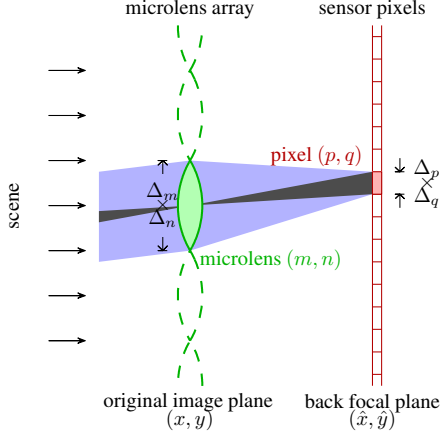


Figure 3. The plenoptic camera consists of a microlens array (green) placed at the original image plane and an imager (red) placed at the back focal plane of the microlens array. A particular value $l_{discrete}[m, n, p, q]$ is the total energy of light (blue) collected by a single pixel (p, q) (filled in red) through a single microlens (m, n) (filled in green). Each microlens aperture is $\Delta_m \times \Delta_n$ in size and each pixel is $\Delta_p \times \Delta_q$ in size. The spatial extent of the pixel causes it to capture a spread of angles in the scene (gray).

For a particular microlens (m, n) , as illustrated in Fig. 3, the image on the back focal plane is the magnitude squared of the impinging scalar field's Fourier transform[16]:

$$I(\hat{x}, \hat{y}) = \left| \iint U(x, y) T(x - m\Delta_m, y - n\Delta_n) \times e^{-j\frac{2\pi}{\lambda f}(x\hat{x} + y\hat{y})} dx dy \right|^2 \quad (23)$$

where $T(x, y) = \text{rect}(x/\Delta_m) \text{rect}(y/\Delta_n)$ is the aperture transmission function of a single microlens and f is its focal length. Using (9) and (10), we can rewrite (23) as:

$$I(\hat{x}, \hat{y}) = l_{obs}^{(T)}(m\Delta_m, n\Delta_n, \hat{x}/f, \hat{y}/f) \quad (24)$$

The total amount of energy captured by the pixel p, q behind microlens m, n can then be written as an integral over the area occupied by a single pixel:

$$l_{discrete}[m, n, p, q] = \iint I(\hat{x}, \hat{y}) P(\hat{x} - p\Delta_p, \hat{y} - q\Delta_q) d\hat{x} d\hat{y} \quad (25)$$

where $P(\hat{x}, \hat{y}) = \text{rect}(\hat{x}/\Delta_p) \text{rect}(\hat{y}/\Delta_q)$ is a function that is 1 inside the pixel and 0 outside. Converting imaging coordinates on the back focal plane to slope coordinates involves setting $u = \hat{x}/f$ and $v = \hat{y}/f$. Applying this transformation to (23) and (25), we obtain:

$$l_{discrete}[m, n, p, q] = \iint l_{obs}^{(T)}(m\Delta_m, n\Delta_n, u, v) \tilde{P}(u - p\Delta_u, v - q\Delta_v) du dv \quad (26)$$

where $\tilde{P}(u, v) = P(uf, vf)$ is the spread of slopes or angles in the scene that a single pixel captures (illustrated in gray in Fig. 3), and $\Delta_u = \Delta_p/f, \Delta_v = \Delta_q/f$ are the new sampling rates due to the coordinate transformation.

The expression in (26) is that of a convolution followed by sampling at intervals of $\Delta_m, \Delta_n, \Delta_u, \Delta_v$ along s, t, u, v . In other words, the discrete light field captured by this plenoptic camera is equivalent to a sampled observable light field with \tilde{P} as a two-dimensional prefilter. Therefore, we can think of the aperture transmission function as the aperture of a single microlens in a plenoptic camera. Equivalently, since the observable light field is itself a convolution, evident from (14), we can think of this discrete light field as a sampled version of the slope-form Wigner distribution with the following four-dimensional prefilter:

$$\hat{W}_T(-s, -t, u, v) \otimes \tilde{P}(u, v) \quad (27)$$

Therefore, analyzing spatioangular tradeoffs in plenoptic capture systems with novel microlens shapes only requires computation of the slope-form Wigner distribution of the aperture transmission function of a single microlens.

There is one thing, however, that we must keep in mind if we are to use this discrete light field in applications where we must add samples together, such as in synthesizing focused images[19]. The observable light field is a function of only the intensity of the scalar field and does not contain phase information. Therefore, rays in the observable light field are assumed to be incoherent with respect to each other, and thus we will not obtain any coherence effects when summing rays.³

For example, in a double slit system like the one illustrated in Fig. 1, we would not obtain interference between the two slits by adding rays in the observable light field if our aperture is too small to cover both slits. This is due to the quadratic nature of the Wigner distribution discussed earlier in Section 2.

4.2. Image Refocusing

Focused images from a light field can be generated by integrating rays that intersect a desired focal surface[19], or through slicing in the Fourier domain[28]. The analogue

³This is not the case in the related fields of synthetic aperture radar and radio astronomy, where the phase is recorded and used to produce interference effects that increase the resolution.

in the optics community is the use of slices of the ambiguity function to simulate defocus[30, 9]. Let us see in more detail how similar the two methods are.




Figure 4. In order to calculate the image formed at a (x, y) plane located z behind the (s, t) plane, we can either use the scalar field measured at the (s, t) plane for a wave optics derivation, or use a light field with (s, t) being the reference plane and the (u, v) plane at infinity for a geometric optics derivation. The Fourier transform of the image formed at the (s, t) plane forms a horizontal slice (red dashed) of both the ambiguity function of the scalar field (a) and the Fourier transform of the light field (b), while the Fourier transform of the image formed at the (x, y) plane forms a tilted slice (blue solid) of both the ambiguity function (a) and Fourier transform of the light field (b).

Suppose we have an optical system, as shown in Fig. 4. The scalar field is measured at the (s, t) plane, and the light field is measured with the (s, t) plane being the reference plane and the (u, v) plane at infinity. Our goal is to produce the image at the (x, y) plane, which is z away from the (s, t) plane along the optical axis.

Using light fields, we can write the image as:

$$I_{geom}(x, y) \propto \iint l(x - uz, y - vz, u, v) dudv \quad (28)$$

where $l(s, t, u, v)$ is the light field. Using the generalized Fourier-slice theorem[28], which equates projection in the original domain with slicing in the Fourier domain, we can rewrite (28) as a slicing operation:

$$I_{geom}(x, y) \propto \iint L(f_x, f_y, z f_x, z f_y) e^{j2\pi(f_x x + f_y y)} df_x df_y \quad (29)$$

where $L(f_s, f_t, f_u, f_v)$ is the 4D Fourier transform of the light field.

Alternatively, using wave optics the image formed at the (x, y) plane is the inverse Fourier transform of a slice of the

ambiguity function[9]:

$$I_{wave}(x, y) \propto \iint A_{\hat{U}}(f_x, f_y, 0, 0) e^{j2\pi(f_x x + f_y y)} df_x df_y \quad (30)$$

where $A_{\hat{U}}(f_x, f_y, \xi, \eta)$ is the ambiguity function of the scalar field $\hat{U}(x, y)$ at the (x, y) plane. This ambiguity function can be written in terms of the ambiguity function of the scalar field $U(s, t)$ at the (s, t) plane through a coordinate transformation[30]:

$$A_{\hat{U}}(f_x, f_y, \xi, \eta) = A_U(f_x, f_y, \xi - z\lambda f_x, \eta - z\lambda f_y) \quad (31)$$

where $A_U(f_s, f_t, \sigma, \tau)$ is the ambiguity function for the scalar field $U(s, t)$ at the (s, t) plane. Substituting (31) into (30), we obtain:

$$I_{wave}(x, y) \propto \iint A_U(f_x, f_y, -z\lambda f_x, -z\lambda f_y) e^{j2\pi(f_x x + f_y y)} df_x df_y \quad (32)$$

This is also a slicing operation, like (29). To highlight the equivalence of (29) and (32), we can rewrite (32) in a form analogous to (29):

$$I_{wave}(x, y) \propto \iint \hat{L}_U(f_x, f_y, z f_x, z f_y) e^{j2\pi(f_x x + f_y y)} df_x df_y \quad (33)$$

where \hat{L}_U is the 4D Fourier transform of the slope-form Wigner distribution and is proportional to the ambiguity function after a coordinate change, i.e.:

$$\begin{aligned} \hat{L}_U(f_s, f_t, f_u, f_v) &= \\ & \iiint \hat{W}_U^{(\lambda)}(s, t, u, v) \\ & \quad \times e^{j2\pi(f_s s + f_t t + f_u u + f_v v)} ds dt du dv \\ &= \iiint W_U\left(s, t, \frac{u}{\lambda}, \frac{v}{\lambda}\right) \\ & \quad \times e^{j2\pi(f_s s + f_t t + f_u u + f_v v)} ds dt du dv \\ &= \lambda^2 A_U(f_s, f_t, -\lambda f_u, -\lambda f_v) \end{aligned} \quad (34)$$

4.3. Cubic Phase Plate Wavefront Coding

Dowski and Cathey have proposed an optical system using a cubic phase plate, which after digital processing produces an image with extended depth of field[11]. This is accomplished by creating an intentionally aberrant point spread function (PSF) that varies little with defocus, then deconvolving with that PSF. The original analysis was conducted using slices of the ambiguity function, but we can conduct a similar analysis using light fields. Let us restrict ourselves to flatland to simplify the notation.

Fig. 5 shows the setup of the optical system. A convex lens with focal length f is placed at the $z = f$ plane in order to create the Fourier transform at the $z = 2f$ plane of a point source at the $z = 0$ plane. We then pass the light through a




Figure 5. In an optical system used for extended depth of field wavefront coding, light from a single point (a) passes through a Fourier transforming convex lens (b), a cubic phase plate (c) and another Fourier transforming convex lens (d) to arrive at an intentionally aberrant point spread function. The PSF at plane (e) varies very little from the PSF at plane (f).

cubic phase plate with thickness $\Delta(x) = \frac{\alpha x^3}{k(n-1)}$ where n is the index of refraction of the plate and $k = \frac{2\pi}{\lambda}$ is the wave number of the light. This results in a phase delay of αx^3 . Finally, we place another lens at $z = 3f$ to create at $z = 4f$ the Fourier transform of the scalar field immediately after the cubic phase plate. This produces an aberrant PSF that is roughly focus-invariant. Hence, deconvolving after image capture will result in an image with extended depth of field.

In [11], Dowski and Cathey derive an expression for the ambiguity function of a scalar field created by the cubic phase plate. They argue that since slices through the ambiguity function vary little with the slope of the slice, the PSF generated for various degrees of misfocus must also vary little with slope, since the ambiguity function is a representation of the Fourier transform of the PSF at various degrees of defocus[9].

We can form a similar argument using a light field derivation. At the $z=0$ plane, suppose we have a single point emitter. Our light field is thus a line embedded in 2D:

$$l_0(s, u) = \delta(s) \quad (35)$$

Propagating to the right by f , passing through a lens of focal length f , and then propagating again to the right by f results in a cumulative ray transfer matrix[17] of:

$$\begin{pmatrix} s_{in} \\ u_{in} \end{pmatrix} = \begin{pmatrix} 0 & -f \\ 1/f & 0 \end{pmatrix} \begin{pmatrix} s_{out} \\ u_{out} \end{pmatrix} \quad (36)$$

Hence, the light field immediately before the cubic phase plate can be written as:

$$l_{2f-}(s, u) = l_0(-uf, s/f) = \delta(-uf) \quad (37)$$

An arbitrary phase plate with phase delay $\phi(s)$ incurs the following transformation on the light field:

$$l_{out}(s, u) = l_{in} \left(s, u - \frac{\partial \phi(s)}{\partial s} \lambda / 2\pi \right) \quad (38)$$

The cubic phase plate in this system has a phase delay of the form $\phi(s) = \alpha s^3$. Therefore, by differentiating this phase

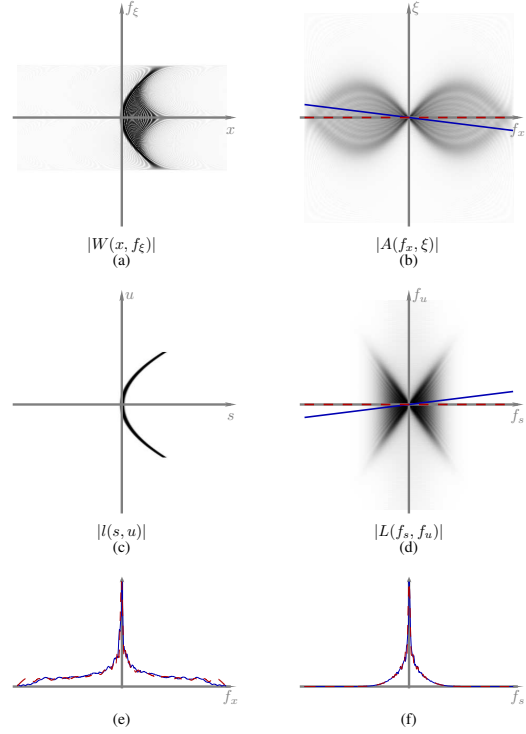


Figure 6. The impulse response, taken at $z = 4f$, of the cubic phase plate system, viewed as a Wigner distribution (a), an ambiguity function (b), a light field (c), and its Fourier transform (d). Horizontal (red dashed) and tilted (blue) slices (e) through the ambiguity function are very similar in magnitude. The same is true for slices (f) through the Fourier transform of the light field. These results were generated numerically and thus only simulate a finitely large cubic phase plate. An analytic derivation for the Wigner distribution is available in Appendix A. The ambiguity function (b) is derived and used in [11], while the light field (c) and its Fourier transform (d) are derived and used in [23, 21].

function according to (38), we obtain a light field immediately after the phase plate that is only nonzero on a parabola:

$$\begin{aligned} l_{2f+}(s, u) &= l_{2f-}(s, u - 3\alpha s^2 \lambda / 2\pi) \\ &= l_0(-f(u - 3\alpha s^2 \lambda / 2\pi), s/f) \\ &= \delta[-f(u - 3\alpha s^2 \lambda / 2\pi)] \end{aligned} \quad (39)$$

Finally, applying (36) for the lens at $z = 3f$, we find that the light field at $z = 4f$ is also zero everywhere except for a parabola, as shown in (c) of Fig. 6:

$$\begin{aligned} l_{4f}(s, u) &= l_{2f+}(-uf, s/f) \\ &= l_0(-s + 3\alpha u^2 f^3 \lambda / 2\pi, -u) \\ &= \delta[-s + 3\alpha u^2 f^3 \lambda / 2\pi] \end{aligned} \quad (40)$$

This result was independently derived in Appendix A of [22] and illustrated in Fig. 1 (i) of [21]. Using (28), the

PSF of this system at $z = 4f$ is a projection along u, v of the light field at $z = 4f$, while the PSF at some other plane $z = 4f + \Delta z$ is the projection along u, v of a shear of the light field. One property of a parabola is that it only translates when sheared. Therefore, the projection of a sheared parabola is a shifted projection of the original parabola[23]. Hence, for two different planes, we have shown that the PSF only varies up to a spatial shift.

The magnitude of the Fourier transform of the PSF, the magnitude transfer function, would be the same for the two planes, since shifts in position only cause linear phase shifts in the Fourier transform. Using the generalized Fourier slice theorem[28], which states that the Fourier transform of a projection of a function is equivalent to a slice of the Fourier transform of the same function, we can also say that different slices of the Fourier transform of the light field of this system have the same magnitude. This is analogous to Dowski and Cathey’s argument that slices of the ambiguity function at different slopes have the same magnitude.

The visual similarities, shown in Fig. 6, between the light field and the Wigner distribution (derivation in Appendix A), and between the Fourier transform of the light field and the ambiguity function are the result of the equivalence we showed in Section 3. Thus, it is possible to analyze the depth-invariance properties of a cubic phase mask using only the light field.

5. Conclusion and Future Work

We’ve shown that analysis using the light field in geometric optics is analogous to analysis using the Wigner distribution in wave optics. Therefore, when reading through the optics literature, the reader can associate the concept of a Wigner distribution with the light field and the concept of an ambiguity function with the Fourier transform of the light field and vice versa.

With this equivalence in mind, we may want to thoroughly explore current research topics in light fields using wave optics. For example, it may be interesting to modify the analysis of scene extraction from various imaging configurations in [21] to include diffraction effects. Furthermore, diffraction effects may alter the benefits and drawbacks in light field capture using dappled photography[31]. Lastly, while we’ve mainly analyzed light field capture and its coherence properties, we have not looked closely at light field generation and its coherence properties, which may be important for simulating certain types of illumination.

It may also be useful to find papers in the optics literature on Wigner distributions that can also be applied to light fields. For instance, there are various optical devices for capturing Wigner distributions[4, 26]. In particular, the device in [26] has no parallel in the computer graphics literature. It uses optical phase conjugators, a nonlinear optical element that inverts the phase of an incoming scalar field,

to produce the Wigner distribution optically. It would be interesting to adapt systems like this for light field capture.

Conversely, it may be useful to adapt some of the work done on light fields to the optics literature. For example, a three-dimensional manifold inside the Fourier transform of the light field is both the only region of interest for focused image generation[28] and the only nonzero region for an isotropically emitting medium[25]. These ideas may make capture of the Wigner distribution for coherent imaging analysis faster under certain situations.

Lastly, it may be useful to compare different solutions to the same problem in the optics and computer graphics literature for performance tradeoffs. For example, plenoptic imaging can also be used to extend the depth of field[29] using the approach in [2], which is different from [11]. Being able to compare the signal-to-noise performance as a function of spatial frequency in the two methods will enable the selection of the right method for the right situation when an extended depth of field image is needed. Furthermore, this comparison may yield further insights into extended depth of field imaging.

Just as the Wigner distribution and the ambiguity function were originally developed for quantum mechanics and radar imaging, respectively, and then later adapted to optical analysis, adaptation of these representations to computer graphics may bear further fruit.

Acknowledgment

We would like to thank Anat Levin, Fredo Durand and Bill Freeman for discussions while developing the ideas in this paper. Furthermore, the work was funded in part by a Stanford Graduate Fellowship from Texas Instruments and National Science Foundation grant number CCF-0540872.

A. Wigner Distribution of a Cubic Phase Plate

Deriving the cubic phase plate system’s output Wigner distribution will allow us to further explore its similarity with the light field. The scalar field right after the phase plate is $e^{j\alpha x^3}$ and its Wigner distribution can be obtained using the derivation given in (8.47)-(8.52) of [10]:

$$W_U^{2f+}(x, f_\xi) = 2\pi \left(\frac{12}{\alpha}\right)^{1/3} \text{Ai} \left(\left(\frac{12}{\alpha}\right)^{1/3} [3\alpha x^2 - 2\pi f_\xi] \right) \quad (41)$$

where $\text{Ai}(x) = \frac{1}{\pi} \int_0^\infty \cos(u^3/3 + xu) du$ is the Airy function. Passing this scalar field through a Fourier transform system results in a 90 degree rotation of the Wigner distribution through the application of ray transfer matrices[17, 5]:

$$W_U^{4f}(x, f_\xi) = 2\pi \left(\frac{12}{\alpha}\right)^{1/3} \text{Ai} \left(\left(\frac{12}{\alpha}\right)^{1/3} [3\alpha \lambda^2 f_\xi^2 f^2 - \frac{2\pi x}{\lambda f}] \right) \quad (42)$$

This Wigner distribution we obtain, shown in Fig. 6 (a), is the convolution of a two-dimensional function:

$$\delta \left(x - 3\alpha \lambda^2 f_\xi^2 f^3 \frac{\lambda}{2\pi} \right) \quad (43)$$

that is only nonzero on a parabola, with an Airy function:

$$2\pi \left(\frac{12}{\alpha}\right)^{1/3} \text{Ai}\left(-\left(\frac{12}{\alpha}\right)^{1/3} 2\pi x\right) \quad (44)$$

The parabola in (43) is the same as the one in (40) after converting slopes to frequencies as per (8).

References

- [1] E. H. Adelson and J. Y. Wang. Single lens stereo with a plenoptic camera. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):99–106, Feb. 1992. 5
- [2] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen. Interactive digital photomontage. In *Proc. ACM SIGGRAPH*, 2004. 9
- [3] N. A. Ahuja and N. K. Bose. Spatiotemporal-bandwidth product of m -dimensional signals. In *IEEE Signal Process. Lett.*, pages 123–125, 2005. 5
- [4] H. O. Bartelt, K.-H. Brenner, and A. W. Lohmann. The wigner distribution function and its optical production. *Opt. Commun.*, 32(1):32–38, Jan. 1980. 9
- [5] M. J. Bastiaans. The wigner distribution function applied to optical signals and systems. *Opt. Commun.*, 25(1):26–30, Apr. 1978. 1, 9
- [6] M. J. Bastiaans. *The Wigner Distribution - Theory and Applications in Signal Processing*, chapter Application of the Wigner distribution in optics. Elsevier Science, 1997. 3, 5
- [7] M. Born and E. Wolf. *Principles of optics, 7th. ed.* Cambridge University Press, Cambridge, UK, 2005. 3
- [8] R. N. Bracewell. *The Fourier transform and its applications, 2nd. ed.* McGraw-Hill, Reading, NY, 1986. 5
- [9] K.-H. Brenner, A. W. Lohmann, and J. Ojeda-Castañeda. The ambiguity function as a polar display of the OTF. *Opt. Commun.*, 44(5):323–326, Feb. 1983. 1, 7, 8
- [10] L. Cohen. *Time-Frequency Analysis*. Prentice Hall, Upper Saddle River, NJ, 1995. 9
- [11] E. R. Dowski, Jr. and W. T. Cathey. Extended depth of field through wave-front coding. *Applied Optics*, 34(11):1859–1866, Apr. 1995. 1, 7, 8, 9
- [12] J. T. Foley and E. Wolf. Radiometry as a short-wavelength limit of statistical wave theory with globally incoherent sources. *Opt. Commun.*, 55(4):236–241, Sept. 1985. 3
- [13] A. T. Friberg. On the existence of a radiance function for finite planar sources of arbitrary states of coherence. *J. Opt. Soc. Am.*, 69(1):192–198, Jan. 1979. 1, 3
- [14] W. D. Furlan, M. Martínez-Corral, B. Javidi, and G. Saavedra. Analysis of 3-d integral imaging displays using the wigner distribution. *J. Display Technol.*, 2(2):180–185, June 2006. 1
- [15] A. Gershun. The light field (translated by P. Moon and G. Timoshenko). *J. Math. and Physics*, 18:51–151, 1939. 1
- [16] J. W. Goodman. *Introduction to Fourier Optics, 3rd. ed.* Roberts and Company Publishers, Greenwood Village, CO, 2004. 2, 4, 6
- [17] K. Halbach. Matrix representation of gaussian optics. *American Journal of Physics*, 32(2):90–108, Feb. 1964. 1, 8, 9
- [18] P. Hanrahan. *Radiosity and Realistic Image Synthesis*, chapter Rendering Concepts, pages 13–40. Academic Press Professional, 1993. 3
- [19] A. Isaksen, L. McMillan, and S. J. Gortler. Dynamically reparameterized light fields. In *Proc. ACM SIGGRAPH*, 2000. 1, 6
- [20] K. Kim and E. Wolf. Propagation law for Walther’s first generalized radiance function and its short-wavelength limit with quasi-homogeneous sources. *J. Opt. Soc. Am. A*, 4(7):1233–1236, July 1987. 3
- [21] A. Levin, W. T. Freeman, and F. Durand. Understanding camera trade-offs through a bayesian analysis of light field projections. In *Proc. ECCV*, Oct. 2008. 8, 9
- [22] A. Levin, W. T. Freeman, and F. Durand. Understanding camera trade-offs through a bayesian analysis of light field projections. Technical Report CSAIL TR 2008-049, MIT, 2008. 8
- [23] A. Levin, P. Sand, T. S. Cho, F. Durand, and W. T. Freeman. Motion-invariant photography. In *Proc. ACM SIGGRAPH*, 2008. 8, 9
- [24] M. Levoy and P. Hanrahan. Light field rendering. In *Proc. ACM SIGGRAPH*, 1996. 1, 3
- [25] M. Levoy, R. Ng, A. Adams, M. Footer, and M. Horowitz. Light field microscopy. In *Proc. ACM SIGGRAPH*, 2006. 5, 9
- [26] Y. Li, G. Eichmann, and M. Conner. Optical wigner distribution and ambiguity function for complex signals and images. *Opt. Commun.*, 67(3):177–179, July 1988. 9
- [27] E. W. Marchand and E. Wolf. Walther’s definitions of generalized radiance. *J. Opt. Soc. Am.*, 64(9):1273–1274, Sept. 1974. 3
- [28] R. Ng. Fourier slice photography. In *Proc. ACM SIGGRAPH*, 2005. 1, 6, 7, 9
- [29] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a hand-held plenoptic camera. Technical Report CTSR 2005-02, Stanford University, 2005. 9
- [30] A. Papoulis. Ambiguity function in fourier optics. *J. Opt. Soc. Am.*, 64(6):779–788, June 1974. 1, 7
- [31] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan, and J. Tumblin. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. In *Proc. ACM SIGGRAPH*, 2007. 9
- [32] A. Walther. Radiometry and coherence. *J. Opt. Soc. Am.*, 58(9):1256–1259, Sept. 1968. 1, 3
- [33] E. Wigner. On the quantum correction for thermodynamic equilibrium. *Physical Review*, 40(5):749–759, 1932. 1
- [34] E. Wolf. Coherence and radiometry. *J. Opt. Soc. Am.*, 68(1):6–17, Jan. 1978. 1, 3
- [35] P. M. Woodward. *Probability and information theory with applications to radar*. Pergamon, London, U.K., 1953. 1
- [36] R. Ziegler, S. Bucheli, L. Ahrenberg, M. Magnor, and M. Gross. A bidirectional light field - hologram transform. *Computer Graphics Forum (Proc. Eurographics EG’07)*, 26(3):435–446, Sept. 2007. 1