

An Eternal Irradiance Camera

JEREMY KLOTZ, Computer Science Department, Columbia University, USA

SHREE K. NAYAR, Computer Science Department, Columbia University, USA




Fig. 1. **An irradiance camera and its applications.** (a) We present a novel camera that measures the total illumination (irradiance) falling on every point of a sphere. Our prototype camera, called FluxCam, has a minimalist design: it uses a very small number of detectors to both measure irradiance and harvest energy. This allows FluxCam to be completely powered by the light falling on it, i.e., it can read out and wirelessly transmit its measurements at a framerate of 30 FPS or more without a battery, cable, or external power supply. FluxCam can be used for (b) estimating egomotion and (c) monitoring a workspace. In these cases, the camera is not only self-powered and untethered but also privacy preserving since it does not capture the visual details needed to identify individuals. FluxCam can also be used as (d) an untethered light probe for diffuse relighting and (e) an omnidirectional pyranometer for analyzing the sky's illumination.

A conventional camera uses millions of pixels to measure radiance from all directions within its field of view. We present an omnidirectional *irradiance camera* that measures the irradiance function—the illumination incident upon every point on a sphere. The irradiance function varies smoothly over the sphere and hence is bandlimited. We have analyzed this function in the frequency domain and have shown that it is well approximated by a weighted sum of the first seven degrees of spherical harmonics. As a result, the irradiance function can be accurately reconstructed from a small number of samples. This implies that an irradiance camera does not need millions of detectors (pixels)—just a handful of measurements suffice. This brings two major benefits. First, the camera consumes such little energy that it can be completely powered by the light falling on its detectors. Second, it does not capture the visual details needed to identify an individual, and hence privacy is preserved. We have built a prototype irradiance camera, called FluxCam, using 49 detectors arranged on the surface of a sphere. In a well-lit indoor environment, FluxCam can read out and wirelessly transmit its measurements at 30 frames per second using energy harvested from the light falling on it (i.e., without a battery, cable, or external power supply). We show how FluxCam can be used as an optical gyroscope for computing rotation, to monitor a workspace, as an untethered light probe for diffuse relighting,

and as an omnidirectional pyranometer for estimating sky conditions and determining the best orientation of a solar panel.

CCS Concepts: • **Hardware** → **Wireless devices; Displays and imagers;** • **Computing methodologies** → **Computer vision.**

Additional Key Words and Phrases: Irradiance Camera, Irradiance Function, Self-Powered Imaging, Optical Gyroscope, Lightweight Vision, Workspace Monitor, Untethered Light Probe, Solar Energy

ACM Reference Format:

Jeremy Klotz and Shree K. Nayar. 2026. An Eternal Irradiance Camera. In *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26)*, December 01–04, 2026, Kuala Lumpur, Malaysia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3829340.3842234>

1 What is an Irradiance Camera?

Consider a scene viewed from a point. The brightness of the scene can be represented as $L(\vec{n})$, where L is the scene radiance in the direction $\vec{n}$. The goal of a conventional omnidirectional camera is to measure the radiance function $L(\vec{n})$. One way to represent such a camera is by arranging detectors (pixels) on a sphere, with a lens above each detector.¹ This architecture is illustrated in fig. 2(a), where each detector measures radiance from a single direction. Since a conventional camera densely samples the radiance function, it can be thought of as a *radiance camera*.

Suppose now that we remove the lenses above the detectors in fig. 2(a), as shown in fig. 2(b). Each detector now measures the total

¹This simplified architecture can be used to represent a wide class of camera designs: a camera with a single lens, a spherical camera with two lenses, a catadioptric camera with a lens and mirror, or a light probe consisting of a camera and mirror ball.

Authors' Contact Information: Jeremy Klotz, Computer Science Department, Columbia University, New York, NY, USA, jklotz@cs.columbia.edu; Shree K. Nayar, Computer Science Department, Columbia University, New York, NY, USA, nayar@cs.columbia.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SA Conference Papers '26, Kuala Lumpur, Malaysia*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2842-6/2026/12
<https://doi.org/10.1145/3829340.3842234>




Fig. 2. **What is an irradiance camera?** (a) A conventional omnidirectional camera uses millions of pixels to measure radiance from every direction. Each pixel can be modeled by a detector that views the scene through a lens. (b) If we remove the lenses, each detector measures irradiance received from a hemisphere. In effect, this camera measures the illumination falling on every point of a sphere. Since it measures the irradiance function, we call it an irradiance camera. (c) The irradiance function in any environment is extremely smooth, and hence an irradiance camera only needs a handful of detectors to measure it. (d) Since so few detectors are required, we can use large detectors that cover the entire surface of the sphere. Large detectors not only produce measurements with a wide dynamic range, but can also be used as solar cells that harvest energy from the incident illumination. This harvested energy can be used to fully power the camera using the light falling on it. As a result, the camera can operate forever, as an eternal camera, without a battery, cable, or external power supply.

irradiance it receives from a hemisphere, which can be written as an integral over the radiance function,

$$E(\vec{n}) = \int_{\vec{s} \in \mathbb{S}^2} L(\vec{s}) \max(\vec{n} \cdot \vec{s}, 0) d\omega, \quad (1)$$

where $\vec{n}$ is the vector normal to the detector, $d\omega$ is the differential solid angle, and the max operator accounts for the fact that a detector only receives light from a hemisphere. Equation (1) defines the irradiance function, which is fundamental to our work. Given that the detectors in fig. 2(b) sample the irradiance function, we refer to this camera as an *irradiance camera*.

The irradiance function at any point in a scene has a simple visual interpretation: it is equivalent (up to a scale factor) to the brightness distribution of an infinitesimally small Lambertian sphere with constant albedo placed at the point. Since each point on the sphere receives light from an entire hemisphere, the irradiance function is smooth, regardless of the complexity of the scene. It does not include the fine-grained details in the corresponding radiance function.

Although it is smooth, we posit that the irradiance function is an expressive signal for a wide class of visual processing tasks. In fact, the irradiance function at every point in space is a complete representation of Gershun’s light field [Gershun 1939], which describes the flow of illumination. When an irradiance camera moves through space, its irradiance function encodes its motion. Furthermore, there exists a large class of vision tasks, called lightweight tasks, that can be solved without sensing radiance functions [Klotz and Nayar 2025b]. An irradiance camera can be used to solve such tasks in an efficient manner.

In nature, several organisms are able to function in their worlds by sensing just the irradiance function [Land and Nilsson 2002]. Figure 3 shows the eye of a limpet and the eye of a flatworm, neither of which has a lens. In each case, the photoreceptors (detectors) are densely packed on a curved surface, similar to the detectors of the irradiance camera shown in fig. 2(b). Each photoreceptor measures irradiance due to a different field of view. Naturally, these eyes do



Fig. 3. **Irradiance cameras in nature.** The eyes of (a) a limpet and (b) a flatworm resemble an irradiance camera. These eyes do not have lenses. Instead, the photoreceptors (detectors) are arranged on a curved surface. Even though these eyes do not sense detailed images, they provide enough visual information for the organisms to function in their environments. Figures adapted from Land and Nilsson [2002].

not form a detailed image; however, they provide the organisms with sufficient visual information to navigate around and survive in their environment.

It has been shown analytically that the irradiance function is smooth (i.e., bandlimited), regardless of the complexity of the radiance function [Basri and Jacobs 2003; Ramamoorthi and Hanrahan 2001b]. Consequently, an irradiance camera does not need millions of detectors (as in fig. 2(b)) to measure the irradiance function—just a handful of detectors suffice (fig. 2(c)). Since so few detectors are needed, an irradiance camera can use large detectors that cover the entire surface of the sphere (fig. 2(d)). These large detectors collect more light, allowing them to not only produce high dynamic range measurements but also harvest enough energy to power the camera.

We have built the prototype irradiance camera shown in fig. 1(a) using 49 detectors arranged on the surface of a sphere. We refer to this prototype as FluxCam since its measurements are proportional to the light flux received by the surface of the sphere. Each detector is a photovoltaic cell that measures irradiance over a wide dynamic range, spanning 2 lx to 13,000 lx. When the photovoltaics are not used for sensing, they harvest energy from the light falling on them to power the camera. In a well-lit indoor environment, FluxCam

can read out and wirelessly transmit its measurements at 30 frames per second (FPS). In a bright outdoor environment, it can function at up to 200 FPS. So long as there is sufficient light to power the camera, it can operate forever, without a battery or power supply. For this reason, we refer to it as an *eternal irradiance camera*.

From the 49 measurements made by the camera, we reconstruct the irradiance function. We have validated the accuracy of our reconstruction method using both simulated and real-world measurements, in indoor and outdoor settings. We have used FluxCam for computing egomotion (fig. 1(b)) and monitoring a workspace by determining which desks are occupied, when a door is open, and which lamps are activated (fig. 1(c)). In these cases, the irradiance camera is not only self-powered and untethered but also privacy preserving since it does not capture the fine-grained visual details needed to identify individuals. In addition, we have used FluxCam as an untethered light probe for diffuse relighting (fig. 1(d)) and as an omnidirectional pyranometer for estimating the current sky conditions and determining the best orientation of a solar panel to maximize its harvested energy (fig. 1(e)).

2 Related Work

Prior work has established that the irradiance function in any environment is bandlimited [Basri and Jacobs 2003; Ramamoorthi and Hanrahan 2001b,c]. Building on this result, we show that the irradiance function can be accurately reconstructed using a very small number of samples (measurements). This enables us to build a camera that consumes very little energy and hence is fully self-powered.

A popular approach for measuring the radiance function is to use a light probe, consisting of a mirror and a camera that observes it from a distance [Debevec 1998]. Subsequent works have extended the mirror probe for measuring high dynamic range (HDR) radiance functions either by adding a diffuse probe [Reinhard et al. 2010] or by attaching diffuse patches to an existing mirror probe [Debevec et al. 2012]. All of these devices capture a detailed representation of incident light, which can be used for rendering objects with arbitrary shapes and material properties. However, the mirror and diffuse probes must be fully visible to the camera and hence cannot be placed at locations at which they are entirely, or even partially, occluded from the camera. Our irradiance camera, on the other hand, is an untethered light probe that measures the complete irradiance function for any location it can be physically placed at.

Although the irradiance function is not as expressive as the radiance function, it has proved useful in real-time rendering pipelines. Ward [1988] proposed irradiance caching as a technique for speeding up the computation of global illumination. Greger et al. [1998] introduced the irradiance volume to represent the irradiance function at every point in space. Ramamoorthi and Hanrahan [2001a] proposed a compact representation of the irradiance function, along with a fast algorithm for rendering diffuse objects. Our irradiance camera is designed to directly measure the irradiance function in real environments.

Prior work has shown that low-dimensional visual signals are sufficient for solving a variety of vision tasks, including human activity recognition [Bobick and Davis 2001; Dai et al. 2015], person detection [Nakashima et al. 2010], and even image classification [Torralba et al. 2008]. In this context, we show that the measurements

produced by an irradiance camera can be used for monitoring a workspace, i.e., determining which desks are occupied, when a door is open, and which lamps are on. We are able to solve these tasks with just 49 measurements, in part, because our detectors have high sensitivity and wide dynamic range. Interestingly, since the camera does not capture fine-grained visual details, it is difficult to identify an individual from the measurements, which is beneficial in applications where privacy needs to be preserved.

In meteorology and solar energy, pyranometers are widely used for obtaining the irradiance measurement corresponding to a single direction [Augustine et al. 2000; Duffie and Beckman 2006]. These measurements, made over time, can be used to obtain coarse estimates of the condition of the sky [Duchon and O'Malley 1999; Long and Ackerman 2000]. Others have used multiple pyranometers in different, fixed orientations to obtain more detailed descriptions of the sky condition [Baltazar et al. 2015; Faiman et al. 1992; Lamm and Adler 1987; Lee 2025]. The irradiance camera is essentially an omnidirectional pyranometer that provides these measurements without any moving parts. We use the camera to demonstrate multiple applications in this context: estimating the sky conditions, decomposing irradiance into sun and sky contributions, and finding the best orientation of a solar panel.

Our work is related to previous work on self-powered light sensing and imaging. A variety of light sensors (pixels) have been proposed that use the incident illumination to power the measurement and wireless transmission [Nayar et al. 2024; Zhang et al. 2020, 2019]. Nayar et al. [2015] developed a 30×40 px² camera that is powered by the light falling on the pixels. This camera harvested enough energy to read out the image, but not enough to wirelessly transmit it. The minimalist camera [Klotz and Nayar 2025b] used a handful of freeform pixels whose shapes are automatically learned for a given task. Since the power consumption of any camera is roughly linear in the number of detectors, the minimalist camera was able to read out and wirelessly transmit using energy harvested from the incident light. Given that the irradiance camera uses very few detectors, it is also powered by the light falling on it. However, rather than designing the camera for a specific task, our goal is to develop a general-purpose device that accurately measures the irradiance function for a variety of downstream tasks.

3 Measuring the Irradiance Function

We begin by analyzing the properties of the irradiance function using spherical harmonics.

3.1 Degrees of Freedom

Prior work by Basri and Jacobs [2003] and Ramamoorthi and Hanrahan [2001c] has shown that the irradiance function (eq. (1)) can be written as the spherical convolution² of the radiance function $L(\vec{n})$ with a kernel given by

$$k(\vec{n}) = \max(\vec{n} \cdot \hat{z}, 0), \quad (2)$$

where $\hat{z}$ is the zenith vector. Figure 4(a) shows the kernel, which is a truncated cosine function. Since the kernel is extremely smooth and broad, it acts as a low-pass filter on the radiance function. Building

²We refer the interested reader to Driscoll and Healy [1994] for an introduction to convolution on the sphere.




Fig. 4. **Spherical harmonic analysis of the kernel.** (a) The irradiance function is the result of convolving the radiance function with the kernel $k(\vec{n})$, which is defined as the cosine of the zenith angle (θ), but clamped to be non-negative. (b) The spherical harmonic coefficients of the kernel as a function of the spherical harmonic degree. The coefficients k_ℓ tend toward 0 as the degree (frequency) increases, as a result of the kernel being a low-pass filter. We see from the plot that k_ℓ is approximately zero for any degree $\ell \geq 7$ (shaded gray region to the right of the vertical red line). Therefore, the kernel can be approximated by a weighted sum of the first seven spherical harmonics. (c) A 1D slice of the kernel through the gray dotted line in (a). The bandlimited approximation computed from the first seven spherical harmonics (red line) closely follows the kernel (black line).

on this prior work, we analyze the kernel’s bandwidth. Figure 4(b) plots the spherical harmonic coefficients k_ℓ of the kernel as a function of the spherical harmonic degree ℓ .³ As can be seen in fig. 4(b), the coefficients k_ℓ decay rapidly and are approximately zero for $\ell \geq 7$, indicating that the kernel is effectively bandlimited. Figure 4(c) shows a 1D slice of the kernel computed using the first seven degrees of spherical harmonics overlaid on a 1D slice of the original kernel. The bandlimited approximation (red line) closely follows the original kernel (black line), indicating that the kernel is well approximated using this bandlimit.

Technically speaking, the kernel is not strictly bandlimited, and hence the accuracy of the approximation depends on the chosen bandwidth. Naturally, a higher bandwidth yields a more accurate approximation of the kernel at the cost of requiring more spherical harmonic bases. This trade-off is illustrated in the first two columns of table 1. Here, the bandwidth L denotes the number of spherical harmonic degrees used in the approximation. In this work, we approximate the kernel using the bandwidth $L = 7$, which captures more than 99.9% of the kernel’s total energy.

Since the kernel is (approximately) bandlimited, the irradiance function also has the same bandwidth. Therefore, we can approximate the irradiance function as a weighted sum of the first seven degrees of spherical harmonic bases:

$$E(\vec{n}) \approx \sum_{\ell=0}^6 \sum_{m=-\ell}^{\ell} E_{\ell m} Y_{\ell m}(\vec{n}), \quad (3)$$

where $E_{\ell m}$ is the spherical harmonic coefficient of degree ℓ and order m , and $Y_{\ell m}$ is the corresponding real spherical harmonic basis function. In general, there are L^2 spherical harmonics in the first L degrees. Therefore, the approximation in eq. (3) is a weighted sum of 49 spherical harmonic bases. However, we know $k_\ell = 0$ for odd degrees $\ell > 1$, and hence the convolution theorem [Driscoll and Healy 1994] implies that the spherical harmonic coefficients $E_{\ell m}$ of those degrees are zero as well. Therefore, eq. (3) reduces to a weighted sum of 31 spherical harmonics (third column of table 1). *In other words, a good approximation of the irradiance function for any environment has just 31 degrees of freedom.*

³Since the kernel is zonal (i.e. symmetric about the north pole), its spherical harmonic coefficients are zero for any order $m \neq 0$. For brevity of notation, we have therefore dropped the dependence on m in k_ℓ .

Table 1. **Impact of the kernel’s approximate bandwidth.** While the first three degrees of spherical harmonics capture 99% of the kernel’s energy (second column) [Basri and Jacobs 2003; Ramamoorthi and Hanrahan 2001c], a higher bandwidth is needed to accurately approximate the irradiance function in real environments. The number of degrees of freedom in the approximated irradiance function (third column) is the number of non-zero spherical harmonic coefficients in the first L degrees. The error in the irradiance function (fourth column) is the error due to the bandwidth, as a percentage of the maximum irradiance in the scene, across 3,299 scenes. In this work, we approximate the kernel (and hence the irradiance function) using the first seven degrees of spherical harmonics (highlighted row).

Bandwidth (L)	Kernel	Irradiance Function	
	Energy Captured (%)	D.O.F.	Approx. Error (%)
1	37.50	1	25.23
2	87.50	4	9.08
3	99.22	9	1.77
5	99.81	18	0.77
7	99.92	31	0.44
9	99.96	48	0.28
11	99.98	69	0.20

3.2 Making the Measurements in the Spatial Domain

One approach for sampling the irradiance function would be to directly measure the spherical harmonic coefficients $L_{\ell m}$ of the radiance function. Then, the coefficients $E_{\ell m}$ could be computed using the convolution theorem.

This approach, however, is impractical for several reasons. First, computing a single coefficient requires a detector that views the scene through a spherical amplitude mask, whose spatially varying transmittance function is given by the corresponding spherical harmonic (see fig. 5(a)). Spherical masks, however, are difficult to fabricate. Second, the detector must be infinitesimally small such that the entire mask is viewed from essentially a single point. Third, a single detector only sees one hemisphere of the mask, and hence two detectors with hemispherical masks are needed to compute the result for the spherical mask. Furthermore, an amplitude mask can only have positive transmittance values, while spherical harmonics have both positive and negative values. Therefore, the measurement for each hemisphere would require the use of two hemispherical masks, one for the positive component and the other for the negative




Fig. 5. **A spherical harmonic sensor?** (a) One approach to sensing the irradiance function would be to directly measure its spherical harmonic coefficients. Each coefficient could be measured using a detector and mask whose transmittance is given by the corresponding spherical harmonic. (b) In practice, this would require four detectors to measure a single spherical harmonic coefficient, since each detector can only receive light from a hemisphere, and a mask can only have positive transmittance values. This approach would require 124 detectors and hemispherical masks, while the irradiance function estimation problem only has 31 degrees of freedom.

component. Hence, the measurement of each spherical harmonic coefficient would require the use of four masks and detectors, as shown in fig. 5(b). In short, measuring 31 coefficients would require 124 hemispherical masks and detectors.

Clearly, directly measuring the spherical harmonic coefficients is impractical. However, our analysis has revealed that a good approximation of the irradiance function only has 31 degrees of freedom. We therefore take a simpler approach of directly sampling the irradiance function in the spatial domain. This can be done by placing bare detectors (without masks) on the surface of a sphere. Given that so few detectors are needed, we can use large detectors that cover the entire surface of the sphere, as shown in fig. 2(d). The large detectors collect more light, allowing them to produce measurements with a high dynamic range and signal-to-noise ratio (SNR), and harvest energy to power the camera.

Although this analysis shows that 31 detectors (samples) are sufficient to reconstruct the irradiance function, we chose to use 49 detectors to ensure that the reconstruction from the measurements is an overdetermined system and hence less sensitive to noise.

3.3 Where Should the Samples Be?

Now that we have decided to sample the irradiance function in the spatial domain, we must determine the sample locations. The problem of evenly arranging points on a sphere was first formalized by Tammes [1930]: Given N points on the surface of a sphere, how should they be arranged to maximize the minimum distance between any two points? This problem, known as the Tammes problem, is well-studied [Conway and Sloane 1999], and various sources provide tabulated solutions that are approximately optimal [Cohn 2024; Sloane et al. 1994]. We use a solution for $N = 49$ to determine the locations of our detectors, which are shown in fig. 6.

3.4 Reconstructing the Irradiance Function

Once we have sampled the irradiance function, we need to reconstruct it. To this end, the sampling process can be written as a linear system by evaluating the irradiance function in eq. (3) at the sample




Fig. 6. **Sampling the irradiance function in the spatial domain.** Shown here is a roughly uniform arrangement of 49 samples (detectors) on the sphere, which is given by an approximate solution to the Tammes problem.

locations:

$$\begin{bmatrix} E(\vec{n}_1) \\ \vdots \\ E(\vec{n}_{49}) \end{bmatrix} = \begin{bmatrix} Y_{0,0}(\vec{n}_1) & \dots & Y_{6,6}(\vec{n}_1) \\ \vdots & \ddots & \vdots \\ Y_{0,0}(\vec{n}_{49}) & \dots & Y_{6,6}(\vec{n}_{49}) \end{bmatrix} \begin{bmatrix} E_{0,0} \\ \vdots \\ E_{6,6} \end{bmatrix}, \quad (4)$$

where $\{\vec{n}_1, \dots, \vec{n}_{49}\}$ is our set of 49 sample locations. Equation (4) forms an overdetermined system with 49 measurements and 31 unknowns (the non-zero spherical harmonic coefficients). We solve for the 31 coefficients using a constrained least squares solver to enforce the constraint that the irradiance function must be non-negative. The supplemental material includes details of this reconstruction method.

3.5 Simulations with Real Environment Maps

We have simulated our approach for sampling and reconstructing the irradiance function using real environment maps from the Laval Indoor [Gardner et al. 2017] and UrbanSky [Klotz and Nayar 2025a] datasets. These datasets together contain 3,299 high dynamic range images (radiance functions) taken in a variety of real indoor and outdoor scenes. Figure 7(a) shows a few of the images from the two datasets, and fig. 7(b) shows the corresponding ground truth irradiance functions.

First, we evaluate the accuracy of the bandlimited approximation to the irradiance function. For each scene, we compute the bandlimited irradiance function using eq. (3) and compare it to the ground truth. The error in each scene is then computed as the root mean square error in the bandlimited irradiance function as a percentage of the maximum irradiance in that scene. The fourth column in table 1 shows the average error across all 3,299 scenes due to the bandlimited approximation for different bandwidths L . Our chosen bandwidth of $L = 7$ yields an average error of 0.44%.

Next, we evaluate the accuracy of the irradiance function reconstructed from the 49 measurements. To simulate an irradiance camera, we assume that each detector produces a measurement proportional to its irradiance. We sample the irradiance functions at the 49 locations (blue dots overlaid in fig. 7(b)). In addition, we emulate the noise characteristics and dynamic range of a real sensor by adding Gaussian noise and applying a saturation function. Finally, we use the samples to reconstruct the irradiance function and compute the error with respect to ground truth. Figure 7(c) shows the reconstructed irradiance functions for each scene in fig. 7(a). Figure 7(d) shows a 1D slice through one of the reconstructions, where the slice corresponds to the gray dashed line in the first reconstructed irradiance function in fig. 7(c). On average, across




Fig. 7. **Simulations with real environment maps.** (a) Example HDR radiance functions. Although the radiance function in each scene is a spherical image, we only show one hemisphere here. (b) Corresponding ground truth irradiance functions, visualized using a colormap. We simulate the irradiance camera by sampling the irradiance function at the 49 sample locations (blue dots). (c) The reconstructed irradiance functions computed from the 49 measurements. (d) A 1D slice of the reconstruction corresponding to the dashed gray line in (c). The 49 measurements are sufficient for accurately reconstructing the irradiance function across diverse real environments.

all 3,299 scenes, the error between the reconstructed and ground truth irradiance functions was 1.41% with respect to the maximum irradiance in each scene. These results show that our approach for sampling the irradiance function yields accurate reconstructions across diverse environments. Please see the supplemental material for simulation details and additional results.

4 FluxCam: Irradiance Camera Prototype

We have built FluxCam, a prototype irradiance camera, using 49 detectors arranged on the surface of a sphere.

4.1 Camera Architecture

Each detector is a $25 \times 20 \text{ mm}^2$ photovoltaic cell (Panasonic AM-5610), one of which is shown in fig. 8(a). The photovoltaic is connected to a custom-designed sensor board that switches the photovoltaic between a transimpedance amplifier (OPA607) for measuring irradiance and an energy harvester (BQ25570) to power the camera. All 49 sensor boards are connected to a custom-designed processor board (fig. 8(b-c)). The processor board includes a microcontroller (STM32L051C6) for communicating with the sensor boards, an RF transmitter (TI CC1310) to wirelessly transmit the measurements, and a supercapacitor to store energy harvested by the photovoltaics.

The camera is typically asleep, with all of the photovoltaics being used for energy harvesting. To measure irradiance, the sensor board switches its photovoltaic into the transimpedance amplifier, waits for the photocurrent to settle, and then digitizes the corresponding voltage using a 16-bit analog-to-digital converter (ADC). Once the measurement is made, the photovoltaic is immediately switched back into the energy harvester. To produce a single frame (49 irradiance measurements), the processor board wakes up, sequentially reads out all 49 sensors as described above, wirelessly transmits the measurements, and then returns to sleep. The supplemental material includes further details of the camera architecture. In addition, the circuit schematics, PCB layout, and firmware are available online [Klotz and Nayar 2026].

4.2 Rotating Detectors to Minimize Camera Size

The detectors (photovoltaics) are placed on the surface of a sphere to sample the irradiance function at the 49 locations shown in fig. 6. Although the detector locations are fixed, each of them can be arbitrarily rotated about the axis that passes through its center and the center of the sphere without affecting its measured irradiance. This raises a practical question: how should the detectors be rotated such that they fit on the smallest possible sphere without overlapping?

As a baseline, consider the smallest sphere that is large enough to avoid detector overlap, regardless of the rotations of the detectors. This case is illustrated on the left side of fig. 8(d), in which the detector rotations are chosen at random, and the sphere is large enough such that the bounding circles (blue) do not overlap. A 1 mm padding (gold border) is added to the dimensions ($25 \times 20 \text{ mm}^2$) of each detector to ensure a minimum separation between detectors in the final design. This conservative approach to packing the detectors on the sphere yields a sphere with a diameter of 13.6 cm.

To obtain a more compact design, we developed a greedy algorithm to optimize the rotations of the detectors. Beginning with the conservatively sized sphere described above, we adjust the rotation of each detector to maximize its minimum distance from its neighboring detectors. Then, we decrement the diameter of the sphere by a small amount (1 mm) and repeat this process. This approach to shrinking the sphere is repeated until the sphere is too small to pack all of the detectors (inclusive of their paddings) without overlap. This method reduced the diameter of the sphere from 13.6 cm to 11.9 cm (fig. 8(d)). To add space for additional surface artifacts (e.g., mounting holes and debugging wires), we slightly increased the sphere’s diameter to 12.8 cm. We then 3D printed this design as four quadrants using stereolithography (see fig. 8(e)).

Figure 8(f) shows the inside of the camera during the assembly process. **The supplemental video shows how FluxCam was assembled.** The processor board is held at the center of the camera, and each photovoltaic fits into a slot on the outside of the shell. Each sensor board is mounted inside the shell (behind the photovoltaics) and has a cable that connects it to the processor board. The antenna is placed outside the camera to maximize the communication range. Figure 1(a) shows the fully assembled camera.

4.3 Radiometric Performance

We have characterized the performance of our sensor using the measurement setup in fig. 9(a). A controlled halogen light source




Fig. 8. **FluxCam: irradiance camera prototype.** (a) Each sensor board consists of a photovoltaic attached to a custom circuit. The photovoltaic is used both for measuring irradiance and harvesting energy. The (b) front side and (c) back side of the processor board show 49× connectors for the sensor boards, readout circuitry, and an RF transmitter. (d) We have optimized the rotations of the detectors to minimize the camera's size. Starting with a large sphere such that the bounding circles (blue) around each detector (dark brown) and its necessary padding (gold) do not overlap, we iteratively adjust the detector rotations and shrink the sphere until the detectors cannot be rotated without overlap. This yields a sphere with a smaller diameter. (e) We 3D printed the final design in four quadrants using stereolithography. (f) The inside of the camera. **The supplemental video includes a time-lapse clip of the camera assembly.**

(Thorlabs OSL2) was used to illuminate the detector head-on, and a lux meter placed next to the detector was used to measure the ground truth illuminance.

The measurements made by FluxCam are weighted by the spectral response of the photovoltaics. In our case, the photovoltaics have a spectral response that closely resembles the photopic response of the human eye [Panasonic 2021]. Due to this spectral response, FluxCam measures illuminance (in lux) and hence can be calibrated using a lux meter. If the spectral distribution of the incident illumination is known, then irradiance can be computed from the measured illuminance using a scale factor that can be determined via calibration.

Figure 9(b) shows the sensor's radiometric response function. The sensor is capable of measuring light levels as high as 13,000 lx, and the response function is linear over the entire dynamic range. Figure 9(c) shows the SNR of the sensor, which is higher than 60 dB for light levels brighter than 750 lx. At just 2.2 lx, the SNR is 2.4 dB. If we use this light level as the minimum detectable signal, this corresponds to a remarkably wide dynamic range of 75.4 dB. This performance can be attributed to the detector's large active area and the absence of any attenuating optics, allowing the detector to collect substantially more light than a conventional pixel. Details of the sensor performance measurements are provided in the supplemental material.

4.4 Self-Powered Operation

As described in Section 4.1, in addition to measuring irradiance, the 49 sensors also harvest energy to power FluxCam. Figure 9(d) shows the power harvested by a single sensor as a function of light level. The plot does not pass through the origin since the energy harvester circuit itself consumes a small amount of power. As a result, the sensor requires a minimum illumination of 71 lx to harvest energy. This threshold is a limitation of the harvester circuit and is independent of the sensor's dynamic range. As described before, all sensors can measure light levels as low as 2.2 lx. In environments for which the average illumination falling on the camera is roughly 1,000 lx, the total energy harvested by all 49 sensors is sufficient to operate the camera at 30 FPS. In the supplemental material, we have included details of the experimental setup and an algorithm to adaptively adjust the camera's framerate based on the harvested power (i.e., the brightness of the environment).

4.5 Accuracy of Reconstructed Irradiance Functions

We have evaluated the accuracy of the irradiance functions reconstructed using FluxCam by comparing them with ones computed using HDR spherical images (taken with a Ricoh Theta Z1) in four indoor and outdoor scenes. Two of these HDR images are shown in fig. 10(a), and their corresponding irradiance functions (ground truth) are shown in fig. 10(b). The irradiance functions reconstructed from the 49 measurements made by FluxCam are shown in fig. 10(c).




Fig. 9. **Sensor performance.** (a) Setup used to measure the radiometric characteristics of the sensor. (b) The radiometric response is linear over the entire dynamic range (up to 13,000 lx). (c) SNR of the sensor. At just 2.2 lx, the SNR is 2.4 dB, which is near the minimum detectable signal. This corresponds to a wide dynamic range of 75.4 dB. (d) Power harvested by the sensor. Due to the harvester’s overhead, the sensor board requires at least 71 lx to harvest power.




Fig. 10. **Reconstructed irradiance functions using FluxCam.** (a) HDR images of two real scenes. (b) The corresponding ground truth irradiance functions. (c) Reconstructed irradiance functions using the 49 measurements produced by FluxCam in each scene. Across four different scenes, the measured irradiance functions attain an average RMSE of 3.93% with respect to the maximum irradiance in each scene. (d) Reconstruction accuracy for the 1D slice corresponding to the dashed gray line in (c).

The reconstruction accuracy of a 1D slice through one of the reconstructed functions can be seen in fig. 10(d). Across the four scenes we measured, the average RMSE in the reconstructed irradiance function, as a fraction of the maximum irradiance in each scene, was 3.93%. The supplemental material includes results for all four scenes.

5 Example Applications

5.1 Estimating Egomotion

In our first application, we show that FluxCam can be used to estimate egomotion. In the case of a pure rotation of the camera, the irradiance function rotates by the same amount, and hence estimating the camera’s rotation is straightforward. Real motion, however, includes translation, causing the irradiance function to undergo a complex, scene-dependent transformation. Therefore, we trained a network to predict the relative rotation between two camera poses, which includes an arbitrary translation between them, using the irradiance measurements at each pose. Training data was generated in simulation by rendering the irradiance of the camera’s 49 detectors in a synthetic urban environment. For each training sample, we rendered the measurements at an initial camera pose, applied a random rotation and a small translation, and rendered the measurements again at the new pose. The camera was positioned at different distances from the ground to include the effects of parallax due to translation. Additionally, the rendering process accounts for the fact that each detector views the scene from a different viewpoint. Using this procedure, we generated two million pairs of irradiance measurements for training. We tested the network using 319,712 pairs of simulated measurements (with rotation and translation), and found the network to yield an average rotation error of 2.25°.

In practice, a camera’s trajectory varies smoothly with time. Therefore, we can refine the estimated camera orientation using temporal filtering. Figure 11(a) shows the estimated camera orientation for two scenarios: rolling FluxCam on the ground (left) and tossing FluxCam in the air (right). In both cases, the network first predicts the relative orientation between frames, after which temporal filtering is used to determine the camera’s (relative) orientation. In effect, FluxCam functions as an untethered optical gyroscope. **Please see the supplemental video.** Details regarding the network architecture and the simulation results are provided in the supplemental material.

5.2 Relighting using an Untethered Light Probe

Mirror probes are widely used to measure the environmental illumination [Debevec 1998; Debevec et al. 2012; Reinhard et al. 2010]. While this approach can be used for relighting objects with complex appearances, the mirror must be fully visible to the camera. In other words, it cannot measure illumination in locations that are fully or partially occluded by other objects. Since FluxCam is untethered, it does not need line of sight with the camera. For example, on the left in fig. 11(b), FluxCam is partially occluded by the vase but measures and transmits the complete irradiance function at that location. Although the irradiance function has less detail than the radiance function, it is useful for diffuse relighting. On the right in fig. 11(b), the measured irradiance is used to render and insert a diffuse object (a light wood vase with grooves) at the location where FluxCam was positioned.

Since FluxCam’s measurements are weighted by the spectral response of the photovoltaics, the measurements do not reveal the illuminant’s color. Therefore, in this example, we assume a white

illuminant and perform relighting in the RGB color space. In the future, FluxCam could be modified to measure spectral irradiance using a larger number of detectors with spectral filters, enabling physically accurate relighting under different illuminants.

5.3 Workspace Monitoring

Many vision tasks only involve high-level inferences about the statistics of objects in a scene. These tasks, referred to as lightweight [Klotz and Nayar 2025b], do not require the detection of fine-grained visual details. Although an irradiance camera produces few measurements, they have high precision and dynamic range and hence encode visual cues about the environment that can be used to solve lightweight tasks. As an example, consider the workspace in fig. 11(c). The goal is to determine which desks are occupied, detect when the door is open, and determine which lamps are on (dashed black boxes in the first image of fig. 11(c)). We placed FluxCam in the center of the room to solve these tasks. For each task, we trained classifiers using 23 minutes of labeled data. The classifier outputs for the test data are overlaid on the images in fig. 11(c), and the inset images show the irradiance functions measured by the camera. The system detects when the door is open with 90.9% accuracy, the occupancy of all three desks with 87.6% accuracy, and the activation of the lamps with 99.9% accuracy. **Please see the supplemental video.** In addition to being self-powered and untethered, the irradiance camera preserves privacy in such a setting since it does not capture the visual details needed to identify individuals.

5.4 Omnidirectional Pyranometer

Pyranometers are widely used in meteorology and solar energy to measure the irradiance of a surface with a fixed orientation [Duffie and Beckman 2006]. The irradiance camera is, in effect, an omnidirectional pyranometer. Consider the HDR fisheye image on the left of fig. 11(d), which was taken on a clear day. We placed FluxCam in this environment to measure the irradiance function (small image outside the fisheye image). These measurements can be used to compute the direction of maximum irradiance (red dot), which corresponds to the orientation of a solar panel that would maximize the harvested energy. For comparison, the direction computed from the fisheye image is shown as a blue dot.

An irradiance camera can also be used to determine the sky conditions. To this end, we fit our irradiance measurements to a low-dimensional analytical sky model [Perez et al. 1993], which allows us to estimate the sky’s brightness and clearness. Furthermore, we can use the fitted sky model to decompose the irradiance into components due to the sun and the sky, which is typically parameterized by the global horizontal irradiance (GHI), diffuse horizontal irradiance (DHI), and direct normal irradiance (DNI) [Duffie and Beckman 2006]. Figure 11(d) shows reconstructed skies under clear and overcast conditions using the Hosek–Wilkie and Perez models, respectively [Hosek and Wilkie 2012; Perez et al. 1993]. In this example, we do not use explicit spectral measurements to fit the analytical sky model. Since the spectral response of FluxCam’s detectors closely resembles the photopic response, it measures illuminance (in lx), which we convert to irradiance (in W m^{-2}) using a luminous efficacy of 109 lm W^{-1} before fitting. For the overcast case,

the irradiance measurements were simulated using an HDR spherical image from the UrbanSky dataset. The supplemental material includes a quantitative evaluation of this approach to sky condition estimation.

6 Discussion

We have introduced the concept of an irradiance camera and demonstrated its potential applications. There are several directions we plan to pursue as future work. First, our current prototype operates at 30 FPS when the average illumination received by it is roughly 1,000 lx. The camera’s power consumption can be further optimized so that it can operate in dimly lit indoor environments. Second, the camera can be miniaturized by using smaller photovoltaics. This would reduce the harvested energy and hence the framerate of the camera, but the lower framerates would be adequate for a variety of applications. Third, while we have shown that the camera can be used to estimate rotation, we believe our approach can be used to estimate the direction of translation as well, making it useful for self-powered egomotion estimation on a variety of robotic platforms, such as drones and wheeled robots. Finally, while our current implementation ensures each detector has a hemispherical field of view, we believe that using apertures to reduce the field of view would produce measurements that are richer in information for certain lightweight tasks. With these improvements, we hope that the irradiance camera could serve as a powerful untethered sensor for a wide variety of vision applications.

Acknowledgments

This work was supported by the Office of Naval Research (ONR) under award number N00014-23-1-2096 and the National Science Foundation (NSF) and Center for Smart Streetscapes (CS3) under NSF Cooperative Agreement No. EEC-2133516. Jeremy Klotz was supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship. The authors are grateful to Behzad Kamgar-Parsi for his support and encouragement. The authors also thank Mikhail Fridberg for help designing the camera electronics, Sylvia Sellán for discussions about Blender, Joaquin Palacios for help with 3D printing, Junbang Liang for discussions about transformers, and Matthew Beveridge for technical feedback.

References

- John A. Augustine, John J. DeLuise, and Charles N. Long. 2000. SURFRAD—A National Surface Radiation Budget Network for Atmospheric Research. *Bulletin of the American Meteorological Society* 81, 10 (Oct. 2000), 2341–2357.
- Juan-Carlos Baltazar, Yifu Sun, and Jeff Haberl. 2015. Improved Methodology to Evaluate Clear-Sky Direct Normal Irradiance with a Multi-Pyranometer Array. *Solar Energy* 121 (Nov. 2015), 123–130.
- Ronen Basri and David W. Jacobs. 2003. Lambertian Reflectance and Linear Subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25, 2 (Feb. 2003), 218–233.
- A.F. Bobick and J.W. Davis. 2001. The Recognition of Human Movement Using Temporal Templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 3 (March 2001), 257–267.
- Henry Cohn. 2024. Table of Spherical Codes. <https://dspace.mit.edu/handle/1721.1/153543>
- J. H. Conway and N. J. A. Sloane. 1999. *Sphere Packings, Lattices and Groups*. Grundlehren Der Mathematischen Wissenschaften, Vol. 290. Springer New York, New York, NY.
- Ji Dai, Jonathan Wu, Behrouz Saghaei, Janusz Konrad, and Prakash Ishwar. 2015. Towards Privacy-Preserving Activity Recognition Using Extremely Low Temporal and Spatial Resolution Cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.



Fig. 11. **Example applications of an irradiance camera.** (a) FluxCam can be used as an optical gyroscope to determine the camera's orientation. Here we show the computed orientation as the camera is rolled across a floor and tossed in the air. (b) The camera can serve as an untethered light probe for diffuse relighting. We placed the camera in a location that is partially occluded by another object (left image) and measured the irradiance function (inset image). This allows us to relight and insert a diffuse object (right image). (c) FluxCam can be used to monitor a workspace. Here we use the camera's 49 measurements to determine which desks are occupied, detect when the door is open, and determine which lights are on (dashed black boxes in the left image). (d) FluxCam can be used as an omnidirectional pyranometer to estimate the best orientation of a solar panel (red dots). Additionally, the measurements can be used to reconstruct the sky, allowing us to quantify the current sky conditions. **Please see the supplemental video.**

- Paul Debevec. 1998. Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. ACM, 189–198.
- Paul Debevec, Paul Graham, Jay Busch, and Mark Bolas. 2012. A Single-Shot Light Probe. In *ACM SIGGRAPH 2012 Talks*. ACM, Los Angeles, California, 1.
- James R. Driscoll and Dennis M. Healy. 1994. Computing Fourier Transforms and Convolutions on the 2-Sphere. *Advances in Applied Mathematics* 15, 2 (June 1994), 202–250.
- Claude E. Duchon and Mark S. O'Malley. 1999. Estimating Cloud Type from Pyranometer Observations. *Journal of Applied Meteorology* 38, 1 (Jan. 1999), 132–141.
- John A. Duffie and William A. Beckman. 2006. *Solar Engineering of Thermal Processes*. John Wiley & Sons, Inc., Hoboken, New Jersey.
- D. Faiman, D. Feuermann, P. Ibbetson, and A. Zemel. 1992. A Multipyranometer Instrument for Obtaining the Solar Beam and Diffuse Components, and the Irradiance on Inclined Planes. *Solar Energy* 48, 4 (1992), 253–259.
- Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gabbaretto, Christian Gagné, and Jean-François Lalonde. 2017. Learning to Predict Indoor Illumination from a Single Image. *ACM Transactions on Graphics* 36, 6 (Dec. 2017), 1–14.
- A. Gershun. 1939. The Light Field. *Journal of Mathematics and Physics* 18, 2 (May 1939), 51–151.
- G. Greger, P. Shirley, P.M. Hubbard, and D.P. Greenberg. 1998. The Irradiance Volume. *IEEE Computer Graphics and Applications* 18, 2 (1998), 32–43.
- Lukas Hosek and Alexander Wilkie. 2012. An Analytic Model for Full Spectral Sky-Dome Radiance. *ACM Transactions on Graphics* 31, 4 (Aug. 2012), 1–9.
- Jeremy Klotz and Shree K. Nayar. 2025a. Minimal Sensing for Orienting a Solar Panel. *Solar Energy* 300 (Nov. 2025), 113833.
- Jeremy Klotz and Shree K. Nayar. 2025b. Minimalist Vision with Freeform Pixels. In *European Conference on Computer Vision (ECCV)*, Vol. 15122. Springer Nature Switzerland, 329–346.
- Jeremy Klotz and Shree K. Nayar. 2026. Project Website. <https://cave.cs.columbia.edu/projects/categories/project?cid=Computational+Imaging&pid=An+Eternal+Irradiance+Camera>
- L.O. Lamm and C.G. Adler. 1987. A New Method for the Determination of Direct Insolation. *Solar Energy* 39, 2 (1987), 109–112.
- Michael F. Land and Dan-Eric Nilsson. 2002. *Animal Eyes*. Oxford University Press.
- Dong-Seok Lee. 2025. A Multi-Directional Pyranometer (CUBE-i) for Real-Time Direct and Diffuse Solar Irradiance Decomposition. *Remote Sensing* 17, 8 (April 2025), 1336.
- Charles N. Long and Thomas P. Ackerman. 2000. Identification of Clear Skies from Broadband Pyranometer Measurements and Calculation of Downwelling Shortwave Cloud Effects. *Journal of Geophysical Research: Atmospheres* 105, D12 (June 2000), 15609–15626.
- Shota Nakashima, Yuhki Kitazono, Lifeng Zhang, and Seiichi Serikawa. 2010. Development of Privacy-Preserving Sensor for Person Detection. *Procedia - Social and Behavioral Sciences* 2, 1 (2010), 213–217.
- Shree K. Nayar, Jeremy Klotz, Nikhil Nanda, and Mikhail Fridberg. 2024. Cricket: A Self-Powered Chirping Pixel. *ACM Transactions on Graphics* 43, 4 (July 2024), 1–11.
- Shree K. Nayar, Daniel C. Sims, and Mikhail Fridberg. 2015. Towards Self-Powered Cameras. In *Proceedings of the IEEE International Conference on Computational Photography (ICCP)*.
- Panasonic. 2021. Photovoltaic Datasheet. https://panasonic.net/electricworks/amorton/assets/pdf/Brochures_Amorton_E_2.pdf
- R. Perez, R. Seals, and J. Michalsky. 1993. All-Weather Model for Sky Luminance Distribution—Preliminary Configuration and Validation. *Solar Energy* 50, 3 (March 1993), 235–245.
- Ravi Ramamoorthi and Pat Hanrahan. 2001a. An Efficient Representation for Irradiance Environment Maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 497–500.
- Ravi Ramamoorthi and Pat Hanrahan. 2001b. On the Relationship between Radiance and Irradiance: Determining the Illumination from Images of a Convex Lambertian Object. *Journal of the Optical Society of America A* 18, 10 (Oct. 2001), 2448.
- Ravi Ramamoorthi and Pat Hanrahan. 2001c. A Signal-Processing Framework for Inverse Rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 117–128.
- Erik Reinhard, Greg Ward, Sumanta Pattanaik, Paul Debevec, Wolfgang Heidrich, and Karol Myszkowski. 2010. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting* (2 ed.). Morgan Kaufmann Publishers/Elsevier, Burlington, MA.
- Neil J. Sloane, R. H. Hardin, and W. D. Smith. 1994. Tables of Spherical Codes. <http://neilsloane.com/packings/>.
- P. M. L. Tammes. 1930. *On the Origin of Number and Arrangement of the Places of Exit on the Surface of Pollen-Grains*. Ph.D. Dissertation.
- Antonio Torralba, Rob Fergus, and William T. Freeman. 2008. 80 Million Tiny Images: A Large Data Set for Nonparametric Object and Scene Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30, 11 (Nov. 2008), 1958–1970.
- Gregory J. Ward. 1988. A Ray Tracing Solution for Diffuse Interreflection. *Computer Graphics* 22, 4 (Aug. 1988), 85–92.
- Dingtian Zhang, Jung Wook Park, Yang Zhang, Yuhui Zhao, Yiyang Wang, Yunzhi Li, Tanvi Bhagwat, Wen-Fang Chou, Xiaojia Jia, Bernard Kippelen, Canek Fuentes-Hernandez, Thad Starner, and Gregory D. Abowd. 2020. OptoSense: Towards Ubiquitous Self-Powered Ambient Light Sensing Surfaces. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (Sept. 2020), 103:1–103:27.
- Yang Zhang, Yasha Iravantchi, Haojian Jin, Swarun Kumar, and Chris Harrison. 2019. Sozu: Self-Powered Radio Tags for Building-Scale Activity Sensing. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (UIST '19)*. Association for Computing Machinery, New York, NY, USA, 973–985.