Minimal Sensing for Orienting a Solar Panel

Jeremy Klotz^{a,*}, Shree K. Nayar^a

^a Computer Science Department, Columbia University, 500 West 120th St., New York, 10027, NY, USA

Abstract

A solar panel harvests the most energy when pointing in the direction that maximizes the total illumination (irradiance) falling on it. Given an arbitrary orientation of a panel and an arbitrary environmental illumination, we address the problem of finding the direction of maximum total irradiance. We develop a minimal sensing approach where measurements from just four photodetectors are used to iteratively vary the tilt of the panel to maximize the irradiance. Many environments produce irradiance functions with multiple local maxima. As a result, simply measuring the gradient of the irradiance function and applying gradient ascent will not work. We show that a larger, optimized tilt between the detectors and the panel is equivalent to blurring the irradiance function. This has the effect of eliminating local maxima and turning the irradiance function into a unimodal one, whose maximum can be found using gradient ascent. We show that there is a close relationship between our approach and scale space theory. We have collected a large dataset of high-dynamic range lighting environments in New York City, called *UrbanSky*. We used this dataset to conduct simulations to verify the robustness of our approach. Finally, we have built a portable solar panel with four compact detectors and an actuator to conduct experiments in various real-world settings: direct sunlight, cloudy sky, urban settings with occlusions and shadows, and complex indoor lighting. In all cases, we show significant improvements in harvested energy compared to standard approaches for controlling the orientation of a solar panel.

Keywords: Solar Panel Orientation, Photodifferential, Minimal Sensing, Scale Space Theory, Urban Lighting Environment, UrbanSky

1. The Orientation of a Solar Panel

Solar panels are widely deployed in open fields, on rooftops, and in urban settings to harvest energy from sunlight. Consider a panel in a desert on a sunny day. Since the sun is the dominant light source, the illumination from the entire sky can be approximated as coming from a single point (the sun). Thus, pointing the panel in the direction of the sun would maximize the irradiance¹ of the panel and hence the energy harvested by it.² In this case, the panel can track the sun as its trajectory is easily determined from the latitude and longitude of the panel, the date, and the time of day [1, 2].

But what is the optimal solar panel orientation on a cloudy day? Now, the sky is an extended light source with a possibly complex radiance function that includes multiple peaks in different directions. In this setting, since a solar panel aggregates light from an entire hemisphere, simply orienting it towards the brightest point in the sky is almost certain to not maximize the irradiance.

The problem gets even more interesting in urban settings. Consider a solar panel in the urban environment shown in fig. 1(a). In this case, the panel only sees the sun for a fraction of the day due to occlusions by nearby buildings. When the sun is occluded, the panel would only be illuminated by a patch of the sky and reflections from buildings and objects around it. As can be seen in fig. 1(a), even within a single day, the lighting may change dramatically due to shadows and reflections of the sun. Many cities have deployed large numbers of solar panels in dense urban environments to power sensors, lamps, and devices such as electric bikes. Currently, these panels are either fixed in orientation or simply track the sun.

Beyond dense urban environments, solar panels are increasingly used indoors to harvest energy from both the indoor and outdoor illumination to power environmental sensors and consumer devices. In this scenario, the illumination can be expected to vary significantly, causing the direction of maximum irradiance to shift dramatically throughout the day. Numerous smart city applications stand to benefit from indoor solar panels that can continuously track the direction of maximum irradiance. The market for indoor panels is growing rapidly and is projected to be \$154 million by 2030 [3].

^{*}Corresponding author

Email addresses: jklotz@cs.columbia.edu (Jeremy Klotz), nayar@cs.columbia.edu (Shree K. Nayar)

¹Unless otherwise specified, by irradiance we mean the total irradiance received by the panel from the entire hemisphere visible to it.

 $^{^2\}mathrm{Due}$ to non-idealities in the energy harvesting system, the harvested energy from a solar panel may be a non-linear function of its irradiance. However, since this function is monotonic, the orientation that maximizes the irradiance also maximizes the harvested energy.

The goal of our work is to find the direction of maximum irradiance in any environment using minimal sensing resources. The first instinct may be to use a fisheye camera placed in the environment to measure the incoming radiance from every direction and compute the direction of maximum irradiance. Traditional cameras, however, not only add cost to the system but are also powerhungry—the image sensor alone consumes hundreds of milliwatts [4], which, in effect, reduces the energy harvested by the panel. While this cost and energy overhead can be ignored in the case of large arrays of panels, it would be prohibitive in the case of smaller stand-alone panels. For applications involving a single small panel, a minimal sensing approach that captures the fewest light measurements and requires negligible processing is highly desirable.

One approach would be to measure the gradient of the incident irradiance along each of the two dimensions of the panel using four photodetectors that are slightly tilted with respect to each other. The measurements can then be used to perform gradient ascent to iteratively tilt the panel towards the direction of maximum irradiance. Unfortunately, the irradiance function of a complex lighting environment, such the urban environment in fig. 1(a), is likely to have multiple local maxima, as shown in fig. 1(b). As a result, such a naïve approach to orient a solar panel would produce sub-optimal panel orientations.

This brings us to our key result. Using Fourier analysis, we show that by using carefully chosen larger tilts between the four photodetectors, we can find the gradient of a function that is the original irradiance function convolved with a box filter of a specific width. Even when the original irradiance function has multiple modes, the convolved (blurred) function is almost certain to have a single mode that is close to the largest mode of the original function. This result is closely related to scale space theory [5, 6]. Therefore, by simply computing finite differences between the measurements produced by four carefully oriented detectors and using them to perform gradient ascent, the panel converges to the orientation that yields maximum irradiance.

To validate our approach in simulation, we have collected a large dataset of high-dynamic range (HDR) lighting environments in New York City. This dataset, called UrbanSky, consists of 1,067 lighting environments under various weather conditions and at different times of day. In each environment in UrbanSky, we capture the illumination using a 360° camera, we measure the global horizontal irradiance using a pyranometer, and we record the date, time of day, GPS location, and current weather conditions. UrbanSky is publicly available online to encourage future work both in solar energy and computer graphics.

We test our approach in real-world settings using a prototype system. Our prototype includes a single solar panel with four compact detectors attached to its periphery, and the panel is mounted on a two-axis actuator. We compare the energy harvested by our prototype with the energy harvested using two widely used strategies: (a) fix-



(a) Radiance Function (b) Irradiance Function

Figure 1: Complex Illumination in Urban Environments. (a) The radiance function of an urban environment, which describes the environmental illumination, at different times of day. As the sun moves through the sky, the illumination varies dramatically due to changing weather, shadows, and reflections of the sun. (b) The irradiance function, which specifies the total irradiance of a solar panel as a function of the panel orientation, varies significantly with the environmental illumination. In urban settings such as this one with complex illumination, the irradiance function often has multiple local maxima (green and red dots). As a result, iteratively tilting a solar panel based on the gradient of the irradiance function would produce sub-optimal orientations. Given any environment, our goal is to orient a solar panel in the direction of maximum irradiance (red dot) using minimal sensing and processing, regardless of the complexity of the illumination.

ing the orientation of the panel and (b) tracking the sun. In diverse real-world settings—direct sunlight, cloudy sky, urban settings with occlusions, shadows, and reflections, and an indoor room with complex lighting—our approach yields significant energy gains.

2. Related Work

Solar panels are either fixed at a specific orientation or mounted on actuators that vary their orientation to increase the harvested energy. With respect to fixed panel orientations [7, 8, 9, 10, 11], a widely used rule of thumb is to orient the panel towards the equator with an angle

from the zenith equal to the panel's latitude [7]. Panels that are mounted on actuators to move throughout the day almost exclusively track the sun [12, 13, 14, 15]. Kelly and Gibson [16] observed that tracking the sun is sub-optimal on a cloudy day and suggested pointing the panel straight up at the sky when the sun is occluded by clouds. All of these prior works only consider outdoor environments with an unobstructed view of the sky. Our work considers a more general problem: we wish to find the panel orientation that yields maximum irradiance in an arbitrary environment. This problem has become highly relevant as panels are now being used in environments that include occlusions, shadows, and multiple dominant sources. We demonstrate significant gains in harvested energy using our approach compared to a panel that is fixed in orientation [7] or one that tracks the sun.

It is widely accepted that solar energy harvesting in urban environments, in particular, is an increasingly important application of photovoltaics. Orienting a solar panel in urban environments is a complex problem since one must consider not only the illumination from the sky but also reflections and shadows from nearby infrastructure. Prior work seeks to find the best *fixed* orientation in urban settings by using knowledge of the scene's 3D structure and material properties to compute the irradiance falling on rooftops and building façades [17, 18, 19, 20, 21, 22]. All of these works account for the complex illumination caused by city infrastructure. The goal of our work is different: given a solar panel in any lighting environment, which may include shadows and reflections in urban settings, we use an actuator to iteratively tilt the panel toward the direction of maximum irradiance using minimal sensor measurements. To encourage future work in the analysis of urban lighting environments, we have released UrbanSky, a dataset of 1,067 outdoor HDR lighting environments captured in New York City.

A variety of methods seek to find the optimal orientation of a solar panel by iteratively tilting the panel based on measurements produced by light sensors. Most closely related to our work are methods that use detectors mounted on the panel at an angle to produce differential measurements [23, 24, 25, 26, 27]. When the panel points toward the sun on a clear day, the detectors are illuminated equally, and the differential measurement is zero. Thus, the tracking algorithm seeks to minimize the differential measurement in order to track the sun. While this approach can orient a panel toward the sun on a clear day, its convergence in arbitrary lighting is not guaranteed. This is because, in all the previous work, the tilt angle between the detectors and the solar panel is chosen in an ad-hoc manner. We provide a detailed analysis of the differential measurements produced by tilted detectors. We first show that the measurement is equal to the derivative of a function that results from blurring the irradiance function with a box filter. When the detector tilt angle is carefully chosen, the blurred irradiance function is unimodal, regardless of the complexity of the illumination. This means that a panel oriented using our approach will converge at, or close to, the direction of maximum irradiance, even when the irradiance function has multiple modes.

Many other visual sensors have also been proposed to iteratively tilt a solar panel toward the orientation that vields maximum irradiance. Shading-based sensors [28, 29, 30, 31, 32, 33, 34, 35] use a vertical wall to cast a shadow onto multiple photodetectors. A tracking system then orients the panel such that the shadow disappears, which occurs when the sun is directly above the panel on a clear day. Position-sensing diodes have also been used to find the direction of the sun [36, 37]. While a panel using either of these methods can be oriented toward the sun on a clear day, it is not guaranteed to find the orientation that yields maximum irradiance in arbitrary lighting. Pineda and Arredondo [38] approach the problem of finding the best panel orientation in a variety of sky conditions by using a large number of detectors with smaller fieldsof-view to sample the irradiance at different orientations. The panel is then oriented in the direction of the detector that produces the largest measurement. Rather than sample the irradiance function, we iteratively tilt a panel toward the direction of maximum irradiance using differential measurements from just four detectors.

3. The Irradiance of a Solar Panel

In this section, we derive an expression for the irradiance of a solar panel as a function of its orientation.

3.1. Total Irradiance from the Visible Hemisphere

The illumination seen by a solar panel can be completely described as a 4D light field [39]. Since the distances of light sources in any environment are typically much larger than the size of the panel, we can assume that the light incident from a specific direction is uniform over the panel's active area. Thus, as shown fig. 2, the environmental illumination that the panel is exposed to can be represented as a 2D radiance function $L(\vec{s})$, where \vec{s} denotes the direction as seen from the center of the panel.

In the context of solar energy harvesting, we are specifically interested in the irradiance of the panel. Let \vec{n} be the vector normal to the panel. Then, the *directional irradiance dE* from the direction \vec{s} and infinitesimal solid angle $d\omega$ is

$$dE(\vec{s}; \vec{n}) = L(\vec{s}) \max(\vec{n} \cdot \vec{s}, 0) \, d\omega, \tag{1}$$

where the max operator accounts for the fact that the panel only receives light from the hemisphere of illumination that is visible to the panel (see fig. 2).

We can now define the *total irradiance* E of the panel as the integral of the directional irradiance (eq. (1)) over the visible hemisphere,

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



Figure 2: Irradiance of a solar panel. The illumination of a panel can be represented as a radiance function $L(\vec{s})$. The panel only receives light from the visible hemisphere, which is determined by the orientation \vec{n} of the panel. The total irradiance of the panel is computed as an integral of the radiance function over the visible hemisphere.

where \mathbb{S}^2 is the set of points on the surface of the unit sphere. $E(\vec{n})$ corresponds to the total power per unit area incident on the panel. The goal of our work is to find, for any given environment, the panel orientation \vec{n} that maximizes the irradiance $E(\vec{n})$.

3.2. Irradiance as a Convolution on the Sphere

From eq. (2), we see that the irradiance function $E(\vec{n})$ can be written as the convolution of the radiance function $L(\vec{n})$ with a kernel defined as

$$k(\vec{n}) = \max(\vec{n} \cdot \hat{z}, 0). \tag{3}$$

Note that the kernel $k(\vec{n})$ is centered about the zenith vector \hat{z} . A detailed derivation that shows that the irradiance function is a convolution on the sphere [40] is given in appendix A. Notice that the kernel is simply the cosine of the zenith angle, but clipped to be non-negative. Given that the kernel is a smooth and broad function, it serves to simply low-pass filter the radiance function corresponding to the environment. An equivalent result in computer graphics shows that Lambertian reflectance low-pass filters the radiance function [41, 42].

Since the irradiance function is the result of convolving the radiance function with a low-pass filter, it is guaranteed to be smooth. Consider the radiance function of a natural environment shown in fig. 3(a). Upon convolving it with the kernel in fig. 3(b), we get the irradiance function in fig. 3(c) which is very smooth. Despite its smoothness, however, the irradiance function may still have multiple modes, and hence multiple local maxima (red and green dots).

4. Differential Light Sensing

Imagine we had a method to measure the gradient of the irradiance function with respect to the solar panel orientation. Then, given any initial panel orientation, a naïve approach would be to iteratively tilt the panel in the direction of the gradient and hope to converge at the direction



Figure 3: Irradiance function as a convolution on the sphere. The radiance function (a) is convolved with the kernel (b) to produce the irradiance function (c). The irradiance function is a smooth version of the radiance function, but could still have multiple modes and hence multiple local maxima (red and green dots).

of maximum irradiance. Unfortunately, this would not always work as we know that although the irradiance function is guaranteed to be smooth, it could have multiple local maxima (fig. 3(c)). In short, naïve gradient ascent would likely get stuck at a local maximum, causing the panel to harvest less energy. In this section, we introduce our approach which uses differential light sensing to iteratively tilt a panel to arrive at the *global maximum* of the irradiance function, irrespective of the initial orientation of the panel or the complexity of the environmental illumination.

4.1. The Photodifferential

Consider a one-dimensional irradiance function $E(\theta)$, where the panel orientation is given by θ . Our sensing method, illustrated in fig. 4, produces a finite difference of the irradiance function $E(\theta)$ using just two photodetectors. The detectors are tilted by $\Delta\theta$ in opposite directions with respect to the panel. Thus, the detectors measure the irradiance function at two different orientations: $E(\theta + \Delta\theta)$ and $E(\theta - \Delta\theta)$. The difference of these two measurements is proportional to a finite difference of the irradiance function. We refer to this difference as the "photodifferential," which is defined as

$$E_D(\theta; \Delta \theta) = \frac{E(\theta + \Delta \theta) - E(\theta - \Delta \theta)}{2\Delta \theta}.$$
 (4)

When the tilt angle $\Delta \theta$ between the detectors and the panel is very small, the photodifferential is equal to the derivative (gradient) of the irradiance function with respect to the panel orientation. As stated earlier, this derivative is not of much use in our context—iteratively tilting the panel in the direction of the derivative will cause the panel to get stuck at a local maximum. In the next section, we will show that the photodifferential produced by detectors with a *large* tilt angle is actually the derivative of a blurred version of the irradiance function.

4.2. Detector Tilt Blurs the Irradiance Function

We now analyze the effect of the detector tilt angle $\Delta \theta$ on the photodifferential using Fourier analysis. Note that the photodifferential (eq. (4)) can be written as the



Figure 4: **Photodifferential sensor.** Two tiny photodetectors on the sides of the solar panel are tilted with respect to the panel by a significant angle $\Delta \theta$. The difference between the detector measurements is a finite difference of the irradiance function, which we refer to as the photodifferential.

result of convolving the irradiance function $E(\theta)$ with the function

$$h(\theta) = \frac{1}{2\Delta\theta} \left(\delta(\theta + \Delta\theta) - \delta(\theta - \Delta\theta) \right).$$
 (5)

The Fourier transform of $h(\theta)$ is

$$H(j\omega) = \frac{1}{2\Delta\theta} \left(e^{j\omega\Delta\theta} - e^{-j\omega\Delta\theta} \right).$$
(6)

Using Euler's formula, $H(j\omega)$ can be written as

$$H(j\omega) = \frac{j}{\Delta\theta}\sin(\omega\Delta\theta).$$
 (7)

By multiplying the numerator and denominator by ω , eq. (7) can be rewritten as

$$H(j\omega) = \begin{cases} j\omega \ \frac{\sin(\omega\Delta\theta)}{\omega\Delta\theta} & \omega \neq 0\\ 0 & \omega = 0 \end{cases}.$$
 (8)

Therefore, we get

$$H(j\omega) = \underbrace{j\omega}_{\text{Derivative Box filter}} \underbrace{\operatorname{sinc}(\omega\Delta\theta)}_{\theta}.$$
(9)

Equation (9) shows that convolving the irradiance function with $h(\theta)$ is equivalent to first blurring the irradiance function with a box filter $b(\theta; \Delta \theta)$, whose width is proportional to the detector tilt angle $\Delta \theta$, and then taking the derivative.³ That is,

$$E_D(\theta; \Delta \theta) = \frac{d}{d\theta} \left(L(\theta) * k(\theta) * b(\theta; \Delta \theta) \right).$$
(10)

Dumont [43] presents a similar result for discrete signals showing that a finite difference with a large step size implicitly filters the signal with a moving average. This brings us to our first observation:

Observation 4.1. The photodifferential $E_D(\theta; \Delta \theta)$ is the derivative of a blurred irradiance function, where the degree of blur is proportional to the detector tilt angle $\Delta \theta$.

4.3. E_D is the Derivative of a Unimodal Function

We now analyze the effect of the box filter $b(\theta; \Delta \theta)$ in eq. (10). Consider the one-dimensional radiance function $L(\theta)$ in fig. 5(a). Convolving it with the kernel $k(\theta)$ produces the irradiance function $E(\theta)$ in fig. 5(b). This function has three modes: a global maximum (red dot) and two local maxima (green dots). In eq. (10), we denote the function to which the derivative is being applied as the blurred irradiance function:

$$E_B(\theta; \Delta \theta) = L(\theta) * k(\theta) * b(\theta; \Delta \theta).$$
(11)

In fig. 5(c), we plot E_B for three different tilt angles: 5°, 20°, and 45°. We know that increasing the tilt angle has the effect of blurring the irradiance function with a wider box filter. Consequently, as the tilt angle increases, the number of local maxima in E_B reduces. When the level of blur is sufficiently large ($\Delta \theta = 45^{\circ}$ for the example radiance function in fig. 5(a)), the local maxima disappear, rendering the blurred irradiance function unimodal. This brings us to our second observation:

Observation 4.2. When the detector tilt angle $\Delta \theta$ is sufficiently large, the blurring induced by the photodifferential creates a blurred irradiance function that is unimodal, regardless of the complexity of the illumination.

4.4. Gradient Ascent using the Photodifferential

Since the photodifferential measures the derivative of a blurred irradiance function E_B that has a single mode, our approach is to iteratively tilt the solar panel in the direction of the photodifferential to find the panel orientation that yields the *global maximum* of E_B . This approach, in effect, is equivalent to applying gradient ascent to E_B . However, we must ensure that the blurring from the detector tilt never introduces a new mode in E_B that was not present in the original irradiance function.

4.5. Can a Large Tilt Angle Introduce New Modes?

The detector tilt angle must satisfy two conditions: (a) The tilt angle should be large enough such that the blurred irradiance function E_B is unimodal for virtually any environmental illumination. (b) The tilt angle should not introduce new modes in the blurred irradiance function. This brings us to a well-known result in scale space theory: Blurring any function with a Gaussian cannot create new modes that were not present in the original function [44, 45, 46].

Recall that the irradiance function $E(\theta)$ is the result of convolving the radiance function $L(\theta)$ with the kernel $k(\theta)$. Figure 6(a) shows the Fourier transform $K(j\omega)$ of the kernel, which is close to zero for frequencies beyond the first zero-crossings, denoted by the two red lines. In other words, the kernel suppresses frequencies in the irradiance function that are outside the frequency band between the red lines. As a result, we are only concerned with

³Observe that $\lim_{\Delta\theta\to 0} H(j\omega) = j\omega$, which means $H(j\omega)$ becomes the continuous-time differentiator in the limit as the finite difference step size approaches 0. This aligns with the intuition that a finite difference computed using a very narrow box filter (infinitesimally small detector tilt) approximates the gradient of the irradiance function.



Figure 5: A large detector tilt angle makes E_B a unimodal function. (a) An example radiance function $L(\theta)$. (b) The irradiance function $E(\theta)$, which is the convolution of the radiance function with the kernel $k(\theta)$, has multiple modes. (c) Increasing the detector tilt angle $\Delta\theta$ has the effect of blurring the irradiance function, thereby eliminating local maxima. In this example, a tilt angle of 45° results in a blurred irradiance function E_B that is unimodal.

frequencies of the irradiance function that lie within this band. Figure 6(b) shows the Fourier transform $B(j\omega)$ of the box filter $b(\theta; \Delta \theta)$ for a wide range of detector tilt angles. Within the frequency band of the red lines, $B(j\omega)$ can be well-approximated as a Gaussian, even for the largest tilt angle of $\Delta \theta = 90^{\circ}$. This brings us to our third observation:

Observation 4.3. Due to the bandwidth of the kernel $k(\theta)$, large detector tilt angles do not introduce new modes in the blurred irradiance function.

5. Simulations using Real Radiance Functions

We have shown that large tilt angles are beneficial since they eliminate local maxima of the irradiance function and do not introduce new modes. In this section, we explore via simulations the range of detector tilt angles for which



Figure 6: The detector tilt angle has the effect of blurring the irradiance function with a Gaussian. (a) The Fourier transform $K(j\omega)$ of the kernel $k(\theta)$ is close to zero beyond its first zero crossings (red lines). This means that the irradiance function only has frequencies within the band between the red lines. (b) The Fourier transform $B(j\omega)$ of the box filter for a wide range of detector tilt angles. Observe that $B(j\omega)$ can be reasonably well-approximated by a Gaussian within the band between the red lines, even for the largest tilt angle of $\Delta \theta = 90^{\circ}$.

the irradiance of the panel (and hence the harvested energy) is at, or close to, the maximum.

5.1. UrbanSky: A Dataset of Urban Lighting Environments

To simulate our approach in real-world lighting environments, we have collected a large dataset of 1,067 HDR lighting environments across New York City. A few of lighting environments in the dataset are shown in fig. 7(a). Figure 7(b) shows the location of each captured environment overlaid on a map of New York City. At each location, we capture the environmental illumination using the 360° camera (Ricoh Theta Z1) shown in fig. 7(c). Since the dynamic range of outdoor scenes is very large, we capture $7 \times$ bracketed images to create a single, HDR panorama. Unfortunately, on a sunny day, the sun's disk will be saturated in all of the captured images since the camera's exposure cannot be lowered beyond a certain point [47]. To faithfully measure the radiance from the sun, we capture a second set of $7 \times$ bracketed images with an OD4 neutral density filter (Kodak Wratten 2) between the camera and the sun, as shown in fig. 7(d). We then use PTGui Pro [48] to merge and stitch the two sets of bracketed images into two $7400 \times 3700 \,\mathrm{px}^2$ HDR panoramas. Finally, we compute the radiance from the sun's disk using the panorama with the neutral density filter and overlay those values onto the panorama without the neutral density filter. This process ensures that none of the measurements of the radiance function are saturated.

Alongside each captured environmental illumination, we also measure the global horizontal irradiance using a



(a) HDR Lighting Environments in UrbanSky



(b) Capture Locations in New York City





(c) 360° Camera (d) Camera with Filter

(e) Pvranometer

Figure 7: UrbanSky: Outdoor HDR Lighting Environments in New York City. (a) The dataset consists of 1,067 HDR lighting environments at different locations, times of day, and under varying weather conditions, a few of which are shown here. Each sphere shows the illumination at a particular location and time of day (tonemapped for visualization). (b) UrbanSky consists of lighting environments captured across New York City. Each red dot on the map indicates the location of a single capture. (c) We used a 360° camera to capture the environmental illumination at each location. This camera captures a set of bracketed images that are merged and stitched into an HDR panorama. (d) To faithfully capture the radiance from the sun without saturation, we place a neutral density filter between the camera and the sun to collect a second set of bracketed images. The neutral density filter is mounted on an acrylic pipe that can be easily slid over the camera without disturbing the camera's position. (e) For each lighting environment, we also measure the global horizontal irradiance using a silicon pyranometer placed directly above the camera. Please see the UrbanSky website to explore and download the dataset.

silicon pyranometer (EKO Instruments ML-01) placed directly on top of the 360° camera (see fig. 7(e)). For each environment in the dataset, we include metadata that specifies the global horizontal irradiance (in W/m^2), the capture date, time of day, GPS location, and current weather conditions. In appendix B, we provide additional details about the pyranometer, the image capture process, and post-processing steps.

UrbanSky focuses specifically on capturing the complex illumination in dense urban settings caused by shadows, reflections, and changing weather conditions. To this end, we have chosen 49 specific locations on the campus of Columbia University at which we repeatedly capture the illumination at least 10 different times under varying weather conditions and at different times of day. We then used feature matching to automatically register the captured illumination at each location. In total, we have captured 526 lighting environments at specific locations at Columbia University, which comprise roughly half of UrbanSky. The remaining half of the dataset is comprised of 541 environments captured across New York City (see fig. 7(b)). The UrbanSky website includes an interactive viewer to explore the data captured at each environment. While there are existing datasets of 360° images of real environments [49, 50], UrbanSky focuses specifically on the illumination in dense urban settings. Furthermore, the

pyranometer's measurement of global horizontal irradiance captures the total power per unit area that would be received by a silicon photovoltaic placed in the scene. We have released the entire dataset online to encourage future work both in solar energy and computer graphics.

5.2. Simulations using UrbanSky

We divided the radiance functions in *UrbanSky* into two sets: one with 281 multimodal irradiance functions and the second with 786 unimodal irradiance functions. In our simulations, we assume the energy harvested by a panel to be proportional to its irradiance. For each of several tilt angles between 0° and 90° , and for each radiance function within one of the two sets, we used gradient ascent to iteratively orient a panel starting from many initial orientations and then found the harvested energy as an average over the converged orientations. This energy is expressed as a percentage of the energy harvested using the optimal orientation (the global maximum). Once this is done for all the radiance functions in the set, and for each detector tilt angle, an average percentage harvested energy is computed.

For the multimodal set, the average percentage energy harvested by the panel as a function of the detector tilt angle is shown in fig. 8(a). A tilt angle of $\Delta \theta = 45^{\circ}$ (denoted by the gray line) harvests 5.8% more energy on average



(c) Avg. Harvested Energy v. Tilt Angle for Unimodal Set

Figure 8: Simulations using real-world radiance functions. (a) The average energy harvested by a solar panel (as a percentage of the energy corresponding to the global maximum), plotted as a function of the detector tilt angle of the set of 281 multimodal irradiance functions. A large tilt angle of $\Delta \theta = 45^{\circ}$ (gray line) increases the harvested energy by 5.8% on average compared to the smallest tilt angle. In general, any tilt angle of $\Delta \theta = 45^{\circ}$ up to $\Delta \theta = 90^{\circ}$ achieves a similar performance. (b) A few radiance functions that show that using a large detector tilt angle increases the harvested energy by up to 41%. (c) The average harvested energy as a function of the detector tilt angle for the set of 786 unimodal irradiance functions. The harvested energy decreases ever so slightly with increasing detector tilt angles, which can be attributed to the global maximum shifting under increasing levels of blur.

than the smallest tilt angle. In general, any tilt angle of $\Delta \theta = 45^{\circ}$ up to 90° performs best, achieving a similar average harvested energy. Figure 8(b) shows the percentage energy harvested as a function of the tilt angle for five radiance functions chosen from the multimodal set. As can be seen from the plots, the use of a large tilt angle results in a gain in harvested energy of up to 41% in some scenes when compared with the smallest tilt. For the unimodal set, the average percentage harvested energy as a function of the detector tilt angle is shown in fig. 8(c). Note that the harvested energy decreases ever so slightly with increasing detector tilt angles. This can be attributed to the fact that the global maximum of the irradiance function shifts as the level of blur increases, causing the panel to converge close to but not exactly at the global maximum.

6. Experiments: Orienting a Solar Panel

We developed the prototype shown in fig. 9(a), which includes a solar panel (Voltaic Systems P107C) mounted on a two-axis actuator $(2 \times \text{Miuezuth RDS3218})$. Four photodetectors (ams OSRAM BPW 34) are attached to the four sides of the panel with a tilt angle of $\Delta \theta = 45^{\circ}$, as seen in fig. 9(b). We chose a tilt angle of $\Delta \theta = 45^{\circ}$ since that angle is large enough to significantly increase the harvested energy in multimodal scenes, but not too large to avoid shifting the global maximum of the irradiance function. Since the illumination of outdoor environments generally varies slowly over the day, our prototype measures the photodifferential once every few minutes and then actuates the solar panel with a fixed step size of 5° along both axes. To avoid oscillations in the panel orientation, the panel is only actuated when the magnitude of the photodifferential is above a threshold. Our prototype also includes a measurement system that continuously monitors the energy harvested by the solar panel at its maximum power point and the energy consumed by the actuator. The harvested energy reported in all of our experiments is the net harvested energy after subtracting the energy consumed by the actuator.

In fig. 9(c), we show results for five different scenarios. In each case, we compare the harvested energy using our approach (Proposed Method) with a solar panel that tracks the sun (Sun Tracker) and another in a fixed orientation (Fixed). The trajectory of the sun tracker is determined by the location (latitude and longitude) of the panel and the date and time of day. The fixed panel is oriented toward the equator, with the angle from the zenith equal to the panel's latitude [7]. The plot in the first row of fig. 9(c) shows the harvested power by each panel as a function of time in direct sunlight. The fisheye image shown on the left corresponds to the time denoted by the dashed line in the harvested power plot. As expected, the panel oriented using our method tracks the sun, and thus it harvests the same amount of energy as the sun tracker and 9.2% more energy than the fixed panel over a one-hour period (see the table in fig. 9(d)). A video of the environmental illumination and moving panels is included on the project website.

On a cloudy day, orienting a solar panel toward the sun is almost certainly not the optimal orientation. In the second row of fig. 9(c), the panels are placed outside on a cloudy day. Since our panel points up toward the clouds, it harvested 11.2% more energy than the sun tracker and 9.5% more energy than the fixed panel over a one-hour period.

When a solar panel is deployed in a dense urban environment, the panel's view of the sun is often obstructed by nearby buildings. This case is shown in the third row of fig. 9(c), where the panels are in the shadow cast by nearby buildings. The panel oriented using our approach, however, orients itself toward the unobstructed portion of the sky. In this example, over a one-hour period, our panel



(c) Orienting a Solar Panel in Various Lighting Environments

Figure 9: **Experiments.** (a-b) A prototype solar panel mounted on a two-axis actuator. A photodetector is attached to each of its sides with a tilt angle of $\Delta \theta = 45^{\circ}$. (c) The performance of our panel (Proposed) is compared with panels using two widely used strategies: one that tracks the sun (Sun Tracker) and another in a fixed orientation (Fixed). Plots of the power harvested by the three methods are shown for five different scenarios: direct sunlight, cloudy sky, shadows cast by buildings, reflections from buildings, and an indoor workspace. The dotted line in each plot denotes the moment at which the fisheye image of the environment (left) and the image of the three panels (middle) were captured. (d) The performances of the three panels are compared for the five scenarios in (c). In the realm of renewable energy, the gains obtained by using the proposed method are significant.

harvested 10.6% more energy than the sun tracker and 12.9% more energy than the fixed panel (see fig. 9(d)).

In urban settings, reflections of sunlight by nearby buildings can have a significant impact on the energy harvested by a panel. In the fourth row of fig. 9(c), the panels are placed in a shaded area that is briefly illuminated by strong reflections from a nearby building. The peaks in the plot of the harvested power correspond to times when the reflections were strong. At such times, our panel orients itself in the direction of the reflections. As reported in fig. 9(d), in this setting, our panel harvested 248.1% more energy than the sun tracker and 69.0% more energy than the fixed panel over a one-hour period. During the strongest reflection, the instantaneous harvested power in our case is more than $10 \times$ that of the sun tracker (bottom two rows of fig. 9(d)). Please refer to the project website for a video of the environmental illumination and solar panels.

In our final example (the fifth row of fig. 9(c)), the panels are placed in an indoor environment. Throughout the day, people working in the space turn on and off different floor lamps. Unlike the outdoors, in indoor spaces, the lighting can vary dramatically over time. Hence, in this case, we programmed our panel to iteratively adjust its tilt until convergence every minute. Compared to a fixed panel oriented toward the center of the ceiling, our method harvested 163.7% more energy over a one-hour period. Please see the project website for a video of the indoor scene.

7. Discussion

We have presented a minimal sensing method for iteratively orienting a solar panel to maximize its irradiance, thereby maximizing the energy harvested by it. Given the low cost and complexity of our approach, we believe it has the potential to have real-world impact. There are two directions we plan to pursue as future work. First, we plan to deploy our method at a large scale by working with companies that install medium- and small-sized panels in cities and indoor environments. Second, we plan to explore the case in which it is preferable not to use any moving parts (actuators). In this case, we wish to find the optimal fixed orientation of a solar panel in an any lighting environment while accounting for the complex, time-varying illumination seen from a particular location. This problem is particularly interesting for locations in which the scene's 3D structure and material properties are unknown. The optimization would need to account for nearby infrastructure, the variation of the sun's trajectory, and the historical weather conditions recorded for the location.

Acknowledgements

This work was supported by the Office of Naval Research (ONR) awards N00014-23-1-2096 and N00014-21-1-2378. Jeremy Klotz was supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship. The authors are grateful to Behzad Kamgar-Parsi at ONR for his support and encouragement. The authors also thank Mikhail Fridberg for his input on the prototype, Changxi Zheng for his feedback on the simulations, Bill Miller for his help with fabrication, Makoto Odamaki for technical discussions about the Ricoh Theta camera, and Jean-François Lalonde for sharing the Laval Indoor and Outdoor HDR datasets. Finally, the authors thank Joanne Chan for her help with collecting data for *Urban-Sky* and Lulu Wang for her help in creating the *UrbanSky* website.

Appendix A. Irradiance as a Convolution

The irradiance function $E(\vec{n})$ is defined in eq. (2) as

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

Here we show that $E(\vec{n})$ can be written as the convolution of the radiance function $L(\vec{n})$ with a kernel $k(\vec{n})$, which is defined in eq. (3) as

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

where \hat{z} is the zenith vector.

When convolving a function defined on the surface of the sphere with a kernel, the kernel is first rotated and then multiplied with the original function. Let $R_{\vec{s}}$ be a rotation operator that rotates the zenith vector \hat{z} to \vec{s} , i.e. $\vec{s} = R_{\vec{s}}\hat{z}$. Equation (A.1) can now be written as

$$E(\vec{n}) = \int_{\vec{s} \in \mathbb{S}^2} L(\vec{s}) \max(\vec{n} \cdot R_{\vec{s}} \hat{z}, 0) \, d\omega. \tag{A.3}$$

Using the fact that $\vec{n} \cdot R_{\vec{s}}\hat{z} = R_{\vec{s}}^{-1}\vec{n} \cdot \hat{z}$, we can rewrite eq. (A.3) as

$$E(\vec{n}) = \int_{\vec{s} \in \mathbb{S}^2} L(\vec{s}) \max(R_{\vec{s}}^{-1}\vec{n} \cdot \hat{z}, 0) \, d\omega. \tag{A.4}$$

Substituting eq. (A.2) into eq. (A.4), $E(\vec{n})$ becomes

$$E(\vec{n}) = \int_{\vec{s} \in \mathbb{S}^2} L(\vec{s}) \, k(R_{\vec{s}}^{-1}\vec{n}) \, d\omega. \tag{A.5}$$

Equation (A.5) shows that the irradiance function $E(\vec{n})$ is the convolution of the radiance function $L(\vec{n})$ with the kernel $k(\vec{n})$.⁴ Both Ramamoorthi and Hanrahan [42] and Basri and Jacobs [41] present a similar derivation showing that Lambertian reflectance has the effect of convolving the radiance function with the kernel defined in eq. (A.2). We refer the interested reader to Driscoll and Healy [40] for a detailed introduction to the convolution of two functions defined on the surface of the sphere.

⁴Strictly speaking, the convolution on a sphere is defined by integrating over the three degrees of freedom of the rotation operator [40]. Even though eq. (A.5) only integrates over two degrees of freedom, adding the third degree of freedom does not change the result in our context. Since the kernel $k(\vec{n})$ is symmetric about the z-axis, $k(\vec{n})$ is invariant to rotations about the z-axis. Thus, integrating over the rotation operator's third Euler angle, ψ , only scales the expression by a constant, leaving the convolutional form unchanged.

Appendix B. Details of the UrbanSky Dataset

For each scene in *UrbanSky*, we use a silicon pyranometer (EKO Instruments ML-01) to measure the global horizontal irradiance. We have chosen the pyranometer such that its spectral response and directional response closely match that of a silicon photovoltaic [51]. Thus, the irradiance measurement corresponds to the power per unit area that would be received by a silicon photovoltaic pointing up in the scene. During the capture process, the neutral density filter in fig. 7(d) is only needed to capture the radiance from the sun's disk. On cloudy days, this filter is not needed since the camera can capture a single set of $7 \times$ bracketed images without any saturation. Thus, we only use the neutral density filter on sunny days. Finally, as a post-processing step, we used EgoBlur [52] to automatically detect and blur faces and license plates in the captured images. We then manually reviewed each image to blur faces and license plates that were not automatically detected by EgoBlur.

References

- J. J. Michalsky, *The Astronomical Almanac*'s Algorithm for Approximate Solar Position (1950–2050), Solar Energy 40 (3) (1988) 227–235.
- [2] J. H. Meeus, Astronomical Algorithms, Willmann-Bell, Incorporated, 1991.
- [3] P. Reports, Global Indoor Solar Cell Market Growth 2024-2030 (Jan. 2024).
- [4] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong, P. Bahl, Energy Characterization and Optimization of Image Sensing Toward Continuous Mobile Vision, in: Proceeding of the International Conference on Mobile Systems, Applications, and Services (MobiSys), Association for Computing Machinery, New York, NY, USA, 2013, pp. 69–81.
- [5] A. P. Witkin, Scale-Space Filtering, in: Readings in Computer Vision, Elsevier, 1987, pp. 329–332.
- [6] T. Lindeberg, Scale-Space Theory in Computer Vision, Springer US, Boston, MA, 1994.
- [7] J. A. Duffie, W. A. Beckman, Solar Engineering of Thermal Processes, John Wiley & Sons, Inc., Hoboken, New Jersey, 2006.
- [8] H. Gunerhan, A. Hepbasli, Determination of the optimum tilt angle of solar collectors for building applications, Building and Environment 42 (2) (2007) 779–783.
- [9] M. Lave, J. Kleissl, Optimum fixed orientations and benefits of tracking for capturing solar radiation in the continental United States, Renewable Energy 36 (3) (2011) 1145–1152.
- [10] W. D. Lubitz, Effect of manual tilt adjustments on incident irradiance on fixed and tracking solar panels, Applied Energy 88 (5) (2011) 1710–1719.
- [11] M. Z. Jacobson, V. Jadhav, World estimates of PV optimal tilt angles and ratios of sunlight incident upon tilted and tracked PV panels relative to horizontal panels, Solar Energy 169 (2018) 55–66.
- [12] H. Mousazadeh, A. Keyhani, A. Javadi, H. Mobli, K. Abrinia, A. Sharifi, A review of principle and sun-tracking methods for maximizing solar systems output, Renewable and Sustainable Energy Reviews 13 (8) (2009) 1800–1818.
- [13] C.-Y. Lee, P.-C. Chou, C.-M. Chiang, C.-F. Lin, Sun Tracking Systems: A Review, Sensors 9 (5) (2009) 3875–3890.
- [14] N. AL-Rousan, N. A. M. Isa, M. K. M. Desa, Advances in solar photovoltaic tracking systems: A review, Renewable and Sustainable Energy Reviews 82 (2018) 2548–2569.

- [15] A. Awasthi, A. K. Shukla, M. M. S.R., C. Dondariya, K. Shukla, D. Porwal, G. Richhariya, Review on sun tracking technology in solar PV system, Energy Reports 6 (2020) 392–405.
- [16] N. A. Kelly, T. L. Gibson, Improved photovoltaic energy output for cloudy conditions with a solar tracking system, Solar Energy 83 (11) (2009) 2092–2102.
- [17] J. Mardaljevic, M. Rylatt, Irradiation mapping of complex urban environments: An image-based approach, Energy and Buildings 35 (1) (2003) 27–35.
- [18] D. Robinson, A. Stone, Solar radiation modelling in the urban context, Solar Energy 77 (3) (2004) 295–309.
- [19] R. Compagnon, Solar and daylight availability in the urban fabric, Energy and Buildings 36 (4) (2004) 321–328.
- [20] J. A. Jakubiec, C. F. Reinhart, A method for predicting citywide electricity gains from photovoltaic panels based on LiDAR and GIS data combined with hourly Daysim simulations, Solar Energy 93 (2013) 127–143.
- [21] N. Jakica, State-of-the-art review of solar design tools and methods for assessing daylighting and solar potential for buildingintegrated photovoltaics, Renewable and Sustainable Energy Reviews 81 (2018) 1296–1328.
- [22] Project Sunroof, Google.
- [23] W. A. Lynch, Z. M. Salameh, Simple Electro-Optically Controlled Dual-Axis Sun Tracker, Solar Energy 45 (2) (1990) 65– 69.
- [24] V. Poulek, M. Libra, New solar tracker, Solar Energy Materials and Solar Cells 51 (2) (1998) 113–120.
- [25] V. Poulek, M. Libra, A very simple solar tracker for space and terrestrial applications, Solar Energy Materials and Solar Cells 60 (2) (2000) 99–103.
- [26] I. Sefa, M. Demirtas, I. Çolak, Application of one-axis sun tracking system, Energy Conversion and Management 50 (11) (2009) 2709–2718.
- [27] Y. Away, M. Ikhsan, Dual-axis sun tracker sensor based on tetrahedron geometry, Automation in Construction 73 (2017) 175–183.
- [28] H. Wang, T. Luo, H. Song, J. B. Christen, On-chip sensor for light direction detection, Optics Letters 38 (22) (2013) 4554.
- [29] H. Wang, T. Luo, Z. Lu, H. Song, J. B. Christen, CMOS selfpowered monolithic light-direction sensor with digitalized output, Optics Letters 39 (9) (2014) 2618.
- [30] H. Wang, T. Luo, Y. Fan, Z. Lu, H. Song, J. B. Christen, A selfpowered single-axis maximum power direction tracking system with an on-chip sensor, Solar Energy 112 (2015) 100–107.
- [31] J.-M. Wang, C.-L. Lu, Design and Implementation of a Sun Tracker with a Dual-Axis Single Motor for an Optical Sensor-Based Photovoltaic System, Sensors 13 (3) (2013) 3157–3168.
- [32] R. Zogbi, D. Laplaze, Design and construction of a sun tracker, Solar Energy 33 (3-4) (1984) 369–372.
- [33] G. L. Haywood, W. J. Haywood, Solar collector and drive circuitry control means (Apr. 1978).
- [34] G. M. Tina, F. Arcidiacono, A. Gagliano, Intelligent suntracking system based on multiple photodiode sensors for maximisation of photovoltaic energy production, Mathematics and Computers in Simulation 91 (2013) 16–28.
- [35] J. M. Quero, A. Guerrero, L. G. Franquelo, M. Dominguez, I. Ameijeiras, L. Castaner, Light source position microsensor, in: ISCAS 2001. The 2001 IEEE International Symposium on Circuits and Systems, Vol. 2, IEEE, Sydney, Australia, 2001, pp. 648–651.
- [36] P. Roth, A. Georgiev, H. Boudinov, Design and construction of a system for sun-tracking, Renewable Energy 29 (3) (2004) 393–402.
- [37] F. Chen, J. Feng, Analogue sun sensor based on the optical nonlinear compensation measuring principle, Measurement Science and Technology 18 (7) (2007) 2111–2115.
- [38] F. Pineda, C. A. Arredondo, Design and implementation of sun tracker prototype for solar module positioning, in: 2012 38th IEEE Photovoltaic Specialists Conference, IEEE, Austin, TX, USA, 2012, pp. 002905–002910.
- [39] M. Levoy, P. Hanrahan, Light Field Rendering, in: Proceed-

ings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '96, Association for Computing Machinery, New York, NY, USA, 1996, pp. 31–42.

- [40] J. R. Driscoll, D. M. Healy, Computing Fourier Transforms and Convolutions on the 2-Sphere, Advances in Applied Mathematics 15 (2) (1994) 202–250.
- [41] R. Basri, D. W. Jacobs, Lambertian Reflectance and Linear Subspaces, IEEE Transactions on Pattern Analysis and Machine Intelligence 25 (2) (2003) 218–233.
- [42] R. Ramamoorthi, P. Hanrahan, 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) (2001) 2448.
- [43] F. Dumont, Analysis of Some Finite Difference Schemes for Slightly Noisy Time Dependent Signals (Mar. 2015).
- [44] J. J. Koenderink, The Structure of Images, Biological Cybernetics 50 (5) (1984) 363–370.
- [45] J. Babaud, A. P. Witkin, M. Baudin, R. O. Duda, Uniqueness of the Gaussian Kernel for Scale-Space Filtering, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 8 (1) (1986) 26–33.
- [46] A. L. Yuille, T. A. Poggio, Scaling Theorems for Zero Crossings, IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-8 (1) (1986) 15–25.
- [47] J. Stumpfel, A. Jones, A. Wenger, C. Tchou, T. Hawkins, P. Debevec, Direct HDR capture of the sun and sky, in: ACM SIG-GRAPH 2006 Courses, SIGGRAPH '06, Association for Computing Machinery, New York, NY, USA, 2006.
- [48] PTGui, New House Internet Services B.V.
- [49] M.-A. Gardner, K. Sunkavalli, E. Yumer, X. Shen, E. Gambaretto, C. Gagné, J.-F. Lalonde, Learning to Predict Indoor Illumination from a Single Image, ACM Transactions on Graphics 36 (6) (2017) 1–14.
- [50] Y. Hold-Geoffroy, A. Athawale, J.-F. Lalonde, Deep Sky Modeling for Single Image Outdoor Lighting Estimation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 6927–6935.
- [51] EKO Instruments, ML-01 Silicon Pyranometer Instruction Manual.
- [52] N. Raina, G. Somasundaram, K. Zheng, S. Miglani, S. Saarinen, J. Meissner, M. Schwesinger, L. Pesqueira, I. Prasad, E. Miller, P. Gupta, M. Yan, R. Newcombe, C. Ren, O. M. Parkhi, EgoBlur: Responsible Innovation in Aria (Sep. 2023). arXiv:2308.13093.