An Eternal Irradiance Camera Supplemental Material

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

S1 SPHERICAL HARMONIC CONVENTIONS

Throughout the paper, we use the real spherical harmonics $Y_{\ell m}$, indexed by degree $\ell \geq 0$ and order $-\ell \leq m \leq \ell$. For any direction $\vec{n}$ with zenith angle θ and azimuthal angle ϕ, the real spherical harmonics are defined as

$$Y_{\ell m}(\theta, \phi) = \begin{cases} \sqrt{2}\, N_\ell^{|m|}\, P_\ell^{|m|}(\cos\theta)\ \sin(|m|\phi), & m < 0, \\ N_\ell^0\, P_\ell^0(\cos\theta), & m = 0, \\ \sqrt{2}\, N_\ell^m\, P_\ell^m(\cos\theta)\ \cos(m\phi), & m > 0, \end{cases} \tag{1}$$

where P_ℓ^m is the associated Legendre polynomial of degree ℓ and order $m \geq 0$, and N_ℓ^m is the normalization constant

$$N_\ell^m = (-1)^m \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}}. \tag{2}$$

We use the Condon–Shortley phase convention in which the $(-1)^m$ phase term is included in P_ℓ^m, which cancels out with the $(-1)^m$ phase term in N_ℓ^m.

S2 RECONSTRUCTING THE IRRADIANCE FUNCTION

S2.1 Reconstruction Method

We reconstruct the spherical harmonic coefficients of the irradiance function by inverting the sampling process in eq. (4) in the paper. Recall from eq. (3) in the paper that only 31 of the 49 coefficients in the first seven degrees are non-zero. We therefore solve only for those 31 coefficients, whose degrees and orders we index as $(\ell_1, m_1), \ldots, (\ell_{31}, m_{31})$. Since irradiance must be non-negative, not every set of spherical harmonic coefficients corresponds to a physically valid irradiance function. To enforce this constraint, we solve for the 31 coefficients from the 49 measurements using constrained least squares:

$$\begin{aligned} \min_x \quad & \|Ax - b\|_2^2 \\ \text{subject to} \quad & Gx \geq 0, \end{aligned} \tag{3}$$

where

$$A = \underbrace{\begin{bmatrix} Y_{\ell_1 m_1}(\vec{n}_1) & \ldots & Y_{\ell_{31} m_{31}}(\vec{n}_1) \\ \vdots & \ddots & \vdots \\ Y_{\ell_1 m_1}(\vec{n}_{49}) & \ldots & Y_{\ell_{31} m_{31}}(\vec{n}_{49}) \end{bmatrix}}_{49\times 31}, \tag{4}$$

$$x = \underbrace{\begin{bmatrix} E_{\ell_1 m_1} \\ \vdots \\ E_{\ell_{31} m_{31}} \end{bmatrix}}_{31\times 1}, \tag{5}$$

$$b = \underbrace{\begin{bmatrix} E(\vec{n}_1) \\ \vdots \\ E(\vec{n}_{49}) \end{bmatrix}}_{49\times 1}, \tag{6}$$

$$G = \underbrace{\begin{bmatrix} Y_{\ell_1 m_1}(\vec{n}'_1) & \ldots & Y_{\ell_{31} m_{31}}(\vec{n}'_1) \\ \vdots & \ddots & \vdots \\ Y_{\ell_1 m_1}(\vec{n}'_K) & \ldots & Y_{\ell_{31} m_{31}}(\vec{n}'_K) \end{bmatrix}}_{K\times 31}. \tag{7}$$

Here, $\{\vec{n}'_1, \ldots, \vec{n}'_K\}$ is a dense set of uniformly distributed directions on the sphere generated using HEALPix [Gorski et al. 2005]. This constrained least squares formulation ensures that the reconstructed irradiance function evaluated at other points on the sphere, not just the measured points $\{\vec{n}_1, \ldots, \vec{n}_{49}\}$, is non-negative. In our implementation, we use $K = 4800$ points. We solve for the spherical harmonic coefficients using the `quadprog` solver in `qpsolvers` [Caron et al. 2026].

The 49 sample locations $\{\vec{n}_1, \ldots, \vec{n}_{49}\}$ are listed in table S1. Code for reconstructing the irradiance function from the 49 samples is available online [Klotz and Nayar 2026].

S2.2 Reconstruction Stability

We analyze the stability of the reconstruction method by examining the condition number of A (eq. (4)). The blue line in Figure S1 plots the condition number of A corresponding to roughly uniform sample arrangements (approximate solutions to the Tammes problem) for different sample counts. For the 49 samples used by our system (vertical red line), A has a condition number of 2.69. As expected, the conditioning improves if more samples are used.

Even though there are 31 unknowns (i.e., A has 31 columns), not every arrangement of $N \geq 31$ samples yields a well-conditioned reconstruction problem. The gaps in fig. S1 denote sample arrangements that are ill-conditioned. This occurs when the specific sample layout exhibits symmetries that produce linearly dependent columns

Authors' addresses: 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.

Table S1. **Sample locations.** The coordinates of the 49 sample locations $\{\vec{n}_1, \dots, \vec{n}_{49}\}$ on the unit sphere. These samples are an approximate solution to the Tammes problem [Sloane et al. 1994].

n	X	Y	Z	n	X	Y	Z
1	0.407	0.511	−0.757	26	0.956	−0.247	0.157
2	−0.569	−0.802	0.179	27	−0.131	0.967	0.217
3	0.307	0.865	−0.396	28	−0.840	0.042	0.540
4	0.654	−0.289	0.699	29	−0.263	−0.761	0.593
5	−0.710	−0.288	−0.643	30	0.403	−0.915	−0.012
6	−0.876	−0.398	0.272	31	0.046	0.213	−0.976
7	−0.484	0.276	0.831	32	−0.096	−0.991	0.097
8	0.512	0.015	−0.859	33	−0.194	0.936	−0.295
9	0.823	−0.223	−0.523	34	−0.591	0.656	−0.470
10	−0.396	−0.866	−0.304	35	−0.583	0.812	0.022
11	0.726	0.668	−0.166	36	0.373	0.921	0.113
12	−0.305	−0.163	−0.938	37	0.507	0.391	0.768
13	−0.801	−0.559	−0.213	38	−0.909	0.416	−0.037
14	0.013	0.401	0.916	39	0.780	−0.601	−0.174
15	−0.289	0.714	0.638	40	−0.080	0.680	−0.729
16	0.286	−0.036	0.958	41	0.221	0.780	0.586
17	−0.597	−0.382	0.705	42	0.125	−0.514	0.849
18	0.169	−0.367	−0.915	43	0.972	0.235	−0.027
19	−0.438	0.325	−0.838	44	0.104	−0.896	−0.432
20	0.807	0.293	−0.513	45	0.230	−0.851	0.472
21	0.861	0.136	0.490	46	−0.988	−0.086	−0.125
22	−0.727	0.534	0.432	47	0.500	−0.617	−0.608
23	0.684	−0.650	0.331	48	0.713	0.610	0.347
24	−0.828	0.200	−0.523	49	−0.236	−0.634	−0.737
25	−0.218	−0.146	0.965				

in A. However, the conditioning for any given number of samples could be improved using a non-uniform sample layout. To illustrate this, we have iteratively adjusted the sample locations using gradient descent to minimize the condition number. The resulting sample layouts (black line in fig. S1) would improve the conditioning compared to the uniform sample layouts from approximate solutions to the Tammes problem. For the layout using $N = 49$ samples, the condition number would decrease from 2.69 to 1.57 using this refined, non-uniform layout. In practice, such non-uniform sample layouts come at the cost of requiring a larger sphere since some samples (detectors) would be more tightly packed.

S3 SIMULATED SENSOR MODEL

For the experiments evaluating the accuracy of the reconstructed irradiance functions in simulation, we model a sensor that saturates at 20,000 lx (corresponding to 3.3 V) and includes 250 µV root mean square (RMS) Gaussian noise. For any scene in which the maximum illumination exceeds 20,000 lx (i.e., the sensor would saturate), we simulate the effect of adding a 3-stop neutral density (ND) filter to all 49 detectors. Then, we model the saturation by clipping the measurements at 3.3 V. Figure S6 shows additional simulation results.

S4 FLUXCAM ARCHITECTURE

Table S2 lists the components used in FluxCam, our prototype irradiance camera. Each photovoltaic is connected through a switch to

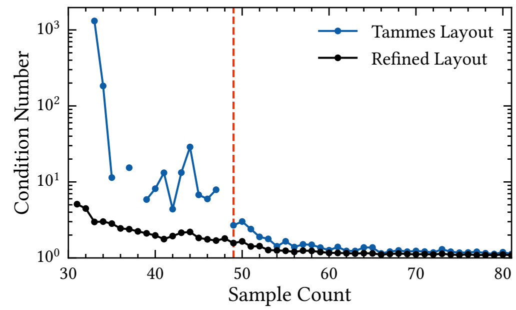

Fig. S1. **Conditioning of the reconstruction problem for different sample arrangements.** The condition number of A using sample arrangements given by approximate solutions to the Tammes problem is shown in blue. The condition number of A after iteratively refining the sample locations is shown in black.

Table S2. **Components in the prototype irradiance camera.** The hardware design (circuit schematics and PCB layouts) is available online [Klotz and Nayar 2026].

Component	Quantity	Description
Photovoltaic	49	Panasonic AM-5610CAR-DGK-T
Amplifier	49	TI OPA607
ADC	49	Analog Devices AD7683
Energy Harvester	49	TI BQ25570
Switch	49	Analog Devices ADG787
Readout Microcontroller	1	STM32L051C6T6
RF Microcontroller	1	TI CC1310F128
Mux	6	Analog Devices ADG706
Supercapacitor	1	107 mF
Antenna	1	ANTX150P116B08683

a transimpedance amplifier (TIA) with a gain of 4700 V/A. The TIA measures the photovoltaic's short circuit current without applying any reverse bias. A 16-bit ADC measures the output of the TIA and sends that measurement to the processor board over an SPI bus. Once the measurement is made, the photovoltaic is switched back into the energy harvester to power the camera. **The circuit schematics, PCB layouts, and firmware are available online [Klotz and Nayar 2026].**

When the photovoltaic is initially switched into the TIA, the photocurrent requires time to settle before reaching its steady state current. Figure S2 shows the transient response of the TIA output immediately after switching the detector from the energy harvester into the TIA. The switch occurs at the moment the ENABLE signal goes high (see fig. S2). The transient takes roughly 20 ms to achieve steady state. This settling time can be set in firmware to trade off measurement fidelity (which benefits from a long settling time) with energy harvesting efficiency (which benefits from a short settling time).

To minimize the power consumption, the microcontrollers on the processor board are typically asleep. In our current implementation, they wake up at a predefined interval (nominally 30 FPS) to read out the measurements from the sensor boards, wirelessly transmit them, and then go back to sleep. At 30 FPS with a 50 µs settling time, the

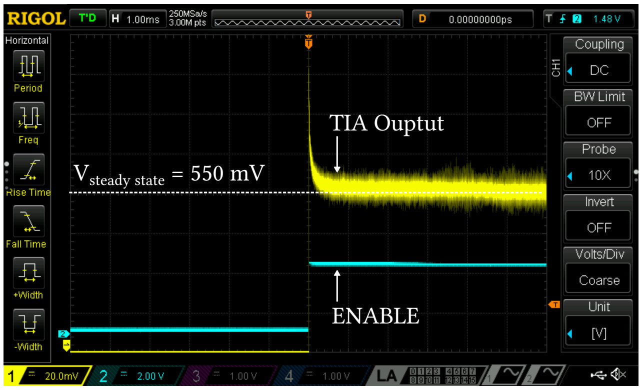

Fig. S2. **Settling time of transimpedance amplifier output.** When the ENABLE signal goes high, the sensor board switches the photovoltaic from the energy harvester into the transimpedance amplifier. The output of the transimpedance amplifier, which is proportional to the photovoltaic's short circuit current, takes roughly 20 ms to settle to its steady state value (550 mV in this case). The choice of when to sample the TIA output represents a tradeoff between producing high fidelity measurements and harvesting more energy.

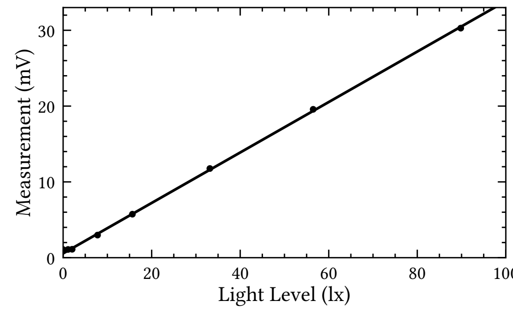

Fig. S3. **Radiometric response function at low light levels.**

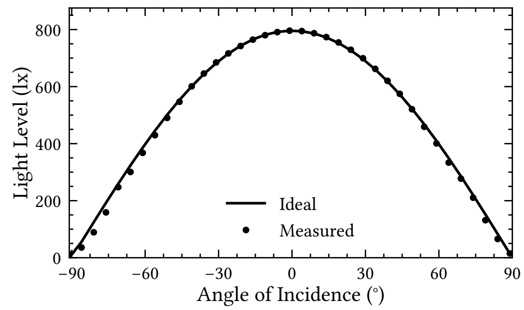

Fig. S4. **Angular response of the detector.** Across incident angles ranging from −80° to +80°, the measurements deviate from the ideal cosine response by 3.1% on average.

processor board consumes 500 µA, and each sensor board consumes roughly 3 µA.

S5 FLUXCAM PERFORMANCE

S5.1 Radiometric Performance

To characterize the radiometric performance and signal-to-noise ratio (SNR), we used a long settling time of 20 ms. In this case, we operated the sensor at 30 FPS and attached a 681 kΩ load to the output of the harvester. To measure the radiometric response at a single light level, we used the average measurement taken over 15 seconds (corresponding to 450 measurements at 30 FPS). Figure S3 shows the radiometric response function at very low light levels, with a line fit to the measurements.

To measure the radiometric response and SNR at low light levels (roughly 1,500 lx and below), we used a Thorlabs OSL2 halogen lamp. To measure the performance at higher light levels, we used a brighter halogen lamp (IBM 663533 B). We applied a constant scale factor to measurements taken under the brighter halogen lamp to account for the spectral mismatch with the dimmer lamp.

Each of the 49 detectors may respond differently to the same light level due to component tolerances. To account for this, we measured the radiometric response of all 49 detectors at two different light levels using a controlled light source. Then, we computed the relative gain for each detector using the two measurements.

In addition, we have measured the angular response of the detector along one dimension by placing the detector on a manual rotation stage. For this experiment, a controlled light source illuminates the bare detector from a distance. Figure S4 shows the measured light level (in lux) as a function of the angle of incidence of the light source. Over incident angles ranging from −80° to +80°, the measurement deviates from the ideal cosine response by an average of 3.1%.

When the detectors are installed in the prototype, the 3D printed shell slightly occludes their field of view, which in turn impacts their angular response. Using the known geometry of the 3D model, we have calculated that the shell blocks rays with an angle of incidence $\theta \geq 78.5°$ from reaching the center of each detector. Given that the foreshortened solid angle is small at these large angles of incidence, the occlusion caused by the shell has little impact on the detector's irradiance.

S5.2 Energy Harvesting

To characterize the energy harvesting performance, we measured the change in energy stored in a capacitor (connected to VSTOR of the BQ25570 energy harvester) without a load attached. The harvested power was computed by dividing the change in stored energy (J) by the time (s). In this case, we operated the sensor at 30 FPS with a short settling time of 50 µs.

S6 ADAPTIVE FRAMERATE ALGORITHM

The power consumption of the irradiance camera is roughly linear with the framerate. Therefore, we could adaptively adjust the camera's framerate so that it only transmits when it has harvested enough energy to do so. This could be implemented as follows: when the irradiance camera wakes up, it reads out all 49 measurements. If the camera has enough stored energy (as measured by the voltage across the supercapacitor), the camera wirelessly transmits those measurements. Then, the camera goes back to sleep for a time that is inversely proportional to the average irradiance measurement over the last 60 seconds. Even if the camera does not have enough energy to wirelessly transmit, this algorithm ensures that it still periodically wakes up to measure irradiance. This allows the camera to adapt to changes in illumination without waiting to harvest enough energy for wireless transmission.

S7 MEASURED IRRADIANCE FUNCTIONS IN REAL ENVIRONMENTS

Figure S7 shows the reconstructed irradiance function using measurements made by the prototype camera in four real environments. We used a Ricoh Theta Z1 to capture an HDR spherical image in each scene, which was used to compute the ground truth irradiance function. To compare the ground truth and measured irradiance functions, we first need to align the poses of the Ricoh Theta Z1 camera and FluxCam. To do this, we begin with a coarse estimate of the relative rotation between the two cameras. Then, we refine this estimate by searching over all rotations within a 20° cone to find the rotation that yields the highest normalized cross-correlation between the ground truth and measured irradiance functions. In this case, we used a long settling time of 20 ms to produce high fidelity irradiance measurements.

S8 ESTIMATING EGOMOTION

S8.1 Simulation Details

We used Mitsuba [Jakob et al. 2022] to simulate the irradiance camera in an urban environment. For each camera location, we rendered the irradiance of its 49 detectors while accounting for their different viewpoints. The irradiance of each detector was simulated using a direct illumination renderer with 16,000 samples. Then, we applied a random rotation and small translation, ranging from 0 m to 0.1 m, which corresponds to a maximum speed of 2 m/s sampled at 30 FPS, and rendered the measurements in this new pose. In total, this process created 1,120,250 samples (pairs of irradiance camera measurements) for training, 160,617 samples for validation, and 319,712 samples for testing.

S8.2 Network Architecture and Training

The network is a transformer, in which each measurement produced by the irradiance camera corresponds to a single token with an embedding dimension of $D = 512$. The inputs to the network are the 49 measurements (in units of volts) at the two poses, and the output of the network is the rotation matrix between the two camera poses. The network was trained by minimizing the relative angle between the predicted and ground truth rotations using the AdamW optimizer [Loshchilov and Hutter 2019] with a learning rate of 1e-4.

S8.3 Results

Over the 319,712 test samples, the average error between the predicted and ground truth rotation was 2.25°, and the median error was 1.17°. Figure S5 shows the predicted rotation error as a function of the irradiance camera's height above the ground and translation distance. The prediction error is higher when the camera is near the ground or translated by a larger distance. This is expected since the impact of translation on the measured irradiance function is amplified by nearby objects and large translations.

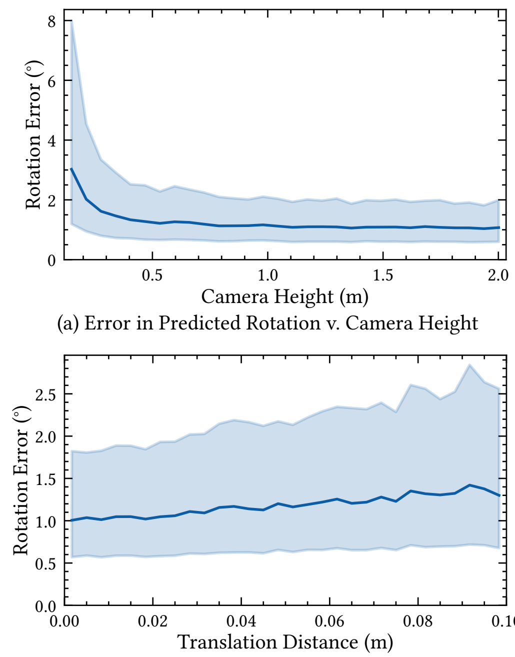

(a) Error in Predicted Rotation v. Camera Height

(b) Error in Predicted Rotation v. Translation Distance

Fig. S5. **Performance of egomotion estimation.** (a) Error in the predicted rotation as a function of camera height above the ground. (b) Error in the predicted rotation as a function of the translation distance. The solid blue line shows the median error, and the light blue region denotes the interquartile range.

S8.4 Estimating Orientation using Real Measurements

To compute the FluxCam's orientation from its measurements, we use the trained network to estimate the relative orientation between individual frames (each frame is a set of 49 irradiance measurements). Rather than estimating the relative orientation between consecutive frames, we use the network to estimate the orientation between two frames that are separated in time. We found empirically that this produced better estimates of the camera orientation than using consecutive frames. To combine the estimated orientations, we first scale the angle of each relative rotation to account for the fact that they were computed over a longer time horizon. Then, we accumulate the relative rotations by applying them in sequence. Finally, we filter the resulting camera orientation in quaternion space using a Savitzky–Golay filter.

S9 ESTIMATING SKY CONDITIONS

S9.1 Experimental Details

We used FluxCam's measurements to estimate the current sky conditions. To do this, we fit a low-dimensional analytical sky model [Perez et al. 1993] to the measurements. Due to the spectral response of FluxCam's photovoltaics, the prototype measures illuminance (in lux). Therefore, we first convert FluxCam's measurements to irradiance (in $\mathrm{W\,m^{-2}}$) using a luminous efficacy of $109\,\mathrm{lm\,W^{-1}}$. This luminous efficacy was computed using the global spectrum in the reference air mass 1.5 standard spectra [ASTM 2021]. To minimize the impact of the ground and nearby buildings on the measured irradiances, we placed FluxCam inside a cylinder with black walls. This cylinder served as an aperture that blocked the irradiance camera's view of the ground and buildings near the horizon. Since the geometry of this shading cylinder is known, we accounted for its impact on each detector's view of the sky.

For any given sky model, we can apply the aperture caused by the shading cylinder to the corresponding radiance function and then compute what the irradiance measurements would be. Using this approach, we fit the Perez sky model (i.e., a model of the sky's radiance function) to the measured irradiances. Additionally, we ignored the measurements produced by the detectors that are oriented sideways or toward the ground.

In order to fit the sky model, we assume that the pose of the irradiance camera in the earth coordinate frame (i.e., its orientation with respect to true north) is known. Additionally, we use the time of day and rough GPS location to compute the sun's position in the sky, which is an input to the sky model.

We have also assumed that the global horizontal irradiance (GHI) is known. This is reasonable since an irradiance camera could be positioned such that one detector has a clear view of the sky, and hence it would directly measure the GHI. In our case, we measured the GHI using a silicon pyranometer (EKO Instruments ML-01). Since GHI is known, the only sky parameter to fit is the diffuse horizontal irradiance (DHI), which specifies the irradiance of a horizontal surface due to the sky, excluding the sun. Note that in the Perez sky model, the direct normal irradiance (DNI) is computed analytically from the GHI and DHI.

S9.2 Validation in Simulation

We validated the accuracy of this approach for measuring the sky conditions in simulation. We used the Perez sky model to generate the sky's radiance function at every hour of a typical meteorological year in Manhattan. We ignored times when the sun's elevation was lower than 12°. Then we simulated the irradiance measurements produced by the camera inside a shading cylinder that blocks objects in the scene with an elevation of 12° or lower. Additionally, this simulation emulates the characteristics of a real sensor that measures irradiance (in $\mathrm{W\,m^{-2}}$) and saturates at 1,200 $\mathrm{W\,m^{-2}}$ (corresponding to 3.3 V). We assume that the voltage produced by the sensor is proportional to its irradiance, and we model sensor noise by adding 1 mV RMS Gaussian noise to the measurements. We also assume that the shading cylinder is completely black, i.e., the radiance from the cylinder is 0. As before, we also assume that the GHI and the camera's pose in the earth coordinate frame are known.

Over 2,494 different skies at different times of day throughout the year, this approach estimated the DHI with a root mean square error (RMSE) of 8.3 $\mathrm{W\,m^{-2}}$ and the DNI with an RMSE of 76.0 $\mathrm{W\,m^{-2}}$. The estimated Perez clearness had an RMSE of 0.28, and the estimated Perez brightness had an RMSE of 0.01.

REFERENCES

ASTM. 2021. G-173-03: Standard Tables for Reference Solar Spectral Irradiances: Direct Normal and Hemispherical on 37° Tilted Surface.

Stéphane Caron, Daniel Arnström, Suraj Bonagiri, Antoine Dechaume, Nikolai Flowers, Adam Heins, Takuma Ishikawa, Dustin Kenefake, Giacomo Mazzamuto, Donato Meoli, Brendan O'Donoghue, Adam A. Oppenheimer, Pavel Otta, Abhishek Pandala, Juan José Quiroz Omaña, Nikitas Rontsis, Paarth Shah, Samuel St-Jean, Nicola Vitucci, Soeren Wolfers, Fengyu Yang, Brian Delhaisse, MeindertHH, rimaddo, urob, shaoanlu, Uriel Sandoval, Ahmed Khalil, Lev Kozlov, Antoine Groudiev, João Sousa Pinto, Roland Schwan, Rohan Budhiraja, Johannes Keustermans, Yubin Wu, Hongpei Li, and Antoine Hoarau. 2026. Qpsolvers: Quadratic Programming Solvers in Python.

K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. 2005. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal* 622, 2 (April 2005), 759–771.

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. 2022. Mitsuba 3 Renderer.

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

Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.

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.

Neil J. Sloane, R. H. Hardin, and W. D. Smith. 1994. Tables of Spherical Codes. http://neilsloane.com/packings/.

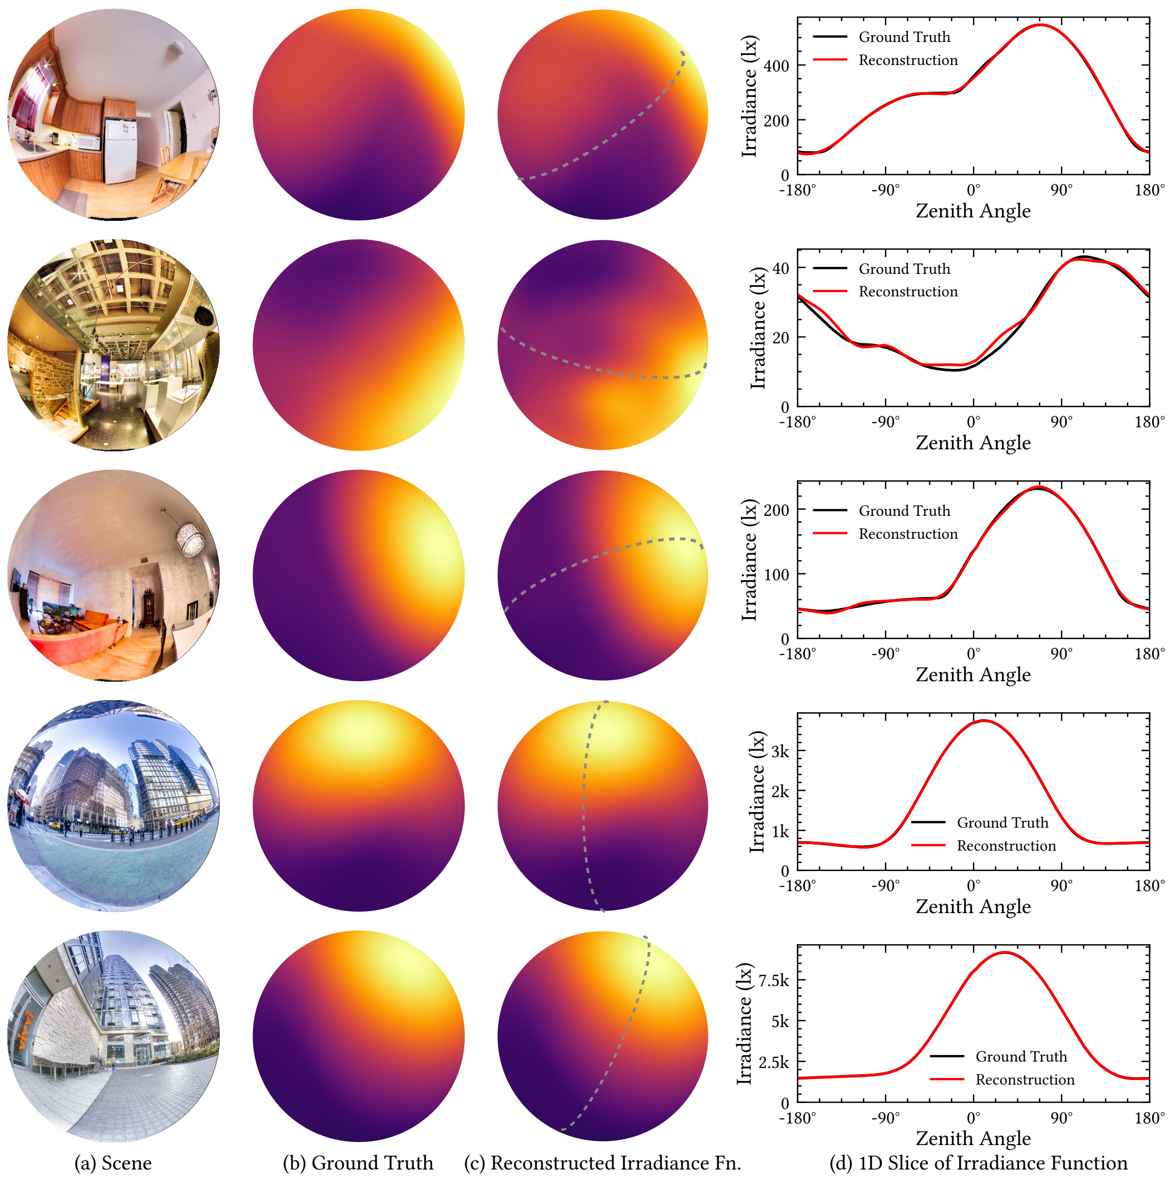

Fig. S6. **Additional simulation results in scenes from the Laval Indoor and UrbanSky datasets.** (a) HDR image of the scene (radiance function). (b) Ground truth irradiance function. (c) Reconstructed irradiance function from simulated measurements. (d) 1D slice of the reconstructed irradiance functions through the dashed gray lines in (c).

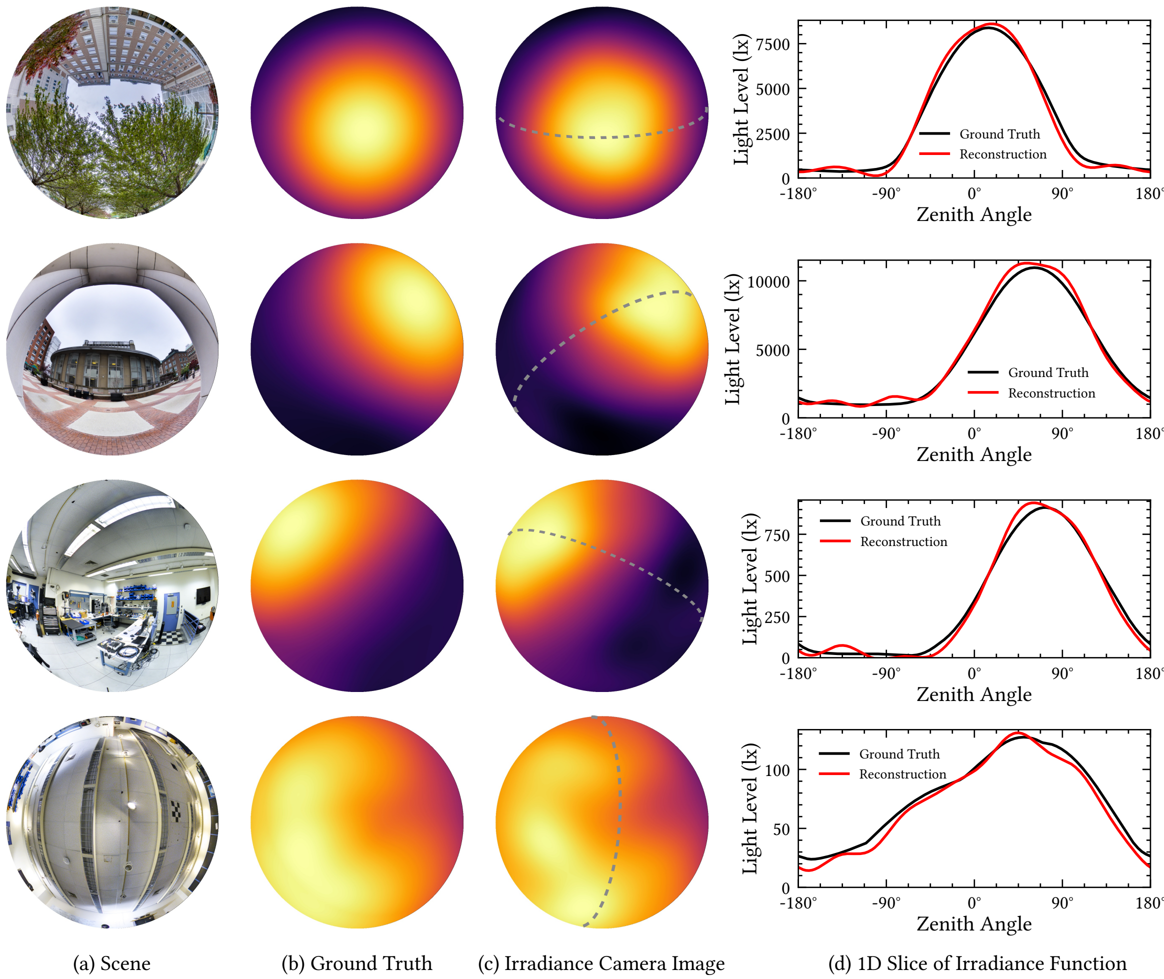

Fig. S7. **Irradiance functions measured by the prototype camera in four environments.** (a) HDR image of the scene captured by a Ricoh Theta Z1. (b) Ground truth irradiance function computed using the HDR fisheye image. (c) Irradiance function measured by FluxCam. (d) A 1D slice through the dashed gray lines in (c).