

Rolling Shutter Imaging on The Electric Grid

Mark Sheinin, Yoav Y. Schechner
Viterbi Faculty of Electrical Engineering
Technion - Israel Institute of Technology

marksheinin@gmail.com, yoav@ee.technion.ac.il

Kiriakos N. Kutulakos
Dept. of Computer Science
University of Toronto

kyros@cs.toronto.edu

Abstract

Flicker of AC-powered lights is useful for probing the electric grid and unmixing reflected contributions of different sources. Flicker has been sensed in great detail with a specially-designed camera tethered to an AC outlet. We argue that even an untethered smartphone can achieve the same task. We exploit the inter-row exposure delay of the ubiquitous rolling-shutter sensor. When pixel exposure time is kept short, this delay creates a spatiotemporal wave pattern that encodes (1) the precise capture time relative to the AC, (2) the response function of individual bulbs, and (3) the AC phase that powers them. To sense point sources, we induce the spatiotemporal wave pattern by placing a star filter or a paper diffuser in front of the camera's lens. We demonstrate several new capabilities, including: high-rate acquisition of bulb response functions from one smartphone photo; recognition of bulb type and phase from one or two images; and rendering of live flicker video, as if it came from a high speed global-shutter camera.

1. Introduction

The electric grid powers light bulbs by alternating current (AC). In response, bulbs flicker quasi-periodically [2, 44]. The flicker is usually too subtle and fast for the naked eye to notice [19]. In photography and videography, flicker has usually been considered a nuisance [13,17,28,38,48], as it leads to spatiotemporal brightness and color artifacts [11]. However, flicker can be useful. Tajbakhsh and Grigat [39] used flicker to classify the electric grid's temporal frequency. More recently, Bianco *et al.* and Sheinin *et al.* [3,36] exploited flicker to extract rich information about the electric grid, including the distribution of electric phases, bulb types and electric perturbations. Moreover, flicker facilitates scene unmixing into reflected contributions of different light sources [36], *i.e.*, passively determining the light transport matrix [6,34] of the scene.

To sense flicker for quantitative analysis, Bianco *et al.* and Sheinin *et al.* [3,36] built computational imaging systems. As in other coded-aperture systems [8,25,31,42], the ACam [36] involves an elaborate electro-optical setup. Moreover, it has to be tethered to the AC grid for real-time synchronization and control. This challenges its ease of use.




Figure 1. (Top) Rolling-shutter imaging flicker of AC lighting results in a vertical wave pattern. We compute global-shutter frames of the scene. (Bottom) Point emitters imaged through an optical star filter. Along the vertical arms of the star point-spread function a rolling-shutter frame reveals wave-like modulation created by AC-induced flicker.

It is desirable to have untethered systems that are off-the-shelf camera devices, *e.g.*, smartphones.

This paper shows that simple, ubiquitous rolling-shutter cameras suffice to quantitatively sense and analyze flicker. A rolling shutter creates an inter-row delay of exposure. This delay is constant and very short, enabling fast and dense temporal sampling of flicker. This effect creates a spatiotemporal wave pattern in raw rolling-shutter images (Figure 1), which we derive in Section 3. Moreover, regularity and periodicity of this pattern lead to a set of scene analysis approaches, suiting different data sizes and scene types (Table 1). Even a single rolling-shutter image yields useful information (Section 4). This includes the types and electric phases of light sources connected to the grid; regions in the scene they illuminate; the bulbs' temporal response to AC; and the time at which the rolling-shutter image was acquired, relative to zero-crossing of the grid voltage.

The results are significantly improved if the input short-exposure rolling-shutter frame is accompanied by an image

Method	Frames	Applies to	Discussed in	Key assumptions
Image deflicker	1	Diffuse surfaces	Section 4.1.1	Spatially-smooth light transport, reflection at each pixel dominated by one source
Single-frame analysis	1	Diffuse surfaces	Section 4.1.2	Spatially-smooth light transport, reflection at each pixel dominated by one source, BRF database
Two-frame analysis	2	Diffuse surfaces	Section 4.2	Reflection at each pixel dominated by one source, BRF database
Star-filter analysis	2	Directly viewed sources	Section 4.3	Image contains spatially separated light sources, BRF database, An optical star-filter attached to the lens
Surface-based	Many	Scenes having a diffuse surface	Section 5.1.1	A wide, diffuse surface is in view
Emitter-based	Many	Any scene	Section 5.1.2	Entire batch of frames
Render global shutter	Many	Any scene	Section 6	After frame assignment by Section 5, one known scene bulb (for AC analysis application only)
Real-time	Stream	Scenes having a diffuse surface	Section 6	Continuous assignment using Section 5, one known scene bulb (for AC analysis application only)

Table 1. Summary of methods considered in this paper. Each method represents a different scene type or a different imaging procedure.

whose exposure time lasts an integer multiple of the flicker cycle (Section 4.2). The latter emulates illumination by direct current (DC), easing signal decomposition to flicker illumination modulation and a flicker independent image. In Section 4.3, we induce the wave pattern spatially in an image of point emitters by mounting a star-filter [29] on the camera (Figure 1).

Section 5 deals with a data sequence of many raw rolling-shutter frames. The sequence is converted to a canonical volume that expresses the spatiotemporal signal of flicker, as if acquired by a very fast global-shutter camera (Figure 2). This conversion can handle temporally disorganized sequences, such as those acquired manually as stills by a smartphone. We use this approach in Section 6 to render the scene at any arbitrary in-cycle canonical time, that had not been sampled during rolling-shutter data acquisition. This capability exploits the periodicity and regularity of flicker, which naturally suits Fourier analysis.

Rolling-shutter cameras were leveraged in computational photography [26,27,37] for three-dimensional sensing [1,5,30], compressed sensing [7,10] and white balance [9]. They have also been used for inter-device communication [12,14,47]. We show experimental results using several rolling-shutter cameras: two off-the-shelf machine vision cameras and two smartphones. This work is related to prior art on use of radiometric measurements for self-calibration [15,18,31,33,43].

2. Theoretical Background

2.1. Imaging on The Electric Grid

We briefly describe the model of illumination under AC power, adopting some notations from [36]. AC voltage is approximately a sinusoid in time. Its frequency f is either 50Hz or 60Hz, depending on the country. Bulbs flicker at double the AC frequency, in a cycle lasting $\Delta = 1/(2f)$ seconds. Each bulb is powered by one of three electric grid lines, having respective AC phases $\phi \in \{0^\circ, 120^\circ, 240^\circ\}$. Any electric outlet or bulb connected to one of these power

lines can be considered as a *reference*, *i.e.*, having $\phi = 0^\circ$.

Events of interest, such as raw image acquisition, take place in *absolute time*, denoted $t^{\text{abs}} \in \mathbb{R}$ (Figure 2a). On the other hand, because flicker is periodic, it suffices to define flicker by an in-cycle *canonical time* $t^{\text{can}} \in [0, \Delta]$ (Figure 2c). We set the origin $t^{\text{can}} = 0$ to be a zero crossing of the AC voltage at a reference bulb. A bulb type is denoted by β . The bulb flicker is its response to AC. The response is expressed as a unit-less bulb response function (BRF), denoted $B_\beta(t^{\text{can}})$. For every color channel the BRF satisfies

$$\frac{1}{\Delta} \int_0^\Delta B_\beta(t^{\text{can}}) dt^{\text{can}} = 1. \quad (1)$$

For the rest of the paper, the time variable t^{can} ranges over the circular domain $[0, \Delta]$. If a bulb is connected to an outlet of phase ϕ , then the bulb’s flicker response is $B_\beta(t^{\text{can}} - \phi\Delta/2\pi)$. BRFs can be measured “in the field” (see Section 7) or obtained from a database such as DE-LIGHT [36].

An image has $R \times C$ pixels, where r indexes rows and c indexes columns. The scene is illuminated by S sources, indexed by s . Let B_s be the BRF of the bulb type to which source s belongs, and ϕ_s be the electric phase to which this source is connected. The intensity of an object illuminated by S sources at pixel (r, c) is

$$i(r, c, t^{\text{can}}) = \sum_{s=1}^S \tau_s(r, c) B_s(t^{\text{can}} - \phi_s\Delta/2\pi). \quad (2)$$

Here $\tau_s(r, c)$ is the response of image pixel (r, c) to illumination by s . It encompasses image radiance factors that depend on the distance of the object from the light source, surface slope, albedo, lens aperture, bidirectional reflectance, inter-reflections and exposure duration T_{exp} . That is, $\tau_s(r, c)$ is the light transport coefficient at this pixel. The transport coefficients in all pixels constitute a matrix τ_s .

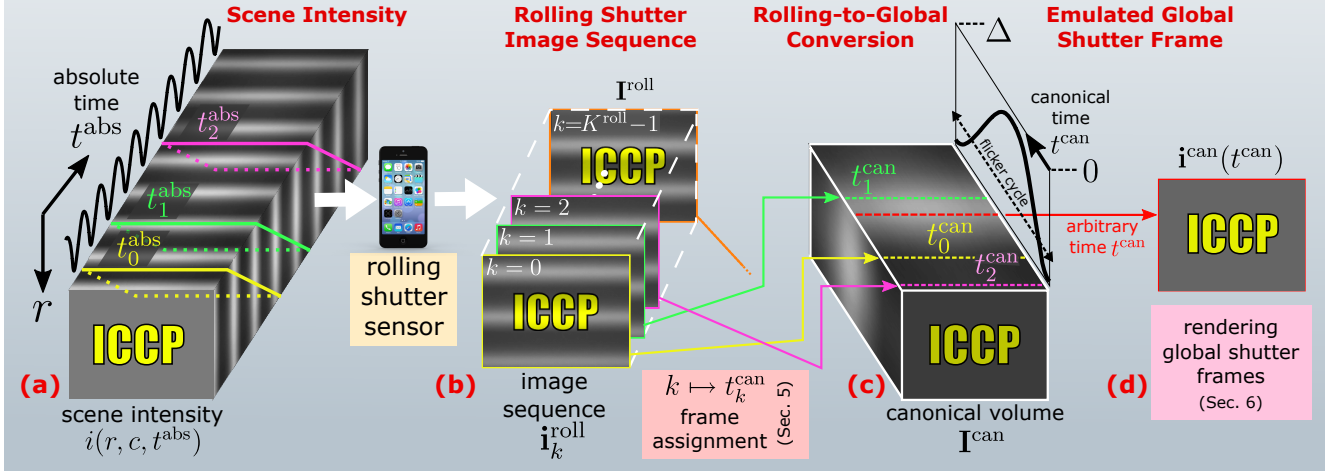


Figure 2. Rolling-shutter imaging under AC lighting. (a) A rolling-shutter camera captures a sequence of frames. In frame k , the topmost row is sampled at absolute time t_k^{abs} . Row r is sampled with constant delay D relative to row $r - 1$. Thus, the space-time volume $i(r, c, t^{\text{abs}})$ is sampled along diagonal slices. (b) The resulting rolling-shutter frames exhibit a wave pattern in each frame and across frames. (c) Each t_k^{abs} corresponds to an in-cycle canonical time $t_k^{\text{can}} \in [0, \Delta]$, where Δ is the flicker cycle. (d) Once we assign canonical times to all frames, we can compute a global-shutter frame of the scene at any arbitrary canonical time t^{can} .

2.2. Rolling-Shutter Imaging

In a rolling shutter, the exposure of each row is delayed by D seconds relative to the preceding row.¹ Let row 0 (topmost row) be acquired at absolute time t_k^{abs} , where $k = 0, 1, \dots$ is the discrete index of a captured rolling-shutter frame. Row r is thus acquired at time

$$t_{k,r}^{\text{abs}} = t_k^{\text{abs}} + rD. \quad (3)$$

Rows thus correspond to distinct temporal samples. Often D is much smaller than the exposure duration T_{exp} . This enables a high temporal sampling rate ($1/D$ Hz), while enough photons accumulate during T_{exp} for a high signal.

3. Rolling-Shutter Model for AC Flicker

We now model image formation of an AC-illuminated scene that is viewed by a rolling-shutter camera (Figure 2b). The topmost row of a rolling-shutter image k is acquired at t_k^{abs} . This t_k^{abs} corresponds to some in-cycle canonical time t_k^{can} . This means that there is a mapping

$$k \mapsto t_k^{\text{can}} \in [0, \Delta]. \quad (4)$$

In Section 5 we describe how this mapping can be derived from image data. For the moment, let this mapping be known. Following Eq. (3), row r of rolling-shutter frame k is acquired at canonical time

$$t_{k,r}^{\text{can}} = t_k^{\text{can}} + rD. \quad (5)$$

¹Some rolling-shutter sensors delay columns rather than rows. Without loss of generality we assume a sensor with inter-row delay.

Let $i_k^{\text{roll}}(r, c)$ be the intensity measured at pixel (r, c) of frame k . Combining Eqs. (2,5), this intensity is equal to

$$i_k^{\text{roll}}(r, c) = \sum_{s=1}^S \tau_s(r, c) B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi). \quad (6)$$

Eq. (6) corresponds to a spatial vertical wave pattern with a period of Δ/D rows (Figures 1,3). To see this mathematically, consider a region Ω for which $\tau_s(r, c)$ is vertically uniform, *i.e.*, independent of r . Then, Eq. (6) satisfies

$$i_k^{\text{roll}}(r, c) = i_k^{\text{roll}}(r + \Delta/D, c). \quad (7)$$

For a sequence of rolling-shutter frames, Eq. (6) expresses a spatiotemporal *wave function* that moves vertically with upward velocity

$$v = 1/D \quad (8)$$

pixels per second. To see this, consider two consecutive rolling-shutter images, k and $k + 1$, whose canonical times are $t_k^{\text{can}}, t_{k+1}^{\text{can}}$, respectively. In an image region Ω for which $\tau_s(r, c)$ is vertically uniform, Eq. (6) satisfies

$$i_k^{\text{roll}}(r, c) = i_{k+1}^{\text{roll}}(r - v(t_{k+1}^{\text{can}} - t_k^{\text{can}}), c). \quad (9)$$

This wave motion is independent of the scene's BRFs.

When the exposure time is Δ (*e.g.*, 0.01 sec in a 50Hz electric grid) or an integer multiple thereof, the image is unaffected by the wave pattern and emulates imaging under DC. Following Eqs. (1,6),

$$\begin{aligned} i^{\text{DC}}(r, c) &= \frac{1}{\Delta} \int_0^{\Delta} \sum_{s=1}^S \tau_s(r, c) B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi) dt^{\text{can}} \\ &= \sum_{s=1}^S \tau_s(r, c). \end{aligned} \quad (10)$$

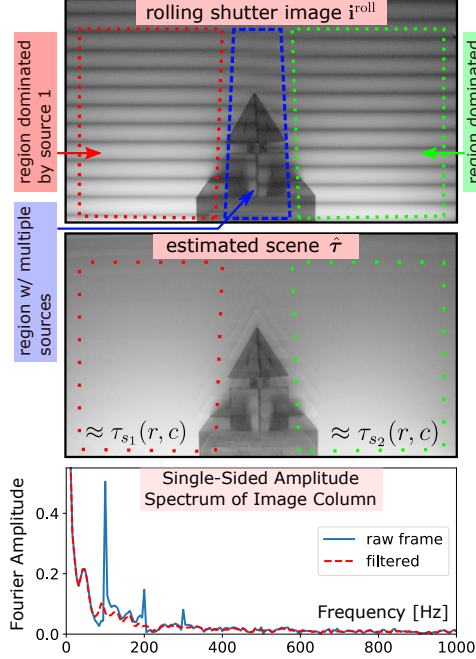


Figure 3. (Top) A rolling-shutter image of a scene illuminated from its sides by two identical bulbs powered by different AC phases. Flicker induces wave patterns. (Middle) Single-frame deflickering using a digital bandstop filter tuned to the flicker frequency and its higher harmonics. (Bottom) Spectrum of a single image column before and after deflicker filtering. In the x-axis, the image spatial frequency (in pixel units) is multiplied by $1/D$ to yield the corresponding temporal frequency.

4. Analysis Using One Image or an Image Pair

We show that from a pair or even a single rolling-shutter image, it is possible to extract useful information about the camera and scene illumination. This information includes (a) a flicker-independent (deflickered) image of the scene; (b) recognition of the types and electric phases of the scene’s light sources; (c) image regions that are dominated by any one source; and (d) the canonical time of the image, relative to the zero-crossing of the AC-grid voltage.

The analysis in this section mainly applies to image regions that are predominately illuminated by a single source. For example, Figure 3(top) has two distinct image regions outlined in red and green that are dominated by one of two sources. In regions where no source dominates (blue region in Figure 3(top)), the estimation degrades. We manually choose one region around a particular pixel (r_0, c_0) , dominated by a single source, to act as a reference whose source is connected to electric phase $\phi = 0$. Moreover, here, bulb type and AC phase analysis requires prior knowledge about BRFs (e.g., DELIGHT).

Section 4.1 discusses feasibility of information extraction using a single image, and thus relies on one more assumption: that τ_s is relatively uniform in pixel neighbor-

hoods dominated by a single source. Section 4.2 obviates this assumption, making the estimation insensitive to τ_s variations, by using a pair rather than a single image.

4.1. Single-Frame Analysis

4.1.1 Single-Frame Deflickering by Filtering

The flicker that affects i_k^{roll} has a known spatial frequency of D/Δ . We leverage this property to decompose i_k^{roll} into a flicker-dependent and flicker-independent component [32, 35]. This operation is motivated by common (non-flicker) intrinsic-image decomposition techniques [4,22,41,46].

Consider an image region illuminated by a single source s . Then, Eq. (6) becomes

$$i_k^{\text{roll}}(r, c) = \tau_s(r, c) B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi) \quad (11)$$

for all pixels in the region. This is a product of a flicker-independent component $\tau_s(r, c)$ and a flicker dependent component $B_s(t_k^{\text{can}}, r) \equiv B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi)$. From Eq. (11)

$$\ln[i_k^{\text{roll}}(r, c)] = \ln[\tau_s(r, c)] + \ln[B_s(t_k^{\text{can}}, r)]. \quad (12)$$

Due to the vertical periodicity of $B_s(t_k^{\text{can}}, r)$, all the signal energy associated with $\ln[B_s(t_k^{\text{can}}, r)]$ resides in specific spatial frequencies: $0, \pm D/\Delta, \pm 2D/\Delta, \dots$. Thus, M non-DC harmonics of $\ln[B_s(t_k^{\text{can}}, r)]$ can be eliminated using a linear filter $V(r)$ whose frequency response blocks frequencies $\pm D/\Delta, \pm 2D/\Delta, \dots, \pm MD/\Delta$. This vertical homomorphic spatially invariant filtering yields

$$\begin{aligned} \hat{\tau}(r, c) &\equiv \exp\left(\ln[i_k^{\text{roll}}] * V(r)\right) = \\ &\exp\left(\ln[\tau_s(r, c)] * V(r) + \ln[B_s(t_k^{\text{can}}, r)] * V(r)\right) = \\ &\exp \bar{B}_s \cdot \exp\left(\ln[\tau_s(r, c)] * V(r)\right) \sim \tau_s(r, c). \end{aligned} \quad (13)$$

Here \bar{B}_s is the temporal mean of $\ln[B_s]$. For any image region dominated by any single source s , the result $\hat{\tau}(r, c)$ approximates $\tau_s(r, c)$ there (Figure 3(middle)).

We used a cascade of digital filters, i.e., $V(r) = V_1(r) * V_2(r) * \dots * V_M(r)$. Each $V_m(r)$ blocks the m^{th} flicker harmonic by a 4th order bandstop Bessel filter, defined by a critical bandstop domain $[D(m/\Delta - \epsilon), D(m/\Delta + \epsilon)]$. In our case $1/\Delta = 100\text{Hz}$, $\epsilon = 8\text{Hz}$ and $M = 8$.

4.1.2 AC Lighting Retrieval from a Single Frame

Image regions mainly illuminated by a single source are modeled by Eq. (11). Different sources may illuminate different image regions. In regions dominated by each source s , we seek: (a) the bulb type $\beta \in \mathbb{B}$ per s , (b) the electric phase ϕ_s , and (c) the canonical time of the image t_k^{can} . The AC phase ϕ_s and t_k^{can} are linearly coupled in Eq. (11). Hence define an unknown

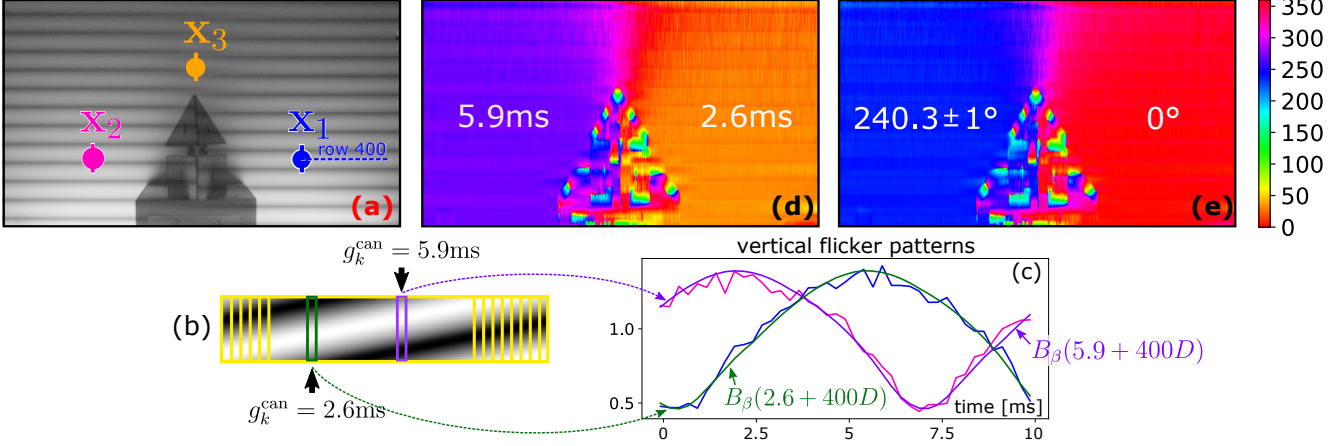


Figure 4. Single-frame analysis. (a) Each scene side is dominated by one of two fluorescent bulbs powered by electric phases 0° and 240° . Per pixel, we search for the AC phase that best fits the flicker wave pattern in a vertical domain centered at the pixel. (b) The bank of $W = 360$ samples of B_β , each corresponding to a different time g_k^{can} . (c) The image signal $i_k^{\text{filt}}[\cdot]/\mu_{r,c}$ for $\bar{\Omega}_{x_1}$ and $\bar{\Omega}_{x_2}$ is plotted in orange and purple, respectively. Superimposed are the time-shifted BRFs yielding the best match according to Eq. (17). (d) The recovered time map $g_k^{\text{can}}(r, c)$. (e) The estimated AC phase per pixel, $\hat{\phi}(r, c)$. Here we used the monochrome IDS camera with $T_{\text{exp}} = 1500\text{us}$.

$$g_k^{\text{can}} = t_k^{\text{can}} - \phi\Delta/2\pi. \quad (14)$$

Suppose we have an approximation of the transport coefficients $\hat{\tau}_s$. Such an approximation can be obtained by the filtering operation of Eq. (13). Alternatively, if the flicker spatial period is small, $\hat{\tau}_s$ may be approximated by a low-pass filter of i_k^{roll} . Then, from Eq. (11),

$$i_k^{\text{filt}}(r, c) \equiv \frac{i_k^{\text{roll}}(r, c)}{\hat{\tau}_s(r, c)} \approx B_s(t_k^{\text{can}} + rD - \phi\Delta/2\pi). \quad (15)$$

Define a vertical image line $\bar{\Omega}_{r,c} \in \Omega$ with length of a single flicker cycle in pixels centered at pixel (r, c) :

$$\bar{\Omega}_{r,c} = \{(r-l, c), \dots, (r, c), \dots, (r+l, c)\}, \quad (16)$$

where $l = \lceil \Delta/(2D) \rceil$. The mean intensity of i_k^{filt} on $\bar{\Omega}_{r,c}$ is $\mu_{r,c}$. Let²

$$F_{r,c}(g_k^{\text{can}}, \beta) = \sum_{(r',c') \in \bar{\Omega}_{r,c}} \left| \frac{i_k^{\text{filt}}(r',c')}{\mu_{r,c}} - B_\beta(g_k^{\text{can}} + r'D) \right|^2. \quad (17)$$

Now, we recognize the bulb type β that dominates pixel (r, c) as well as the value of g_k^{can} there using

$$\{\hat{g}_k^{\text{can}}(r, c), \hat{\beta}(r, c)\} = \arg \min_{g_k^{\text{can}} \in [0, \Delta], \beta \in \mathbb{B}} F_{r,c}(g_k^{\text{can}}, \beta). \quad (18)$$

Once $\hat{g}_k^{\text{can}}(r, c)$ has been estimated for all pixels, we decompose it into an associated per-pixel electric-grid phase $\phi(r, c)$ and frame canonical time t_k^{can} . Let pixel (r_0, c_0) be

²In Eq. (17) normalizes i_k^{filt} by $\mu_{r,c}$, so that the subsequent analysis is insensitive to any potential residual scale factor.

strongly dominated by illumination from the reference AC outlet. Hence, set $\hat{\phi}(r_0, c_0) = 0^\circ$. Thus, following Eq. (14),

$$\hat{t}_k^{\text{can}} = \hat{g}_k^{\text{can}}(r_0, c_0). \quad (19)$$

Based on Eqs. (14,19),

$$\hat{\phi}(r, c) = 2\pi[\hat{t}_k^{\text{can}} - \hat{g}_k^{\text{can}}(r, c)]/\Delta. \quad (20)$$

It is expected that in areas that are strongly dominated by a single source, $\hat{\phi}(r, c)$ should be either 0° , 120° , or 240° .

Demonstration on a real scene. In Figure 4a, pixels x_1, x_2 are in image regions dominated exclusively by either of the two bulbs. Vertical image lines $\bar{\Omega}_{x_1}, \bar{\Omega}_{x_2}$ are defined around each pixel according to Eq. (16). The corresponding values $i_k^{\text{filt}}[\bar{\Omega}_{x_1}]/\mu_{x_1}, i_k^{\text{filt}}[\bar{\Omega}_{x_2}]/\mu_{x_2}$ are plotted in Figure 4c, where row indices are converted to canonical time using:

$$t_{k,r'}^{\text{can}} = (r' - r + l)D, \quad (r', c) \in \bar{\Omega}_{r,c}. \quad (21)$$

We implement Eq. (18) by searching for the optimal g_k^{can} on a discrete grid having W samples, i.e. $g_k^{\text{can}} = 0, \Delta/W, 2\Delta/W, \dots$. The estimation precision is thus Δ/W . In the example of Figure 4, we manually set $W = 360$. For x_1 and x_2 , the optimal respective values of g_k^{can} are 2.6ms and 5.9ms. Because x_1 is chosen to correspond to the reference electric outlet, Eq. (19) sets $\hat{t}_k^{\text{can}} = 2.6\text{ms}$. Figure 4e maps the result to electric grid phases using Eq. (20), $\forall(r, c)$. Accordingly, the left-hand side of the image is dominated by illumination which is consistent with electric phase $\phi = 240^\circ$.

4.2. Two Frames: Transport-Invariant Analysis

Section 4.1 assumed an approximately constant τ_s per region Ω . We compensated for deviation from this assump-

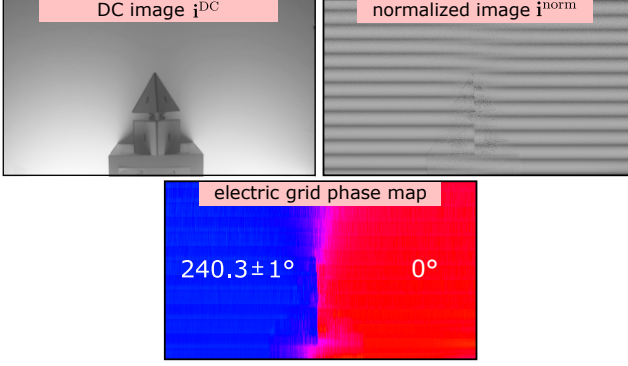


Figure 5. AC analysis using two frames. Normalizing a rolling-shutter frame (Fig 4a) by a DC frame [top-left] yields i^{norm} . This improves bulb information recovery. Compare the estimated electric phase map to the one in Figure 4e.

tion by roughly estimating τ_s from a single image. However, we can retrieve τ_s directly by using a *second* image. Specifically, in addition to i_k^{roll} , let us capture i_k^{DC} using Eq. (10). In image regions dominated by s ,

$$i^{\text{DC}}(r, c) = \tau_s(r, c) + \sum_{s' \neq s} \tau_{s'}(r, c) \approx \tau_s(r, c). \quad (22)$$

Then, from Eqs. (11,22) derive

$$i_k^{\text{norm}}(r, c) = \frac{i_k^{\text{roll}}(r, c)}{i^{\text{DC}}(r, c)} = B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi). \quad (23)$$

Thus, in a region dominated by a single source, the image i_k^{norm} is invariant to the scene's light transport τ_s , expressing $B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi)$, as desired. Consequently, we use i_k^{norm} instead of i_k^{filt} in Eq. (17), and then apply Eqs. (18-20). Figure 5 shows the result of this process.

4.3. Optical Star-Filter for Emitter Analysis

Consider scene regions having strong emitters. This is essentially the situation in wide-field cityscapes (Figure 1). It is generally encountered in direct measurements and analysis of AC bulbs. In images of such scenes, there is no region Ω with uniform τ_s , as required in Sections 4.1 and 4.2. Hence, we *create* Ω *optically, simultaneously around all sources*. The simplest way to induce Ω is to defocus the camera. Then, each light source creates a two-dimensional circular defocus point spread function (PSF), which provides support for Ω , (see Figure 6). However, defocus blur limits spatial resolution.

Ideally, we need an optical PSF that is vertical and narrow having a weak central point so as to minimize interference with adjacent sources. These desiderata are practically addressed by mounting an *optical star-filter* on the camera lens. Figure 1 shows a cityscape scene viewed by a rolling-shutter camera through a *Star-4* filter aligned with the sensor's axes. Most of the light sources are bright and spatially

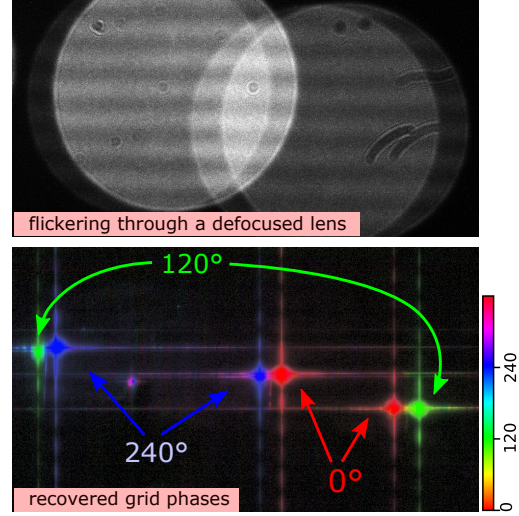


Figure 6. (Top) Imaging the same scene as in Figure 1 (bottom) using a severely defocused lens and a rolling-shutter camera. The PSF reveals the flicker-induced wave pattern but degrades the spatial image content. (Bottom) A star filter enables bulb and phase recognition while also preserving most image details.

resolved, despite minor overlap caused by the star PSF. The vertical arms of the PSF play the role of Ω , modulated by $B_s(t_k^{\text{can}} + rD - \phi_s \Delta / 2\pi)$ for each source s . We acquired two star-filtered images: i_k^{roll} (Figure 1) and i^{DC} (Eq. 10). We then followed the process of Section 4.2 to recognize the bulbs and recover their electric phases (Figure 6 (bottom)).

5. Canonical Volume from an Image Sequence

We now deal with a long sequence of K^{roll} raw rolling-shutter frames. They express a data volume $\mathbf{I}^{\text{roll}} = [i_0^{\text{roll}}, i_1^{\text{roll}}, \dots, i_{K^{\text{roll}}-1}^{\text{roll}}]$. We seek a canonical volume \mathbf{I}^{can} , whose spatial domain is discrete (pixels), and whose time domain is continuous. The canonical volume expresses the spatiotemporal flicker signal. To achieve this, each input index k should be mapped to an in-cycle canonical time t_k^{can} according to Eq. (4). Consequently, the output intensity $i^{\text{can}}(r, c, t_k^{\text{can}})$ at each pixel (r, c) can be thought of as the output of a very fast light detector that captures K^{roll} consecutive samples in $[0, \Delta]$ (Figure 2). Once the assignment in Eq. (4) is computed, we can re-sample and discretize the canonical time domain t^{can} to emulate a global-shutter sequence using any method of choice.

Assigning according to Eq. (4) faces three challenges:

(a) The raw temporal order of the rolling-shutter frames is often meaningless. Consider the frames captured in absolute times $t_1^{\text{abs}} < t_2^{\text{abs}} < t_3^{\text{abs}} \dots$. This ordering is generally not maintained in the corresponding canonical times t^{can} , *i.e.*, we may have $t_2^{\text{can}} < t_{32}^{\text{can}} < t_3^{\text{can}} < t_1^{\text{can}} \dots$. This is largely because our rolling-shutter camera's clock is unsynchronized and untethered to the AC. This problem is exac-




Figure 7. Histogram showing $(t_{k+1}^{\text{abs}} - t_k^{\text{abs}}) - t^{\text{FR}}$, where t^{FR} corresponds to the declared setting of the camera. The histogram is based on 15000 frames of an IDS UI-348xLE camera.

erbed when taking a sequence of still images at random times, *e.g.*, by a cellphone.

(b) Even when the frames are well sorted, the time interval between them is not a fixed fraction of the flicker cycle, *i.e.*, $|t_{k+1}^{\text{abs}} - t_k^{\text{abs}}| \neq |t_{k+1}^{\text{can}} - t_k^{\text{can}}|$. This is mainly due to jitter of the AC grid [36] and is exacerbated by jitter of the camera electronics (Figure 7).

(c) Temporal shift. Suppose the entire canonical time axis is cyclically shifted globally by δ . This shift does not affect the temporal order of frames or the interval between them, and is inconsequential to many applications. However, such a shift does affect the relation between canonical time and the AC zero crossing and must be eliminated to accurately relate image measurements to the grid’s behavior. We solve these challenges below.

5.1. Canonical Time up to a Global Time Shift

Here we describe how to compute the canonical time of each frame up to an unknown shift both for scenes having a significant large diffuse surface, and for emitter scenes. Because in this section we do not solve for a global shift, we determine all canonical times relative to the first frame, whose canonical time is denoted t_0^{can} .

5.1.1 Canonical Times from Diffuse Surfaces

Suppose there is a large *diffuse* surface in the scene, such as a wall, reflecting light towards the camera. Recall that in this case, the rolling-shutter sequence creates a spatiotemporal wave whose vertical motion and velocity are independent of BRFs (Section 3). Consequently, the wave’s motion enables estimation of canonical time differences, *e.g.*, $t_k^{\text{can}} - t_0^{\text{can}}$, *without prior knowledge* of bulb BRFs.

Similarly to Section 4.1, we consider a vertical image region Ω , where the corresponding diffuse surface is mainly illuminated by a single source s . We choose Ω manually in an image region that clearly exhibits a flicker wave pattern. The imaging model is thus represented by Eq. (11). As described in Section 4.2, shift analysis is simplified if it relies on images $\mathbf{i}_k^{\text{norm}}$ having vertical flicker that is invariant to spatial variations of τ_s in the scene.

Irrespective of the BRFs in the scene, the wave pattern

in $\mathbf{i}_k^{\text{norm}}$ has a spatial period of Δ/D pixels, *i.e.*, fundamental angular spatial frequency $2\pi D/\Delta$. We can then rely on simple Fourier analysis. The fundamental mode of the pattern has phase

$$\theta_k = \text{angle} \left\{ \sum_{(r',c') \in \bar{\Omega}} \mathbf{i}_k^{\text{norm}}(r',c') \exp(j2\pi r' D/\Delta) \right\}, \quad (24)$$

not to be confused with electric grid phase. Here angle is the argument of a complex number and $\bar{\Omega}$ is an image line in Ω . Comparing to the pattern phase θ_0 of $k = 0$, the canonical time assignment of frame k satisfies

$$t_k^{\text{can}} - t_0^{\text{can}} = \frac{\theta_k - \theta_0}{2\pi} \Delta. \quad (25)$$

Recall that for $\mathbf{i}_k^{\text{norm}}$ we require \mathbf{i}^{DC} . However, there is no need to capture \mathbf{i}^{DC} if we compute assignment of for an entire batch of frames, whose sampling times t_k^{can} span the interval $[0, \Delta]$ uniformly. We can approximate it using

$$\mathbf{i}^{\text{DC}}(r,c) \approx \frac{1}{K^{\text{roll}}} \sum_k \mathbf{i}_k^{\text{roll}}(r,c) \approx \tau_s(r,c). \quad (26)$$

See Figure 8 for an example result.

5.1.2 Canonical Times Without Surface Reflection

In outdoor nocturnal scenes, sometimes no reflecting object is reliably sensed in the short exposures needed, due to low photon counts away from emitting bulbs. Thus only bulbs are effectively observed. Contrary to the situation in Section 5.1.1, here temporal assignment cannot rely on a wave pattern that is spatially observed on a surface. We now show that canonical time estimation is still possible.

The method described here applies to general scenes. There is no need for a star filter. Moreover, since this method does not make use of the rolling-shutter model, it can apply to global shutter sensors as well. However, unlike Section 5.1.1 where assignment is done for each frame individually, here, we rely on the entire batch of captured frames to assign canonical time to each frame. Moreover, we require that our image batch’s random samples span the interval $[0, \Delta]$ uniformly.

We make the following observation. In any scene, if the canonical times of two rolling-shutter frames are very close, then the pixel intensities in these frames are very close as well, *i.e.*, if $|t_k^{\text{can}} - t_{k'}^{\text{can}}| \ll \Delta$, then $\|\mathbf{i}_k^{\text{roll}} - \mathbf{i}_{k'}^{\text{roll}}\|$ is small. This pair-wise *affinity* between closely-timed images leads to an overall mutual ordering between all rolling-shutter frames, *i.e.*, expressing how far t_k^{can} and $t_{k'}^{\text{can}}$ are, even for dissimilar frames in the set.

Any frame $\mathbf{i}_k^{\text{roll}}$ is a point in a vector space of dimension $R \times C$. For a specific scene, the set of all K^{roll} rolling-shutter

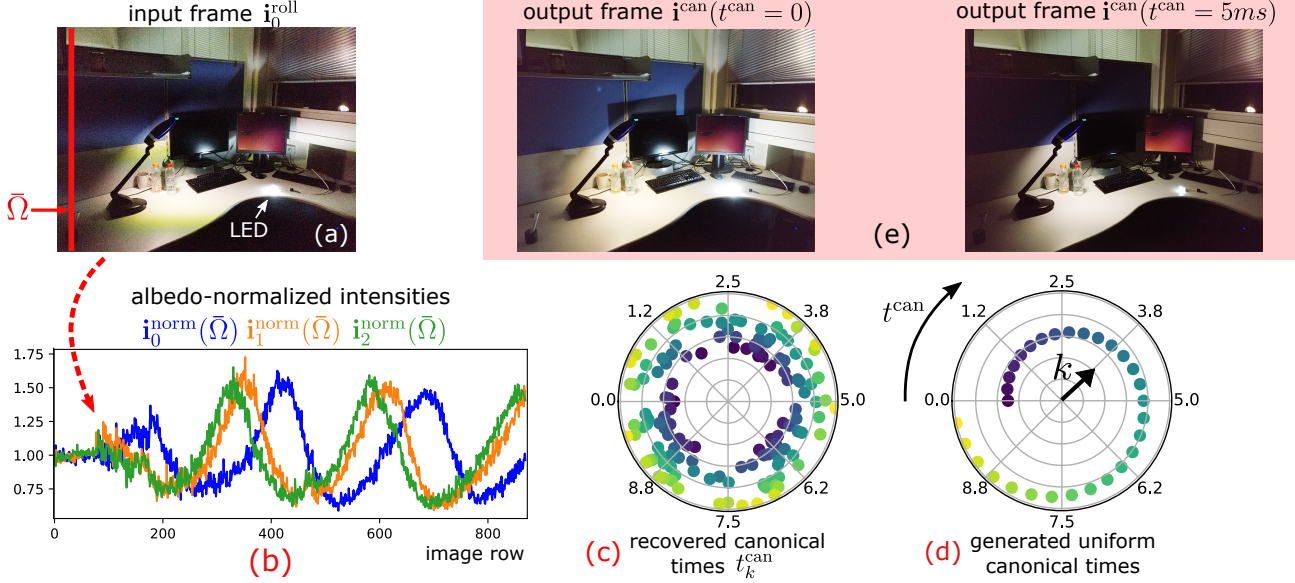


Figure 8. An experiment using a OnePlus cellphone with exposure set to 1/500sec. We manually triggered 138 rolling-shutter frames. (a) The first frame. (b) The vertical signal along $\bar{\Omega}$. (c) The recovered canonical times. In the polar plot, the radii and color denote k while the angle represents $t_k^{\text{can}} \in [0, 10\text{ms}]$. Note that t_k^{can} is distributed randomly. We render 30 uniformly spread global shutter frames spanning the flicker cycle. (d) The times of the rendered frames. (e) Two rendered frames at times 0 and 5ms. See [45] for video results.

frames $\{\mathbf{i}_k^{\text{roll}}\}$ resides on a low-dimensional manifold in this space. Moreover, when $|t_k^{\text{can}} - t_{k'}^{\text{can}}| \ll \Delta$, it follows that $\mathbf{i}_k^{\text{roll}}$ and $\mathbf{i}_{k'}^{\text{roll}}$ lie close to each other on this manifold [21]. Furthermore, the underlying phenomena that drive changes in the images are temporally periodic. For this reason, the manifold has a *ring topology*.

Hence, location on the ring encodes the canonical time t_k^{can} of any rolling-shutter frame. Thus, to estimate t_k^{can} relative to t_0^{can} , we need to measure how far on the ring $\mathbf{i}_k^{\text{roll}}$ is relative to $\mathbf{i}_0^{\text{roll}}$. A ring can be expressed as a curve in a two dimensional space. Dimensionality reduction from $R \times C$ into two dimensions, the ring structure and the embedding of $\mathbf{i}_k^{\text{roll}}$ on the ring emerge from a *diffusion map* [40]. Let us denote the dimensionality reduction of the frames from $R \times C$ to two dimensions by an operator

$$\mathbf{A} = \mathcal{DM}_{RC \rightarrow 2}\{\mathbf{I}^{\text{roll}}\}. \quad (27)$$

The diffusion maps algorithm requires specifying the pairwise affinity of frames. We used a Gaussian kernel whose width was manually adjusted per scene to yield a ring. As seen in Figure 9, the resulting ring is generally skewed. We remove this skew using singular value decomposition (SVD). First, define the ring's center of mass

$$[a_0 \ a_1] = \mathbf{1}_{K^{\text{roll}}} \mathbf{A}^\top / K^{\text{roll}}. \quad (28)$$

where $\mathbf{1}_{K^{\text{roll}}}$ is a row vector of length K^{roll} , all of whose elements are 1's. Here \top denotes transposition. Then, the ring is centered by

$$\tilde{\mathbf{A}} = \mathbf{A} - \begin{bmatrix} a_0 \mathbf{1}_{K^{\text{roll}}} \\ a_1 \mathbf{1}_{K^{\text{roll}}} \end{bmatrix}. \quad (29)$$

Using SVD, decompose $\tilde{\mathbf{A}}$ to

$$\tilde{\mathbf{A}} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}, \quad (30)$$

where $\mathbf{\Sigma}$ is a 2×2 diagonal positive matrix. Now, any rolling-shutter frame k is represented by column k of \mathbf{V} , where $\mathbf{V} \in \mathbb{R}^{2 \times K^{\text{roll}}}$. As seen in Figure 9, now the ring manifold is not skewed. Let column k of \mathbf{V} be $[v_0(k) \ v_1(k)]^\top$. Then, using the argument of a complex number in lieu of a temporal phase-angle, set

$$\theta_k = \text{angle}\{v_0(k) + jv_1(k)\}. \quad (31)$$

We use Eq. (25) to convert temporal phase-angle θ_k to canonical time.

5.2. Temporal Alignment to the AC

So far we assumed that the relation between t_0^{can} and the AC zero crossings is unknown. To recover this relation, we rely on two priors: a database \mathbb{B} of BRFs [36] and knowledge of the bulb type that affects one pixel (r_0, c_0) in the image. We treat that bulb as a reference whose AC phase is ϕ is zero. Then we find the temporal shift δ that best aligns the BRF B_β to the time-varying normalized intensity at (r_0, c_0) :

$$\hat{\delta} = \underset{\delta}{\text{argmin}} \sum_k |i_k^{\text{norm}}(r_0, c_0) - B_\beta(t_k^{\text{can}} - t_0^{\text{can}} + \delta)|^2 \quad (32)$$

Once $\hat{\delta}$ is estimated, we set $t_0^{\text{can}} = \hat{\delta}$ and thus determine t_k^{can} for all frames relative to the AC zero crossing.

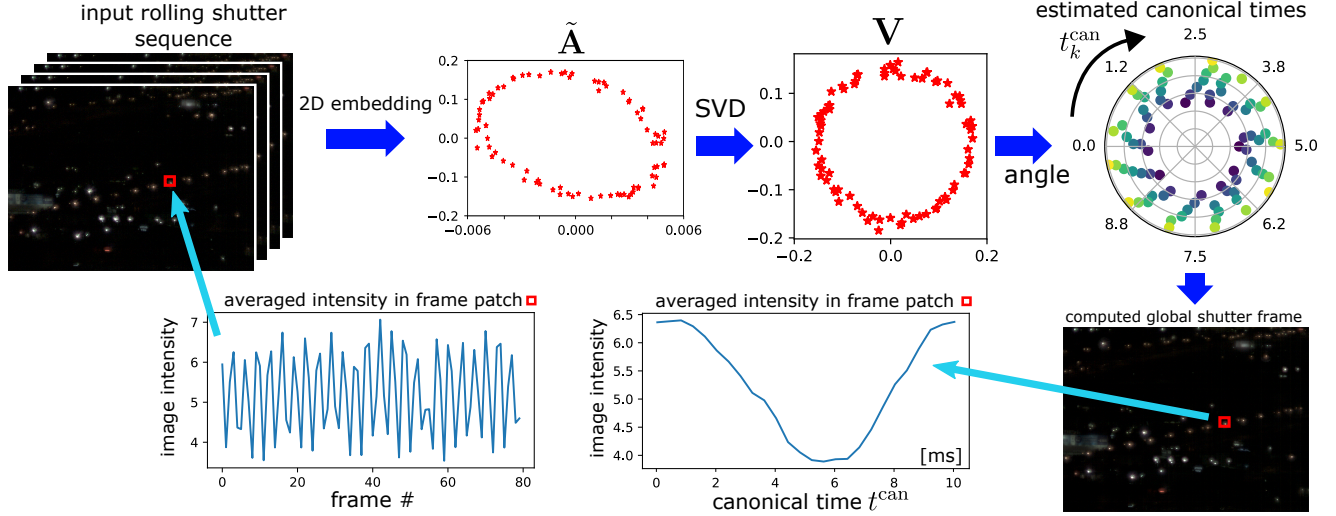


Figure 9. A machine vision color camera (IDS UI-155xLE) with exposure time set to 1788us captured 80 frames of the Haifa bay area from afar. By computing a 2D embedding of the frames we obtain their canonical times. See [45] for video comparison.

6. Computing Global-Shutter Frames

Section 5 described temporal assignment of rolling-shutter frames to canonical times $\{t_0^{\text{can}}, \dots, t_{K^{\text{roll}}-1}^{\text{can}}\}$. Following this assignment we can now compute canonical images, as if taken by a global shutter at an arbitrary t^{can} . The canonical volume \mathbf{I}^{can} represents a signal which is temporally periodic. Hence for each pixel (r, c) , the canonical intensity function $i^{\text{can}}(r, c, t^{\text{can}})$ is represented by a *Fourier series*: the fundamental temporal frequency is $1/\Delta$, and there are higher harmonics. We use this representation to compute \mathbf{I}^{can} in two steps:

- Use the K^{roll} raw rolling-shutter frames to estimate the Fourier series coefficients of the 1D temporal signal observed at each pixel (r, c) .
- Use the per-pixel Fourier coefficients to compute the canonical volume \mathbf{I}^{can} and/or any of its slices for any t^{can} .

Here are the mathematical details. Per pixel, the real-valued Fourier series of the canonical signal is

$$i^{\text{can}}(r, c, t^{\text{can}}) = p_0(r, c) + \sum_{m=0}^M \left[p_m(r, c) \cos \frac{2\pi m t^{\text{can}}}{\Delta} + q_m(r, c) \sin \frac{2\pi m t^{\text{can}}}{\Delta} \right]. \quad (33)$$

There are $2M + 1$ real-valued Fourier coefficients p_m, q_m for each pixel (r, c) . To make the estimation well-posed, $2M+1 \leq K^{\text{roll}}$. Fortunately, BRFs are usually smooth. So, intensity variations due to flicker can be represented with just a few Fourier coefficients.

Concatenate these coefficients to form a vector

$$\mathbf{p}(r, c) = [p_0(r, c), p_1(r, c), \dots, q_1(r, c), \dots, q_M(r, c)]^{\top}. \quad (34)$$

Define a row vector

$$\mathbf{h}(t) = \left[1, \cos \frac{2\pi t^{\text{can}}}{\Delta}, \dots, \cos \frac{2\pi M t^{\text{can}}}{\Delta}, \sin \frac{2\pi t^{\text{can}}}{\Delta}, \dots, \sin \frac{2\pi M t^{\text{can}}}{\Delta} \right]. \quad (35)$$

Then, rendering of a global-shutter image for any time t^{can} is enabled by expressing Eq. (33) as

$$i^{\text{can}}(r, c, t^{\text{can}}) = \mathbf{h}(t^{\text{can}}) \mathbf{p}(r, c). \quad (36)$$

Let us estimate $\mathbf{p}(r, c)$ from a sequence of rolling shutter frames. The canonical time assigned to row r of frame k is given by Eq. (5). Intensities in frame k are given by

$$i_k^{\text{roll}}(r, c) = i^{\text{can}}(r, c, t_{k,r}^{\text{can}}) = \mathbf{h}(t_{k,r}^{\text{can}}) \mathbf{p}(r, c). \quad (37)$$

Define a matrix

$$\mathbf{H}_r = \left[\mathbf{h}(t_{0,r}^{\text{can}})^{\top}, \dots, \mathbf{h}(t_{k,r}^{\text{can}})^{\top}, \dots, \mathbf{h}(t_{K^{\text{roll}}-1,r}^{\text{can}})^{\top} \right]^{\top}. \quad (38)$$

Let column vector $\mathbf{d}_{r,c}$ of length K^{roll} denote the intensities at pixel (r, c) , *i.e.*, $\mathbf{d}_{r,c}[k] = i_k^{\text{roll}}(r, c)$. Then,

$$\mathbf{d}_{r,c} = \mathbf{H}_r \mathbf{p}(r, c). \quad (39)$$

Computing $\mathbf{p}(r, c)$ is essentially a least-squares Fourier series estimation:

$$\hat{\mathbf{p}}(r, c) = (\mathbf{H}_r^{\top} \mathbf{H}_r)^{-1} \mathbf{H}_r^{\top} \mathbf{d}_{r,c}. \quad (40)$$

Live Rolling-to-Global Shutter Conversion

When the field of view includes a diffuse surface, we can estimate global-shutter images in real-time. This yields




Figure 10. Light transport decomposition. (Left) A scene is illuminated by four different bulbs. In a standard long exposure, illumination from all sources is mixed together. (Top right) We capture 120 rolling-shutter frames (not shown) and estimate a 25-frame global-shutter sequence of the scene. We then apply source unmixing as described in [36]. BRFs used here for unmixing were extracted as described in Section 4.2. (Bottom right) Ground-truth single-source images were captured by turning on each source individually.

a live-streaming algorithm that can continuously compute global-shutter frames. To achieve this, we develop the following *update* process:

1. In a ramp-up stage, K^{roll} increases with every new input frame, until reaching a pre-set maximum value $K^{\text{roll}} = K_{\text{max}}^{\text{roll}}$. After this ramp-up, for each new frame streaming in, the oldest frame in the queue is deleted.
2. Use the K^{roll} frames in the queue to estimate the Fourier series coefficients of each pixel.
3. When a new rolling-shutter frame is acquired, compute its canonical time t_k^{can} , using the method of Section 5.1.1. Then, *update* the Fourier series coefficients for all pixels.
4. Using the Fourier series, the canonical volume \mathbf{I}^{can} can be rendered for any t^{can} .

Temporal smoothness of BRFs allows using a small number of Fourier coefficients. Hence even during ramp-up, when K^{roll} is small, we obtain a good estimate of \mathbf{I}^{can} .

7. More Experimental Details

We used two rolling-shutter machine-vision cameras: IDS UI-155xLE (color) and IDS UI-348xLE (monochrome). In addition, sequences were captured using a OnePlus 3 smartphone mounted on a tripod and triggered manually, as well as an iPhone 7. The exposure time was set per scene depending on its brightness. We used a HOYA 52mm cross-screen filter to capture emitter scenes. For our smartphone cameras, we estimated D using a still image of a smooth surface. Then, we applied a spatial discrete Fourier transform on a long vertical line $\bar{\Omega}$, and detected the spatial frequency having the strongest signal energy.

The emulated global-shutter images resulting from Section 6 enable source separation as in [36]. An example is shown in Figure 10. Here, the reference pixel for estimating the shift δ relative to the AC zero crossing was in a region dominated by an LED (listed as LED2 in the DELIGHT [36]). As seen in the figure, the unmixing is consis-




Figure 11. We placed a white printer paper over an iPhone 7 camera lens and pointed the camera at a bulb. The paper diffused the bulb’s illumination, destroying all scene-dependent spatial image variations and making τ_s uniform. The resulting image captured the bulb response function *curve* at high temporal resolution, for three color bands. To reduce noise, we averaged across columns.

tent with the ground truth.

We acquired BRFs in separate, dedicated images in which the camera’s inter-row delay D is minimal. We did this in two ways: The first way was to place a paper diffuser in front of a smartphone lens as in Figure 11. This allowed us to measure BRFs up to an unknown shift δ . The second, which allows measurement of the exact BRF, is to image a white wall without a diffuser, similar to the setup in Figure 4. This wall was illuminated from two sides by two bulbs (a similar setup is seen in Figure 4a). One bulb was the one for which we needed to measure the BRF. The other bulb was a known reference (LED2 of DELIGHT) which provided δ . More can be seen in [45], including videos and a demonstration of Section 6.

8. Comparison to the ACam [36]

The electronic control of the ACam of [36] is replaced in our current work by computational post processing that assigns t^{can} to each frame and row. This enables a rolling-shutter camera to operate asynchronously when capturing an image sequence. In contrast, data acquisition in the ACam is intermittently suspended during operation by a

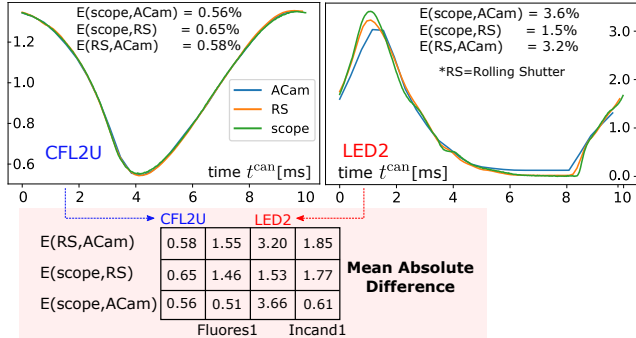


Figure 12. BRF curves measured by three methods: ACam, rolling shutter (RS) using the method of Section 4.2, and a photodiode. We compute the Mean Absolute Difference, normalized by BRF amplitude, between each pair of methods, per bulb. The curves are consistent. Because ACam [36] yields only 26 BRF samples, its measurement of the LED2 BRF is somewhat inferior.

controller enslaved to AC zero crossings. Moreover, as typically the inter-row delay is much smaller than a flicker period, a rolling shutter offers very dense temporal sampling of flicker. The ACam of [36] is essentially a global-shutter system, which does not exploit the spatial domain for temporal sampling. These characteristics offer a speed advantage to a rolling shutter system.

Unlike our rolling-shutter system, the ACam has relay optics, which limit the system’s overall numerical aperture and complicate the system’s robustness. Consequently, the ACam needed significantly longer exposures than our rolling-shutter system. Experiment times in [36] ranged between 26-400sec, depending on scene illumination strength, while yielding 26 samples of a flicker cycle. In comparison, the experiment of Figure 9, which shows our dimmest scene, needed less than 7sec to yield 80 samples of a flicker cycle. Moreover, in Section 4, one or two frames lasting at most 0.2secs provided similar information about AC phases (in uncluttered areas), for which the ACam of [36] required a full set of images. The quality of our rolling shutter measurements is shown in Figure 12. It shows that the quality of the systems is comparable, while rolling shutter offers an advantage when high temporal sampling is needed.

The ACam still has advantages, though. Being a coded aperture system, ACam can integrate optically an arbitrarily large number of flicker cycles. This enables superior dynamic range and sensitivity to low-light conditions, with which rolling-shutter readouts are currently incompatible.

9. Discussion

Our approach is much simpler than a tethered coded-aperture system [36]. It enables practical use of flicker for computer vision methods that rely on illumination and reflection modulation [20,23], *e.g.*, photometric stereo, shape from shadows and spacetime stereo. Bulb recognition can




Figure 13. In the scene of Figure 4a, pixel x_3 resides in a region affected strongly by two bulbs. Therefore, the vertical flicker signal around x_3 (in blue) is a mixture of their BRFs. Here we show that this is approximately an equal-weight mixture (red curve), thereby demonstrating crude single-image unmixing.

assist in color balancing. Flicker may assist smartphones in separating window reflections [36]. Scene separation into single-source images can be done using a sequence of rolling-shutter images in real-time (Figure 10). We hypothesize that a single rolling-shutter image can yield coarse image source unmixing (see motivation in Figure 13). Some present and future LED bulbs may deviate from the model of Section 3, thus will require new imaging models.

Acknowledgments: We thank Yoash Levron for useful discussion, Vadim Holodovsky and Michael Fisher for help with experiments. YYS is a Landau Fellow—supported by the Taub Foundation. This research was supported by the Israel Science Foundation (Grant 542/16) and conducted in the Ollendorff Minerva Center. Minerva is funded through the BMBF. KNK gratefully acknowledges the support of the Natural Sciences and Engineering Research Council of Canada under the RGPIN and SGP programs, as well as the support of DARPA under the REVEAL program.

References

- [1] S. Achar, J. R. Bartels, W. L. Whittaker, K. N. Kutulakos and S. G. Narasimhan. Epipolar time-of-flight imaging. In *ACM TOG*, 36(4), 2017.
- [2] N. Andersson, M. Sandström, A. Berglund, and K. Hansson. Amplitude modulation of light from various sources. *Lighting Research and Technology*, 26(3):157–160, Sept. 1994.
- [3] F. B. Bianco, S. E. Konnin, C. Mydlarz and M. S. Sharma. Hypertemporal imaging of NYC grid dynamics. In *Proc. ACM Intern. Conf. on Sys. for Energy-Efficient Built Env.*, 2016.
- [4] B.P. Bogert, M.J.R. Healy, and J.W. Tukey. The quefrency analysis of time series for echoes: Cepstrum, pseudo-autocovariance, cross-cepstrum, and saphe cracking, In *Time Series Analysis*, M. Rosenblatt, Ed., (15)209243, 1963.
- [5] Y. Dai, H. Li and L. Kneip. Rolling shutter camera relative pose: generalized epipolar geometry. In *Proc. IEEE CVPR*, 2016
- [6] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *ACM SIGGRAPH*, 145–156, 2000.
- [7] J. Gu, Y. Hitomi, T. Mitsunaga and S. Nayar. Coded rolling shutter photography: Flexible space-time sampling. In *Proc. IEEE ICCP*, 2010.

- [8] Y. Hitomi, J. Gu, M. Gupta, T. Mitsunaga, and S. K. Nayar. Video from a single coded exposure photograph using a learned over-complete dictionary. In *Proc. IEEE ICCV*, 287–294, 2011.
- [9] Z. Hui, A. C. Sankaranarayanan, K. Sunkavalli and S. Hadap. White balance under mixed illumination using flash photography. In *Proc. IEEE ICCP*, 2016.
- [10] J. Holloway, A. C. Sankaranarayanan, A. Veeraraghavan, and S. Tambe. Flutter shutter video camera for compressive sensing of videos. In *Proc. IEEE ICCP*, 2012.
- [11] E. Hsu, T. Mertens, S. Paris, S. Avidan, and F. Durand. Light mixture estimation for spatially varying white balance. In *ACM SIGGRAPH*, 2008.
- [12] K. Jo, M. Gupta, and S. K. Nayar. DisCo: Display-camera communication using rolling shutter sensors. *ACM TOG*, 35(5), article 150, 2016.
- [13] H. R. V. Joze and M. S. Drew. Exemplar-based color constancy and multiple illumination. *IEEE TPAMI*, 36(5):860–873, 2014.
- [14] H. Y. Lee, H. M. Lin, Y. L. Wei, H. I. Wu, H. M. Tsai and K. C. J. Lin. Rollinglight: Enabling line-of-sight light-to-camera communications. In *Proc. ACM 13th Annual Inter. Conf. on Mobile Systems, Appl., and Serv.*, 167-180, 2015.
- [15] A. Levis, Y. Y. Schechner, and R. Talmon. Statistical tomography of microscopic life. In *Proc. IEEE CVPR*, 2018.
- [16] J. P. Lewis. Fast normalized cross-correlation. In *Vision Interface*, 10(1), 120-123, 1995.
- [17] B. Li, W. Xiong, W. Hu, and B. Funt. Evaluating combinational illumination estimation methods on real-world images. *IEEE TIP*, 23(3):1194–1209, Mar. 2014.
- [18] A. Litvinov and Y. Y. Schechner. Radiometric framework for image mosaicking. *JOSA A* 22(5):839-848, 2005.
- [19] M. G. Masi, L. Peretto, R. Tinarelli, and L. Rovati. Modeling of the physiological behavior of human vision system under flicker condition. In *IEEE Int. Conf. Harmonics & Quality of Power*, 2008.
- [20] S. U. Mehta, K. Kim, D. Pajak, K. Pulli, J. Kautz and R. Ramamoorthi. Filtering environment illumination for interactive physically-based rendering in mixed reality. In *ESGR*. 2015.
- [21] H. Murase and S. K. Nayar. Visual learning and recognition of 3-D objects from appearance. In *Int. J. Computer Vision*, 14(1), 5-24, 1995.
- [22] A. V. Oppenheim and R. W. Schaffer. From frequency to quefrency: A history of the cepstrum. *IEEE Signal Proc. Magazine*, 21(5):95-106, 2004.
- [23] M. O’Toole, S. Achar, S. G. Narasimhan, and K. N. Kutulakos. Homogeneous codes for energy-efficient illumination and imaging. *ACM SIGGRAPH*, 2015.
- [24] M. O’Toole, J. Mather, and K. N. Kutulakos. 3D shape and indirect appearance by structured light transport. *IEEE TPAMI*, 38(7):1298–1312, 2016.
- [25] D. Reddy, A. Veeraraghavan, and R. Chellappa. P2C2: Programmable pixel compressive camera for high speed imaging. In *Proc. IEEE CVPR*, 329–336, 2011.
- [26] V. Rengarajan, A. N. Rajagopalan and R. Aravind. From bows to arrows: rolling shutter rectification of urban scenes. In *Proc. IEEE CVPR*, 2016.
- [27] V. Rengarajan, Y. Balaji and A. N. Rajagopalan. Unrolling the shutter: CNN to correct motion distortions. In *Proc. IEEE CVPR*, 2017.
- [28] V. Reno, R. Marani, M. Nitti, N. Mosca, T. D’Orazio and E. Stella. A powerline-tuned camera trigger for ac illumination flickering reduction. *IEEE Embedded Systems Letters*, 9(4), 97-100, 2017.
- [29] M. Rouf, R. Mantiuk, W. Heidrich, M. Trentacoste and C. Lau. Glare encoding of high dynamic range images. In *Proc. IEEE CVPR*, 2011.
- [30] O. Saurer, K. Koser, J. Y. Bouguet and M. Pollefeys. Rolling shutter stereo. In *Proc. IEEE ICCV*, 2013.
- [31] Y. Y. Schechner and S. K. Nayar. Uncontrolled modulation imaging. In *Proc. IEEE CVPR*, 197-204, 2004.
- [32] Y. Y. Schechner and N. Karpel. Attenuating natural flicker patterns. In *MTS/IEEE OCEANS*, 1262-1268, 2004.
- [33] Y. Y. Schechner. Self-calibrating imaging polarimetry, In *Proc. IEEE ICCP*, (2015).
- [34] P. Sen, B. Chen, G. Garg, S. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch. Dual photography. In *ACM SIGGRAPH*, 745–755, 2005.
- [35] Y. Swirski and Y. Y. Schechner, 3Deflicker from motion. In *Proc. IEEE ICCP*, 2013.
- [36] M. Sheinin, Y. Y. Schechner, and K. Kutulakos. Computational imaging on the electric grid. In *Proc. IEEE CVPR*, 2017.
- [37] S. Su and W. Heidrich. Rolling shutter motion deblurring. In *Proc. IEEE CVPR*, 2015.
- [38] T. Tajbakhsh and R.-R. Grigat. Illumination flicker correction and frequency classification methods. In *Proc. SPIE Electronic Imaging*, (6502):1-10, 2007.
- [39] T. Tajbakhsh and R. R. Grigat. Illumination flicker frequency classification in rolling shutter camera systems. In *Proc. SIP*, (7):282-287, 2007.
- [40] R. Talmon, I. Cohen, S. Gannot and R.R. Coifman. Diffusion maps for signal processing: A deeper look at manifold-learning techniques based on kernels and graphs. *IEEE Signal Processing Magazine*, 30(4), 75-86, 2013.
- [41] M.F. Tappen, W.T. Freeman, E.H. and Adelson. Recovering intrinsic images from a single image. In *In NIPS*, 1367-1374, 2003.
- [42] A. Veeraraghavan, D. Reddy, and R. Raskar. Coded strobing photography: compressive sensing of high speed periodic videos. *IEEE TPAMI*, 33(4):671–686, 2011.
- [43] D. Veikherman, A. Aides, Y. Y. Schechner, and A. Levis. Clouds in the cloud. In *Proc. ACCV*, 659-674, 2014.
- [44] M. Vollmer and K.-P. Moellmann. Flickering lamps. *European Journal of Physics*, 36(3), 2015.
- [45] *Rolling Shutter Imaging on the Electric Grid: Webpage*. [Online] (2018). Available at: <http://www.marksheinin.com> <http://webee.technion.ac.il/~yoav/research/RScam.html> <http://www.dgp.toronto.edu/RScam>
- [46] Y. Weiss. Deriving intrinsic images from image sequences. In *Proc. IEEE ICCV* (2):68-75, 2001.
- [47] G. Woo, A. Lippman, and R. Raskar. VRcodes: Unobtrusive and active visual codes for interaction by exploiting rolling shutter. In *IEEE ISMAR*, 59-64, 2012.
- [48] Y. Yoo, J. Im, and J. Paik. Flicker removal for CMOS wide dynamic range imaging based on alternating current component analysis. *IEEE Trans. on Consumer Electronics*, 60(3):294–301, 2014.