

Uncontrolled Modulation Imaging

Yoav Y. Schechner

Dept. Electrical Engineering
Technion - Israel Inst. Technology
Haifa 32000, ISRAEL
yoav@ee.technion.ac.il

Shree K. Nayar

Dept. Computer Science
Columbia University
New York, NY 10027
nayar@cs.columbia.edu

Abstract

To obtain high dynamic range or hyperspectral images, multiple frames of the same field of view are acquired while the imaging settings are modulated; images are taken at different exposures or through different wavelength bands. A major problem associated with such modulations has been the need for perfect synchronization between image acquisition and modulation control. In the past, this problem has been addressed by using sophisticated servo-control mechanisms. In this work, we show that the process of modulation imaging can be made much simpler by using vision algorithms to automatically relate each acquired frame to its corresponding modulation level. This correspondence is determined solely from the acquired image sequence and does not require measurement or control of the modulation. The image acquisition and the modulation work continuously, in parallel, and independently. We refer to this approach as computational synchronization. It makes the imaging process simple and easy to implement. We have developed a prototype modulation imaging system that uses computational synchronization and used it to acquire high dynamic range and multispectral images.

1 Modulation Imaging

The quality of acquired scene information is often enhanced by capturing a sequence of images of the scene and combining the information within them. Each image of the sequence corresponds to a different modulated state of one or more of the imaging parameters. We concentrate here on two modulation types. One type modulates the exposure settings, with the aim of extending the radiometric dynamic range of the system. At each exposure setting, a different range of intensities can be measured reliably. Fusing the data from all the exposures [6, 14] results in a single high dynamic range (HDR) image.¹ Several methods have been developed to enable the display of HDR images on low dynamic range media [5, 8, 9, 19, 25]. A problem with prior acquisition methods of this sort has been the control of the acquisition process: The exposure parameters have to be set at a new level before a new frame can be acquired.

¹Other approaches for HDR imaging use nonlinear components [17, 27] and systems [16] or global camera motion [1, 3, 22].

The other type of modulation we deal with filters into the detector a narrow wavelength band. Taking multiple images at different modulation settings results in multispectral/hyperspectral images. This imaging mode is extremely useful in several applications [21], including material and object recognition [7, 24], color analysis and color constancy [12, 13, 26], biomedical imaging [18], remote sensing [24] and astronomy [2, 28]. Widefield multispectral images are typically acquired by sequentially changing the central wavelength of the measured spectral band [4, 7, 11, 12, 18, 28]. As in the case of HDR imaging, existing acquisition methods for widefield multispectral imaging require precise control of the coupling between image acquisition and the state of the spectral filter. After each image is acquired, the spectral band is changed to one out of several discrete levels. During this process, image acquisition is put on hold.

In many applications, it is beneficial to have uncontrolled modulation of the imaging system. In such a scheme, no hardware is needed to synchronize the camera with the modulation device. Both work continuously, in parallel, and without a stepped sequencing mechanism. Such a scheme can be realized using systems that are much simpler than controlled systems; there is no need for a controller and feedback channels. In addition, the system can use continuous modulation, where a filter undergoes a simple periodic transformation. In comparison with existing “stepped” approaches, this is simpler to implement.

Even when the modulation is uncontrolled during acquisition, we need accurate knowledge of the modulation state corresponding to each acquired image. Without this knowledge, we cannot recover the spectral distribution of the object’s radiance in multispectral imaging, or an image with known radiometric response in the case of HDR imaging.

In this paper we show that we can determine the modulation state at each image *based only on the acquired image sequence*. No auxiliary recording of the modulation state is required. By simply analyzing the acquired image sequence, we determine a-posteriori at which wavelength band or exposure setting each of the images was taken. We call this method *computational synchronization*. We develop algorithms for computational synchronization for





Figure 1. A system for demonstrating the algorithms of uncontrolled modulation imaging. It uses a monochrome image detector. Adjacent to the detector (before the lens), there are two spatially varying filters: one with varying transmittance, the other with a varying spectral response.

HDR imaging and hyperspectral imaging. We demonstrate our algorithms on image sequences acquired using a prototype uncontrolled modulation imaging system we built.

2 A Prototype Implementation

The emphasis of the work is algorithmic. However, to demonstrate our ideas, we implemented a simple image acquisition system. We point out that the algorithms we develop are not limited to this specific implementation. The system uses a monochrome video camera having a linear radiometric response. We placed spatially varying filters right in front of the CCD detector, between it and the lens, as depicted in Fig. 1. One filter is a circular variable neutral density filter. Its transmittance depends on its rotation angle. The other is a linear variable interference filter. The transmitted spectral band varies across this filter. Both filters are driven by simple DC motors (Fig. 1). When the motors are turned on, the state of the filters change with time. This causes temporal modulation of the light transmittance and/or spectral band. The modulations are *continuous*, and

uncontrolled. There is no synchronization between the state of the filters and image acquisition.

3 Image Analysis for High Dynamic Range

3.1 Estimation of the Modulation

We alleviate the need of controlling the temporal modulation of the transmittance. Rather than controlling the transmittance, we estimate it from the sequence itself. Let F_k be the transmittance of the density filter at frame k , and I be the intensity of light falling on the detector (irradiance), when the transmittance of the filter is maximal ($F = 1$). When the detector is not saturated, the intensity in frame k is $v_k = F_k I$. Then, this image point is measured without saturation in frame p , with intensity value $v_p = F_p I$. Assuming the scene radiance is constant between frames, these measurements should satisfy

$$F_k v_p - F_p v_k = 0 \quad . \quad (1)$$

Measuring the intensities of various image points \vec{x} in several frames provides many such linear equations, which the filter modulation should satisfy. This set of equations can be written as $\mathbf{V}\mathbf{F} = 0$.

In addition, we impose temporal smoothness on the estimated modulation, by penalizing for 2nd order variations in F . The smoothest modulation would satisfy $\mathbf{L}\mathbf{F} = 0$, where \mathbf{L} is the 2nd order derivative operator (See [22]). Overall, we obtain an overconstrained system of equations. The least squares solution of this system of equations is

$$\hat{F} = \arg \min_F (F^t \mathbf{M}^t \mathbf{M} F) \quad , \quad (2)$$

where $\mathbf{M} = [\mathbf{V}^t \beta \mathbf{L}^t]^t$. Here β is a parameter that weights the penalty for unsmooth solutions relative to the penalty for disagreement with the data. The solution is derived using singular value decomposition. The solution is scaled by setting $\max \hat{F} = 1$.

We demonstrated this method on several sequences, samples of which are shown in Fig. 2. There are millions of corresponding measurements in each sequence. From them, we obtained about 40,000 equations as Eq. (1), based on randomly picked measurement pairs. The intensity values selected were unsaturated and also not dark. The self-calibrated modulations for each of the sequences are plotted in Fig. 3. Note that each sequence has a different transmittance variation in time. The reason is that the motor in each experiment was run at a different speed. In addition, the mechanical coupling of the motor to the filter was not tight. Nevertheless, the self-calibration method we described bypasses the need for motion control.

3.2 Fusing the Measurements

We now describe the method we used for estimating the HDR intensity. Let a measured intensity value at frame k be v_k with uncertainty Δv_k , corresponding to an estimated transmittance \hat{F}_k . Compensating the readout for the filter's



Figure 2. Raw sample images from the “Church” and “Face” video streams. The video streams were acquired while the filter transmittance changed. These changes were uncontrolled.

attenuation, the scene point’s intensity is $I_k = v_k / \hat{F}_k$. Estimates are combined using a weighted average,

$$\hat{I} = W^{-1} \sum_k w_k I_k, \quad (3)$$

where $W \equiv \sum_k w_k$ is a normalization factor. The weights are set by

$$w_k = \frac{1}{\Delta I_k^2} = \left| \frac{\partial I_k}{\partial v_k} \Delta v_k \right|^{-2} = \left(\frac{\hat{F}_k}{\Delta v_k} \right)^2, \quad (4)$$

so that I_k carries a smaller weight, if its uncertainty ΔI_k is large. The estimates corresponding to higher light transmittance \hat{F}_k are more reliable and are thus more influential. We assume the readout uncertainty to be $\Delta v_k = 0.5$, since the readout values are integers. Measurements suspected as saturated are treated as having high uncertainty, thus their corresponding Δv_k is set to be a very large number.

The images from the sequences, of which samples are shown in Fig. 2, were fused into HDR intensity images using this method. To display the image content on standard, low dynamic range media, we processed the HDR images in the two steps. First, we used the HDR image to render several standard images of the scene (which include satu-

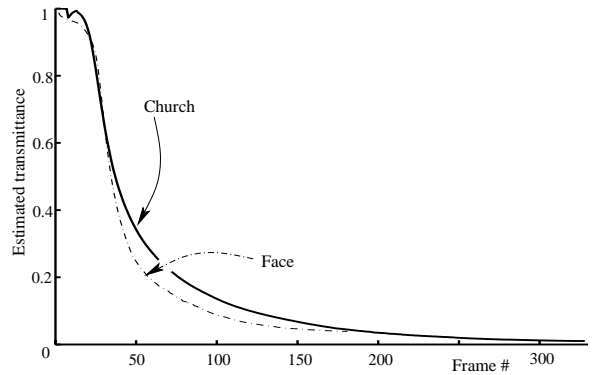


Figure 3. The estimated transmittance as a function of time (frame index), for two video streams. The transmittance was calculated based on the raw images, bypassing the need for imaging synchronization.



Figure 4. Images derived from HDR intensity information. The HDR information was calculated from raw images, samples of which are shown in 2. The calculation used the estimated temporal transmittance plotted in Fig. 3.

ration), each of which simulating imaging at a distinct exposure. These images have a better quality than the raw frames, since no blooming effect takes place around simulated saturation points. In addition, these processed images are the result of data gathered in numerous frames, and thus are much less noisy. Then, we fused these images into a single high-content image, using the method described in Ref. [5]. The results are displayed in Fig. 4.

4 Hyperspectral Imaging

4.1 Common Approaches

There are several approaches for obtaining hyperspectral images. Our method belongs to the most common category, which is based on a widefield approach. In the widefield category, the entire field of view (FOV) is imaged at once. Multiple images of the scene are taken, each of which through a filter having a different narrow wavelength band.² There are several realizations of this approach. One real-

²There are two additional categories for hyperspectral imaging systems. Pushbroom systems scan the FOV in narrow spatial swaths. Transform systems multiplex spectral components, rather than scanning individual spectral bands. Those categories are disjoint to the work presented here.




Figure 5. Raw sample images from several video streams. During video acquisition, the sensed spectral band changed in an uncontrolled manner. Note the strong variations of the flower appearance [left], where in some wavelengths they appear bright while in others they are very dark. Contrast polarity also flips in the intensity of the fluorescent bulb [right].

ization is based on a filter wheel, over which discrete filters are mounted. The wheel is rotated by a discrete angle between image acquisitions. Related realizations are based on continuous variation of the wavelength passband across a filter [4, 18, 28] which moves in discrete steps. Additional widefield systems are based on electronic tunable filters as liquid crystal tunable filters (LCTF) [11], acousto-optical tunable filters (AOTF) [7, 12], and Fabry-Perot etalons [2].

4.2 Our Method

The modulation of the spectral pass-band is obtained in our system by periodically moving a variable interference filter. Sample images taken using the system are shown in Fig. 5. Some of our sequences had dramatic effects. Consider the flowers on the left side of Fig. 5. There are very strong variations of flower appearance during spectral modulation: in some wavelengths they appear bright while in others they are very dark. Strong variations also appear in the “Bulbs” sequence (right side of Fig. 5), where a fluorescent bulb radiates intensively only in several narrow bands.

To extract quantitative spectral information from such data, we need to know which wavelength band each measurement corresponds to. Thus, the spectral modulation state needs to be known at each frame. Nevertheless, controlling the modulation state in conjunction to the image acquisition complicates the system. We thus introduce in this section a method for determining the filtered passband at

each frame, *based on the acquired images*. The resulting computational synchronization alleviates the need for on-line control or monitoring of the spectral modulation.

The key for this blind estimation lies in creating an *oscillatory modulation*. In this scheme, the wavelength passband shifts back and forth periodically, thus each band is sampled twice during each modulation cycle. In the particular system we built, this was achieved by a crankshaft, which converted motor rotation to oscillatory linear translation.

4.3 State Dynamics

Let λ_{\max} and λ_{\min} be the maximum and minimum wavelengths of the scanned spectrum, respectively. At acquisition time t , the central wavelength of the filter’s passband is

$$\lambda(t) = (\lambda_{\max} + \lambda_{\min})/2 - [(\lambda_{\max} - \lambda_{\min})/2]f(t) \quad , \quad (5)$$

where $f(t)$ is the modulation function which determines the wavelength band observed at each frame. We require that $f(t)$ be oscillatory with period T :

periodic: For each t , $f(t) = f(t + T)$

oscillatory: For each t_1 exists a unique $\Delta t_1 > 0$, such that

$$f(t_1) = f(t_1 + \Delta t_1), \quad \text{while } \Delta t_1 < T \quad . \quad (6)$$

We stress that the oscillations *do not need to be harmonic* (i.e. a cosine function). Indeed, in the system we use for demonstration, $f(t)$ is not a cosine function.

Define a function $f_{\text{baseline}}(\theta)$, which is similar to $f(t)$, but shifted and scaled in time so its minimum is at $\theta = 0$, and its period is 2π . We may thus write

$$f(t) = f_{\text{baseline}}[(2\pi t/T) - \theta_0] \quad , \quad (7)$$

where θ_0 is the *phase* of the function $f(t)$. In general, we know $f_{\text{baseline}}(\theta)$. The reason is that it can usually be pre-calibrated or derived in closed form. For example, the modulation may be created by a motor rotation, as we indeed do in the experiments described in Sec. 4.5. In this case, we know for each angle θ of the motor what the corresponding wavelength is, although we do not know (yet) what is the angle θ at any given time. The same happens when the spectral modulation is determined by a periodic voltage (or another signal), as is the case in tunable filters: it is known in advance which wavelength corresponds to any voltage, but we need to find out which voltage is used at any time.

It follows from Eqs. (5,7), that in order to estimate the central wavelength at time t , we need to estimate two global parameters: the modulator’s period T and the phase θ_0 . The underlying assumption we make here is that θ changes linearly with t , e.g., the motor rotates at a constant rate.

4.4 Estimating the Spectral Modulation

In this section we show a way to estimate the global parameters T and θ_0 from the image sequence. Consider the spectrum of a scene point corresponding to a single image pixel, e.g., the function shown in Fig. 6a. Then, consider

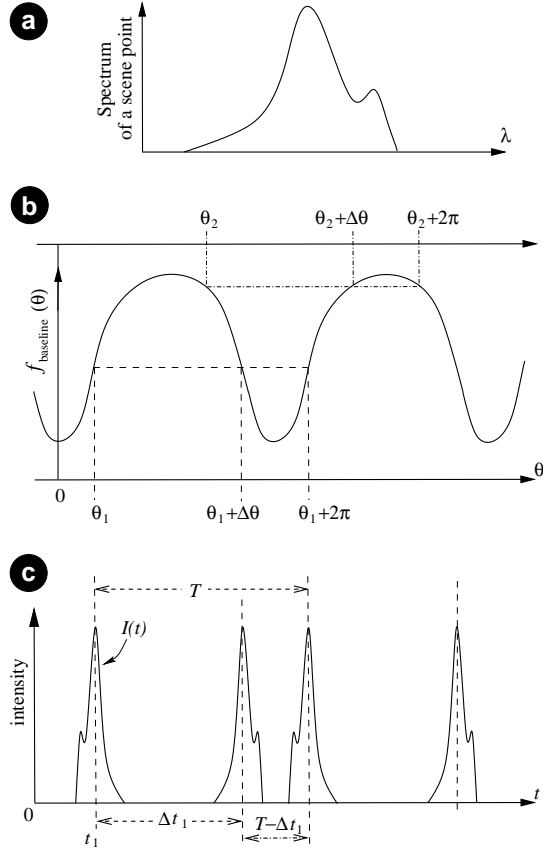


Figure 6. (a) Spectrum of a scene point. (b) The modulation function is periodic and oscillating. (c) The spectrum is scanned leading to intensity readout $I(t)$. Oscillations cause two scans of each wavelength in each period. A feature in the spectrum appears at times t_1 , $t_1 + \Delta t_1$. This pattern repeats at the modulation period T . The measured Δt_1 and T indicate the phase of the modulation. This indication has a simple ambiguity seen in (b): both phases θ_1 and θ_2 can account for the measured Δt_1 and T .

the oscillatory function $f_{\text{baseline}}(\theta)$ depicted in Fig. 6b. As θ changes in time, the measured intensity $I(t)$ varies in correspondence to the spectrum³. The maxima and minima of $I(t)$ are easy features to detect, and they correspond to extremum points of the spectrum. Since $f(t)$ is periodic, the sequence of scanned wavelengths is repeated in the same order and rate at $t + T$. Thus, $I(t)$ is replicated at $t + T$ (Fig. 6c). The estimation of T is thus simply obtained from the temporal period of the measured intensity.

We are now left with the problem of estimating θ_0 . Recall that the modulation oscillates *back and forth*. Consider Fig. 6b and Eq. (6). For every θ_1 , there exists $\Delta\theta$, such that $0 < \Delta\theta < 2\pi$, while $f_{\text{baseline}}(\theta_1) = f_{\text{baseline}}(\theta_1 + \Delta\theta)$.

³The intensity $I(t)$ resembles a temporally *distorted* (unevenly stretched) version of the spectrum. The reason is that in general $f(t)$ is not piecewise linear, thus λ and t are related in a nonlinear fashion.

Corresponding to the opposite directions of the modulation variation at these two phases,

$$\text{sign} \left[\frac{\partial f_{\text{baseline}}(\theta)}{\partial \theta} \right]_{\theta=\theta_1} = -\text{sign} \left[\frac{\partial f_{\text{baseline}}(\theta)}{\partial \theta} \right]_{\theta=\theta_1+\Delta\theta} \quad (8)$$

Since θ varies linearly with t , there is a time instance t_1 and a time lag Δt_1 corresponding to θ_1 and $\Delta\theta$. The same wavelength is sampled at t_1 and $t_1 + \Delta t_1$, but it is scanned in a mutually opposite order:

$$\begin{aligned} \text{sign} [\partial f(t)/\partial t]_{t=t_1} &= -\text{sign} [\partial f(t)/\partial t]_{t=t_1+\Delta t_1} \\ \Rightarrow \text{sign} [d\lambda(t)/dt]_{t=t_1} &= -\text{sign} [d\lambda(t)/dt]_{t=t_1+\Delta t_1} \end{aligned} \quad (9)$$

Therefore, if the measured intensity $I(t)$ has a maximum value at t_1 , then a *time-reflected* replicate⁴ of $I(t)$ exists with its corresponding maximum at $t_1 + \Delta t_1$ (See Fig. 6c). It is often easy to detect these corresponding maxima. Therefore, from the acquired data we can measure the time lags Δt_1 and $T - \Delta t_1$ depicted in Fig. 6c. As we show next, these measurements facilitate the required phase estimation.

From the measurements we empirically estimate the ratio

$$r \equiv \frac{\Delta t_1}{T} = \frac{\Delta\theta}{2\pi} . \quad (10)$$

Thus

$$\Delta\theta = 2\pi r . \quad (11)$$

As seen in Fig. 6b, in each modulation period T there are exactly two phases θ_1 and θ_2 , for which the modulation state is replicated, $f_{\text{baseline}}(\theta_k) = f_{\text{baseline}}(\theta_k + \Delta\theta)$ where $k = 1$ or 2 . Since we know a priori the function $f_{\text{baseline}}(\theta)$, we can easily derive the solutions θ_k and $f_{\text{baseline}}(\theta_k)$.

The values $f_{\text{baseline}}(\theta_k)$ correspond to two specific wavelengths λ_k :

$$\lambda(\theta_k) = (\lambda_{\text{max}} + \lambda_{\text{min}})/2 - [(\lambda_{\text{max}} - \lambda_{\text{min}})/2]f(\theta_k) . \quad (12)$$

To appreciate the significance of this observation, note that this derivation enables us to *estimate the wavelength of a source measurement, although we have not used any control mechanism to synchronize the filter modulation with the measuring sensor (camera)*. Yet, the estimate is prone to a simple ambiguity, since there are two possible solutions, $\lambda(\theta_1)$ and $\lambda(\theta_2)$, as depicted in Fig. 6b. This ambiguity is easily removed using a qualitative cue derived from a second scene point of a different color, as we detail in [23].

To estimate T and Eq. (10), a little more than a full period of the scan should be completed. During that period, the spectrum is practically scanned three times. For example, the wavelength at t is scanned also at $t + \Delta t_1$ and $t + T$. This repetition can be exploited. Each of scan can be performed in different settings of the imaging system, in order to extend the information in other imaging dimensions. For

⁴Similarly to explanation in footnote 3, this replicate is temporally distorted (unevenly stretched in time) relative to $I(t)$. This is due to the different nonlinear behavior of $f(t)$ in different time instances. This distortion is in addition to the temporal reflection and translation.

example, the scan can be performed in different exposure settings to obtain HDR data. In fact, this is what we chose to do in the experiments described in Sec. 4.5. Alternatively, the scan can be performed in different polarization settings to determine the polarization state [15, 30], or different focus settings to obtain depth from defocus [10, 20, 29].

4.4.1 Symmetric Modulation

If the oscillatory modulation is symmetric in time, then its phase can be calculated in closed form. It is reasonable to assume that symmetric modulation will most often be used. For example, in electronic tunable spectral filters the modulating voltage can be increased at the rate of its decrease. In the mechanical system like the one we describe in Sec. 4.5 (see Fig. 1), the wavelength is symmetric with respect to the motor angle θ . In symmetric modulation,

$$f_{\text{baseline}}(\theta) = f_{\text{baseline}}(-\theta) = f_{\text{baseline}}(2\pi - \theta) . \quad (13)$$

We seek $\theta_k > 0$ such that

$$f_{\text{baseline}}(\theta_k) = f_{\text{baseline}}(\theta_k + 2\pi r) , \quad (14)$$

considering Eqs. (10,11). Note that

$$f_{\text{baseline}}(-\theta) = f_{\text{baseline}}(2\pi - \theta) = f_{\text{baseline}}[\theta + (2\pi - 2\theta)] . \quad (15)$$

Comparing Eqs. (13,14,15), the first solution satisfies $2\pi r = 2\pi - 2\theta_1$, thus

$$\theta_1 = \pi(1 - r) . \quad (16)$$

Similarly,

$$f_{\text{baseline}}(-\theta) = f_{\text{baseline}}(4\pi - \theta) = f_{\text{baseline}}[\theta + (4\pi - 2\theta)] , \quad (17)$$

leads to

$$\theta_2 = \pi(2 - r) . \quad (18)$$

Since $0 < r < 1$, then $0 < \theta_1 < \pi < \theta_2 < 2\pi$.

To conclude, we can determine the phase of the modulation cycle in closed form (either θ_1 or θ_2) based on the empirical measurement of Δt_1 and T . Interestingly, this phase estimate does not depend at all on the nature of $f(\theta)$: the modulation can be piecewise linear, a cosine or an arbitrary function. As long as $f(\theta)$ is oscillatory and symmetric, its phase is determined by Eqs. (10,16,18).

4.5 Experiments

In our implementation, depicted in Fig. 1, we use a linear variable interference filter (LVF). The central wavelength of the LVF passband varies linearly across the filter. Let $\lambda(x_{\text{internal}})$ be the central wavelength of the filter's passband, where x_{internal} is a position along the filter, measured in the internal coordinate system of the filter.

$$\lambda(x_{\text{internal}}) = (\lambda_{\text{max}} + \lambda_{\text{min}})/2 + [(\lambda_{\text{max}} - \lambda_{\text{min}})/l]x_{\text{internal}} \quad (19)$$

where l is the length of the LVF.

Consider Fig. 7. As the motor rotates, the LVF translates horizontally. Thus, the position of the center of the LVF (where $x_{\text{internal}} = 0$) changes with time. We denote this position by $x_0(t)$ in the external coordinate system of the




Figure 7. A system having a periodic, but not harmonic, spectral modulation. A crankshaft of length A coupled to ring of radius R transforms the rotation of a motor located at x_{motor} to the linear motion of an LVF. The motor angle θ corresponds to the LVF part affecting the imaging system existing in $x = 0$. A point of the filter is parameterized by the internal coordinate x_{internal} relative to the filter's middle, whose coordinate x_0 varies in time. The filter's middle is at a constant distance B from the crankshaft.

camera lens. We define the origin of this coordinate system x at the lens optical axis. In particular, $x_0 = 0$ indicates a state in which the center of the filter is on the optical axis. In this coordinate system, a general position on the filter is $x = x_0 + x_{\text{internal}}$. From Eqs. (5) and (19),

$$x_0(t) = lf(t)/2 . \quad (20)$$

The LVF moved using a crankshaft mechanism, depicted in Fig. 7. We know for each angle θ of the motor, what the corresponding modulation $f_{\text{baseline}}(\theta)$ is:

$$x_0 = lf_{\text{baseline}}(\theta)/2 = x_{\text{motor}} + \left(\sqrt{A^2 + R^2 \sin^2 \theta} - R \cos \theta \right) + B \quad (21)$$

where x_{motor} is the position of the motor axis, measured in the coordinate system of lens. The radius of the crankshaft is R , while A is the length of its connecting rod. Here B is the horizontal distance between the filter's center and the connecting point of the crankshaft rod to the filter holder (See Fig. 7). Obviously, Eq. (21) is not a pure cosine function. Nevertheless, it is periodic, oscillatory, and symmetric with respect to θ .

As indicated in Fig. 5, we placed a few LEDs of different colors in the periphery of the FOV. Although we *did not know* their spectra, they determined the global modulation parameters. In each sequence, we measured the intensity of the LEDs as a function of time. Fig. 8 plots part the measurements taken from a video sequence a couple of minutes long. The dotted line plots the intensity⁵ of a yellow LED. By matching the peaks of this plot, we determined T and Δt_1 , as described in Sec. 4.4. The solid line plots

⁵In the specific system we used, the translation range of the filter $2R$ significantly exceeded the length of the filter l . For this reason, a significant number of frames were taken outside the effective range of the spectral filter. At these frames the light throughput increased. For this reason the LED intensity plots in Fig. 8 are not unimodal.

the intensity of a green LED. There is a slight temporal difference between the peaks corresponding to the two LEDs. This difference removes the solution ambiguity described in Fig. 6b, i.e., it uniquely determines whether θ_1 is the modulation phase, or whether it is θ_2 (See Ref. [23]).

We used the multiple wavelength scans to capture the scene in multiple exposures, using the system’s circular variable density (Fig. 1). For this reason, peaks in Fig. 8 corresponding to different oscillations have different intensities. This operation was very useful in scenes having a high dynamic range of intensities. For instance, the right hand side of Fig. 5 shows several images of the “Bulbs” sequence. The 2nd and 3rd images were taken at a similar wavelength band but at different exposure settings, revealing details of dim as well as bright objects.

Following the procedure described in Secs. 4.3 and 4.4, we used the measurement of T and Δt_1 to determine the global parameters of $f(t)$. Based on that and Eqs. (5,21), we determined the central passband wavelength at each frame. Then, we fused measurements of corresponding bands which were taken at different exposure settings, as described in Sec. 3. This resulted in HDR multispectral images. Spectra of selected object points are plotted in Fig.9. Note that some gray-level values far exceed 255, which is the limit set by our 8-bit camera. The reason we manage to derive such values stems from the HDR information extracted from the multiple exposure data.

4.6 Spatial Inhomogeneities

For the sake of explanation simplicity, the derivations starting at Sec. 4.3 have assumed that the temporal modulation is spatially invariant. This assumption is easily alleviated. To see this, observe that Eqs. (5,20) refer to the central wavelength on the optical axis, that is, at the center of the FOV. When the spectral modulation is achieved by shifting a LVF, pixels left or right of the center will sense passbands with shifted central wavelengths. This *does not* affect the procedure of computational synchronization. A global shift of the wavelengths at some pixel corresponds to an addition of a constant to the function $f_{\text{baseline}}(\theta)$. As seen in Fig. 6b, the DC component of $f_{\text{baseline}}(\theta)$ does not influence at all the phases θ_k and $\Delta\theta$. The measured T and r are invariant to spatial shifts in the image, and so is the end result.

Nevertheless, knowing the global parameters of the wavelength modulation is just one step. Eventually, we are interested in knowing, for each pixel, which wavelength band it sampled at any of the frames. Thus, for building the multispectral image cube, we need to account for shifts of the wavelength band which vary across the FOV. If the interference filter resides close to the imaging detector array whose dimensions are known, the wavelength shift is known in closed form. It is then easy to account for that when building the image cube. If the filter resides away from the detector, the shift may be empirically calibrated.




Figure 8. Computational synchronization in an experiment. Features as intensity peaks indicated the period T . Measuring an intermediate appearance of a peak at Δt_1 indicated the phase of the modulation. Two spectral features, e.g., peaks corresponding to different colored objects, had different values of Δt_1 . This was used to alleviate the ambiguity mentioned in Fig. 6b. The measured corresponding peaks have different intensities because of changes in the transmittance (exposure) settings per scan.

5 Discussion: Natural Extensions

During the modulation process, some objects in the FOV may move. This has always been a problem faced by modulation imaging. It is worth noting that in the method we proposed the temporal modulation is gradual and continuous. We therefore believe that tracking corresponding scene points is much easier in our setting, than in cases where the modulation is done in large discrete steps. In future work, we plan to explore ways of making modulation imaging robust to scene and camera motions.

Our computational spectral synchronization was based on matching a few prominent spectral features. Greater robustness can be achieved using standard registration algorithms, such as those used frequently in vision to register images. Temporal registration will optimize the autocorrelation of the temporal intensity signal, by seeking the best temporal distortions and shifts. If this can be achieved, we will not have to rely on scene points with prominent spectral peaks (e.g. LEDs); all the scene points, irrespective of their spectra, will be able to contribute to the synchronization.

Acknowledgments

The work was supported by the US-Israel Binational Science Foundation (BSF). Yoav Schechner is a Landau Fellow - supported by the Taub Foundation, and an Alon Fellow. His work is supported by the Ollendorff Minerva Center. Minerva is funded through the BMBF. Shree Nayar is supported by NSF ITR Award 00-85864.




Figure 9. Experimental results. Spectral plots extracted from video streams, samples of which are shown in Fig. 5. The correspondence between intensities and wavelengths was achieved using computational synchronization. The HDR intensity values in the plots are obtained thanks to the multiple transmittance (exposure) settings used.

References

- [1] M. Aggarwal and N. Ahuja, "High dynamic range panoramic imaging," Proc. Int. Conf. Comp. Vision, Vol. I, 2-9 (2001).
- [2] I. K. Baldry and J. Bland-Hawthorn, "A tunable echelle imager," Pub. Astron. Soc. Pacific **112**, 1112-1120 (2000).
- [3] F. M. Candocia, "Jointly registering images in domain and range by piecewise linear comparometric analysis," IEEE Trans. IP **12**, 409-419 (2003).
- [4] J. Blommaert, R. Siebenmorgen, A. Coulais, K. Okumura, S. Ott, M. Sauvage and J.L. Starck, *The ISO handbook vol.III: CAM - the ISO camera* (European Space Agency, 2001).
- [5] P. J. Burt and R. J. Kolczynski, "Enhanced image capture through fusion," Proc. Int. Conf. Comp. Vis. 173-182 (1993).
- [6] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," Proc. SIGGRAPH 369-378 (1997).
- [7] L. J. Denes, M. Gottlieb, B. Kaminsky and D. F. Huber, "Spectro-polarimetric imaging for object recognition," Proc. SPIE **3240**, 8-18 (1998).
- [8] F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," ACM TOG **21**, 257-266 (2002).
- [9] R. Fattal, D. Lischinski, and M. Werman, "Gradient domain high dynamic range compression," ACM TOG **21**, 249-256 (2002).
- [10] P. Favaro, A. Mennucci and S. Soatto, "Observing shape from defocus," Int. J. Computer Vision **52**, 25-43 (2003).
- [11] N. Gat, "Imaging spectroscopy using tunable filters: A review," Proc. SPIE **4056** Wavelet Applic. VII, 50-64 (2000).
- [12] J. Hallikainen, J. P. S. Parkkinen, and T. Jaaskelainen "Color image processing with AOTF," Proc. SCIA 294-300 (1989).
- [13] M. Hauta-Kasari, K. Miyazawa, S. Toyooka and J. Parkkinen, "Spectral vision system for measuring color images," JOSA A **16**, 2352-2362 (1999).
- [14] S. Mann and R. W. Picard, "On being 'Undigital' with digital cameras: extending dynamic range by combining differently exposed pictures" IS&T Annual Conf., 422-428 (1995).
- [15] D. Miyazaki, M. Kagesawa and K. Ikeuchi, "Transparent surface modeling from a pair of polarization images," IEEE Trans. PAMI **26**, 73-82 (2004).
- [16] S. K. Nayar and V. Branzoi, "Adaptive dynamic range imaging: Optical control of pixel exposures over space and time," Proc. Int. Conf. Comp. Vision, 1168-1175 (2003).
- [17] W. Ogiers, "Survey of CMOS imagers," IMEC report P60280-MS-RP-002, Issue 1.1, part 1 (1997).
- [18] A. Papadakis, E. Stathopoulos, G. Delides, K. Berberides, G. Nikiforidis and C. Balas, "A novel spectral microscope system: application in quantitative pathology," IEEE Trans. Biomedical Eng. **50**, 207-217 (2003).
- [19] A. Pardo and G. Sapiro, "Visualization of high dynamic range images," IEEE Trans. IP **12**, 639-647 (2003).
- [20] Y. Y. Schechner and N. Kiryati, "Depth from defocus vs. stereo: How different really are they?" Int. J. Computer Vision **89**, 141-162 (2000).
- [21] Y. Y. Schechner and S. K. Nayar, "Generalized mosaicing: Wide field of view multispectral imaging," IEEE Trans. PAMI **24**, 1334-1348 (2002).
- [22] Y. Y. Schechner and S. K. Nayar, "Generalized mosaicing: High dynamic range in a wide field of view," Int. J. Computer Vision **53**, 245-267 (2003).
- [23] Y. Y. Schechner and S. K. Nayar, "Uncontrolled Modulation Imaging," Tech. Report (EE Dept., Technion 2004).
- [24] D. Slater and G. Healey, "Material classification for 3D objects in aerial hyperspectral images," Proc. Computer Vision and Pattern Recognition, vol. 2, 268-273 (1999).
- [25] D. A. Socolinsky, "Dynamic range constraints in image fusion and realization." Proc. IASTED Int. Conf. Signal and Image Process, 349-354 (2000).
- [26] H. M. G. Stokman, T. Gevers and J. J. Koenderink, "Color Measurement by Imaging Spectrometry," Computer Vision and Image Understanding **79**, 236-249 (2000).
- [27] N. Tabiryan, and S. Nersisyan, "Liquid-crystal film eclipses the sun artificially," Laser Focus World **38/5**, 105-108 (2002).
- [28] K. I. Uchida, K. Sellgren and M. Werner, "Do the infrared emission features need ultraviolet excitation?" The Astrophysical J. **493**, L109-L112 (1998).
- [29] M. Watanabe and S. K. Nayar "Rational filters for passive depth from defocus," Int. J. Comp. Vis. **27**, 203-225 (1998).
- [30] L. B. Wolff and T. E. Boulton "Constraining object features using a polarization reflectance model," IEEE Trans. PAMI **13** 635-657 (1991).