

An Introduction to Data-Efficient Spectral Imaging

Mitchell R. Rosen, Mark D. Fairchild, and Noboru Ohta
Munsell Color Science Laboratory, RIT
Rochester, New York

Abstract

Until now, descriptions in the literature of capture systems for multi-channel visible-spectrum imaging (MVS) have been dominated by a pair of designs. The first of these designs is extremely general. It acquires images through a large number of filtered channels each with a narrow spectral bandpass. Disadvantages are associated with low light throughput per channel and a torrent of image data. A second popular approach offers higher light throughput and lower data demands because it depends upon a limited number of wide-bandwidth channels chosen with prior knowledge of scene contents. Unfortunately, this second approach is customized to scene contents so it is not a general solution. Hybrid approaches are introduced here that combine the generality of a many-narrow-channels solution with the lower data demands of a customized, wide-band system.

Tailored Versus General MVS Systems

Multi-channel imaging systems used for spectral reconstruction require the capture and storage of far more data than that needed for current RGB systems. Implementations have tended to fall into one of two camps. The most robust systems are those separating incoming radiation into narrow bins using interference filters¹ or tunable filters². The other approach to MVS has been using a small number of wide-bandwidth channels and associated reconstruction transforms customized to anticipated scene objects.^{3,4}

For any particular object within a scene there is a minimum set of channels that could be used for accurate spectral reconstruction of that object and an associated transform that would produce that reconstruction from the captured channels.⁵ The number of channels in this optimum set could be as few as two or could be many more, depending upon the specific spectral characteristics of the channels and the object. Figure 1 illustrates how such an optimally constrained system is constructed and works. An important observation is that a system outfitted for one object's optimally minimum set of filtered channels is not guaranteed to deliver any level of accuracy for any other object.

Although imaging applications are of primary interest to this discussion, it is interesting to note that there is a non-imaging application for which the approach of using a minimum number of channels to reconstruct reflectance or transmittance spectra has been well known for many

years. Within the photographic industry, status sets of densitometer filters have been specified such that when taking densitometry measurements on a specific type of film material, if the appropriate status set of filters is used, the material contents of a piece of film can be estimated.⁶ Knowledge of physical models of film systems has allowed photoscientists to accurately estimate the reflectance or transmittance spectra of a colored patch of film given the three filtered channel values that are returned as densitometer measurements.




Figure 1. Constructing and using a spectral imaging system designed for the optimally minimal set of filtered channels. 1) Materials from a typical scene measured using a spectrophotometer and analyzed. 2) System constructed. Based on that analysis as well as knowledge of the system optics, spectral sensitivity of the imager and spectral transmittances of available filters, the smallest set of filters that can be used to reconstruct spectra from anticipated scene materials is chosen. 3) The transform from channel signals back to spectra is derived, where dc_i is the digital count associated with the i^{th} channel. 4) Scene is captured using optimally minimal set of filtered channels. 5) Spectra is reconstructed pixel by pixel based on predetermined transform. See reference 7.

As mentioned above, such a highly customized approach requiring just the right set of filters for a particular film type is not a general solution. If one is going to make hundreds, thousands or millions of spectral

measurements on a new type of film, then perhaps it would make sense to design a new customized set of filters for a densitometer for this new film type. Thus, the three numbers returned from the densitometer combined with the physical model or *transform* will yield low error spectral reconstruction. But, if no large investment is warranted in measurement efficiency, then using a general solution makes sense. In this case, the general solution would be that of measuring with a spectrophotometer. Of course, that requires a complete set of values being taken, far more than three that would come from the densitometer.

Returning to the topic of imaging applications, an analogous situation is encountered as that which is seen in the conflict between the densitometer and the spectrophotometer for film measurements. When one wishes to image an arbitrary scene and reproduce the spectra, but there has been no pre-analysis, no customized system, a general solution is sought. A multi-channel imaging system with a large set of narrow-band filtered channels is such a general spectral imaging system. See Figure 2. When those filtered channels sufficiently sample across the visible spectrum, arbitrary imaged spectra can be easily reconstructed. On the other hand, data overload, speed of capture and light starvation become significant problems for these more general systems.

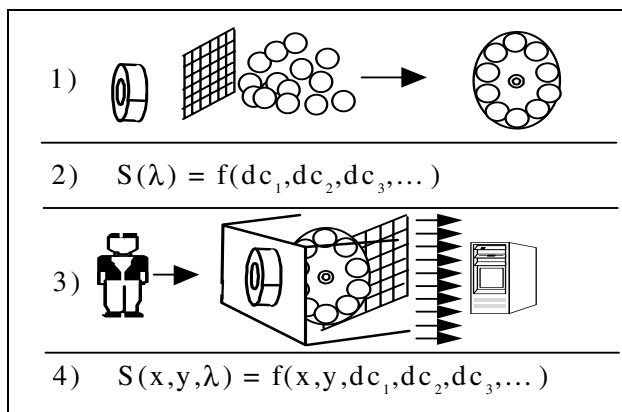


Figure 2. A general spectral imaging system. 1) System constructed. Based on knowledge of the system optics, spectral sensitivity of the imager and spectral transmittances of available filters, a large set of narrow-band filters that span and well sample the visible spectrum are chosen. 2) The transform from channel signals back to spectra is derived, where dc_i is the digital count associated with the i^{th} channel. 3) Scene is captured using full set of filtered channels. 4) Spectra is reconstructed pixel by pixel based on predetermined transform.

The tension between the two systems is now clear. One either lacks generality or must confront tremendous bandwidth, speed and sensitivity requirements. Since a highly tailored system will capture fewer data but still exhibit high accuracy when reconstructing spectra, it seems to follow that there must be a large quantity of redundancy within the data captured by the general

system. Redundancy in a data stream is the essence of what a compression scheme will seek out and exploit. It may seem that a good compression algorithm will solve the drawbacks of the general imaging solution. While compression can and should be used to reduce the eventual storage demands it only partially addresses the drawbacks described above. Depending on where bottlenecks arise, post-processing may not answer problems associated with speed of imaging and it certainly will not address problems of light starvation. Further, and perhaps most important, the imaging chain that includes this post-processing compression still needs at some point to collect the full high-spatial high-spectral resolution data stream from the capture system. See Figure 3.

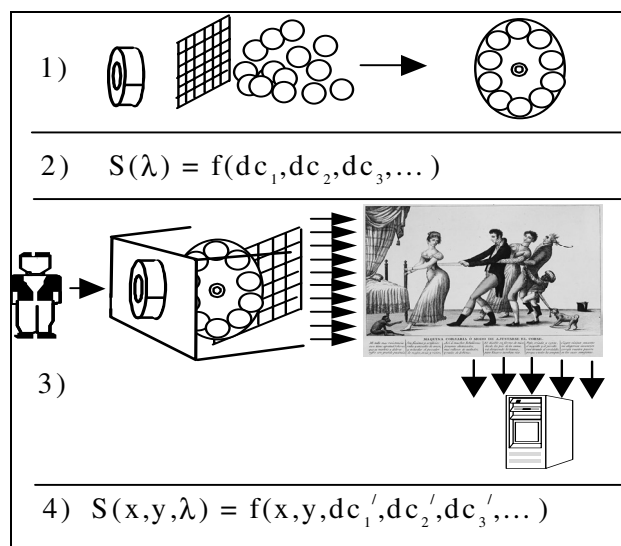


Figure 3. A general spectral imaging system with compression. 1) System constructed. Based on knowledge of the system optics, spectral sensitivity of the imager and spectral transmittances of available filters, a large set of narrow-band filters that span and well sample the visible spectrum are chosen. 2) The transform from channel signals back to spectra is derived, where dc_i is the digital count associated with the i^{th} channel. 3) Scene is captured using full set of filtered channels. Compression takes place in post-processing. It may take place on-board the camera or it may take place on an external computer prior to storage. 4) Digital counts are decompressed and spectra is reconstructed pixel by pixel based on predetermined transform, where dc'_i is the decompressed digital count associated with the i^{th} channel.

Avoiding high-spatial, high-spectral resolution becomes the goal of a data-efficient system. Compression after data capture may be helpful, but compression prior to data capture solves the drawbacks of the general system described above. It can be said that systems that use limited sets of filtered channels are examples of general systems with compression taking place prior to data capture. Unfortunately, as has been noted, such systems have applicability only to scenes for which they are optimized. The problem becomes a

question of how to perform pre-capture compression on a general system or, stated differently, how to configure a limited filter system in-the-field.

Hybrid Configurable MVSII Systems

It would be desirable to combine the best of the two approaches, the general with the highly tailored. If one could take a robust system such as the one described in Figure 2 and use it to automatically synthesize a tailored system such as the one described in Figure 1, the problem would be solved. By logically replacing Step 1 in Figure 1 with the system found in Figure 2, it is possible to design such a system. See Figure 4.

Conceptually, there are two subsystems in a hybrid system. The first subsystem consists of a general system. The second subsystem is a configurable tailored system. Some system designs might not physically separate the two subsystems. They may, in fact, work with the same optics, the same imager, maybe even the same filters. On the other hand, there may be very good reason to separate the two subsystems. In that case, it is likely that the general subsystem will be ancillary to the tailorable subsystem, and thus may be low in spatial and possibly even temporal resolution. It may also be low quality in other ways, because it will not be used to capture final images.

The example in Figure 4 anticipates that the two subsystems are physically separated. Step 1 represents the construction of the general subsystem. Note that in this example, the imager is relatively low in spatial resolution. Step 3 shows the general subsystem capturing portions of the scene. Here the small imager can be used to probe the entire scene with fat pixels or small areas in the scene with sharp pixels. Regardless of the fact that this subsystem may be low in resolution spatially and temporally, it has very high spectral resolution. Its purpose is to derive spectra or average spectra of the scene at-large or of important places in the scene. Once these spectra are derived after Step 4, an analysis takes place that determines what would be a qualifying set of filters from all filters available to the tailorable subsystem that can be used by the configured main subsystem to reconstruct spectra found by the general subsystem. Error tradeoffs and other considerations will be taken into account. A transform, f_g , from the chosen filtered channels to spectral estimates is derived and stored away.

Figure 4's Step 5 is the essential aspect of the data-efficient approach. It differentiates the hybrid system from the previous systems, showing Figure 1 and 2's systems to be static in relation to this system. Step 5 shows the major subsystem being configured. This should not be confused with Step 2 of Figure 1. In Figure 1, filters are chosen and the system is constructed. If new needs arise, a new system must be built. The system here, shown in Figure 4, is configurable. If new conditions are encountered, a new analysis takes place (Steps 3 and 4) and a new configuration is chosen. Once a set of filters is put in place, then Steps 6 through 8 mimic the final steps in Figure 1 for the reconstruction of scene spectra.

The system in Figure 4 will continually return to Step 3. This is but one of many ways in which a configurable, data-efficient system can be implemented.




Figure 4. A hybrid spectral imaging system. 1) General system constructed. Based on knowledge of the general system's optics, imager spectral sensitivity and available filters spectral transmittances, a large set of narrow-band filters that span and well sample the visible spectrum are chosen. 2) The transform, f_g , from general channel signals back to spectra is derived, where gdc_i is the digital count associated with the i^{th} channel for the general system. 3) Portions of the scene are captured using full set of filtered channels. Note that the imager is low in spatial resolution. 4) Spectra is reconstructed pixel by pixel based on predetermined transform, f_g and analyzed. 5) Tailored system configured. Based on the Step 4 analysis as well as knowledge of the tailored system's optics, imager spectral sensitivity and available filters spectral transmittances, a small set of filters that can be used to reconstruct spectra from scene portions is chosen. 6) The transform, f_t , from tailored channel signals back to spectra is derived, where tdc_i is the digital count associated with the i^{th} channel for the tailored system. 7) Scene is captured using optimally minimal set of filtered channels. 8) Spectra is reconstructed pixel by pixel based on predetermined transform, f_t .

Data-efficient Spectral Imaging – A Studio Camera Example

The needs of different applications call for different system implementations. For example, real-time processing demands would be relatively relaxed for a studio environment where unmoving objects were being spectrally imaged. Here, setup shots could be taken, analysis performed and the system placed into optimal configuration with potential small pauses being acceptable. Data storage constraints would likewise be relatively loose but requirements associated with spectral accuracy would likely be relatively high.




Figure 5. Three chip configurable studio spectral camera. A single capture system is used for both analysis and final image capture. Note that each imager has an identical mosaic of three filters on it. In between the camera optics and the imagers sit solid state tunable filters. These filters have very narrow bandpass and can be controlled to change their central passing wavelength. By switching through a series of filter settings over an extremely short time, the tunable filter can emulate any arbitrary filter. Thus, the system has great flexibility in choosing among almost any combination of nine filtered channels.

Unlike the system shown in Figure 4, the studio camera example described here does not have two physically separated imaging subsystems. This design entails a single image capture chain, used for both analysis and configuration. The camera is a three-chip device, but, unlike today's typical three-chip systems, each imager has a mosaic of three filters on it. See Figure 5. Each imager has identical filters on it. In addition, in series with each imager is a solid-state tunable filter. Temporal modulation of the narrow-band tunable filters are used to emulate wideband filters. While current tunable filter technology would make exposure times shorter than 0.5 sec difficult to achieve, it is reasonable to assume that faster tunable filters will be available in coming years. The emulated filters will further shape the three filtered channels on the imagers. Thus, by having each tunable filter emulate a different wide-band filter, a total of 9 filterings is available to the system.

As in the earlier example illustrated in Figure 4, the first steps involve analysis of the scene. Consider the child shown in Figure 6. When she first sits for her portrait, the system shown in Figure 5 can be used to perform a complete spectral analysis of her face, hair and clothes. Because there are three imagers and three tunable filters, the data capture for this analysis stage can be performed very quickly. Central pixels will be analyzed as-is, but to avoid data-overload, non-center pixels will be drastically subsampled. This will produce a spectral image with very fat pixels at the edges but high resolution pixels in the center. All low spatial frequency aspects of the scene will participate in the data analysis and high spatial frequency in the center will also be important to the results.

The optimal set of nine filters from those that can be configured will be chosen for portrait imaging. There are obviously restrictions to which nine filters may be chosen. This is because there is already a fixed relationship between the three filters on the chip. The tunable filter will attenuate their spectral shape.



Figure 6. Portrait shots that could be taken by the configurable studio spectral camera.

Between portraits (a) and (b) of Figure 6, there will be no reason to update the filters and transform between them. Objects within the scene are not changing spectrally. The sweater is changed by the time portrait (c) is taken, so a new analysis might take place. If reanalysis is a manual aspect of the system, it may be decided by the photographer that introducing spectral estimation error for the sweater might not be sufficient reason to update the filter choices. If reanalysis is automatically initiated, the system could use low-overhead methods to determine if a new analysis is warranted.

Although the studio camera captures only 9 channels of data, it is configured based on much higher spectral resolution. It, thus, delivers the same spectral quality of a general system consisting of 10's of channels.

Data-efficient Spectral Imaging – A Video Camera Example

A video application that captures 60 fields every second would have drastically increased demands on real-time analysis over the studio example. The system would require high speed update and constraints on data storage would be extremely tight. Much lower relative expectations would be acceptable for accuracy of spectral reconstruction.

For this example, it is reasonable to include two physically separated subsystems in the system design. The main subsystem can be constantly gathering video based on the most recent configuration update while the analysis subsystem determines the best configuration for constantly changing conditions. As objects come into view and go beyond view of the camera, configurations become inaccurate and need updating.

Figure 7 demonstrates such a video system. This system is not designed to deliver the low error on spectral reconstruction for which the studio camera was built. Here approximate results are good enough. A tunable filter in the analysis subsystem gives high spectral resolution. Spectral data are gathered from the central portion of the main system's view. Five degrees of freedom are afforded the system. It may choose any one of five filters on a filter wheel to place in the optical path of the main system imager. A mosaic of five additional filtered channels are found on the main imager. A transform that estimates spectra from the five doubly filtered channels is stored along with the main system digital counts.

The algorithm used to determine how to update the main system is biased toward not changing the current filter on the filter wheel. As long as quality is within a pre-determined tolerance, the filter stays fixed and only the transform will be updated. When errors go beyond the tolerance, a new filter is chosen.




Figure 7. Dual imaging system configurable video spectral camera. A general capture system (at top of figure) is used to spectrally analyze the scene. This subsystem consists of a solid state tunable filter in front of a monochrome imager. Low spatial resolution images are captured in sequence. This makes for a smear in the time-domain, but does not affect results very much since they tend to be relatively high in error anyway. The main subsystem consists of an imager with a five filter mosaic on it. There are five supplementary filters on a filter wheel that can be placed in the optical path. The analysis system chooses the best supplementary filter and the best transform to derive spectra from the five captured channels.




Figure 8. Analysis subsystem looking in the center of this scene will do the best job reconstructing grass and tree. When the children first emerge at the end of this scene, they will likely have low accuracy associated with their spectral reconstruction.




Figure 9. For this scene, the children are in the center of the camera's view and the analysis subsystem takes their spectra into account when determining the updated configuration and transform.

Data-efficient Imaging Spectrometers

A variety of analysis protocols could be used by an MVSI system to choose among available imaging configurations to find that which gives the best trade-off between spectral accuracy, speed of acquisition and data quantity constraints. Two configurable system examples were

discussed above, one with a large number of potential filtered channels and another that has a more limited number of potential channels. Each has use within a different application.

A spectral studio camera was considered. There, a single imaging chain was used in two ways. First, a low spatial resolution image was captured with high spectral resolution. Analysis was applied to determine the best set of nine wide-bandwidth filtered channels that would deliver highly accurate spectral estimates for the scene. The auto-synthesized filters were then used in-line with the same imaging setup, this time delivering high spatial resolution. The result is a data-efficient spectral capture of the scene where highly accurate spectral estimates could be reconstructed from the captured channels. The system configuration for this first example was appropriate for applications where spectral quality is of paramount concern such as when imaging portraits or fine-arts paintings.

The second simulated data-efficient spectral imaging system would more likely be found in a digital cinema application. Here a physically separated secondary imaging subsystem was used to capture a small central region of the imaged scene. Although this subsystem is spatially and temporally of low resolution, it has very high spectral resolution. A spectral error tolerance is constantly compared to see if the current filtration/transform configuration of the main imaging system is delivering sufficient spectral accuracy. When the error tolerance is no longer within range, a new transform is derived. If that is still out-of-tolerance, a new filter configuration will be chosen and the appropriate transform derived. This system will have more inherent spectral error than the one first described but will deliver spectral estimates within a pre-determined tolerable error while still adhering to strict speed requirements and limited storage budget.

Conclusions

It has been shown that the majority of current spectral image capture approaches come from two ends of the potential design spectrum: from brute-force data-hungry narrow-band to highly customized non-general wide-band. The concept of data-efficient designs based upon hybrid approaches has been introduced with two different system designs illustrated. While the systems do have high spectral resolution capabilities, they avoid ever taking data that are high in both spectral and spatial resolution at the same time. By posing solutions to many of the roadblocks for MVSI, these self-configurable systems may be the key to moving spectral imaging out of the laboratory and into commercial use.

Acknowledgements

Acknowledgement is given to Lawrence Taplin and Roy Berns who suggested the temporal modulation of tunable filters for emulating arbitrary filters. As usual, Francisco Imai is appreciated for his help. Thanks to Marissa and Peter Balonon-Rosen for appearing in Figures 6, 8 and 9.

References

1. F. Konig and W. Praefcke, The practice of Multispectral Image Acquisition, International Symposium on Electronic Image Capture and Publishing, 34 – 41 (1998).
2. F. Schmitt, H. Brettel and J. Hardeberg, Multispectral Imaging Development at ENST, International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives, 50 – 57 (1999).
3. F. Imai, N. Tsumura, H. Haneishi and Y. Miyake, Principal Component Analysis of Skin Color and Its Application to Colorimetric Color Reproduction on CRT Display and Hardcopy, JIST, 442-430 (1996).
4. F. Imai and R. Berns, Spectral Estimation Using Trichromatic Digital Cameras, International Symposium on Multispectral Imaging and Color Reproduction for Digital Archives, 42-49 (1999).
5. H. Haneishi, T. Hasegawa, N. Tsumura and Y. Miyake, Design of Color Filters for Recording Artworks, IS&T's 50th Annual Conference, 369-372 (1997).
6. P. Kowaliski, Sensitometry of Color Films and Papers, Chapter 18, T. James, ed., The Theory of the Photographic Process, 4th Edition, Macmillan Publishing Co., Inc, New York, 1977, pg. 525.
7. S. Quan, N. Ohta, R. Berns and N. Katoh, Optimal Design of Camera Spectral Sensitivity Functions Based on Practical Filter Components, Proceedings of The 9th Color Imaging Conference, 326-331 (2001).

Biography

Mitchell Rosen is a Senior Color Scientist with the RIT Munsell Color Science Laboratory. He received his BS in Computer Science from Tufts University and his MS in Imaging Science from RIT. He teaches color systems and color management courses in the Center for Imaging Science at RIT. He is also a candidate in the RIT Imaging Science doctoral program. From 1989 to 1998 he was a member of Polaroid's Image Science Laboratory. His research at the Munsell Lab grows out of recognition that widespread efforts to capture spectral images have introduced new challenges for image processing and new opportunities for improved quality in color reproduction. He is a member of IS&T.