

Experimental characterization of the quality of image reconstruction from a chromotomographic system

Kyle J. Dufaud, Michael R. Hawks, and Ryan Tervo

Department of Engineering Physics, Air Force Institute of Technology, Wright-Patterson AFB,
OH USA 45433

ABSTRACT

A fieldable hyperspectral chromotomographic imager has been developed at the Air Force Institute of Technology to refine component requirements for a space-based system. The imager uses a high speed visible band camera behind a direct-vision prism to simultaneously record two spatial dimensions and the spectral dimension. Capturing all three dimensions simultaneously allows for the hyperspectral imaging of transient events. The prism multiplexes the spectral and spatial information, so a tomographic reconstruction algorithm is required to separate hyperspectral channels. The fixed dispersion of the prism limits the available projections, leading to artifacts in the reconstruction which limit the image quality and spectrometric accuracy of the reconstructions. The amount of degradation is highly dependent on the content of the scene. Experiments were conducted to characterize the image and spectral quality as a function of spatial, spectral, and temporal complexity. We find that in general, image quality degrades as the source bandwidth increases. Spectra estimated from the reconstructed data cube are generally best for point-like sources, and can be highly inaccurate for extended scenes. In other words, the spatial accuracy varies inversely with the spectral width, and the spectral accuracy varies inversely with the spatial width. Experiment results also demonstrate the ability to reconstruct hyperspectral images from transient combustion events.

Keywords: hyperspectral imaging, tomography, tomosynthesis, chromotomography

1. INTRODUCTION

1.1 Background

Hyperspectral imaging (HSI) has proven to be an important tool for a variety of intelligence and remote sensing applications. Attempting to capture three dimensions of information (2 spatial, 1 spectral) on a two-dimensional focal plane array detector requires some tradeoffs. Conventional HSI systems solve this by removing information from the scene to capture two dimensions at a time. The full data cube is built up of many individual frames of camera data. For example, push-broom (spatial scanning) systems image the scene through a narrow slit so that each frame of data contains one dimension of spatial information which is then dispersed (typically using a grating) to spread the spectral information across the second dimension of the detector array. The slit is reoriented to capture different parts of the scene in successive frames. Alternatively, spectral scanning systems see the full spatial scene in each camera frame, but require multiple frames to develop spectral information through tunable filters or interferometers.

This creates issues for changing scenes. Conventional systems cannot handle transient or rapidly changing events well. Spectral scanning systems will see the event in at most a few bands, and spatial scanning systems might not even notice the event happened at all.

Chromotomographic systems are similar to slit-less spectrometers in that the entire scene is viewed simultaneously through a dispersive element—a prism in this case. The prism multiplexes the spatial and spectral information. Traditional slit-less spectrometers are therefore only useful for very simple scenes like point sources, where the structure of the scene prevents the spectra of different object points overlapping on the detector. In chromotomography the spectra of different points in an extended scene can overlap, but we rotate the prism and capture multiple frames which are used to reconstruct the data cube. While this method still requires

Send correspondence to M.R.H.: E-mail: michael.hawks@afit.edu, Telephone: 1-937-255-3636 x4828

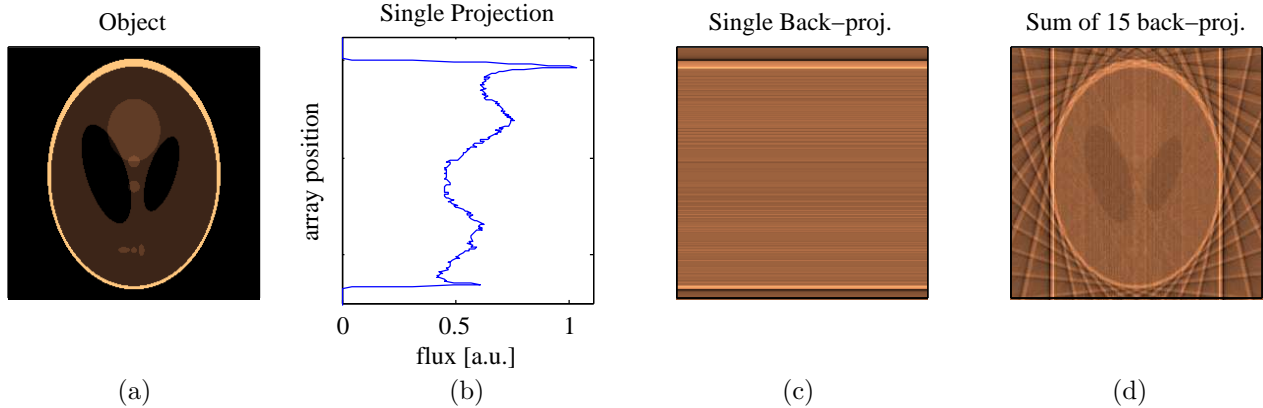


Figure 1. Example to illustrate projections and the back-projection algorithm. Here, the x-rays are assumed to be incident from the left. They pass through the 2-D object (a), resulting in a 1-D measured projection along the axis perpendicular to the ray direction (b). To reconstruct the image, this projection is back-projected by applying the profile uniformly across the entire reconstruction space (c) along an axis parallel to the original x-rays. As more back-projections at different angles are added, the reconstruction approximates the object (d).

multiple frames (and therefore time), the instrument captures all the data (all bands and all spatial points) in every frame. This improves situational awareness and provides an ability to measure transient events, as well as boosting instrument throughput.

1.2 Analog to Medical Tomography

Tomography reconstructs N -dimensional data structures based on a series of $(N-1)$ -dimensional projections. In medical tomography, we reconstruct three spatial dimensions from a collection of two-dimensional images. Okamoto¹ and Bulygin² were among the first to propose using the same idea to reconstruct a three-dimensional hyperspectral data cube from two-dimensional projections. The optical systems used to create these projections vary^{3,4} but the present instrument closely resembles the concept first described by Mooney.⁵

In medical tomography, the three-dimensional object being reconstructed is typically a body part which is held still between an X-ray source and imager. The imager records the total x-ray flux transmitted through the object—there is no depth information available in the projections. By looking at the object from multiple angles, a full three-dimensional picture is formed. The reconstruction relies on back-projection, which is illustrated in figure 1. Because the projection has lost all depth information, back-projection assumes the observed intensity could be due to objects located at any point along the line of x-ray travel, with equal probability. The observed intensity profile is therefore uniformly back-projected along that line. As back-projections from multiple angles are added, the shape of the original object is recovered.

In chromotomography, the process is very similar, but the object being reconstructed is a hyperspectral data cube. If we think of the object as a “stack” of monochromatic images, the camera records an image (one projection) that is the sum of all these monochromatic images, where each has been shifted by a distance proportional to its wavelength. This is equivalent to a projection through the data cube along an axis that is tilted by an angle ψ relative to the spectral dimension. For a more detailed discussion of this, see Brodzik⁶ or Hagen.⁷ As the prism is rotated about the optic axis, the projection direction sweeps out a cone of angle ψ .

This represents one of the fundamental limits of chromotomography. In medical tomography, the source/detector assembly is rotated in a full circle around the object. Because the projections are measured in a plane normal to the beam, the projection planes therefore fill in the 3-dimensional sample space completely (in the limit of infinite number of projections). Chromotomography on the other hand, has projections that describe a cone. Even in the limit of infinite projections, the central slice theorem⁸ tells us there is a portion of the space that is never sampled.




Figure 2. A schematic illustration of a CT imaging system. Lenses L_1 and L_2 collect and collimate light. The DVP disperses the colors, and L_3 re-images onto the FPA.

1.3 Objective

The primary goal of this project is to evaluate how the quality of the reconstructed data depends upon the content of the scene. In the limit of point-like sources, the instrument is expected to perform as a slitless spectrometer. As the amount of spatial and spectral complexity increases, we expect the back-projection artifacts to degrade the reconstruction. Of particular interest is the accuracy of the spectra, as previous studies in this area have emphasized reconstructed images. The second goal was to demonstrate the ability to recover hyperspectral data from a transient event.

2. EXPERIMENT

Figure 2 shows a simplified diagram of the CT optics. The lenses used here were all commercial multi-element photography lenses. The direct vision prism (DVP) is a type of Amici prism,⁹ meaning it is a dispersive prism that allows one wavelength to pass through undeviated. The DVP used for this project is made from two triangular prisms joined by optical contact bonding. Prism 1 is made of SF L6 and prism 2 is LaSF N30, constructed so that rays (from an on-axis object) enter the prism normal to the first surface and are incident 61 degrees from normal at the interface between the two materials. Because the indices of refraction are nearly equal at 547nm, that ray is undeviated.

The DVP is mounted in a hollow-shaft motor that causes it to rotate about the optic axis. An encoder measures the angle of the prism, ϕ . The Phantom v5 camera sends a trigger pulse at the start and stop of every integration time, and the experiment control computer logs the encoder angles at each time.¹⁰

2.1 Methodology

To study the effects of the scene content, we used 3 different spatial scenes, illuminated with 6 different light sources. The scenes were a 25 μm pinhole, the AF resolution target, and the cover of Hsieh’s text book⁸ on tomography. Spectral sources included Hg and Ne pen lamps, a quartz-tungsten-halogen lamp, and a variety of spectral filters used to modify the lamp spectrum.

To test the ability to capture transient events, we observed a transient fireball created by blowing a small “puff” of coffee creamer over a propane flame. The cloud of small particles created an irregularly shaped fireball that spread rapidly from the ignition source to a maximum diameter of approximately 1 m, then dissipated. The fireballs lasted about 1 second from start to finish. An Hg pen lamp was also placed within the field of view to have a clear reference source to aid reconstruction.

3. THEORY

3.1 Reconstruction

The nature of back-projection takes the measured projections and applies them uniformly along the line of projection. In CT, this line runs through the data cube at an angle that depends mainly on the prism dispersion. As we look at reconstructed images in a single wavelength, we see the information from all projections, each shifted by an amount determined by the wavelength and the prism dispersion. In other words, a monochromatic point source in the object produces a line through the reconstructed cube, which we would see as a point which shifts across the image as we look at different wavelengths.

The projection measured when the prism is rotated at angle ϕ is given by

$$g(x, y, \phi) = \int f(x - \delta x, y - \delta y, \lambda) d\lambda, \quad (1)$$

where $f(x, y, \lambda)$ represents the original object as viewed at wavelength λ , and $(\delta x, \delta y)$ represents the chromatic shift of that wavelength. The shift depends on the prism dispersion, $\theta(\lambda)$, and prism orientation, ϕ by

$$(\delta x, \delta y) = f_3 \theta(\lambda) (\cos \phi, \sin \phi), \quad (2)$$

where f_3 is the focal length of the imaging lens between the prism and detector. The effect of this shift can be viewed as convolution by a shifted point-spread function (PSF). If we take the 2-D (spatial) Fourier transform of both sides, the convolution becomes a product and the shift becomes a phase term, or

$$G(u, v, \phi) = \int F(u, v, \lambda) \exp[-2\pi i(u \delta x(\lambda) + v \delta y(\lambda))] d\lambda, \quad (3)$$

where we have used upper case letters to represent the transformed variables and (u, v) for the spatial frequencies. To reconstruct an estimate $f'(x, y, \lambda)$ of the original hypercube, we begin by shifting the recorded data back by a distance proportional to the wavelength and summing over all angles (projections). In fourier space, this means we simply apply the opposite phase,

$$F'(u, v, \lambda) = \sum_{\phi} G(u, v, \phi) \exp[2\pi i(u \delta x(\lambda) + v \delta y(\lambda))] d\phi. \quad (4)$$

Transforming returns our estimated hypercube,

$$f'(x, y, \lambda) = \sum_{\phi} g(x + \delta x(\lambda), y + \delta y(\lambda), \phi). \quad (5)$$

The shifts can also be done in (x, y) space without transforming by simply shifting the images. Because the data is discretely sampled, we often end up with shift distances that are not an integer number of pixels, so the space-domain shifting requires some finesse with sampling and interpolation. The Fourier-space phase shifts are simpler to implement.

Equation 4 implies information from all projections will be present at every wavelength in the reconstruction, but its position in the image will depend on the wavelength. In practice, we reconstruct a single image, $f'(x, y, \lambda_r)$ at a time. Each projection is shifted by a distance $-f_3 \theta(\lambda_r)$ in a direction determined by the prism orientation when the projection (image) was captured. For an object point emitting light at wavelength $\lambda_o = \lambda_r$, this will shift all the projections so that all images of that object overlap at one point. The brightness at this point is then multiplied by the number of projections, so that the “correct” reconstructions should be much brighter than the artifacts. When $\lambda_o \neq \lambda_r$, the back-projection of that point will trace out a circle of radius

$$r(\Delta\lambda) = f_3 (|\theta(\lambda_r) - \theta(\lambda_o)|) \quad (6)$$

in the image reconstructed at λ_r . An example of this is shown in figure 3. We refer to these circles as off-wavelength artifacts. Assessing the impact of these artifacts is the main focus of the present work.

There are a number of more sophisticated reconstruction algorithms (e.g. Brodzik⁶), but they all use back-projection as the basis. The method of back-projection then filtering¹¹ was also tested in this project, but few results of that will be presented here for brevity. The filtering method requires the user to select a filter parameter, μ , but there is no established method for selecting intelligent values. Picking μ by trial-and-error can, in some cases, lead to real improvements in the reconstruction, but the improvements we see are generally small compared to the issues we will present below, so the full side-by-side comparison is not discussed here.



Figure 3. Reconstructions at 4 wavelengths to illustrate off-wavelength artifacts. The target is a pinhole in front of a white light source with a 20nm bandpass filter centered at 540nm. The wavelengths of reconstruction are (from left to right) 500,520,540, and 575 nm. The in-band reconstruction shows an image of the pinhole, but out-of band reconstructions create circles of radius $f_3 (|\theta(\lambda_r) - \theta(\lambda_o)|)$. The dimmer (off-center) circular artifacts are the result of uncorrected hot pixels in the camera.




Figure 4. Hg spectra as estimated from a single pixel in the CT reconstruction of a pinhole (thick line) and an AF resolution target (dashed line), compared to the truth spectrum as measured with a grating spectrometer (thin black line). The off-wavelength image artifacts caused by the spatial extent of the resolution target put significant energy onto this pixel, corrupting the spectral data.

3.2 Spectral Estimation

In a conventional hyperspectral data cube, one can inspect the spectrum at some point (x_i, y_i) in the object simply by plotting the estimated intensity $f'(x_i, y_i, \lambda)$ for all λ . In chromotomographic imaging, the off-wavelength artifacts resulting from back projection can, in many cases, contribute significant energy that corrupts the direct spectral estimation. An example of this is shown in Figure 4. This problem is generally minor for point sources, but as figure 4 shows this can be a significant problem for an extended source. Consider an object of radius ρ centered at a point (x_o, y_o) . According to equation 6, we would expect the light from that source to contribute energy to the spectrum of pixel (x_o, y_o) whenever $r(\Delta\lambda) < \rho$. That means larger objects create wider spectral lines. Figure 5 shows the anticipated spectral linewidth for objects of various sizes. The shape of this curve will depend on the prism and the focal lengths used, but the overall trend would apply to all CT systems.

Because of these artifacts, a different method for estimating the source spectrum is often needed for extended targets. For the targets measured in this experiment, we note that images are generally clear and sharp when reconstructed at wavelengths that are present in the source, but blurred at other wavelengths. Plotting image quality as a function of wavelength often produces a closer estimate of the source spectrum. We refer to these plots as pseudo-spectra. Figure 6 shows a comparison of several different pseudo-spectrum estimation methods. The edge and gradient techniques look for sharp edges in the image, and so are only useful for some scenes. The only difference between them is that “gradient” averages the image gradient over pixels in a user-selected region, while “edge” counts the number of pixels where the gradient exceeds a given threshold. The cross-correlation




Figure 5. The expected spectral line width predicted for objects of various size. The boxed numbers represent the object size (in pixels on the detector).




Figure 6. A comparison of spectral estimation methods from the reconstruction of the textbook cover illuminated by an Hg pen lamp. The figure on the left is for back-projection and the figure on the right is back-projection then filtering. The thick black line is the conventional spectrum, $f(x_i, y_j, \lambda)$ averaged over a 2×2 pixel region. The other lines are pseudo-spectra from the edge (solid green), gradient (dashed green), MSE (solid blue) and cross-correlation (dashed blue) methods. All have been independently normalized to unit peak.

and mean-squared-error (MSE) methods both compare the full image against a truth image measured by the same system with the prism removed. They are therefore not useful for true hyperspectral imaging, since the entire data cube is required to generate a single spectrum. The cross-correlation metric plots the peak value of the cross-correlation of the reconstructed image with the true image.

4. RESULTS

This section presents examples of the resulting image reconstructions and spectral estimates for several scenes. Examples are chosen to illustrate specific key points, but a more exhaustive presentation of results is available in source.¹²

4.1 Image Quality

In general, we find image quality for in-band reconstructions to be good, though there is some dependence on the spectral content of the source. The off-wavelength artifacts often create out-of-band images the look like out-of-focus images. Figure 7 shows an example of this. Figure 8 shows the modulation transfer function (MTF) of the system as calculated from images of the AF resolution target illuminated with an Hg lamp. This was calculated twice, once from raw images taken with the DVP removed from the system, and again from CT reconstructed




Figure 7. Reconstructed images of the AF Resolution Target (panel a) illuminated with an Hg-pen lamp. For reconstructions on a strong Hg wavelength (b), the reconstructed image is clear. Ideally, reconstructions at other wavelengths (c,d) should be blank, but in reality we see off-wavelength artifacts. Filtering can help this, but it is imperfect.




Figure 8. The MTF of the system as measured with the DVP removed (black squares) and from CT reconstructions (red circles). MTF is estimated from observed contrast in AF resolution target illuminated by a Hg pen lamp.

images at 546 nm. One would expect the CT reconstructions to have a lower MTF than undispersed images, but this data shows that is not always the case. This may be due to the fact that noise is averaged over many frames in the reconstruction. The lenses used in this experiment showed significant chromatic aberrations, so the increase may also be due to the prism shifting the more heavily aberrated 435 nm light away from the in-focus 546nm light used for MTF estimation. To avoid the effects of uneven illumination and vignetting from corrupting the contrast measurements, the target was shifted between collections to center a different group of bars (different spatial frequency) in the field of view. In other words, each point on the graph represents a separate reconstruction (or image, for the squares) considering a single bar group.

Figure 9 shows a comparison of in-band images of the resolution target illuminated with different light sources. The sources used to illuminate the target are listed in table 1. Here one can see that the spectral content of the scene affects image quality. All reconstructions shown here are for the wavelength at the peak of the spectrum (or band center if it is not peaked). There is a clear trend of decreasing image quality with increasing bandwidth. The Neon lamp spectrum is composed of very narrow emission lines, but the red end of the spectrum contains many lines that are not resolvable or barely resolvable with this instrument. It therefore behaves like a broadband source for these purposes. The overhead fluorescent light gives a much better reconstruction than the white light in spite of having similar bandwidth, because a large fraction of the energy is in a few narrow, bright peaks. The long-pass filter, combined with the response cutoff of the detector, is roughly equivalent to an ± 40 nm band pass.




Figure 9. Reconstructed images of the AF resolution target illuminated by 8 different light sources, to show dependence of image quality on spectrum. Sources are listed in Table 1

Table 1. Sources used to illuminate target in Fig 9. Filters (rows d-g) were used with the white light source.

Panel	Source	Reconstruction λ shown [nm]
(a)	Hg pen lamp	546.0
(b)	Ne pen lamp	703.0
(c)	Overhead fluorescent lamp	546.0
(d)	Band pass filter (640 ± 6 nm)	642.0
(e)	Band pass filter (540 ± 10 nm)	540.0
(f)	Band pass filter (525 ± 25 nm)	535.0
(g)	Long pass filter (665 nm)	680.0
(h)	white light (Fiber-Lite)	600.0




Figure 10. The pseudo-spectrum of the resolution target, illuminated with an Hg lamp. The spectrum is estimated using the gradient method. The central peak correlates to the bright 546 nm emission line, but the smaller side-lobes show areas of false contrast where the shift error matches the spacing between bars.




Figure 11. Left: Estimated spectra of the 641 ± 6 nm light source based on reconstructions of a pinhole (thin black line), the AF resolution target (blue lines) and the text book cover (red lines). The solid lines are spectra from large, bright regions and the dotted lines are from smaller bright features. The thick black line is the known filter transmission. Center and Right: single reconstructed images of the book and resolution target showing the areas used to make the spectra at left. All spectra are based on 2×2 pixel averages.

4.2 Spectral Accuracy

As has been previously discussed, the nature of back-projection means that energy from all wavelengths of the source will be present in all wavelength of the reconstruction. The pseudo-spectrum estimation methods described in section 3.2 are an attempt to avoid this. Figure 10 shows an example of another type of artifact. This plot shows the pseudo-spectrum of an Hg lamp as estimated from the CT reconstruction of an AF resolution target. This plot was made using the gradient method, which uses image quality as a proxy for brightness in estimating a spectrum. The large peak at 546 nm shows an actual spectral peak and corresponds to a very sharp, clear image at that wavelength. This occurs when the shifted back-projections line up perfectly. The smaller side-lobes in the spectrum actually show regions of false contrast where the shift error, $r(\Delta\lambda)$, matches the distance between bars in the target region used to compute the gradient metric. This effect turns the three-bar pattern of the object into an image with four bars, which can be seen in panel (c) of figure 7 in the largest (lowest spatial frequency) bars.

In general, spectral errors can be attributed to the off-wavelength artifacts. This is demonstrated in figure 11, which shows several spectra of the same source, estimated from reconstructed hypercubes of different targets. The pinhole gives the most accurate spectrum, as expected. Figure 11 shows a correlation between the spatial extent of a structure in the scene and the width of the spectrum estimated from pixels near that structure, as discussed in section 3.2.




Figure 12. Left: Reconstructed image from the transient fireball, reconstructed at 770 nm displayed in false color (where blue = darkest to red = brightest). The two separate fireballs are clearly resolved spatially. A fainter "halo" of vertical bars also demonstrates the off-wavelength artifacts from the 546 nm line in the Hg lamp. Right: the direct spectrum averaged over a 2x2 pixel region in the smaller fireball (dashed line) and a pseudo-spectrum using the edge method (thin line), compared to the truth spectrum (thick line) measured with a grating spectrometer.

4.3 Transient

The reconstruction results of the transient fireball are shown in 12. The transient event reconstructed here showed two separated fireballs, likely due to the coffee creamer clumping in the nozzle. The larger fireball ignited first, and is near its peak diameter in this reconstruction. As momentum carried it past the propane flame, a second puff ignited behind it. The spectra shown in 12 are from the smaller of the two fireballs. The truth spectrum, measured with a grating spectrometer shows very distinct emission lines from Potassium and Sodium, as well as a dimmer broadband thermal emission. The atomic emission lines can also be seen in the reconstructed spectra, but have been significantly broadened by the off-wavelength artifacts. There is also some error in the locations of the peaks in this data. All spectra here were independently normalized to unit peak.

5. CONCLUSIONS

Deriving performance metrics for a CT imaging system is a difficult and complex question because the quality of the reconstruction depends on the content of the scene being imaged. System performance is best for point-like sources on a dark background. For extended scenes, the spatial extent leads to overlapping artifacts in the reconstruction that can corrupt spectral estimation. Images reconstructed at wavelength corresponding to spectral peaks are generally clear, but images can be degraded in broadband sources. Spectral estimation of point-like sources is quite good, but larger objects can cause significant degradation in the spectrum. Pseudo-spectra constructed by plotting image quality as a function of wavelength appear to better represent the source spectrum in cases examined here.

The system used here was successful in reconstructing images and spectra of a transient combustion fireball, but the spectral resolution and accuracy were both poor. Improvements to the optics (to eliminate chromatic aberration) and more advanced reconstruction algorithms will be explored in future work.

REFERENCES

- [1] Okamoto, T. and Yamaguchi, I., "Simultaneous acquisition of spectral image information," *Opt Lett* **16**(16), 1277–1279 (1991).
- [2] Bulygin, F., Vishnyakov, G., Levin, G., and Karpukhin, D., "Spectrotomography—a new method of obtaining spectrograms of 2-d objects," *Opt Spectrosc (USSR)* **71**(6), 561–563 (1991).
- [3] Descour, M. and Dereniak, E., "Nonscanning no-moving-parts imaging spectrometer," *Proc of the SPIE* **2480**, 65–77 (1995).

- [4] Bernhardt, P. A., “Direct reconstruction methods for hyperspectral imaging with rotational spectrotomography,” *J. Opt. Soc. Am. A* **12**, 1884–1901 (Sep 1995).
- [5] Mooney, J. T., “Angularly multiplexed spectral imager,” *Proc of the SPIE* **2480**, 65–77 (1995).
- [6] Brodzik, A. and Mooney, J., “Convex projections algorithm for restoration of limited-angle chromotomographic images,” *J. Opt. Soc. Am. A* **16**(2), 246–257 (February, 1999).
- [7] Hagen, N. and Deriniak, E., “Analysis of computed tomographic imaging spectrometers. i. spatial and spectral resolution,” *Applied Optics* **47**(28), F85–F95 (2008).
- [8] Hsieh, J., [*Computed Tomography*], SPIE Press, Bellingham, Washington (2007).
- [9] Hagen, N. and Tkaczyk, T., “Compound prism design principles i,” *App Opt* **50**(25), 4998–5011 (2011).
- [10] Su’e, C. B., *Characterization of a Hyperspectral Chromotographic Imaging Ground System*, Master’s thesis, Air Force Institute of Technology (2012).
- [11] Deming, R. W., “Chromotomography for a rotating-prism instrument using backprojection, then filtering,” *Opt. Lett.* **31**, 2281–2283 (Aug 2006).
- [12] Dufaud, K., *An experimental evaluation of image quality for various scenarios in a chromotomographic system with a spinning prism*, Master’s thesis, Air Force Institute of Technology (2014).