

Analysis of computed tomographic imaging spectrometers. I. Spatial and spectral resolution

Nathan Hagen^{1,*} and Eustace L. Dereniak²

¹Fitzpatrick Institute for Photonics, Duke University, Durham, North Carolina 27708, USA

²College of Optical Sciences, University of Arizona, Tucson, Arizona 85721, USA

*Corresponding author: nhagen@optics.arizona.edu

Received 7 March 2008; revised 22 April 2008; accepted 30 May 2008;
posted 9 June 2008 (Doc. ID 93420); published 25 July 2008

Computed tomographic imaging spectrometers measure the spectrally resolved image of an object scene in an entirely different manner from traditional whisk-broom or push-broom systems, and thus their noise behavior and data artifacts are unfamiliar. We review computed tomographic imaging spectrometry (CTIS) measurement systems and analyze their performance, with the aim of providing a vocabulary for discussing resolution in CTIS instruments, by illustrating the artifacts present in their reconstructed data and contributing a rule-of-thumb measure of their spectral resolution. We also show how the data reconstruction speed can be improved, at no cost in reconstruction quality, by ignoring redundant projections within the measured raw images. © 2008 Optical Society of America

OCIS codes: 110.4234, 300.6190, 110.3000, 120.6200.

1. Introduction

Computed tomographic imaging spectrometry (CTIS) is a technique by which a scene is viewed directly through a 2D grating (as in Fig. 1), producing a number of prismatically dispersed images of the object on the focal plane. This is much like placing a pair of crossed gratings in front of one's eyes (as in Fig. 2), except that the CTIS imaging sensor is typically not spectrally selective like the human eye. Inserting a field stop in the system limits the instrument's field of view and allows the prismatic images to be spatially separated (producing data like those shown in Fig. 3).

Looking at any one of these prismatic images, we see that the data are what we would get with a slit spectrometer instrument if the slit were opened very wide (and the grating tilted to the appropriate angle), allowing the object's spatial and spectral information to overlap. If only one diffraction order is measured, the result is a disastrous loss of spectral

resolution, but if we design the system to measure multiple diffraction orders, then we can use the non-redundant information present within them to simultaneously recover the spectral resolution and obtain the spatial distribution as well. Thus, the complete spectrally resolved image (the data cube) can be obtained in a snapshot. This process of estimating a 3D distribution from a set of 2D images is a familiar problem from computed tomography, from which we recognize that each of the prismatically dispersed images is a projection of the data cube, as illustrated in Fig. 3.

The historical development of the CTIS disperser element has been to find ways of steadily increasing the number of diffraction orders (going from left to right in Fig. 4). Descour's original design involved using sets of cosine gratings, such as the 3×3 pattern produced by a pair of cosine gratings crossed at 90° [1], or the hexagonal pattern produced by crossing three cosine gratings at 60° to one another (Fig. 4, first and second panels) [2]. Descour *et al.* expanded the range of dispersion patterns available by designing efficient computer-generated holographic

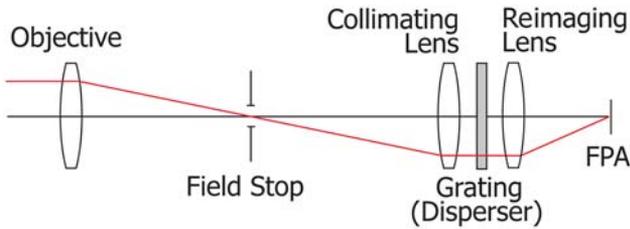


Fig. 1. (Color online) CTIS instrument views the target scene through a 2D grating. The field stop limits the field of view, such that the dispersed diffraction orders are spatially separated.

dispersers capable of producing diffraction orders in both the x and y directions, creating the 5×5 and 7×7 patterns shown in Fig. 4, third and fourth panels [3]. This substantially increased the number of projections measurable by CTIS, packing them all within the same physical dimensions of the focal plane array. CTIS instruments in current use typically work with the 5×5 diffraction pattern on the focal plane array (Fig. 4, third panel), so that they sample the data cube with 25 projections. Additionally, computer-generated holographic design techniques are flexible enough to channel light away from the zeroth order (where the detector array is prone to saturate) and into the dispersed orders, which are otherwise light starved by comparison. This prevents a loss of dynamic range in the measured data.

In a given CTIS design, there is a simple trade-off between the spatial resolution and the spectral resolution. In terms of sampling rate, the former is given by the size, in pixels, of the field stop image on the focal plane array (FPA)—the white square at the center of Fig. 3 (left). The latter is given by the length of the longest projection, minus the size of the field stop. (The actual spectral resolution of CTIS instrument is a complicated issue, which we discuss at length in Section 6 below.) For a 2048×2048 FPA, a typical CTIS design will image the field stop to a 100×100 region, giving a projection along an outer diagonal that is 626 pixels long. The above trade-



Fig. 2. (Color online) Nondescript scene taken with an RGB camera and a pair of crossed gratings inserted in front of the lens.

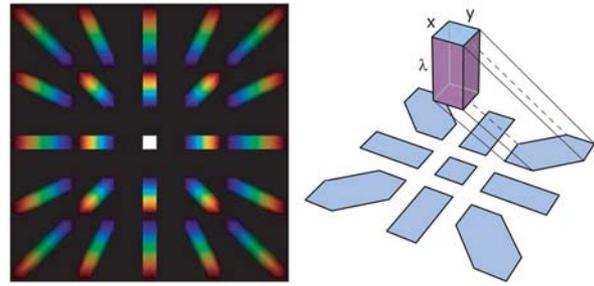


Fig. 3. (Color online) The 5×5 CTIS design produces a pattern of 25 projections onto the detector array, each of which is a prismatic version of the zero-order image. Each of these prismatic images produces a footprint on the detector array that is a tomographic projection through the object scene's 3D data cube.

off between spatial and spectral resolution can be expressed approximately by the relation (see Fig. 5)

$$L \approx \frac{3\sqrt{2}}{14}(Z - D), \quad (1)$$

where L is the maximum possible spectral sampling, Z is the length in pixels of the FPA (e.g., 2048), and D is the spatial sampling (e.g., 100).

2. Fourier Slice Representations

An important means of visualizing the characteristics of a CTIS instrument and of comparing various dispersion patterns is the central slice theorem ([4], p. 205). This theorem states that for a 2D projection of a 3D volume, the Fourier transform of the projection is equivalent to evaluating the Fourier transform of the 3D volume along a plane passing through the origin in Fourier space. (The orientation of the evaluation plane is determined by the angle of projection.) This result allows us to draw a plane within the “Fourier cube” (i.e., the Fourier transform representation of the data cube) corresponding to the frequencies sampled by a given projection. Doing this for all of the projections of a dispersion pattern gives a Fourier slice diagram representation of that particular CTIS design. While the Fourier slice theorem is strictly true only for continuous mappings, it remains a useful approximation for sufficiently sampled discrete systems.

Note that the spectral dimension of the data cube is of course not a physical dimension, and so its physical size can be defined arbitrarily. The choice made here is to define the data cube as being a perfect cube

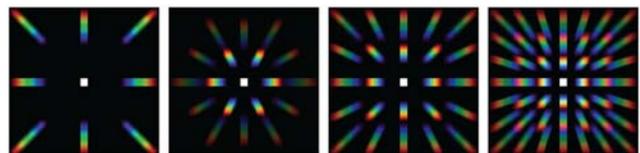


Fig. 4. (Color online) Past CTIS dispersion patterns: from left to right, 3×3 pattern produced by a pair of crossed cosine gratings, 3XCos pattern produced by three cosine gratings oriented at 60° to one another, 5×5 pattern, 7×7 pattern.

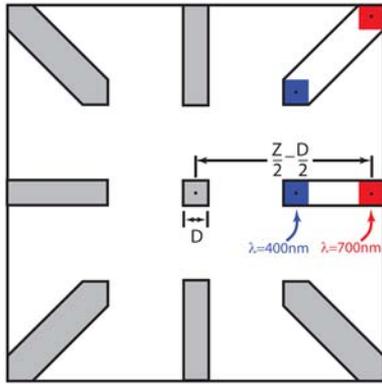


Fig. 5. (Color online) Illustration of the spatial–spectral sampling trade-off in a given CTIS design. For a focal plane array of dimension Z and a data cube dimension of D (this latter giving the spatial sampling), the longest wavelength is diffracted to the edge of the array: a distance of $(Z/2) - (D/2)$ in pixel units. If we assume a spectral range of 400–700 nm, then the shortest wavelength image is diffracted by 4/7 of the same distance. The spectral dispersion distance is therefore $(3/7)[(Z - D)/2]$. Along an outer diagonal order, this distance is increased by $\sqrt{2}$

in size (the physical dimensions of the cube along x , y , and λ axes are identical). Typically, CTIS samples the spectrum at a greater rate than for the spatial information of the scene, so that an individual voxel takes on a flattened pillbox appearance by this definition.

The Fourier slice diagrams for the various CTIS dispersion patterns are shown in Fig. 6. Evident in each of them is a large polygonal double-cone region, called the “missing cone,” which is not sampled by the CTIS system. (For the corner designs, this cone is extended such that two octants in Fourier space are entirely unsampled.) Naturally, the shape of these regions affects the reconstruction quality by altering the shape and size of the system blurring kernel.

3. CTIS Voxel Spread Function

The presence of a missing cone in the Fourier cube indicates that the voxel spread function (VSF) of a CTIS system will have peculiar characteristics. Here the VSF is described as the image of a point source (that is, one pointlike both spatially and spectrally) not on the CTIS FPA but in the reconstructed data cube—the VSF is thus a 3D analog of the 2D point spread function.

An instrument that has characteristics largely similar to CTIS instruments, but which does not operate in snapshot mode, is a chromotomographic hyperspectral imaging spectrometer (CTHIS) [5,6]. The CTHIS is capable of measuring as many projections as desired, a quality that can be used to advantage in analysis here, since a well-sampled CTHIS measurement has a missing cone that is a true right-circular double-cone [Fig. 7(a)]. In the CTHIS Fourier cube, any planar slice containing the h axis (the reciprocal of the λ axis) will contain four triangular regions, containing sampled and unsampled regions [Fig. 7(b)]. Taking the Fourier transform of the slice gives the

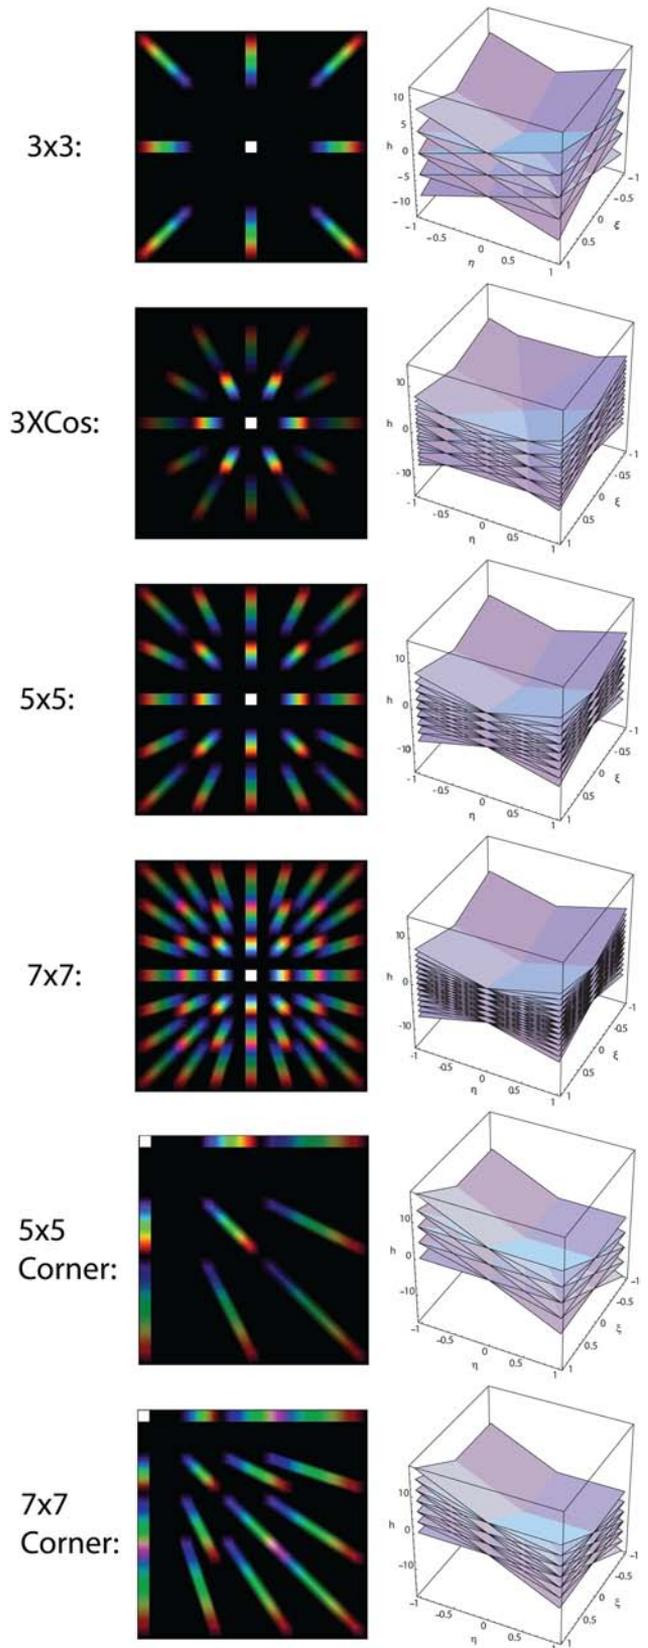


Fig. 6. (Color online) Left, various CTIS dispersion patterns and, right, their corresponding Fourier slice diagrams.

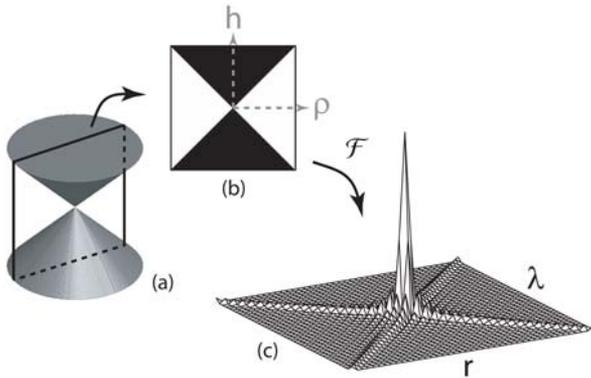


Fig. 7. (Color online) (a) For a CTHIS instrument taking a large number of projections, the missing cone becomes a true right-circular double cone. (b) In the corresponding Fourier cube, any planar slice containing the h axis will have the form shown. (Black indicates the missing cone region; white indicates the sampled regions.) (c) Corresponding VSF. Here (ρ, h) and (x, λ) are Fourier conjugate coordinates.

VSF in terms of cylindrical radial coordinates $\mathbf{r} = (x, y)$ and the λ coordinate [Fig. 7(c)].

Figure 7(c) indicates that any pointlike object creates a signal within the data cube that is a maximum at the appropriate voxel, but which also produces weak but prolonged tails in the form of a double-cone (the shape shown in Fig. 7(c) rotated about the λ axis). The theoretical treatment of the missing cone provides a useful perspective, but the small number of projections available in a snapshot CTIS instrument means that the continuous approach gives only a rough approximation to the true VSF. Examining data cubes reconstructed from simulated and experimentally fabricated point sources gives an empirical measure of the true 3D VSF. Figure 8 shows several spectra in the VSF, determined by imaging a point source in the center of the field of view (here the field of view is mapped to a 25×25 pixel region on the CTIS FPA). The displayed spectra are taken from various spatial positions within the data cube, progressively farther from the point-source image. These show a close correspondence to the VSF illustrated in Fig. 7(c).

Figure 9 illustrates the same data set from another perspective, showing the horizontal slices through the VSF at various wavelength locations. It is here that we see the discrete number of projections of a CTIS instrument (relative to that of an ideal CTHIS measurement) significantly alter the shape of the VSF, and we begin to see a 5×5 pattern reminiscent of the dispersion pattern used. This becomes even more apparent when we look at the backprojected point-source image, shown in Fig. 10. The gray-scale map of the backprojected images indicates that the tails of the backprojection VSF are much stronger than in the final reconstruction. The reconstruction algorithm has thus filtered the backprojection, greatly improving the fidelity of the estimated data cube.

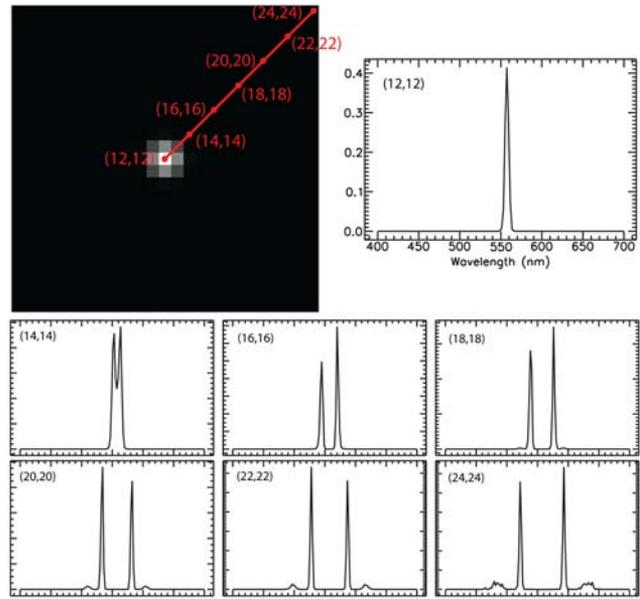


Fig. 8. (Color online) Spectra from different spatial positions in the reconstructed $25 \times 25 \times 145$ data cube for a point source with $\lambda = 557$ nm. The locations sampled in the data cube are indicated. Note that each of the spectra shares the same abscissa, ranging from 400 to 700 nm, and that the vertical axes (in arbitrary irradiance units) have a different scale in each plot: as we move further out from the center of the data cube, the peaks show a drop in amplitude by about a factor of 10 from one plot to the next.

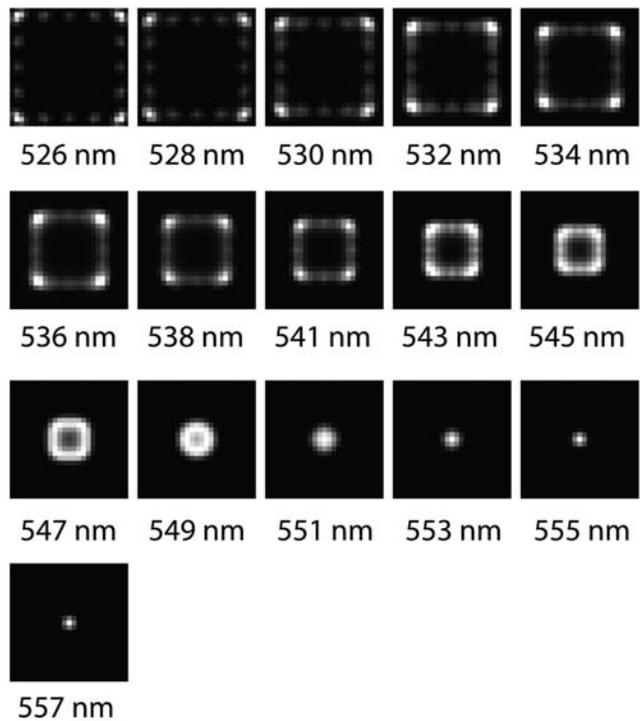


Fig. 9. Several planes shown from the reconstructed data cube for a point source at $\lambda = 557$ nm. The gray-scale mapping on these figures is logarithmic to accentuate the low-irradiance pixels. Each image is also scaled to have its brightest pixel map to white.

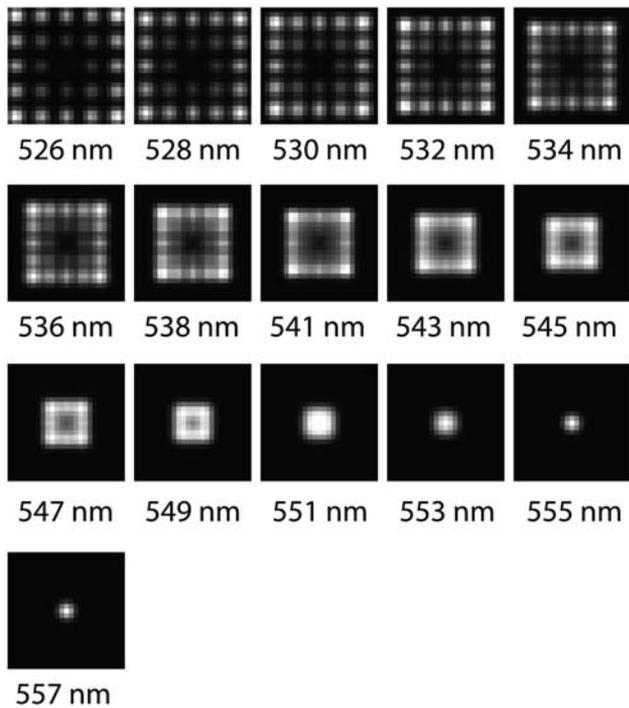


Fig. 10. Several planes shown from the backprojected data cube for a point source at $\lambda = 557$ nm. The gray-scale mapping on these figures is linear. Each image is also scaled to have its brightest pixel map to white.

Figures 8–10 all show simulated data; corresponding plots from experimental data are quite similar, but the presence of noise hides some of the salient features.

4. First-Order Metrics for Comparing the Various Patterns

Having a set of CTIS dispersion patterns, we would like to have a way to characterize each CTIS design in terms of what we want: spectral and spatial resolution. Resolution by itself is a vague quantity until we define how we plan to measure it. In the case of CTIS, it is also an object-dependent quantity, as we will see below. The first step toward defining what we are looking for is to create a set of metrics that we know correlate to desirable characteristics in a CTIS instrument. These provide a rapid and simple means of comparing designs.

There are at least six important metrics commonly used for characterizing the various CTIS designs.

1. *Spatial dimension, D* : the number of pixels in x , and also in y , in the data cube if we assume that the data cube is square. [This definition coincides with the D used in Eq. (1).]

2. *Spectral dimension, L* : the number of spectral samples as expressed by Eq. (1). If spectral resolution is particularly important for a given CTIS system, then the number of spectral samples is a metric closely related to the realizable spectral resolution of the system.

3. *Number of projections, P* : increasing the number of projections improves reconstruction quality. (Medical imaging systems typically use hundreds of projection angles for performing reconstructions.)

4. *Fill factor*: The ratio of illuminated to unused pixels on the FPA. To maximize its ability to measure the object data cube, a good CTIS design will maximize the use of space on the FPA (and minimize the number of unused pixels).

5. *Blur factor*: the average number of voxels mapping to a given illuminated pixel on the CTIS FPA. This is a rough measure of the amount of blurring incurred when performing a backprojection operation (an essential step in the reconstruction algorithms). Increased blurring of spatial and spectral components of the data cube reduces the theoretically achievable spatial and spectral resolution.

6. *Missing cone angle, ϕ* : the half-angle of the missing cone inside the Fourier cube. While this metric is dependent on both the spectral sampling L and the spatial sampling D , the missing cone angle provides another convenient way of comparing CTIS designs.

Table 1 presents a set of CTIS dispersion patterns, together with their respective design parameters and metrics.

5. Simulation Results

A CTIS simulator was developed to provide a more objective means of comparing and assessing CTIS

Table 1. Comparison of Linear CTIS Patterns^a

Design Type	Design Parameters		Design Metrics			
	D	P	L	Fill	Blur	ϕ
Full patterns						
3Xcos	75	19	558	0.18	111	7.66
5 × 5	75	25	570	0.24	81	7.50
7 × 7	75	49	570	0.42	70	7.50
3Xcos	100	19	550	0.25	141	10.28
5 × 5	100	25	526	0.31	102	10.76
7 × 7	100	49	526	0.54	112	10.76
3Xcos	200	19	522	0.57	236	20.95
5 × 5	200	25	349	0.57	147	29.80
7 × 7	200	49	349	0.88	157	29.80
Corner patterns						
5 × 5 Corner	75	9	1261	0.15	103	
7 × 7 Corner	75	16	1261	0.25	107	
5 × 5 Corner	100	9	1208	0.19	133	
7 × 7 Corner	100	16	1208	0.33	138	
5 × 5 Corner	200	9	996	0.37	231	
7 × 7 Corner	200	16	996	0.61	247	
Outer patterns						
5 × 5 Outer	75	17	570	0.16	80	7.50
7 × 7 Outer	75	25	570	0.29	66	7.50
5 × 5 Outer	100	17	526	0.22	107	10.76
7 × 7 Outer	100	25	526	0.37	83	10.76
5 × 5 Outer	200	17	349	0.50	180	29.80
7 × 7 Outer	200	25	349	0.61	135	29.80

^a D , dimension of the zero order; P , number of projections; L , pixels of spectral dispersion. The patterns are assumed to be dispersed onto a 2048 × 2048 FPA.

designs. While the six metrics $D, P, L, \text{fill}, \text{blur},$ and ϕ provide a simple means of comparing the various CTIS patterns, they can be ambiguous. For example, if one design has a larger P but a smaller fill than another design, Table 1 provides no way of knowing which of the two metrics is more important. A more objective measure is necessary. The simulation is able to create data cubes, images, and system matrices of a variety of CTIS systems and to perform complete reconstructions. All of the simulations below are performed for a 512×512 FPA and a data cube with spatial dimensions 25×25 (the data cube's λ dimension is determined by the specific dispersion pattern used). The reconstruction algorithm used is that of Eq. (3) of [3].

The first step toward providing an objective measure is to realize that any comparison is bound to be object dependent. Although this complicates comparisons, we can use object data cube models that closely correspond to useful objects for a CTIS to measure, and thus we can gain a rough idea of how well we can expect the instrument to perform in practice.

Allowing the simulator to generate an image and then perform a reconstruction on this data allows us to compare the result against the input. The normalized root mean square error (RMSE) ϵ of the reconstructed data cube is defined by

$$\epsilon = \left(\frac{1}{\sum_{m=1}^M f_m} \sum_{m=1}^M (f_m - \hat{f}_m)^2 \right)^{1/2},$$

where M is the number of voxels in the true data cube \mathbf{f} (whose elements are f_m), the reconstructed estimate of which is $\hat{\mathbf{f}}$. The RMSE, though not a truly objective measure of performance, provides a good starting point for assessing the capability of each CTIS design.

To use the RMSE measure, we first need to construct model data cubes for the simulator. CTIS objects can take on a variety of forms, so no one single data cube can be taken as representative, and thus a variety of models were used (see Table 2). For the simulations, each of the data cube models was taken to

Table 2. Data Cube Models Used in the CTIS Simulation^a

1.		Uniform: a data cube consisting of voxels whose values are all 1s	9.		Mod5-5-0: a λ -uniform data cube modulated along the x and y axes, with 5 modulation periods in both directions
2.		Point: a point source object modeled as a 3D Gaussian of width parameters $\text{width}(x, y, w) = (2 \text{ nm}, 1 \text{ pix}, 1 \text{ pix})$	10.		Mod8-8-0: a λ -uniform data cube modulated along the x and y axes, with 8 modulation periods in both directions
3.		Blob: a "Gaussian ball" object: a 3D Gaussian of width $\text{width}(x, y, w) = (8 \text{ nm}, 4 \text{ pix}, 4 \text{ pix})$	11.		Mod5-5-5: a data cube modulated equally along all three axes, with 5 modulation periods in each direction
4.		Column: a λ -uniform object with spatial Gaussian profile of width $\text{width}(x, y) = (1 \text{ pix}, 1 \text{ pix})$	12.		Mod8-8-8: a data cube modulated equally along all three axes, with 8 modulation periods in each direction
5.		Wide Column: a λ -uniform object with spatial Gaussian profile of width $= (4 \text{ pix}, 4 \text{ pix})$	13.		Point+Bkgd: a point-source white fluorescent lamp spectral object placed within a spatially uniform background of a white tungsten-halogen lamp
6.		Column_Mod10: a spatial Gaussian profile of width $= (1 \text{ pix}, 1 \text{ pix})$ modulated along the λ axis, with 10 modulation periods	14.		UA logo: a RGB image of the UA logo expanded into a data cube using standard RGB spectra
7.		Column_Mod20: a spatial Gaussian profile of width $= (1 \text{ pix}, 1 \text{ pix})$ modulated along the λ axis, with 20 modulation periods	15.		Starfield: a crowded field of Gaussian points, with varying broadband spectra (as might be found in an astronomical image)
8.		Mod0-0-10: a spatially uniform data cube modulated along the w axis with 10 modulation periods	16.		Checkerboard: a lambda-uniform data cube with 3×3 blocks of 1s and 0s.

^aFigures show projections of the data cube in two ways: (left) a perspective or orthographic projection and (right) a color-mapped panoramic projection.

have fixed spatial dimensions of 25×25 pixels, whereas the spectral dimension is taken to depend on the CTIS dispersion pattern being used. The normalization by $\sum f_m$ in the RMSE should prevent any given pattern from being penalized by this approach. The detector was taken to be a noiseless 512×512 pixel array, and the PSF for each projection to be a Gaussian of $\text{FWHM} = 1$ pixel.

Figure 11 shows the performance for various CTIS dispersion patterns on the 16 data cube models. The reconstruction algorithm uses only 20 iterations for the simulation results, even though the data cube estimate for the more complex objects is generally not yet near convergence there. (This is done to keep the simulation close to typical operating conditions for a physical CTIS system.) The RMSE values shown are for the final iteration. The feature to note in the plots is the spread in RMSE among the various CTIS designs for a given data cube model. For example, for the uniform model (no. 1) all designs perform well, but for the point-source model (no. 2) the two corner designs perform much worse than do the three full designs. For the more complicated data cube models, such as UA Logo (no. 14), Starfield (no. 15), and Checkerboard (no. 16), it appears that even for a noiseless CTIS system we cannot expect to achieve better than 25% mean error for data cube voxels, regardless of the dispersion pattern chosen. For the uniform model, this corresponds to a mean modulation transfer function [7] (MTF) of about 0.75 (see Table 3). This is explained in more detailed in Section 6 below.

It is important to note that the values plotted in Fig. 11 should be treated as guidelines rather than as absolute criteria of CTIS reconstruction quality. Since the various CTIS designs do not sample the wavelength dimension of the data cube at the same rate, a single RMSE measure cannot provide an objective comparison across all designs. Figure 12 shows example reconstructions from a 5×5 CTIS design on a 512×512 FPA.

A. Corner Patterns versus Full Patterns

When seeing CTIS for the first time and hearing about the trade-off in spatial and spectral resolution in the system [expressed by Eq. (1)], almost invariably the initial reaction is: Why not place the zero-order into the corner of the image and allow the diffraction orders to spread across the full extent of the focal plane array? (This configuration is shown in Fig. 6.)

Previous research has dealt with this idea [8], concluding that since the asymmetric patterns extend the missing cone to create an even larger, skewed, missing region within the Fourier cube, these designs must be untenable. However, since the standard designs already have a missing region and thus sample Fourier space nonuniformly anyway, this argument seems weak. Do we really need to sample the Fourier cube symmetrically in order to achieve a good reconstruction? The corner patterns sample some regions of Fourier space which are inaccessible to the symmetric designs, so that it seems plausible that they may have better performance when those frequencies are the important ones.

The simulation results shown in Fig. 11 (left) indicate that while the corner patterns generally perform worse than the comparable symmetric patterns, the difference is smaller than might be expected. And there are cases where their performance is marginally better, such as for the Mod0-0-10 data cube model (no. 8) and the Mod8-8-8 model (no. 12)—the data cubes where spectral resolution is more important than spatial. For objects possessing spatial cylindrical symmetry, however, the corner patterns' performance is disastrous, reflecting the lack of symmetry in their sampling of the data cube.

The excellent performance of the 3×3 pattern in Fig. 11 (left) may be surprising, given that the number of measured data cube projections is only nine. This is unfortunately an artifact of working with noise-free data that does not translate directly into noise-corrupted experimental systems. Because of

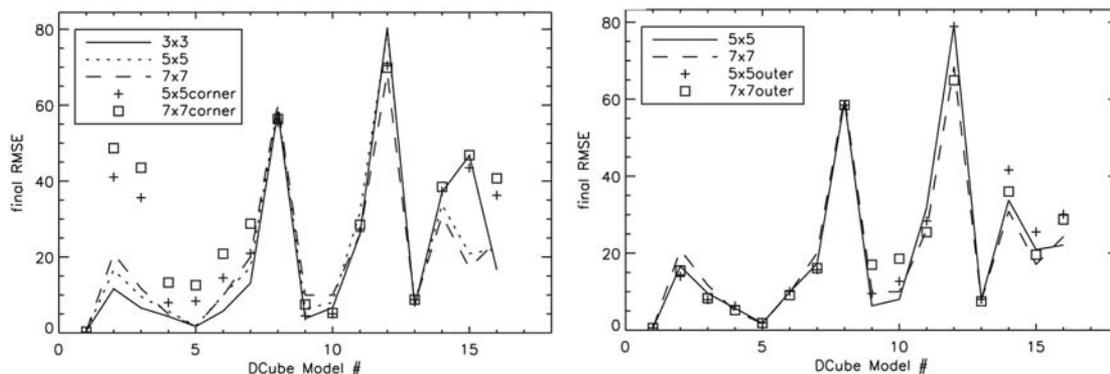


Fig. 11. Best RMSE for the 16 data cube models. Left, comparing the corner patterns with the symmetric full patterns; right, comparing the outer patterns with the full patterns. The data cube models are numbered according to: (1) uniform, (2) point, (3) blob, (4) column, (5) wide column, (6) column_Mod10, (7) column_Mod20, (8) Mod0-0-10, (9) Mod5-5-0, (10) Mod8-8-0, (11) Mod5-5-5, (12) Mod8-8-8, (13) point + Bkgd, (14) UA_Logo, (15) starfield, (16) checkerboard. In the plots above, the most important feature to notice is where the curves diverge from one another at a given data cube model, showing that one pattern performs substantially better than the others. (Note that higher values in the plots indicate poorer reconstruction performance.)

Table 3. Comparing MTF Values with the RMSE Measure for the Two Spectrally Modulated Data Cube Models

MTF (avg)	RMSE (Uniform)	RMSE (Point Source)
1.0	0%	≤40%
0.9	7%	50%
0.8	14%	115%
0.7	21%	122%
0.6	28%	128%
0.5	35%	132%
0.4	42%	136%
0.3	49%	138%
0.25	53%	140%
0.2	56%	145%
0.1	63%	150%
0.0	≥70%	≥155%

the lack of space to show this effect here, we plan to illustrate this behavior in the future.

B. Outer Patterns versus Full Patterns

Some of the projections in a given rectangular dispersion pattern have the same azimuthal projection angle and differ only in the length of spectral spread. In a 5×5 dispersion pattern, the 8 projections forming the inner ring thus share the same azimuthal angle with 8 of the 16 projections in the outer ring of diffraction orders. The primary difference between any two projections lying along a radial line extending from zero order is that the outer order has a greater spectral sampling, and we may wonder whether we can replicate the data of the inner order by merely averaging pairs of points in the outer projection along the radial direction. This implies that there may be some redundancy in the multiplexing and that the added information contributed by the inner projection may be negligible. To a degree, the simulation appears to support this idea.

To simulate this case, we can use the simulation to choose only the outer ring of projections and ignore the inner diffraction orders (see Fig. 13). Performing reconstructions on the same 16 data cube models allows us to compare the outer and full patterns (the results are given in Fig. 11, right). The comparison shows that there is actually little difference in performance between the two. For the Mod5-5-0, Mod8-8-0, Mod5-5-5, and Mod8-8-8 models (nos. 9–12), the outer patterns display worse performance than do the full patterns, but for almost all of the other data cube models, the outer patterns perform slightly better than their full counterparts. And this comes with the added benefit of a faster reconstruction. For the 5×5 pattern, ignoring 8 of the 25 projections allows us to reduce the size of the system matrix used in reconstruction to $17/25 = 0.68$ of the full size, providing a boost of 32% in reconstruction speed with no cost in reconstruction quality. In fact, likewise removing the zero-order image from the center of the FPA data appears to have no significant effect on the reconstruction quality, allowing an even smaller

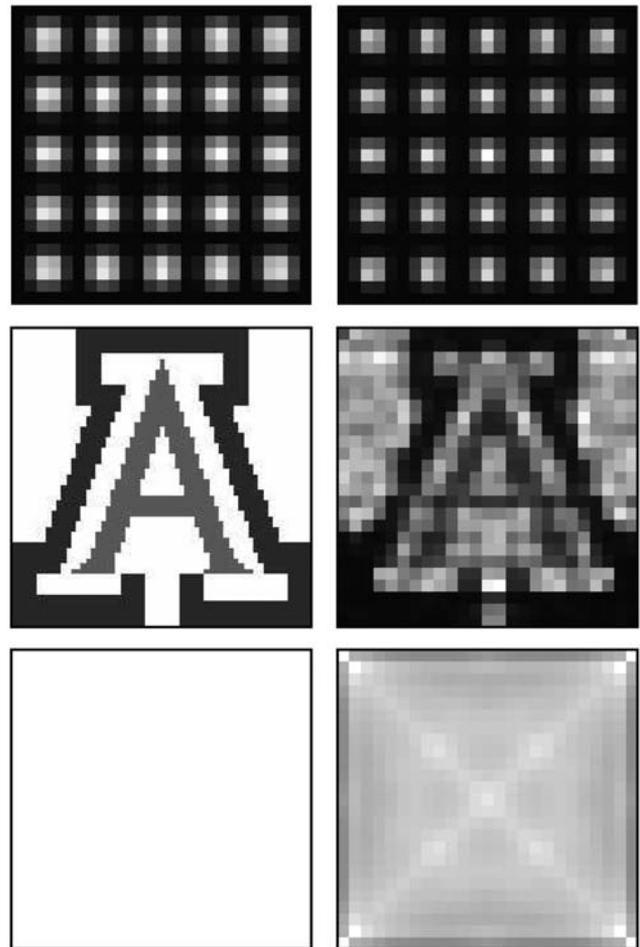


Fig. 12. Left, panchromatic images of three example input data cubes. Right, panchromatic images of the corresponding reconstructed data cubes, for the models (top) Mod5-5-5, (middle) UA_Logo, and (bottom) uniform.

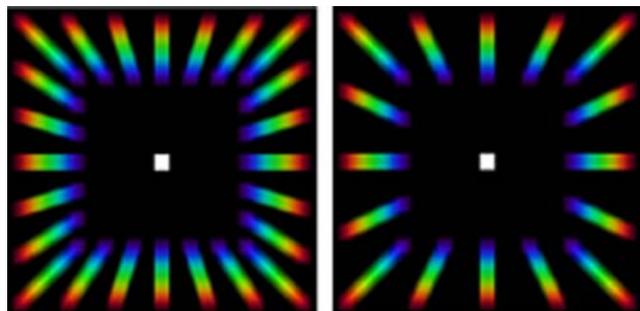


Fig. 13. (Color online) Outer dispersion patterns from (left) 5×5 and (right) 7×7 patterns.

system matrix ($16/25 = 0.64$). Thus, not only may the CTIS benefit by ignoring the inner set of diffraction orders, but when designing the disperser for a given system, it is possible to improve the signal-to-noise ratio by suppressing light going to the inner ring of projections and instead to channel light into the outer ring.

6. Specifying the Spectral Resolution of CTIS

A frustrating feature of a CTIS instrument is the fact that its spatial and spectral resolution are not constants but rather object-dependent quantities. That is, a CTIS instrument can demonstrate its maximum spectral resolution when viewing a point-source object, but for a spatially uniform object the achievable spectral resolution is far lower. Moreover, even when considering only a single object, the resolution is not a constant for all voxels within the data cube (generally, the corners of the cube show a much higher resolution than the middle regions of the cube). The simulation provides a convenient means of studying these complex effects.

One way of determining the spectral resolution of CTIS is to create a spatially pointlike object having a spectrum consisting of two monochromatic wavelengths. As the distance between the two wavelengths is altered, we can examine the reconstructed

spectrum and use a criterion to estimate when the two wavelengths are successfully resolved. A more typical use of a CTIS instrument, however, would be to view broadband objects, in which case the above approach is inappropriate.

For broadband objects, a convenient approach is to use a spectral transfer function, defining a point below which the modulations in the spectra are attenuated to such an extent that they are no longer distinguishable from the noise. This is the approach used below. Moreover, since the resolution is object dependent, it is probably best to define the resolution at the two extremes mentioned above. The resolution of any given object can then be bracketed as being somewhere between these two.

Figure 14 shows vectorized plots of an example input data cube, and its reconstructed counterpart, for a spatially uniform broadband object whose spectrum is modulated (i.e., the Mod0-0-10 model).

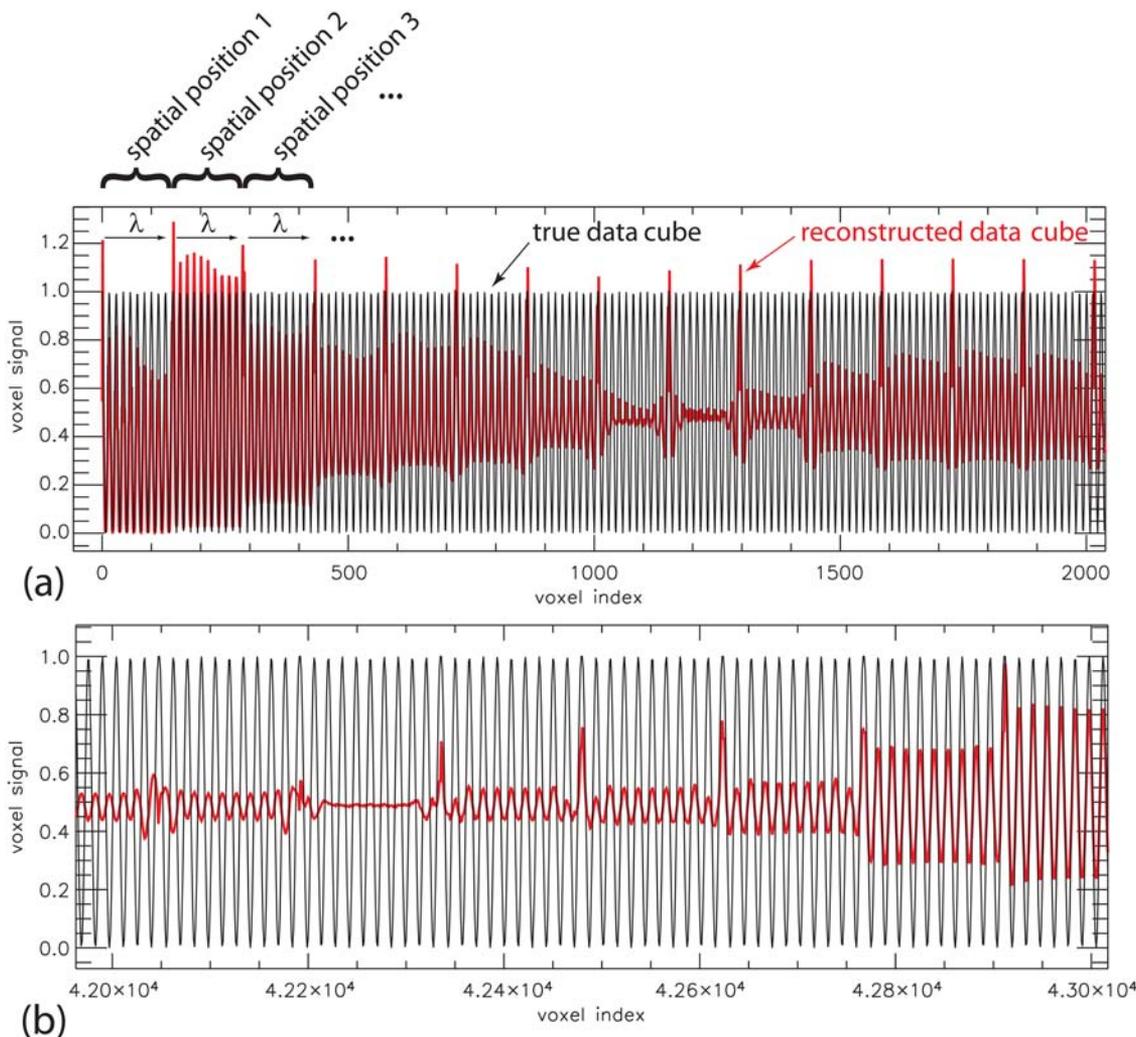


Fig. 14. (Color online) Sections of the vectorized Mod0-0-10 data cube (of dimensions $25 \times 25 \times 145$), where every 145 voxel section (containing 10 periods of modulation) represents the spectrum at a single spatial location in the data cube. (a) Snapshot from the $(x, y) = (0, 0)$ corner of the data cube, moving to increasing values of x from left to right; (b) snapshot from the middle $(x, y) = (12, 12)$ region in the data cube. Note that the flat region of the red (central) curve of (b) is the exact middle spatial location in the data cube, and that the reconstructed voxels on the left-hand side show reverse contrast.

The dimensions of the data cube are $(x, y, \lambda) = 25 \times 25 \times 145$, where the 145 wavelength samples result from the spectral spread in the 5×5 design's outer diagonal orders. The data cube vectorization is performed as follows: the first vertical column in the data cube (the spectrum corresponding to the lower left spatial position) is extracted and then concatenated with the spectrum at the second spatial position, the third, and so on. This provides a convenient means of visualizing the reconstruction.

The modulation contrast decrease from the input is evident in the reconstructed data cube, as is the change in transfer function from a corner toward the middle of the cube. The RMSE of the reconstruction here is 53%, which for this object corresponds to an average MTF of about 0.25. Of course, the choice of where to declare a fringe resolvable depends on the noise in the reconstructed data. If we sidestep noise considerations for the moment, then we can tentatively define the above situation as being at the limit of resolvability. (The definition is arbitrary, but provides a useful simple rule that has worked for our experimental systems.) Thus, for this data cube model, resolvability is reached when the RMSE falls below 53%.

The opposite extreme for determining spectral resolution is that of a spectrally modulated point-source object. If we take a look at a section of the input data cube and its reconstruction (Fig. 15), we can see that the spectral modulations are largely intact, equivalent to an MTF of about 0.95.

If we again use the criterion that the resolvability limit is where modulations in the data cube on average experience a fringe contrast loss of no more than 75%, then we can increase the modulation frequency in the object's spectrum, and we find that the point-source spectral fringes can be resolved up to 52 modulations across the spectrum. For use with a 400–700 nm spectral range, this corresponds to a spectral resolution of 5.8 nm.

Since both data cube models (uniform and point source) use spectra that are pure sinusoids, the nor-

malized RMSE values are easily compared with the MTF (i.e., fringe contrast). A comparison between the two is given in Table 3. The boldface row of the table indicates the cutoff point for resolvability. The surprisingly large RMSE numbers for the point source are due almost entirely to the overestimation of fringe contrast in the central pixel of the object. The huge discrepancy in RMSE values for the same fringe contrast in the signal of interest is an example of the limitations of the RMSE as a figure of merit.

The simulation results above are taken from a CTIS design using a 512×512 detector array, giving the number of resolvable modulation periods as 10 and 52 (i.e., 30 and 5.8 nm resolution) for the uniform and point-source data cube models, respectively. Performing the same simulation procedure for a 2048×2048 detector array, we find that the results for choosing a MTF of 0.25 as the cutoff give resolvable modulation periods of 30 and 200 (i.e., 10 and 1.5 nm resolution), respectively. Thus, for a 5×5 CTIS instrument, we can use a rule of thumb that the point-source spectral resolution is roughly equivalent to the spectral dispersion across 3 pixels along one of the outer diagonal orders (or 2 pixels across one of the other horizontal or vertical orders). Likewise, the spatially uniform object's spectral resolution is roughly equivalent to 20 pixels along an outer diagonal order.

The mean-MTF-based approach taken here to define spectral resolution is not only useful for defining a rule of thumb, it is also ideally suited to establishing a design criterion for a CTICS instrument. "CTICS" is an acronym adopted for a CTIS system incorporating channeled spectropolarimetry, a technique whereby polarization information is encoded into modulations of the spectrum [9]. When designing the polarization components for a CTICS system, it is necessary to estimate the resolvable modulation frequency in order to choose retarder thicknesses. The rule of thumb given here for spectral resolution can therefore replace more *ad hoc* means of estimating this information.

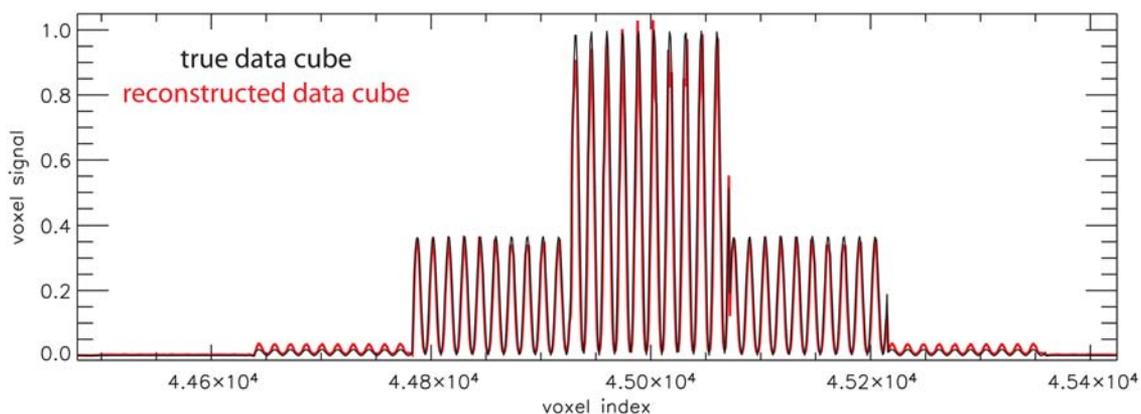


Fig. 15. (Color online) Section from the middle region of the vectorized Column_Mod10 data cube, where every 145 voxel section (containing 10 periods of modulation) represents the spectrum at a single spatial location.

7. Summary and Comments

Here we see that a CTIS system is capable of reconstructing a data cube that is spatially uniform but spectrally modulated. This is a data cube whose non-zero elements in the Fourier domain lie entirely within the missing cone region. If we take the missing cone idea at face value, such a data cube could not be reconstructed at all. This works only up to a point, of course; as the object data moves deeper into the missing cone region, the measured signal becomes increasingly attenuated and accurate reconstruction increasingly difficult. The fact that for discrete systems the missing cone is not a region of zero sampling, but rather one where the sampled signals are severely attenuated, has been noted before [10], but the simulation clearly confirms the concept.

Moreover, we find that for the above situation (object no. 8) the reconstruction quality is better than for some object models whose data lie entirely outside the missing cone (such as object no. 12). Thus, it is a mistake to think of the missing cone as the dominant mechanism for resolution loss. Rather, it is more useful to think in terms of (1) the multiplexing of data cube voxels on the measured image and (2) the blurring introduced by the VSF emerging from each voxel in the measured scene.

An unexpected result from the simulation is that the 5×5 and 7×7 patterns appear to be improved by deleting the inner set of diffraction orders, which has the side benefit of speeding up the reconstructions. This provides an easy, if marginal, means of

improving existing reconstruction speed with little effort.

References

1. T. Okamoto, A. Takahashi, and I. Yamaguchi, "Simultaneous acquisition of spectral and spatial intensity distribution," *Appl. Spectrosc.* **47**, 1198–1202 (1993).
2. M. Descour and E. Dereniak, "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results," *Appl. Opt.* **34**, 4817–4826 (1995).
3. M. R. Descour, C. E. Volin, E. L. Dereniak, T. M. Gleeson, M. F. Hopkins, D. W. Wilson, and P. D. Maker, "Demonstration of a computed-tomography imaging spectrometer using a computer-generated hologram disperser," *Appl. Opt.* **36**, 3694–3698 (1997).
4. H. H. Barrett and K. Myers, *Foundations of Image Science* (Wiley, 2004).
5. J. M. Mooney, V. E. Vickers, M. An, and A. K. Brodzik, "High-throughput hyperspectral infrared camera," *J. Opt. Soc. Am. A* **14**, 2951–2961 (1997).
6. A. K. Brodzik and J. M. Mooney, "Convex projections algorithm for restoration of limited-angle chromotomographic images," *J. Opt. Soc. Am. A* **16**, 246–257 (1999).
7. H. H. Barrett and K. Myers, *Foundations of Image Science* (Wiley, 2004), p. 311.
8. J. D. George, "Designing a non-scanning imaging spectrometer," Ph.D. dissertation (University of Arizona, 2001).
9. D. Sabatke, A. Locke, E. L. Dereniak, M. Descour, J. Garcia, T. Hamilton, and R. W. McMillan, "Snapshot imaging spectropolarimeter," *Opt. Eng.* **41**, 1048–1054 (2002).
10. K. C. Tam and V. Perez-Mendez, "Tomographical imaging with limited-angle input," *J. Opt. Soc. Am.* **71**, 582–592 (1981).