3-D Reconstruction from Few Projections: Structural Assumptions for Graceful Degradation

University of Waterloo Technical Report CS-2014-07

Michael Cormier, Daniel J. Lizotte, and Richard Mann
David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario N2L 5K5
Email: m4cormie@uwaterloo.ca

***Abstract*—We present a spatial-domain method for the reconstruction of a three-dimensional density distribution from one or more projections (images formed by integration of density along lines of sight) and using the three-dimensional reconstruction to explain features of the two-dimensional images. The advantages of our proposed method are that it degrades gracefully down to a single image, that it uses linear equations and constraints (allowing the use of convex optimization), that it is amenable to three-dimensional structural biases, and that ambiguity can be expressed precisely (it is possible to "know what we don't know"). Previously described methods have some, but not all, of these properties.**

I. Introduction

We examine the problem of reconstructing a density distribution from projections in a way that degrades gracefully to a single image, using symmetry and other structural properties to reduce or eliminate ambiguity. Projection, in the sense that we consider here, is a process of image formation by integration along the line of sight. Using structural constraints allows reconstruction even from a single viewpoint, which would otherwise provide insufficient information for a reconstruction. Our framework allows a wide variety of structural constraints to be defined, including various types of symmetry, smoothness, and compactness. It also allows any remaining uncertainty to be precisely expressed for each model. This makes our framework more flexible and more robust than existing methods for single-projection reconstruction (discussed in V).

A 3-D reconstruction is useful on its own but it can also be used to explain aspects of a 2-D image. Subtracting the projection of the 3-D reconstruction from the original image results in a residual image that reveals aspects of the original image that cannot be explained under the structural assumptions used for reconstruction.

Reconstruction from projections is useful in many domains. One such domain is astronomy; images of galaxies, for example, can be modelled as projections of "luminosity density." Since galaxies are so distant, only one viewing angle is available to use in reconstruction, which makes it impossible to use existing techniques; on the other hand, there are strong symmetry assumptions that can be made about galaxies. Radiographs (X-rays) are a very different type of image which can also be modelled as a projection process in which light is absorbed rather than emitted. Computed tomography allows three-dimensional reconstruction with many X-ray images from different angles, but is not effective with a small number of images. Our framework allows reconstruction with as little as a single image if appropriate structural assumptions can be made.

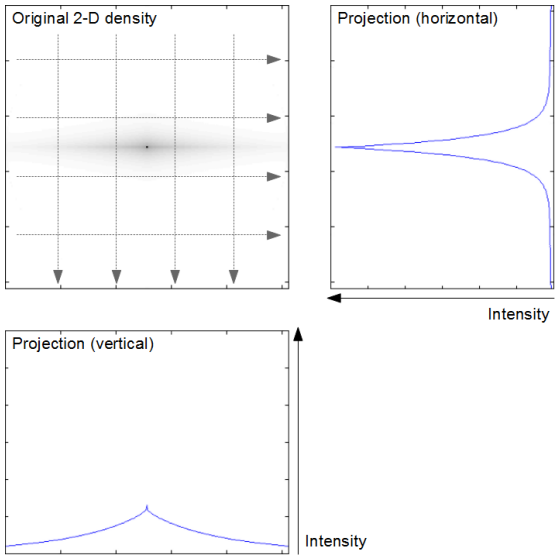

Fig. 1. Example of projection of a two-dimensional function along two directions

II. Projection & Reconstruction

In this paper, projection refers to integration of density along a line of sight. This is distinct from orthographic or perspective projection *geometry*, which determine the directions of lines of sight; we assume parallel lines of sight (*i.e.* orthographic projection geometry) in our definitions of the projection process. A projection is defined along a specific direction (the direction of the lines of sight); for notational convenience, we define the z-axis of the viewer's coordinate system to be parallel to the direction of projection. Given a three-dimensional density

function D, its two-dimensional projection P along the z-axis can be defined as follows:

$$P(x,y) = \text{Proj}[D](x,y) = \int_{-\infty}^{\infty} D(x,y,z)\mathrm{d}z. \quad (1)$$

This process is shown in Figure 1. Projection can also be defined in a discrete form, where the image is discretized as a grid of pixels and the density distribution is discretized as a grid of voxels. Assuming that the size of the density distribution is finite, a discrete projection process can be defined as follows:

$$P(x,y) = \text{Proj}[D](x,y) = \sum_{z=z_{min}}^{z_{max}} D(x,y,z) \quad (2)$$

where z_{min} and z_{max} represent the limits of non-zero density along the z-axis. In either formulation, it is clear that the projection process is linear; this will be critical to our reconstruction framework. Although this formulation is discrete, the density at an intermediate point can be defined using interpolation.

For images of objects like galaxies that emit light, the value of a pixel in the image is the sum of light emitted along a line of sight, which corresponds directly to projection. For images formed by light absorption, like X-ray images, a pixel value is the product of the proportion of light passing through each voxel on the line of sight. This does not correspond to projection as stated, but taking the logarithm of the transmission proportions and pixel values converts the product operation to a sum; this modified problem is a case of projection.

Reconstruction is the problem of inverting the projection process. Given a projection P, the objective is to produce a reconstructed density distribution $\hat{D}$. The optimal $\hat{D}$ must be determined based on the evidence of the projected image P, since the original density distribution is not directly accessible. An objective function representing the image reproduction error can be defined based on the sum of squared errors:

$$E(\hat{D}) = \sum_{x,y} \Big(P(x,y) - \text{Proj}[\hat{D}](x,y)\Big)^2 \quad (3)$$

The optimal reconstruction is therefore

$$\underset{\hat{D}}{\text{argmin}} \Big(E(\hat{D})\Big). \quad (4)$$

If the image is both noise-free and compatible with the model used, $E(\hat{D}) = 0$; otherwise reproduction error is inevitable.

III. Our Framework

We use a spatial-domain method for reconstruction which is based on a discrete voxel representation of the density distribution. The projection process is linear, and each pixel value is a linear combination of voxel values along the line of sight; this implies that projection can be represented by a matrix operation. Let $\mathbf{p}$ and $\mathbf{d}$ be vector representations of P and D, respectively, generated by placing pixel (or voxel) values in the vector in raster-scan order. Then a *projection matrix* G_P can be defined such that

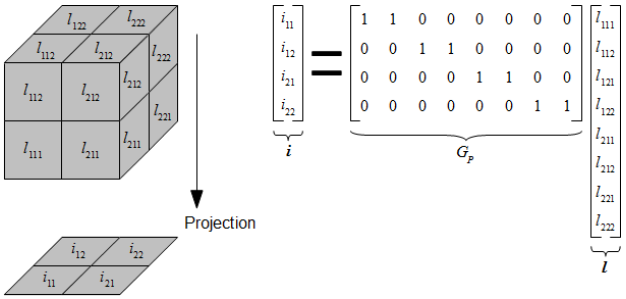

Fig. 2. Three-dimensional projection with a matrix formulation.

$$\mathbf{p} = G_P\mathbf{d}. \quad (5)$$

Assuming that D is $n \times n \times n$ voxels and P is $n \times n$ pixels, G_P is $n^2 \times n^3$. It is the matrix formulation of the reconstruction problem that we use to derive our framework. Figure 2 shows an example of a three-dimensional projection process and the corresponding matrix formulation.

If there are multiple images, G_P can be extended to describe the relationship between the volume and all of the images simultaneously. The meaning of the volume vector $\mathbf{d}$ remains the same; the image vector $\mathbf{p}$ must now represent each pixel value in *every* image. The projection matrix G_P must be altered to represent the contribution of each voxel to each pixel value in the new image vector. These modifications can be implemented by simply vertically concatenating the image vectors and projection matrices for each of the original images.

The matrix formulation of the projection problem hints at a way to solve the reconstruction problem: $\hat{\mathbf{d}} = G_P^{-1}\mathbf{p}$. The "unpacked" $\hat{D}$ corresponding to $\hat{\mathbf{d}}$ would minimize the sum of squared image reconstruction error. This requires that G_P be invertible. For a single image, the projection matrix has n^3 columns and n^2 rows; it represents an underdetermined system of equations, with infinitely many solutions. Multiple images can provide more evidence upon which to base the reconstruction; if there are enough images from the appropriate viewing angles, G_P is invertible and this method is viable. In particular, if the total number of images is greater than the number of voxels, this corresponds to the algebraic reconstruction method (ARM) in computed tomography [1].

A. Ambiguity

Since the correct 3-D reconstruction is ambiguous when there are fewer independent constraints than voxels, it is important to precisely characterize this ambiguity. This can be done using the properties of the projection matrix. Ambiguity means that there are multiple distinct density distributions which project to the same image; thus it can be characterized by defining the set of all possible solutions. We now show how the null space of the projection matrix can be used to characterize ambiguity:

Theorem 1. *If* $\mathbf{d}_1$ *and* $\mathbf{d}_2$ *are density distributions, then* $G_P\mathbf{d}_1 = G_P\mathbf{d}_2$ *if and only if* $\mathbf{d} = \mathbf{d}_1 - \mathbf{d}_2$ *is in the null space of* G_P.

Proof. Assume that $\mathbf{d}$ is a member of the null space of G_P. Then $G_P\mathbf{d} = 0$. It follows that

$$\begin{aligned} G_P(\mathbf{d}_1 - \mathbf{d}_2) &= \mathbf{0} \\ G_P\mathbf{d}_1 - G_P\mathbf{d}_2 &= \mathbf{0} \\ G_P\mathbf{d}_1 &= G_P\mathbf{d}_2 \end{aligned} \tag{6}$$

Similarly, if $G_P\mathbf{d}_1 = G_P\mathbf{d}_2$, it follows that

$$\begin{aligned} G_P\mathbf{d}_1 - G_P\mathbf{d}_2 &= \mathbf{0} \\ G_P(\mathbf{d}_1 - \mathbf{d}_2) &= \mathbf{0} \\ G_P\mathbf{d} &= \mathbf{0}. \end{aligned} \tag{7}$$

Thus, $\mathbf{d} = \mathbf{d}_1 - \mathbf{d}_2$ is in the null space of G_P. □

Since, for any $\mathbf{d}$ in the null space of G_P, we have $\mathbf{p} = G_P\mathbf{d}^* \Rightarrow \mathbf{p} = G_P(\mathbf{d}^* + \mathbf{d}$, the null space of G_P characterizes the inherent ambiguity in that given any one "base" solution $\mathbf{d}^*$ the null space expresses the entire range of possible solutions. If the null space of G_P contains only the zero vector, there can be only one solution; in this case, the problem is unambiguous. Otherwise, infinitely many distinct solutions exist. The dimensionality of the null space can be used to quantify the severity of ambiguity. Depending on the viewing angle, number of images, and choice of structural model, a practical problem may be either ambiguous or unambiguous. In general, it is desirable to reduce ambiguity, either by reducing the number of unknowns (*e.g.* by imposing symmetry constraints) or by increasing the number of equations (by supplying additional images or applying a structural bias).

If the problem is ambiguous, must must still be find $\mathbf{d}^*$. The Moore-Penrose pseudoinverse of the projection matrix, G_P^+, can be used; $\hat{\mathbf{d}} = G_P^+\mathbf{p}$ is the solution with the smallest L_2-norm $\|\hat{\mathbf{d}}\|$ such that $\|\mathbf{p} - G_P\hat{\mathbf{d}}\|^2$ is as small as possible.

B. Symmetry

Symmetry constraints reduce ambiguity by reducing the number of unknowns. Other constraints in which a subset of voxels must have the same density value work similarly; for convenience, we refer to all constraints of this type as "symmetry constraints".

Symmetry constraints can be enforced strictly by reparameterizing the density distribution such that a single variable is used to represent the density at all points in a set of points where the density must be identical. This corresponds to a change in the basis used to represent the space of density distributions. This new basis can be denoted β_{model}, and is defined such that $\mathbf{l}_{\text{voxel}} = \beta_{\text{model}}\mathbf{l}_{\text{model}}$. This can be integrated into the projection process:

$$\begin{aligned} \mathbf{i} &= G_P\mathbf{l}_{\text{voxel}} \\ &= G_P\beta_{\text{model}}\mathbf{l}_{\text{model}} \end{aligned} \tag{8}$$

If each element of the basis is symmetric, every linear combination must also be symmetric. For reasons of efficiency it is often desirable to calculate $G_P\beta_{\text{model}}$ more directly, but the matrix multiplication representation is convenient for analysis.

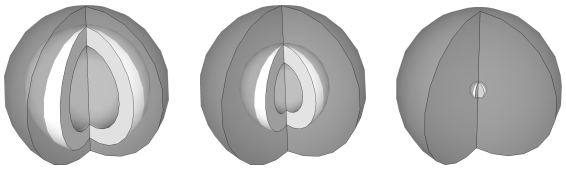

Fig. 3. Examples of spherically symmetric basis vectors.

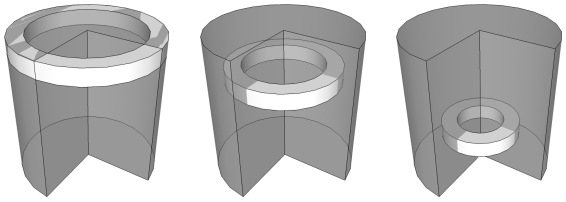

Fig. 4. Examples of cylindrically symmetric basis vectors.

For many types of symmetry, the viewing angle is significant. In these cases, β_{model} is also dependent on the orientation of the density distribution relative to the viewer. For example, it may represent a cylindrically symmetric density distribution viewed at an angle of 30° from the axis of symmetry.

1) Spherical Symmetry: A restrictive but often physically plausible type of symmetry is spherical symmetry. Spherical shells are constrained to have a uniform density; each basis vector corresponds to one of these shells, some of which are shown in Figure 3.

2) Cylindrical Symmetry: Cylindrical symmetry is a very flexible model but considerably reduces the space of solutions consistent with an image. The basis of a cylindrically symmetric density distribution is a set of rings of uniform density value, each of which has a radius and a position along the axis of the cylinder; some of these basis vectors are shown in Figure 4. The appearance of a cylindrically symmetric distribution depends on the angle between the direction of projection and the axis of symmetry, which we denote θ.

Other types of structural model can be defined by defining other sets of equal-density points. The circular ring basis used for cylindrically symmetric functions might be replaced with square rings or some other shape.

C. Structural Bias

Ambiguity can also be addressed by introducing a structural bias. This imposes a penalty term in the objective function which encourages some specific properties in the solution:

$$E(\hat{D}) = \sum_{x,y} \Big(P(x,y) - \text{Proj}[\hat{D}](x,y)\Big)^2 + S(\hat{D}) \tag{9}$$

where S represents the penalty term. This allows a compromise between reconstruction accuracy and the desired properties of the solution. Some structural biases correspond to regularization—a penalty proportional to the L_2-norm of the solution vector, for example. Many structural biases can be expressed by adding constraint rows to G_P. It is important to

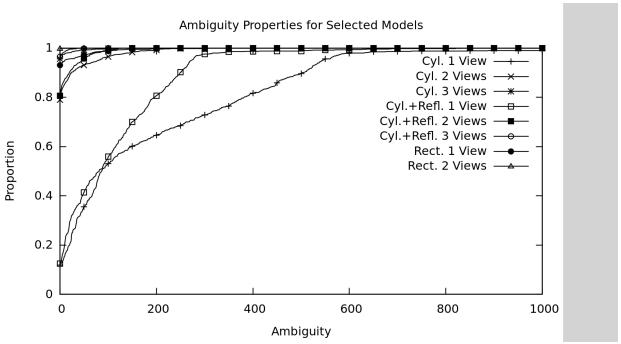

Fig. 5. Proportion of randomly-generated viewing scenarios with ambiguity no greater than the horizontal axis value for a $63 \times 63 \times 63$ volume with 250047 voxels. Cylindrical, cylindrical and reflective, and rectangular symmetry are shown for 1, 2, and 3 randomly-selected viewing angles. Rectangular symmetry with 3 views is omitted because there is so little ambiguity with 2 views.

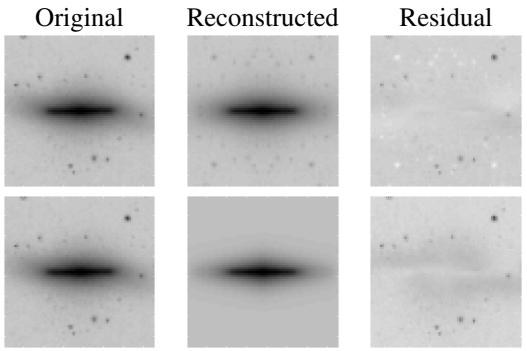

Fig. 7. Telescopic image of a galaxy (left column), the reconstructed image (middle column), and the corresponding residual (right column); the upper row shows reconstruction without a structural bias, and the lower shows reconstruction with a bias encouraging concentration to the equatorial plane.

choose an appropriate bias; an inappropriate one guides the solution away from, rather than toward, plausible reconstructions.

IV. Experimental Results

We test the performance of our framework in three ways: first, we examine ambiguity properties under different symmetry assumptions and with different numbers of images; second, we show experimental results on synthetic data; finally, we test the reconstruction method on real data.

A. Ambiguity Properties

We examine ambiguity properties using the dimensionality of the null space to measure ambiguity under different structural models. Figure IV-A shows the degree of ambiguity under three structural models: cylindrical symmetry, cylindrical symmetry with reflective symmetry across the equatorial plane, and a square-ring variant of cylindrical symmetry we call rectangular symmetry. For multiple viewing angles, fine uniform samples of all combinations of viewing angles is too computationally intensive. Instead we use random sampling of a range of viewing angles to generate Figure IV-A. This graph corresponds to the cumulative probability that a randomly-selected viewing scenario is less ambiguous than each value on the horizontal axis. These tests show that low or zero ambiguity is typical, especially with three images. The size of the null space of G_P is found by counting the number of zero eigenvalues (since the calculations are numerical, extremely small eigenvalues are taken to be zero).

B. Synthetic Data

Synthetic data is useful because truly symmetric objects are difficult to find in practice. It is through synthetic data that the performance of our framework under ideal circumstances can be tested. Figure 6 shows the results of reconstruction of a ring with and without symmetry constraints and a structural bias towards a small L_2-norm for one projection at $45°$ and three projections at using projections at $30°$, $45°$, and $60°$. The results clearly show the advantages of using a structural model. For the case of reconstruction from three views, the problem is unambiguous under the assumption of cylindrical symmetry and reconstruction is exact; adding structural bias slightly degrades the result in this case. In practice, an L_2-norm-minimizing bias has been found to be effective on noisy or asymmetric data (*e.g.* in Section IV-C). In all cases, the original image is reconstructed well.

C. Real Data

One example of image data formed by projection is an image of a galaxy. The galaxy can (in many cases) be modelled as a continuous, transparent, light-emitting substance; thus, the value of a pixel in an image of the galaxy is the sum of the "luminosity density" of the galaxy along a line of sight. Since the angular diameter of a galaxy is small, it is reasonable to assume that the lines of sight are approximately parallel. We use an image of NGC 4452 from the Sloan Digital Sky Survey III [2] as an example case.

The reconstructions shown in Figure 7 assume cylindrical symmetry with reflective symmetry across the equatorial plane; one uses no structural bias, and one uses a structural bias to encourage concentration of luminosity toward the equatorial plane. Both methods are effective, but the biased version is more robust with respect to extraneous objects like stars outside of the galaxy. The residual clearly shows the small deviations from cylindrical symmetry that are difficult to notice in the original image.

A second practical example of projection is an X-ray image. We use a CT scan of a roughly conical chocolate from the NIH/NIGMS Center for Integrative Biomedical Computing CT Scan Database to generate our test images. Since the raw imaging data is not available, we simulate the imaging process using the real volumetric radiopacity data provided in the original file. The raw density data was converted to normalized absorption proportions; the important aspects of the data for our purposes are its structural properties, not the specific density of this particular piece of chocolate.

Figure 8 shows an original projection and its reconstruction

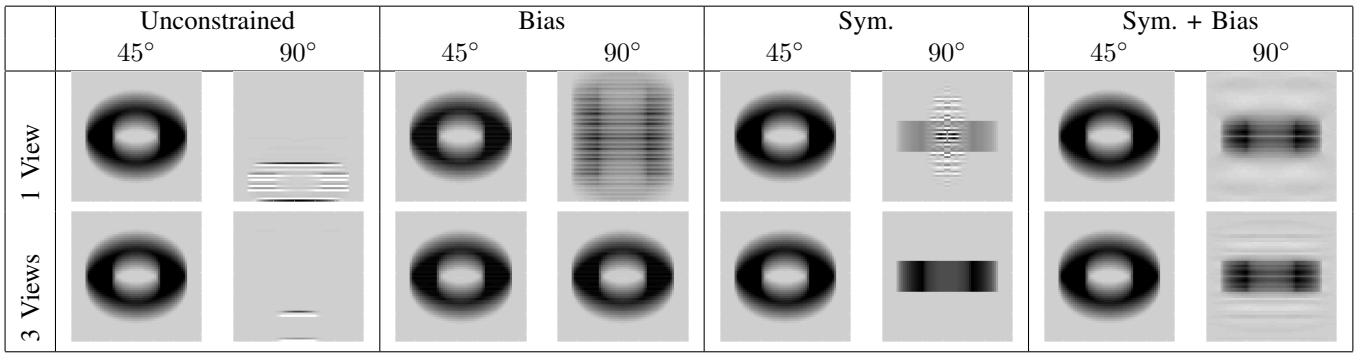

Fig. 6. Reconstruction of synthetic data for 1 and 3 projections, with and without symmetry constraints and structural bias.

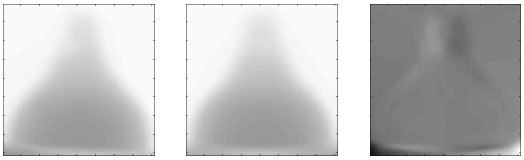

Fig. 8. Reconstruction (middle) of a chocolate from a single projection taken at 90° to the axis of symmetry (left), with residual (right, not to scale).

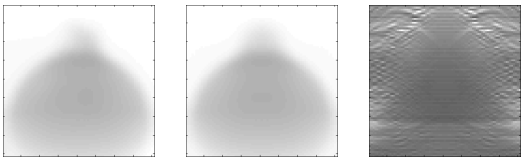

Fig. 9. Reconstruction of a chocolate from a single projection taken at 45° to the axis of symmetry. The original image is on the left, a projection of the reconstruction at the original viewing angle is in the center, and a projection of the reconstruction at 90° to the axis of symmetry is on the left

assuming cylindrical symmetry. The reconstruction is very accurate, and the residual clearly shows the asymmetries in the density distribution. Figure 9 shows a reconstruction with a less-favourable direction of projection; although the reconstruction projects to a reasonable image, overfitting is clear when the reconstructed density distribution is viewed from another angle. A structural bias can be used to reduce the degree of overfitting; Figure 10 shows the effects of a strong bias against a large L_2-norm. Artefacts caused by asymmetries in the original density distribution are not completely eliminated, but the reconstructed density distribution is substantially improved, and the general structure is much more readily visible.

Reconstruction can also be performed with multiple images. With only a few images, overfitting has been found to be a distinct possibility in cases where the true distribution is not perfectly symmetric, even when the problem appears to be unambiguous. This is because while there may be no density distribution that projects to an image of zeros, there may still be some that produce images that are zero nearly everywhere

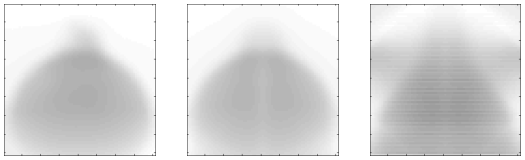

Fig. 10. Regularized reconstruction of a chocolate from a single projection taken at 45° to the axis of symmetry. The original image is on the left, a projection of the reconstruction at the original viewing angle is in the center, and a projection of the reconstruction at 90° to the axis of symmetry is on the left

or nearly zero everywhere. As in the case of a single image, however, structural biases can be used to address the problem. Figure 11 shows the results of reconstruction using projections at 30°, 45°, and 60° assuming cylindrical symmetry and using a strong bias against a large L_2-norm. The results here are much more accurate than the results using a single image; the only strong artefact occurs where the presence of the platform the object rests on causes the density distribution to be strongly asymmetric. As can be clearly seen, a high-quality reconstruction can be obtained using only three image with very simple structural assumptions.

V. Related Work

Most existing single-view three-dimensional reconstruction techniques are designed for images of opaque objects, rather than images formed by projection. For example, shape-from-shading [3] and shape-from-silhouette [4] techniques have been used with single images. These images are very different from images formed by projection. In a projected image, each pixel is influenced by every part of the object along the corresponding line of sight, but this information must be disentangled. In an image of an opaque object, a pixel provides information about a single patch of the surface of the object. Different techniques must be used for reconstruction.

Existing work done with projections typically assumes the availability of projections from a wide range of angles. This is the case for computed tomography (CT), which is the basis of CT scan medical imaging [5], [1]. The projection-slice theorem, which is critical to CT scanning, states that the Fourier transform of a projection is a "slice" through the Fourier

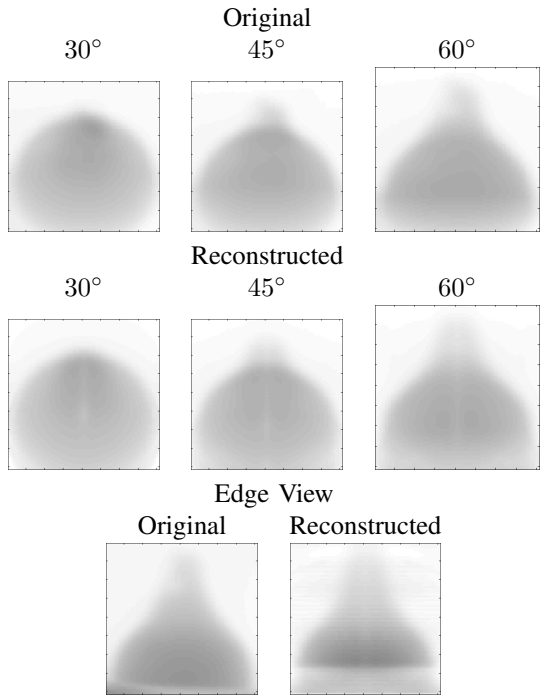

Fig. 11. Regularized reconstruction of a chocolate from three projections taken at 30°, 45°, and 60° to the axis of symmetry. Note that the only significant artefact occurs at the base of the image, where there is strong asymmetry in the original density distribution.

transform of the density distribution, at an angle determined by the angle of projection [6]. CT methods using the projection-slice theorem take the Fourier transforms of projections at many different angles to reconstruct the Fourier transform of the density distribution; from this, the reconstructed density distribution can be readily obtained [5], [1]. The algebraic reconstruction method (ARM) uses the inverse of a projection matrix describing the formation of many images [1]. Another method for CT reconstruction is back-projection [5], [1], in which the original density distribution is generated by "extruding" each projection back along the direction of projection and compensating for the differing numbers of intersecting projections at different distances from the center. CT is very useful and widely applicable, but requires many projections from a range of angles, which are not always available. Our framework uses a single image but requires stronger structural assumptions; it is complementary to CT methods, rather than directly competing with them.

Zaroubi *et al.* [7] describe a method for reconstructing cylindrically symmetric density distributions from a single projection using the projection-slice theorem, emphasizing astronomical applications. The assumption of cylindrical symmetry reduces ambiguity in the reconstruction problem, but does not eliminate it in most practical cases. In the Fourier domain, in fact, it can be shown that ambiguity exists for any direction of projection that is not perpendicular to the axis of symmetry. The frequency-domain method is a very different approach from ours, but it is less flexible because it strictly assumes rotational symmetry whereas our framework allows a wide range of symmetry properties to be defined and exploited.

VI. Conclusion

We have presented a framework for 3-D reconstruction of a density distribution from a small number of projections (as few as one), using structural assumptions to reduce or eliminate ambiguity, and for determining which aspects of an image can be explained by a given structural model. Our framework allows a wide variety of structural assumptions to be made, resulting in great flexibility. It is complementary to reconstruction methods that require many images (such as CT) and to single-image reconstruction methods for images that are not formed by projection.

While projection is not the most common image formation process encountered in day-to-day life, it occurs in very important settings, including X-ray images and images of transparent, light-emitting structures such as galaxies. Our framework is applicable to any of these image classes, and is especially important if only one viewpoint is accessible.

Extensions to the framework such as methods for handling images that are blurred or missing areas and reconstruction of multiple distinct objects are potentially interesting areas for further research.

Acknowledgment

This work was made possible in part by software from the NIH/NIGMS Center for Integrative Biomedical Computing, P41-RR12553-15.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is http://www.sdss3.org/.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the Univ. of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon Univ., Univ. of Florida, the French Participation Group, the German Participation Group, Harvard Univ., the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins Univ., Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State Univ., New York Univ., Ohio State Univ., Pennsylvania State Univ., Univ. of Portsmouth, Princeton Univ., the Spanish Participation Group, Univ. of Tokyo, Univ. of Utah, Vanderbilt Univ., Univ. of Virginia, Univ. of Washington, and Yale Univ.

References

[1] B. Gustafsson, "Mathematics for computer tomography," *Physica Scripta*, vol. T61, pp. 38–43, 1996.

[2] C. P. Ahn, R. Alexandroff, C. Allende Prieto, S. F. Anderson, T. Anderton, B. H. Andrews, É. Aubourg, S. Bailey, E. Balbinot, R. Barnes, and et al., "The Ninth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDSS-III Baryon Oscillation Spectroscopic Survey," *The Astrophysical Journal Supplement*, vol. 203, p. 21, Dec. 2012.

[3] H. Saito and N. Tsunashima, "Superquadrics parameter estimation from shading image using genetic algorithm," in *20th International Conference on Industrial Electronics, Control, and Instrumentation*, vol. 2, 1994, pp. 978–983.
[4] C. Zhang and B. Yuan, "Representation and reconstruction of three-dimensional objects using nonlinear deformable superquadric models," in *IAPR Workshop on Machine Vision Applications*, 1992, pp. 171–174.
[5] R. N. Bracewell, *Fourier Analysis and Imaging*, 1st ed. Kluwer Academic/Plenum Publishers, 2003.
[6] ——, "Strip integration in radio astronomy," *Australian Journal of Physics*, vol. 9, no. 2, pp. 198–217, 1956.
[7] S. Zaroubi, G. Squires, Y. Hoffman, and J. Silk, "Deprojection of rich cluster images," *The Astrophysical Journal Letters*, vol. 500, no. 2, 1998.