PRIOR KNOWLEDGE ANALYSIS OF INTERFACES IN CT DATA

James E. Youngberg

Perceptics Corporation
Suite 201
125 West 10600 South
Salt Lake City, Utah 84070

INTRODUCTION

This paper describes an approach to prior knowledge analysis of component interfaces in data measured using X-ray computed tomography (CT). In this approach, geometric details, including the presence or absence of component separations, are estimated more precisely than is normally permitted by the CT reconstruction interval. In the following section, the problem of interface analysis in CT data is presented. The next section discusses prior knowledge analysis and lists the assumptions necessary to make this approach possible. The final two sections discuss limitations inherent in prior knowledge analysis, and results obtained using both simulated and measured CT data.

BACKGROUND

Since early in this decade, considerable progress has been made in applying x-ray CT to the imaging of industrial objects. This has been evident in the defense propulsion industry, where a variety of solid-propellant rocket motors, nozzles, and nozzle components have been successfully imaged using machines ranging from low-energy medical imagers to a few which operate at 16 MeV. During this period, new designs have been characterized by improved sensitivity, finer resolution, and reduced artifact levels.

Despite imaging hardware successes, however, some expectations associated with these projects have not been realized. This is due to the nature of the objects being imaged, to limitations of the imagers, and to a general failure to anticipate the resulting problems with image processing solutions. Man-made objects typically consist of assemblages of homogeneous components. Inspections of such objects focus on (1) the presence and proper orientation of each component, (2) the internal uniformity of each component, and (3) the condition of the interfaces among components. X-ray CT properly addresses the first need when components are larger than the resolution of the imager. It satisfies the second need, so long as the detailed shape of material inclusions, cracks or discontinuities is not required. But because interfaces are very narrow features they are frequently seriously under-resolved, preventing adequate response to the third need.

This problem is illustrated in Fig. 1(a), which shows, somewhat simplified, the meaning of pixels surrounding an idealized bondline. For this example, the simplifying assumption is made that the pixel value is equal to the average density of the material in a square grid region around the pixel center. The oblique line in this diagram represents an abrupt density discontinuity (i.e. an interface) between two homogeneous materials. The discontinuity's orientation and location are arbitrary and unknown. The problem is to estimate the location and orientation of this discontinuity (interface) from individual pixel values. Since the each pixel along the interface assumes some value between the

Review of Progress in Quantitative Nondestructive Evaluation, Vol. 9
Edited by D.O. Thompson and D.E. Chimenti
Plenum Press, New York, 1990

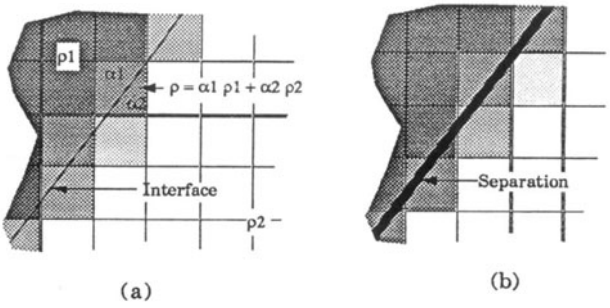

Fig. 1. (a) Simple example showing a representation of a discontinuity between materials having densities, ρ1 and ρ2. The measured density, ρ, of interface pixels is a weighted sum of the contributing materials. (b) A similar example, showing the ambiguous effect of a separation.

values of the adjacent components, visual analysis finer than the sampling interval is impossible.

This example can easily be extended to illustrate the impossibility of measuring (or even of detecting) separations which are significantly narrower than the pixel width. In Fig. 1(b) an additional discontinuity has been added parallel to the original discontinuity. The density between the two discontinuities is zero, representing the narrow gas-filled region which occurs in a separation. The important thing to recognize in this example is that the separation results in only a slight depression in the value of the pixels which intersect the interface. Visually, this depression could easily be mistaken for a slight shift in the *location* of the interface, rather than a separation. Thus, the general nature of the anomaly is uncertain, as is its specific geometry.

Information about details finer than the imaging device resolution is not only difficult to see: it is actually *absent* from an appropriately sampled image. In other words, *no amount of conventional image processing (contrast stretching, pseudocoloring, filtering, edge enhancing, etc.) can ever successfully define features which are narrower than twice the pixel width.* It is true that such features may be *detectable* by reason of the intensity modulations they cause, but as illustrated above, their nature can never be unambiguously known. Even detection is seriously jeopardized when feature dimensions are small compared to the pixel dimension and when they occur at an interface. This is especially true in real-world CT data which has random noise and artifacts. Thus, even though a pixel's value may reflect shifting or missing material within its volume element, the exact nature of the anomaly is gone unless information can be added to the problem.

Fortunately, in the case of many man-made objects, there is additional information available which can be combined with image measurements to enable refined detection and definition of anomalous component features.

As an example of the use of prior knowledge to improve the accuracy of a sampled measurement, study again Fig. 1(a). Consider the simplified problem of determining the position and location of the discontinuity from pixel values, assuming no noise or blurring. If only a single pixel value is known, or if nothing is known about the shape of the discontinuity, or if the densities of the homogeneous materials are not known, there are many possible solutions. On the other hand, if the discontinuity is known to be straight and the component densities are known, then only two pixel values are required to uniquely determine the interface location.

This example suggests how prior knowledge, combined with basic CT image information, can be used to deduce subject detail which is otherwise unavailable from the image itself. However, several complexities distinguish this simplified case from the real problem. First, CT pixel values are *not* simple density averages of rectangular volumes: They are three-dimensional weighted averages, whose weighting function has a

complicated form which is only approximately known. Second, noise distorts the pixel values. Other complicating effects include multiple interfaces within a pixel, irregular interface shapes, and component non-homogeneity. Nevertheless, when prior knowledge assumptions are carefully identified and their limitations are understood, they lead to significant improvements in the extraction of information from CT imagery.

ESTIMATION BASED ON PRIOR KNOWLEDGE

Two classes of prior knowledge are available which pertain to the subject and the imager, respectively. We first examine available prior knowledge, then discuss its use in an estimation algorithm.

The first class of prior knowledge assumptions has to do with the nature of the image subject. This class includes (1) axisymmetry, and (2) component homogeneity. Axisymmetry requires uniformity of value for all points at a particular distance from an axis. Although no man-made object is *strictly* axisymmetric, many parts can be considered to possess local axisymmetry. Component homogeneity requires that each mechanical component of the image subject have uniform density throughout. Parts containing delaminations, porosity, or inclusions do not meet this criterion, although such features, when small relative to the resolution of the CT imager can be grouped with random noise as long as the overall component density is accurately known.

The second class of assumptions pertains to the imaging system. We can assume that, at least locally, a CT imager is a linear, shift-invariant system which has been sampled and to which uncorrelated noise has been added. This assumption carries with it a rich intuitive and analytical heritage which facilitates an understanding of the effects of various assumptions, and which provides a mathematical framework for an analysis algorithm.

Figure 2 depicts an estimation algorithm based on the prior knowledge assumptions discussed above. In this algorithm, a geometric description of the input object drives a model of the CT imager, which implements the imager-related assumptions identified above. The output of the model is a set of synthetic measurements over an annular region of limited angular extent. The algorithm adjusts portions of the geometric description to minimize the mean squared error between the model output and the actual data. This algorithm amounts to a minimization procedure for the quantity,

$$\chi^2 = \sum_{i,j} [\, I_{i,j} - E_{i,j}(p) \,]^2 \qquad (1)$$

with respect to some set of parameters, p. The parameters represented by p might include, for example, the radial displacement of the interface from nominal, and the width of any separation. In Equation 1, $I_{i,j}$ is the measured image, and $E_{i,j}$ is the estimated image. The points (i,j) are constrained to lie within some annular analysis sector which is fixed for the minimization.

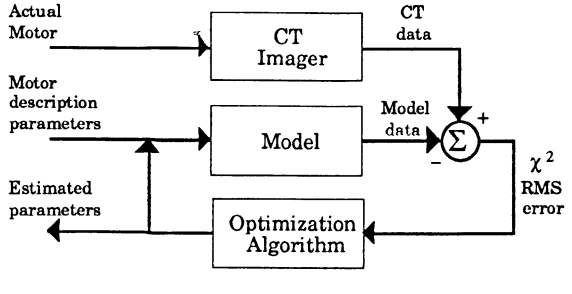

Fig. 2. Block diagram of the prior knowledge estimation algorithm.

LIMITATIONS OF PRIOR KNOWLEDGE CT INTERFACE ANALYSIS

The approach described in the previous section exploits the prior knowledge assumptions discussed earlier. It is entirely objective, potentially producing numeric results by which automatic defect detection can be driven. However, because its output is numeric rather than pictorial, its limitations and accuracy must be understood. The ability of the prior knowledge estimation algorithm to correctly determine the locations of the various interfaces depends on the validity of stated assumptions and on the precision with which various parameters are measured. There are two main approaches to understanding the algorithm's limitations. They are theoretical and empirical. Either approach must take into account the effects of noise, resolution, and model bias.

Theoretical limitations

The theoretical approach has been pursued by Jaffey [1], yielding an expression which quantifies the fundamental limit in the estimation algorithm's ability to unambiguously define a separation in CT as a function of various basic imaging system and imaged object parameters. These results use σ_g, the standard deviation of the gap estimate in the case where no gap is actually present, as an approximation of the reliability of the prior knowledge estimator. This quantity is a function of several parameters:

$$\sigma_g = 1.1\, A(r,\theta) \left| \frac{d_1 - d_2}{d_1 d_2} \right|^{1/2} \sigma_n^{1/2}\, \Delta_r^{1/4}\, \sigma_r^{3/4} \qquad (2)$$

Here $A(r,\theta)$ is a scale factor between 1.0 and 1.3 which takes into account effects arising from the geometry of the sampling grid; d_1 and d_2 are the true densities of the components on either side of the interface; σ_n is the standard deviation of the uncorrelated Gaussian CT noise; Δ_r is the average radial interval between samples in the analysis sector; and σ_r is a measure of the width of the CT radial point response function.

Under conditions typical in solid propulsion CT, where the resolution is 2.5 mm (0.10 in), the reconstruction intersample distance is 0.74 mm (0.03 in), and the noise is 1.5 percent, σ_g = 0.16 mm (0.006 in). Note that this result is most sensitive to the parameters, $|d_1 - d_2|/d_1 d_2$ and σ_r. On the basis of the parameters considered in the above analysis, automatic detection of separation features to the level of one tenth the sampling dimension seems not to be ruled out.

The effects of other parameters on the performance of the estimator are less easily modelled. Such effects were investigated empirically, as described in the following section.

Experimentally-Determined Limitations

A set of synthetic images was created to enable characterization of the estimator algorithm in the presence of failed prior knowledge assumptions. The general nature of this image set is depicted in exaggerated form in Fig 3, which show a section of a right cylindrical object containing a graded set of tapered separations between components whose densities differ in a manner typical of solid propulsion interfaces. Each synthetic slice was created using the typical parameters mentioned previously. When the synthetic slices are examined visually, the tapered separations occurring around 180°, 215°, 270°, and 325° are virtually invisible, even in the absence of noise.

Each synthetic test image was made to probe one or more of the following specific aspects of the estimator's behavior: robustness in noise, effect of an improperly-estimated point response function, behavior when cupping is present, robustness with improperly estimated component densities, and behavior with non-axisymmetric (wavy) interfaces.

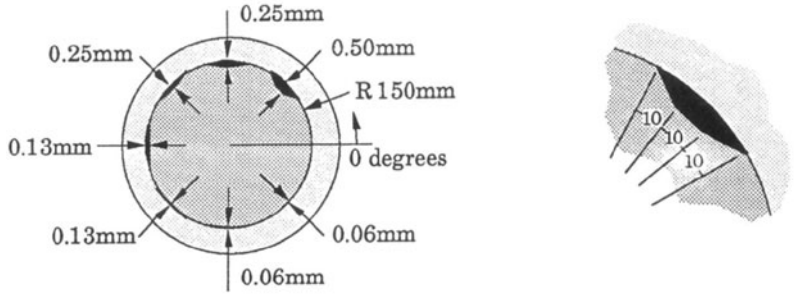

Fig. 3. Synthetic test object for initial estimator testing.

Typical results are plotted in Figure 4(a). In this graph, the trace representing the separation series experiment is prominent. This trace shows the near-ideal performance of the estimator in locating tapered separations less than one tenth the intersample distance (or two percent of the transaxial point response function's 2.5 mm width). Other traces, grouped at the bottom of the graph, show the estimator's performance in the presence of an ideal interface, with one percent gaussian noise, and with interface nonaxisymmetries. Close examination of these results shows that the largest error resulting from any of these conditions was less than 0.03 mm (0.001 in).

In order to assess the effects of incorrectly estimated prior knowledge parameters, these were tested in the absence of simulated gaps. The results are shown in Figure 4(b). Clearly, errors in estimating the radial line spread function are of little consequence. While errors in estimating the center can result in slightly more significant errors, it is possible to estimate the effective center of a local to precision which eliminates this problem. The most serious error results from incorrectly-estimated component density. Fortunately, the resulting signal has variance consistent with other perturbations. Unfortunately, the apparent constant bias shown in this test is less well-behaved in tests containing gaps. Because local density variations can result from several causes (including ray artifacts, cupping, component inhomogeneity, for example) this sensitivity represents a potential weakness in the prior knowledge estimation approach. The problem may yield to approaches which control the variations. It may also be possible to include the component densities as estimated parameters. The most promising approach seems to be direct estimation of component densities in the vicinity of the analysis sector (though this would be difficult for components which are thinner than the CT interpixel distance).

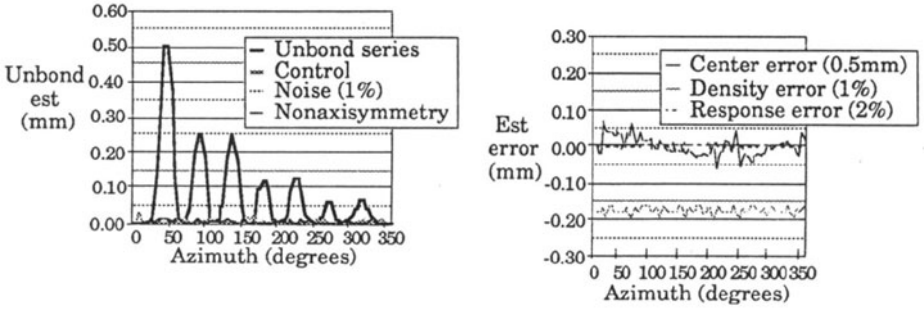

Figure 4. (a) Controlled experiments showing the effects various conditions on the performance of the estimator (assuming an axially-independent interface radius). (b) Effects of various incorrectly estimated a priori parameters.

A set of experiments to investigate the effects arising from tapered interfaces was performed. These experiments demonstrate that tapers effect the detection of the separations only slightly. Effects of other perturbations were somewhat increased, but were still limited to about ten percent of the CT interpixel distance.

Finally, experiments combining multiple failed prior knowledge assumptions were performed. Estimator performance suffered slightly, but errors generally remained within one-tenth the CT interpixel distance.

Results Using Measured Data

Finally, we processed a measured CT image using the prior knowledge estimation algorithm described above. This image was measured from a CT characterization object containing multiple artificial separations, induced to simulate typical chamber/insulator and insulator/propellant separation conditions. A sketch of this object is shown (not to scale) in Figure 5. The CT scan was made using geometric parameters similar to those used in simulated scans used for this program. The CT intersample distance was 0.37 mm (0.015 inch). The CT reconstruction exhibited the usual cupping phenomenon, which was estimated and corrected prior to performing prior knowledge estimation.

The median filtered result of prior knowledge estimation of the inner interface is shown in Figure 6(a). In this figure, all six artificial defects are clearly detected. The size of each defect agrees with the object's design to within a few hundredths of a millimeter (a few thousands of an inch). Note that false indications are limited in amplitude to roughly 0.02 mm (0.001 in), which is *5 percent* of the reconstruction spacing and less then *1 percent* of the radial line spread function width.

The analogous result corresponding to the outer interface is plotted in Figure 6(b) Here, as with the inner interface, all six defects were detected, and the estimated size of each defect is consistent with the test object's design. The estimate has a false indication of 0.08 mm (0.003 in), and three others 0.02-0.05 mm (0.001-0.002 in). Further analysis suggests that these false alarms are directly attributable to CT reconstruction-related periodic density variations. Thus, they can probably be reduced in practice by local estimation of component densities. Non-periodic errors are similar to those of Figure 7.

In an earlier section, the importance of providing correct a priori component densities was discussed. In estimates performed with the test object, this necessity was reconfirmed. Density errors of only one percent were sufficient to cause estimator errors of roughly 0.15 mm (0.006 in).

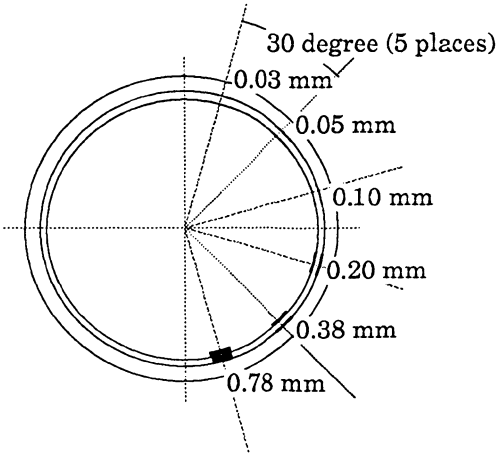

Figure 5. Layout of the CT characterization test object, showing the placement of the ten artificial separations.

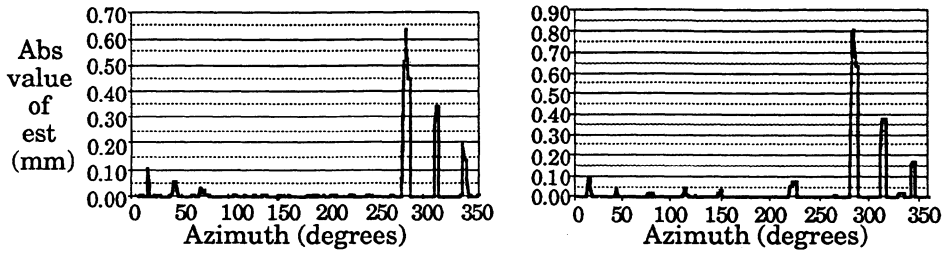

(a) (b)

Figure 6. Result of running the prior knowledge estimator on the (a) inner and (b) outer interfaces of the test object. Note that all six artificial defects are clearly detected.

CONCLUSION

We have characterized actual performance of a prior knowledge interface analysis algorithm using (1) carefully-controlled synthetic input and (2) data measured on an actual CT machine. These results agree generally with analytical predictions and provide a realistic view of the conditions under which prior knowledge interface analysis might routinely be used, and what problems are involved in achieving this goal. From our experiments, we conclude that (1) random noise at levels typical of industrial CT imagery does not prevent model-based estimation from reliably detecting separations as small as one tenth the pixel dimension, (2) small failures of the roundness (axisymmetry) assumption are properly tracked by the estimation algorithm, providing that such variations are gradual with respect to the analysis sector width, (3) errors in estimating the width of the radial line spread function are of little consequence when limited to a few percent of actual width, and that (4) inaccuracy of the a priori densities for the various components, whether resulting from actual object variations, cupping in the CT reconstruction, ray artifact, or reconstruction filter overshoot, results in serious (but avoidable) estimator errors. Taking the above into account, we conclude that prior knowledge interface analysis can detect component separations as small as one tenth the CT intersample distance.

REFERENCES

1. Jaffey, S.M., et al, Improving D5 Motor Separation Detection Using Computer Processing of HECT Tomograms, (Internal report prepared by the Digital Image Processing Laboratory, Lockheed R & D Division, LMSC), April 1988.