

Institut für Computergraphik und
Algorithmen

Technische Universität Wien

Karlsplatz 13/186/2

A-1040 Wien

AUSTRIA

Tel: +43 (1) 58801-18601

Fax: +43 (1) 58801-18698

Institute of Computer Graphics and
Algorithms

Vienna University of Technology

email:

technical-report@cg.tuwien.ac.at

other services:

<http://www.cg.tuwien.ac.at/>

<ftp://ftp.cg.tuwien.ac.at/>

TECHNICAL REPORT

Variance Comparison for Volumetric Industrial CT Data

Muhammad Muddassir Malik, Christoph Heinzl and M. Eduard Gröller

Institute of Computer Graphics and Algorithms,
Vienna University of Technology

TR-186-2-08-06

April 2008

Variance Comparison for Volumetric Industrial CT Data

Muhammad Muddassir Malik, Christoph Heinzl and M. Eduard Gröller

Institute of Computer Graphics and Algorithms,
Vienna University of Technology

April 7, 2008

Abstract

This paper proposes a novel technique for the direct comparison of a surface model with the corresponding industrial CT volume. We do not require the generation of a mesh from the CT scan and instead perform comparison directly with the raw volume dataset. Our technique uses the information from the surface model to locate corresponding points in the CT dataset. We then compute various comparison metrics to perform distance analysis and normal analysis. The metrics are presented to the user both visually as well as quantitatively. The comparison techniques are divided into two groups namely geometry-driven comparison techniques and visual-driven comparison techniques. The geometry-driven techniques color code the datasets and render distance glyphs to provide an overview, while the visual-driven techniques can be used for a localized examination and for determining precise information about the deviation between the datasets.

1 Introduction

Comparison of two almost identical datasets is very important for the continuously rising demands of quality control in industrial engineering. Recently there has been a whole body of work in the area of variance comparison between two surface models. Originally, the basic intent for the comparison of two surface models was to measure the differences introduced during the simplification of mesh datasets. The high number of vertices and edges are hard to process in real time due to the limited processing power available in hardware. Therefore it is necessary to simplify the datasets by reducing the number of triangles in planar areas and keeping the resolution of the mesh in curved regions. This initiated research to simplify mesh datasets in such a way that

the rendering speed is maximized while the mesh distortion is limited. Distortions introduced through mesh simplification led to research on the variance comparison between mesh datasets.

In the manufacturing industry, it is necessary to produce mechanical parts as close to the computer aided design model (CAD) of the part as possible. Engineers use CAD tools like AutoCAD, Catia, ProEngineering etc. for designing, which are purpose built tools for engineering drawings. The CAD model is considered to be the ground truth during the manufacturing process. To verify the accuracy of the production process, manufactured parts are scanned with an industrial computed tomography (CT) machine. The volumetric dataset obtained from the CT scan is then compared to the CAD model of the part (called surface model henceforth). The variance comparison between the two datasets is supposed to clearly identify erroneous regions in order to improve the production process.

The most common method for variance comparison is to generate an iso-surface mesh from the CT scan and to compare the two surface models. This may not be the ideal case. First, the generation of a mesh from the CT dataset requires a surface extraction algorithm, typically involving reconstruction artifacts. Here we observe a loss of important high frequency information. Additionally, if we also try to simplify the mesh dataset, then there is a further deviation from the ground truth. Second, mesh generation for a given iso-value may not be interactively possible during the comparison process. Therefore, the need to do a comparison with a higher or lower resolution mesh may lead to a delay in the examination process. Third, a CT dataset goes beyond a surface model and has information about the interior of the mechanical part as well. Losing this information limits the examination possibilities of the CT dataset.

Figure 1 shows a designed CAD model in (a), direct volume rendering (DVR) of the industrial CT scan in (b) and an iso-surface mesh extracted from the CT scan

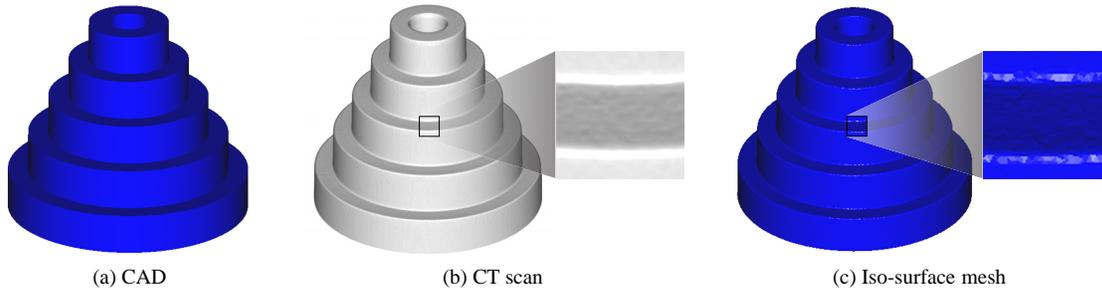


Figure 1: (a) CAD model of test-part-1. (b) Direct volume rendering of test-part-1, manufactured according to (a) and scanned with an industrial CT machine. (c) Iso-surface mesh extracted from the volumetric dataset in (b). Zoom-ins of the volumetric dataset and the iso-surface mesh are shown in (b) and (c) respectively. Reconstruction artifacts are visible in (c). The surface model of test-part-1 consists of 200,000 triangles and its volumetric dataset has a resolution of 561x559x436.

in (c). In figure 1(c), all the internal information of the volumetric dataset is lost. Areas marked with black rectangles in figure 1(b) and (c) are shown as zoom-ins. We observe reconstruction artifacts in figure 1(c).

In this paper we present a novel system to perform variance comparison directly between the surface model (which is the ground truth) and the raw dataset obtained from the industrial CT scan. As we compare the surface model directly with the volumetric dataset, we successfully avoid reconstruction artifacts (section 4). The proposed method is implemented on the Graphics Processing Unit (GPU) and thus provides an interactive variance comparison. We avoid delays in the examination process by embedding the complete comparison pipeline in a single system (section 4).

Our system uses DVR for visualizing the raw datasets. Techniques for the exploration of the volumetric datasets can be easily integrated into the proposed system. We thus combine volume visualization and variance comparison in our system, which are both working interactively but independently. The visualization techniques help the user both in searching for the deviations and in precisely and quantitatively viewing the deviations. We do not assume any prior information about the shape or topology of the datasets being compared. The geometry-driven and visual-driven variance comparison techniques are detailed in sections 3.2 and 3.3 respectively. We include an iterative closest point algorithm (ICP) for dataset registration (section 3.1) in our system. The impact of the errors introduced by the registration in the measurement process is analyzed in section 4.

2 Previous Work

Large numbers of triangles are inefficient to render and also hard to stream over a network. Subsequently, algorithms are proposed to simplify meshes [3], [13]. Mesh simplification distorts the original shape to some extent and therefore techniques are proposed to measure the differences between the two meshes.

Metro [4] is a general tool designed to compare two surfaces. One of the surfaces is scan converted to a set of points and then the distance between each point and the other surface is measured. Aspert et al. [1] propose to use the Hausdorff distance for measuring differences which is computationally and memory wise efficient. Pichon et al. [11] propose to use the gradient of the Laplacian equation to locate corresponding points on the two surfaces. The corresponding points are then used to measure distances between the surfaces. They argue that the Hausdorff distance fails to consider the variation in shape. We calculate uncertainties in the datasets to cater for the shape variations.

A variety of metrics and visualization techniques are proposed by Zhou and Pang [16] to measure mesh differences and present the information visually. The combination of metrics and visualization methods intends to help users test and calibrate various mesh simplification algorithms and find what suits their specific requirements.

Turk presents an algorithm [13] to create multiple levels of detail from a surface model. Samples are randomly distributed over the surface and subsequently a relaxation process spreads them uniformly. In the final step the user can choose a surface from a discrete set of surfaces to perform comparison. The generation of a surface model and variance comparison are two separate

processes. The method proposed by Turk can be used to avoid delays in the variance comparison process by generating multiple surface models with different levels of detail in advance.

Weigle and Taylor [15] investigate visualization methods for distance and local shape comparison. Their study shows that glyphs are better in conveying deviation information between surfaces than color coding alone. They use intersecting surfaces with known alignment for their study.

The above mesh comparison techniques do not need registration. As the simplified mesh is extracted from the original mesh, both meshes are perfectly aligned. In the case of a surface to a CT dataset comparison, where both datasets originate from a separate process, there is a need to perform registration. We have included an iterative closest point algorithm [2], [10] for dataset registration. A survey of registration algorithms specific to medical datasets is provided by Wang et al. [9].

There has been some recent work on the variance comparison between a surface model and an industrial CT dataset. These methods however introduce a preliminary step to the variance comparison process, where an iso-surface mesh is generated from the CT dataset. Heinzl et al. [6] propose a technique for generating a feature preserving mesh from a CT dataset. They use filtering to suppress noise and a watershed segmentation to create a binary dataset. In the final step a surface model is created using elastic surface nets, which is then used for variance comparison. There is a whole body of work in the area of surface extraction, which is considered to be out of scope for this work.

Geomagic Quality [12] is a well-known software product, used for quality control in industrial engineering. A surface model and an iso-surface mesh of the volumetric dataset are inputs to this tool. It uses a point-set to point-set or a best fit alignment algorithm for registration and then performs distance analysis between the two datasets. Results are presented as a color coded surface. Methods for extracting an iso-surface mesh from a volumetric dataset [14], [5], [6], [8] have to be used in a preprocessing step for performing comparisons using Geomagic Quality. Geomagic Quality works independently from the surface extraction process and therefore assumes no reconstruction artifacts during the comparison process, although such errors are introduced in the pre-processing step.

3 Variance Comparison

Our variance comparison system is divided into geometry-driven and visual-driven comparison tech-

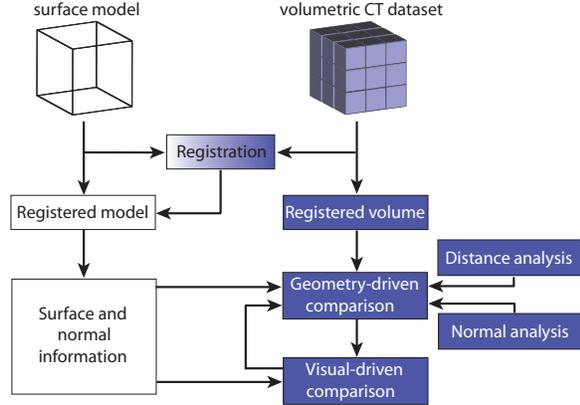


Figure 2: Variance comparison, system overview. Geometry-driven comparison techniques (distance and normal analysis) color code the datasets and render glyphs. Visual-driven comparison techniques provide localized variance information in a graphical and quantitative way.

niques. Geometry-driven techniques provide an overall visualization of the variance between the surface model and the volumetric dataset. Visual-driven techniques are used on top of the geometry-driven comparison techniques for a user guided analysis and for obtaining precise quantitative information. The use of graphics hardware makes all this possible in real time.

An overview of our system is shown in figure 2. The output of the rigid registration process is a transformation matrix. The transformation matrix is applied to the surface model which is transformed to closely match the orientation and scaling of the volumetric dataset. Both types of comparison, i.e., geometry-driven and visual-driven comparison techniques, query the registered surface model for the necessary information but work completely independent from each other. The results of the query, meta data from the CT dataset, and the chosen visualization technique are used to compute quantitative data and to produce images. Images generated through geometry-driven and visual-driven analysis techniques are displayed in separate windows. The user can employ the geometry-driven and visual-driven analysis techniques simultaneously.

Geometry-driven comparison techniques consist of distance analysis and normal analysis. Distance analysis calculates the deviation between the surface model and an interface in the volumetric dataset as Euclidean distances. It also measures the uncertainty of the evaluated distance. Normal analysis precisely locates differences in curvature and compares the surface smoothness of the two datasets.

We provide ray-profile analysis and an innovative magic lens as building blocks of the visual-driven comparison. Ray-profile analysis visually presents the data and deviations at a user specified location and also displays the information quantitatively. The magic lens is used for two purposes. First, it is used to zoom-in/out of the volumetric dataset. Second, it extracts the variance between datasets at a user specified neighborhood and displays the deviations using glyphs, i.e., 3D box plots.

3.1 Surface Model to CT Dataset Registration

The iterative closest point algorithm (ICP) performs rigid registration and produces a transformation matrix as output. The output matrix transforms (translation, uniform scaling and rotation) the surface model (moving dataset) to closely orient it to the CT dataset (fixed dataset).

The algorithm iteratively produces a matrix to transform the moving dataset. After each iteration, the mean square error between the datasets is calculated (equation 1). The iteration continues until the mean square error becomes stable, i.e., there is no change in error from one iteration to the next. At this point, the composite matrix of all the transformations and the final error of the registration process are returned. In equation 1, n is the number of reference points, s_i and m_i depict points on the fixed and moving datasets respectively, q_R and q_T represent quaternions for rotation and translation respectively and $R(q_R)$ is the rotation matrix formed from q_R .

$$f(error) = \frac{1.0}{n} \sum_{i=1}^n \|s_i - R(q_R)m_i - q_T\|^2 \quad (1)$$

Specification of points in the volume data is a two step process. In the first step, the user employs a transfer function to volume render the CT dataset. The points can then be chosen by simply clicking on the rendered image. Our system casts a ray along the viewing direction from each point specified by the user. Edge detection is performed on the profile of the opacity values encountered along the ray, to locate a point in the volumetric dataset. Registration is not the major scope of our work. We performed it with high accuracy (see section 4). More automatic registration techniques have not been investigated but might be possible.

3.2 Geometry-Driven Variance Comparison

Geometry-driven variance comparison is composed of distance analysis and normal analysis methods. Both of these analysis methods require the specification of a corresponding point in the CT dataset for each surface point on the surface model. Starting from a surface point we have to locate the corresponding point in the volumetric data. The search direction is approximately along the surface normal. In high curvature areas the search should be extended to nearby directions as well to ensure robustness.

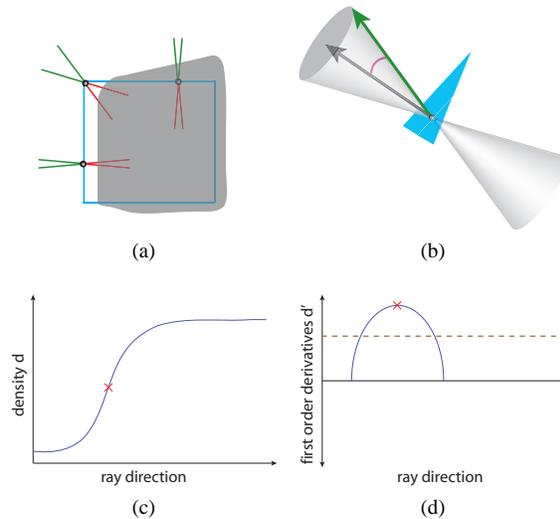


Figure 3: (a) Pairs of red and green lines depict the space in which we search for a corresponding point in the volume data (gray object) for each surface point (black sphere) on the surface model (blue rectangle). The double cone in (b) has an opening angle equal to twice the search-angle. Multiple rays are cast, starting from the surface point (apex of the double cone) and directed towards the base of the double cone. Edge detection is performed along the density profiles of each ray. A density profile and the first order derivative of a density profile are illustrated in (c) and (d) respectively. The dashed brown line in (d) shows a threshold for the first order derivative. The first peak or valley with absolute value above the threshold is considered an interface point (red cross).

Consider the blue rectangle and the gray object in figure 3(a) to be a surface model and a volumetric dataset respectively. Black spheres represent surface points and a pair of red and green lines originating from each surface point indicate the space in which we search for a corresponding point in the volume data. The space is

larger for surface points in high curvature regions (see the surface point at the corner in figure 3(a)).

For each triangle of the surface model we evaluate a facet normal and the three vertex normals. The angle between the facet normal and each of the vertex normals is computed and the maximum of the three angles is stored. We call the maximum angle evaluated as search-angle henceforth. The search-angle indicates the local curvature of the surface model. In areas of high curvature, a large search-angle will be calculated whereas the search-angle will approach zero in planar areas of the surface model. In figure 3(b) we indicate the search-angle as a red arc between the facet normal (black arrow) and one of the vertex normals (green arrow) of the blue triangle. Using the search-angle we can construct a double cone with the opening angle set to twice the search-angle. The double cone is depicted in figure 3(b) with the apex of the double cone placed on the surface of a triangle.

We then extract the spatial locations and the normal vectors for a set of uniformly distributed surface points on the triangles of the surface model. At each surface point the apex of a double cone is placed and the cone axis is oriented along the surface normal. The triangle therefore bisects the double cone at its apex (figure 3(b)). We call the nappe of the double cone that lies on the front face of the triangle as outside nappe, while the nappe on the back face of the triangle is called inside nappe. The double cone defines a region in which we can search for an interface point in the volumetric dataset. An interface point found inside the double cone will be associated with the surface point of the triangle for further computations.

In order to search for an interface point in the volume data, we start from the surface point and traverse the volume data along various rays distributed inside the double cone. The rays originate from the surface point and are directed towards the two bases of the double cone. The density profile of each ray is used to identify the interface point as the position with highest/lowest gradient magnitude (first order derivative is a maximum/minimum and the second order derivative is zero). The gradient magnitude must be greater than a user specified threshold to be considered an interface point. Thresholding is necessary to filter out noise. The interface point with minimum distance to the surface point is stored for further processing. A density profile of a ray is illustrated in figure 3(c). The graph of the first order derivative of such a density profile is drawn as a blue curve in figure 3(d). The dashed brown line shows a threshold for the first order derivatives. The first peak or valley with absolute derivative above the

threshold is considered an interface point in the volumetric dataset. The interface point is depicted as a red cross in figure 3(c) and (d).

As we find an interface point in the volumetric dataset, we store its spatial location, the nappe (inside or outside) in which the interface point was found and the gradient. The information extracted from the surface model and the CT dataset provide all the required parameters to evaluate the metrics for distance analysis and normal analysis.

3.2.1 Distance Analysis

The computationally intensive step of finding a corresponding interface point in the volume data for each surface point on the surface model has already been done. The distance analysis shows the deviation between the datasets as Euclidean distances. For two datasets without any deviations, the spatial locations of the surface points and the corresponding interface points should be exactly the same. We compute the differences between the spatial locations on the surface model and their corresponding interface points in the CT dataset. We also have information about the nappe of the double cone in which the interface point was found. Using this information we color code the dataset for distance analysis.

Figure 4(a), (b) and (c) show test-part-1, test-part-2, and test-part-3 respectively. The surface model of test-part-2 consists of 152,054 triangles and its volumetric dataset has a resolution of 408x351x355. The surface model of test-part-3 has 11,424 triangles and the resolution of the volume data is 256x256x256. The test-parts are rendered using distance analysis with distances measured in millimeters. The color scale used for color coding is shown on the right of figure 4. The distance has positive sign if the interface point is found in the inside nappe of the double cone.

Investigations have shown glyphs to be better for visualizing variances than color coding alone [15]. Figure 5 shows test-part-4 rendered using our distance analysis technique. The surface model of test-part-4 consists of 25,880 triangles and its volumetric dataset has a resolution of 329x527x181. We render distance glyphs on the zoom-in of the user specified area (black rectangle). The arrow of the distance glyph is aligned with the normal vector of the surface and the diameter of the disc is proportional to the base of the double cone. The color of the disc indicates if the deviation was found in the inside nappe (yellow), outside nappe (blue) or no deviation was recorded (white).

We only consider the minimum distance between the surface model and the interface of the volumetric dataset

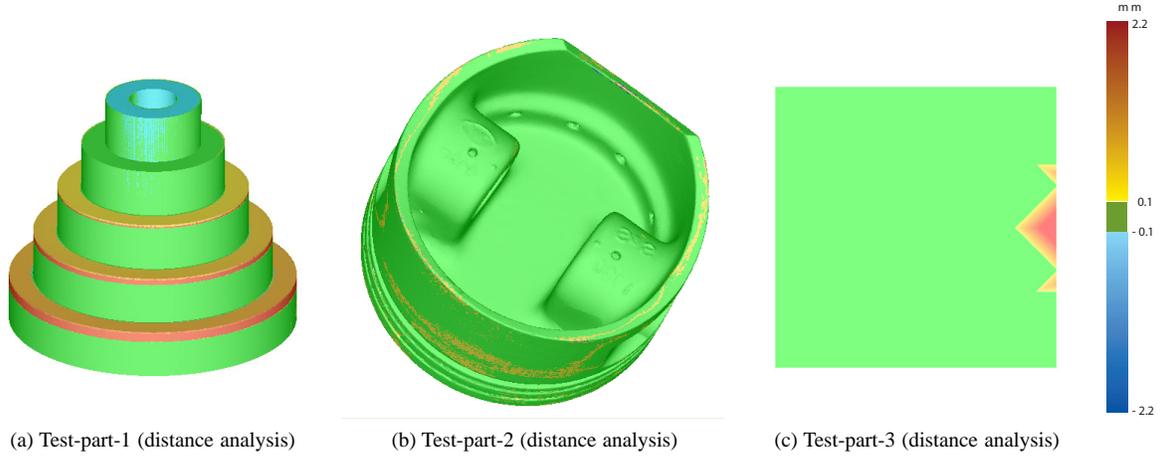


Figure 4: Test-part-1 (a), test-part-2 (b) and test-part-3 (c) rendered using distance analysis. The distances are measured in millimeters and the view port is set to 512x512.

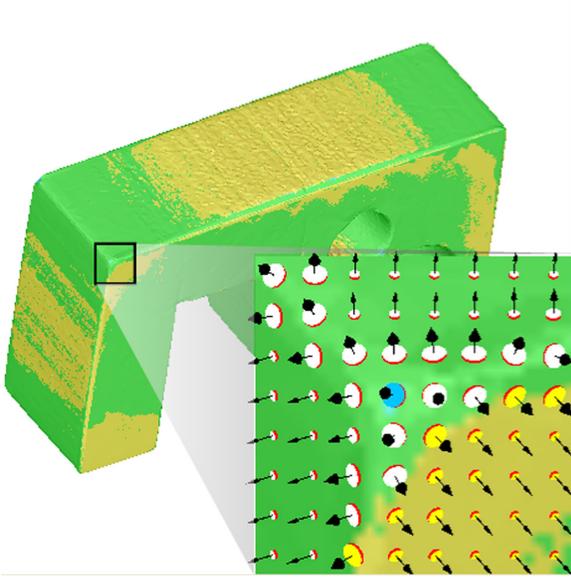


Figure 5: Distance glyphs are rendered on the zoom-in of test-part-4. The color scale is the same as in figure 4.

for distance analysis. The technique does not take the interface shape into consideration. The results have uncertainty in high curvature regions which needs to be highlighted. For uncertainty measurement we look for the maximum distance from the surface point to the interface in the volume data. The search for the maximum distance is conducted in the neighborhood of the ray along which the minimum distance was found. The neighborhood for searching the maximum distance has a radius of one voxel. We choose a radius of one voxel, as

the search space should be smaller than the smallest feature in the dataset and any feature less than the size of a voxel is not detectable in the volumetric dataset anyway.

Figure 6(a) illustrates the uncertainty measurement process. Let us assume that the closest interface point was found along the ray which starts from the surface point (cone vertex) and extends towards the black sphere depicted on the base of the cone. In the neighborhood around that ray (brown cone), we search for an interface point with maximum distance to the surface point. The difference between the minimum and the maximum distance from the surface point to the interface in the volume data is considered the uncertainty of the distance measurement.

The uncertainty in the case of test-part-4 is shown in figure 6(b). It becomes apparent that areas of high curvature and rough surfaces which are highlighted using a dotted and a dashed oval respectively, have higher uncertainty.

3.2.2 Normal Analysis

Normal analysis is proposed as an efficient method to compare surface smoothness. Normal analysis compares the orientation of the normal vectors extracted from the surface model with the gradients obtained from the CT dataset. The angle between the normal vector and the gradient indicates the variance in the curvature of the surface model and the interface of the CT dataset. Normal analysis is easy and efficient to compute given that the surface points and the corresponding interface points are already evaluated.

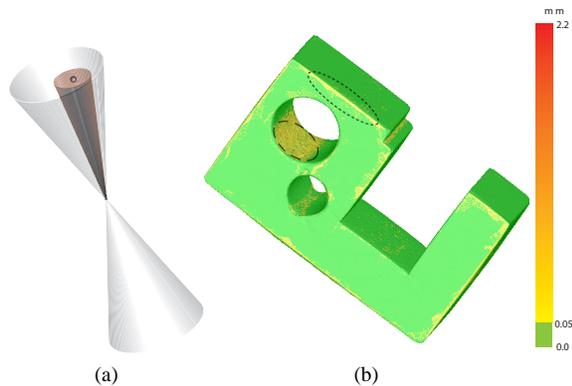


Figure 6: The maximum distance to the interface point is searched in the neighborhood (brown cone) of the ray (cone vertex to black sphere) along which the minimum distance to the interface point was recorded. Uncertainty rendering for test-part-4 is shown in (b). A dotted and a dashed oval highlight areas of high curvature and rough surfaces respectively.

The type of variance shown by normal analysis may pass undetected by distance analysis. Consider the black surface in figure 7(a) to be on a surface model with the normal vector indicated by a black arrow. The interface of the volume data (blue surface) overlaps the surface model in the area marked with a red oval. The distance analysis will report no deviation in such a case, however, there is a deviation in the orientation of the two datasets as the normal vector and the gradient do not point in the same direction. Such deviations can be emphasized correctly using normal analysis. Normal analysis will report a constant deviation along the entire surface in this example.

Figure 7(b) shows test-part-1 rendered using normal analysis. Normal analysis detects deviations at the edges and the rough surfaces. As the volumetric dataset is generated from an industrial process, it does not match the smoothness and exactness of the surface model, especially at the edges. The zoom-in in figure 7(b) shows that the top of test-part-1 has a rough surface. This roughness is not visible using distance analysis in figure 4(a). The color scale can be changed dynamically by the user. We emphasized the response of the normal deviation by appropriately setting the color scale as shown in figure 7(b).

3.3 Visual-Driven Variance Comparison

Visual-driven variance comparison techniques are grouped into ray profile analysis and magic lens dis-

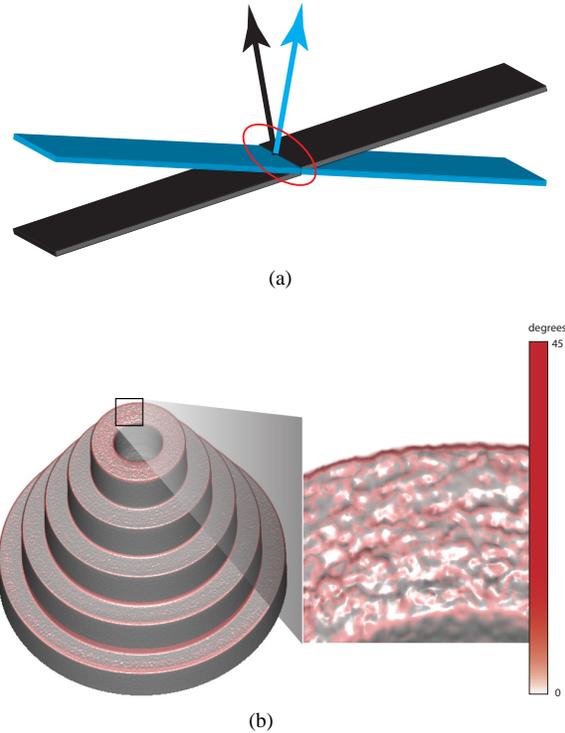


Figure 7: (a) Normal analysis emphasizes deviations in the orientation of the surface model and the interface of the volume data. (b) Test-part-1 rendered using normal analysis. The zoom-in shows roughness at the top of the dataset. The roughness is not recognizable using distance analysis.

plays. Ray profile analysis displays the accurate deviation between the datasets both as a 2D plot and as quantitative numbers. A magic lens is used to zoom-in/out of the dataset and to view the variance graphically.

Ray profile display is generated by plotting the first derivative of the density values encountered by the ray. The location of the surface point is marked on top of the graph. The peaks and valleys in the graph show the interface points. The user can see the deviation between the surface points extracted from the surface model and the interface points evaluated from the ray profile. This provides precise information about the deviations in the datasets.

Figure 8 shows two ray profiles generated from test-part-4 locations marked with two black crosses. The vertical red lines depict the surface points on the surface model. The blue graph shows the first order derivatives of the density values encountered by the ray, along which the interface point in the volume data was found. The peaks in the blue graphs are the edges detected in the volume dataset. The horizontal distance between a

peak and the vertical red line indicates the deviation between the datasets.

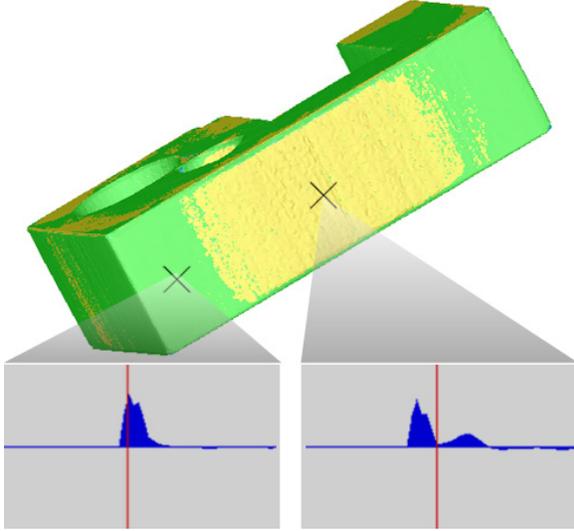


Figure 8: Two ray profiles are extracted from the locations marked with crosses on test-part-4. The ray profile on the left shows no variance whereas the ray profile on the right reports a deviation of 0.2 mm . The horizontal difference between the red line (surface point) and the peaks in the graph (interface points) depicts the deviation between the datasets.

In the ray profile on the left in figure 8, we observe that the surface point (red line) and the interface point (peak) overlap and thus there is no deviation between the two datasets. The ray profile on the right in figure 8 however shows a deviation of the surface model from the volumetric dataset as there is a horizontal difference between the red line and the nearest peak. Our system reported a deviation of 0.2 mm .

A magic lens can provide a precise graphical view of the variance by means of 3D box plots. The 3D box plots are rendered in a user specified area and each 3D box plot shows the minimum, the maximum, the mean and the standard deviation of the variance between the two datasets at each location. The box plots are oriented along the normal vectors of the surface model and the diameter of the box plots is directly proportional to the base of the double cone in which the interface point was searched (see figure 9). They therefore encode distance values, uncertainty and the dependent variables (normal vector and the base of the double cone).

Figure 10 shows 3D box plots over a user specified area (black rectangle) on test-part-4. The length and diameter of a cylinder depict the standard deviation of the variance between datasets and the base of the double

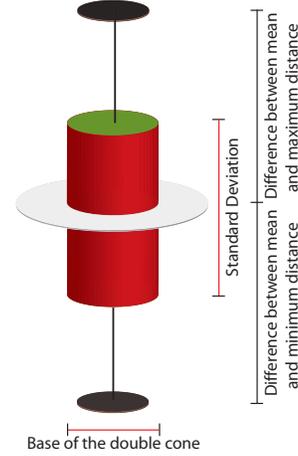


Figure 9: 3D box plot: combined visualization for distance analysis and uncertainty.

cone respectively. The white ring around the cylinder depicts the mean distance between the surface point and the interface of the volume data. The black discs at the bottom and on top of the box plot show the minimum and the maximum distances respectively. A 3D box plot similar to a plane disc indicates no difference in the minimum, maximum and the mean distance recorded between the two datasets. The measurement is most certain in areas where flat 3D box plots are rendered. The difference between the mean and the minimum distance and the maximum and the minimum distance indicate the reliability of the distance analysis.

4 Results

We implemented our variance comparison system on a Pentium 4, 3.4 GHz CPU and an NVidia GeForce 8800 graphics board. The system renders the volumetric data and the surface model side by side (volume view, surface view). Both the views are synchronized and any operation (rotation, scaling, ray profile analysis etc.) on one view modifies the other view accordingly.

We maintain a central queue for the events performed in the synchronized mode. An operation initiated on one of the datasets, also pushes an event into the central event queue and releases a signal. The other dataset pops the event from the queue and executes it. We implement first-come, first-served scheduling for the central queue. As we perform all the rendering (of surface model and of the CT dataset) and computationally intensive operations on the GPU, the complete system operates in real time and the convoy effect is successfully avoided.

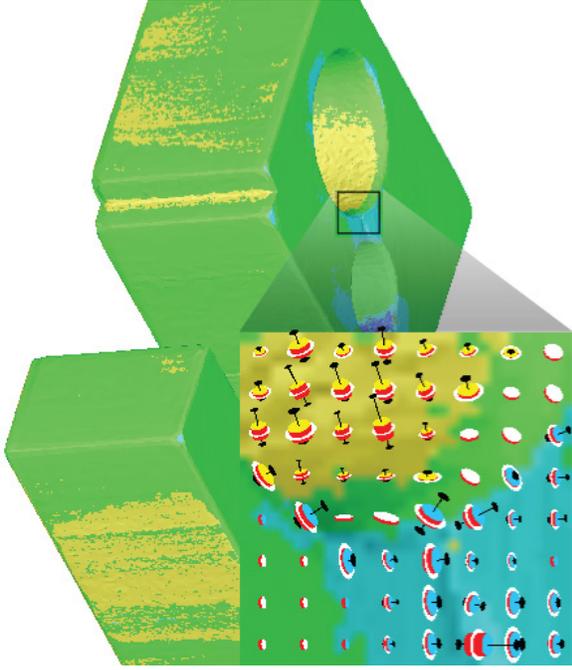


Figure 10: 3D box plots over a user selected area of interest (black rectangle). Box plots containing long cylinders (large standard deviation) and large differences between minimum and maximum distances depict areas of higher uncertainty. The size of the box plots can be scaled dynamically.

The industrial computed tomography of mechanical parts consists of fabrication artifacts and measurement errors. Quality assurance engineers are primarily concerned about the fabrication artifacts as they are introduced during the manufacturing process and should be minimized for a high quality product. Measurement errors are introduced into the volumetric dataset by the measurement process. Both types of errors are present simultaneously in any industrial CT scan. We can however minimize reconstruction artifacts and perform registration with high accuracy. This ensures that the reconstruction and the registration process do not have any serious impact on the measurement of the fabrication artifacts.

We evaluate our registration algorithm by successively performing registration between test-part-4 and a feature preserving mesh [6] of test-part-4. We use a feature preserving mesh for testing purposes so that the fabrication artifacts and the measurement errors are minimized and we can monitor the registration error. We measure the mean square error between the mesh and the test-part-4 (see figure 11) and recorded an average registration error of 0.0152 mm . We specified 6 control points

on both the datasets and the registration algorithm converged in 3.5 iterations on average.

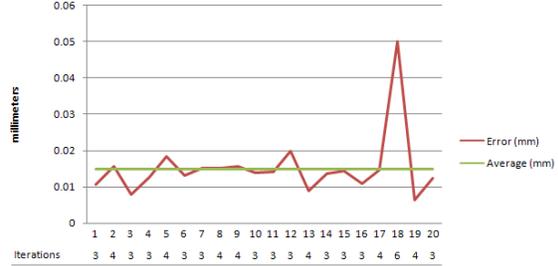


Figure 11: Mean square error produced by point-set to point-set registration on test-part-4 ($60 \times 100 \times 30 \text{ mm}$). Graph shows the results of 20 such experiments.

Experiment number 18 produced a very high error compared to the rest of the experiments. The ICP registration algorithm requires user interaction and the large error in experiment 18 is produced due to a bad specification of control points. The maximum fabrication and measurement artifact found in test-part-4 is 1.93 mm and the mean variance recorded is 0.27 mm . The average registration error introduced by the ICP algorithm is considerably lower than the mean and the maximum deviation in the dataset.

Reconstruction artifacts are introduced while extracting a mesh from a volumetric dataset. We use a synthetic dataset with known fabrication artifacts to evaluate our technique. Measurement and registration errors are not present in a synthetic dataset and this provides a good opportunity to analyze the effect of reconstruction artifacts on the measurement of fabrication artifacts. Figure 12(a) shows a surface model of a cube dataset and (b) shows a volume datasets with known fabrication artifacts. Fabrication artifacts are marked with an oval. The surface model consists of 12,288 triangles and the volume data has a resolution of $256 \times 256 \times 256$.

We generated a feature preserving mesh [6] of the volumetric dataset and compared it with the surface model using Geomagic Qualify (see figure 12(c)). Reconstruction artifacts are visible in the deviations shown by both zoom-ins. A deviation is also reported at the vertical edge of the mesh (zoom-in at the bottom) even though there is no variance there. The deviation at the edges is purely caused by the reconstruction artifacts. Figure 12(d) shows the variance comparison using our system. Our system correctly calculates no deviation on the edges (zoom-in at the bottom). The fabrication artifacts in the volumetric dataset are also reported correctly (zoom-in on top). The color coding is smooth and we do not observe any reconstruction artifacts.

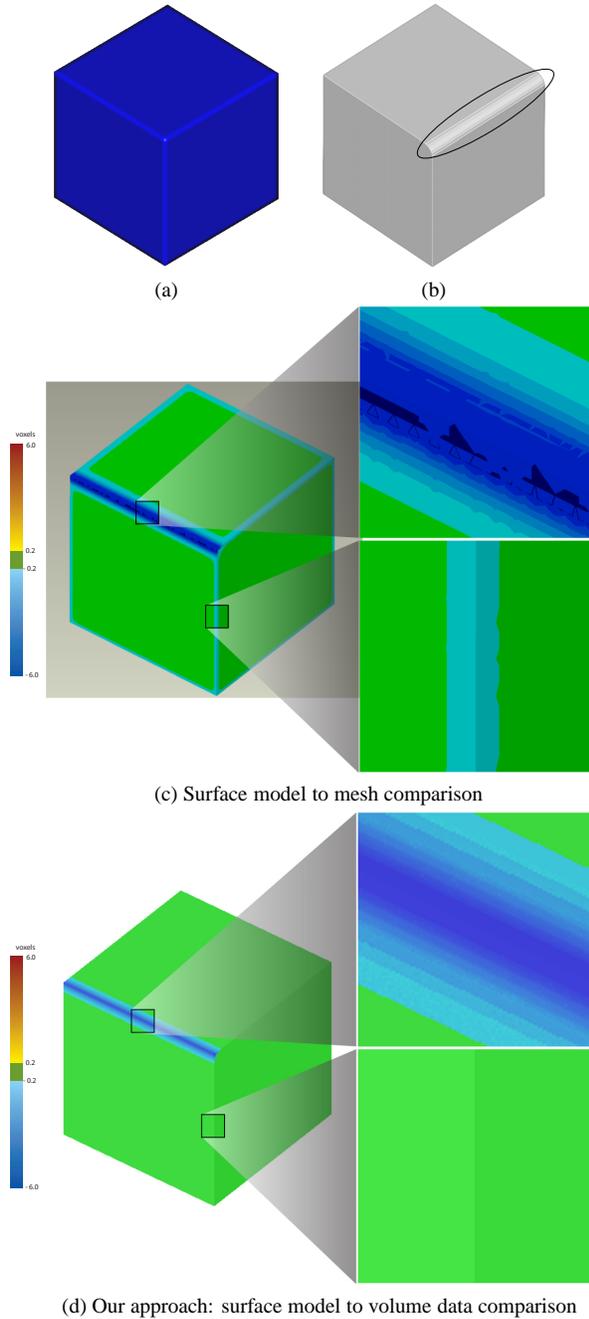


Figure 12: (a) Surface model. (b) Volume data with known fabrication artifacts. Artifacts are highlighted with a black oval. (c) Variance comparison between a surface model and a feature preserving mesh of (b). (d) Direct variance comparison between a surface model and the volume data.

The comparison of the maximum and average deviation evaluated by Geomagic Quality and our technique is

given in table 1. Our method calculates the deviation very close to the original. Geomagic Quality reports the maximum deviation close to the ground truth but the average deviation has a large error. Reconstruction of a mesh from the volume data introduces artifacts distributed over the entire mesh. This is why the average error reported by Geomagic Quality is very small compared to the original deviation. As we avoid reconstruction artifacts, our calculations are more accurate.

Table 1: Maximum and average deviation (voxels) reported by Geomagic Quality and our system.

	Ground truth	Geomagic	Our technique
Maximum	8.485	7.95	8.91
Average	3.42	0.195	3.51

The earlier solutions proposed for variance comparison are divided into two major steps. For instance, Heinzl et al. [6] propose a robust surface detection pipeline for effective variance comparison. They extract a feature preserving mesh from the raw dataset and then perform comparison between the surface model and the extracted mesh. The mesh extraction part consists of a four step pipeline. In the first three steps, an anisotropic diffusion filter, a gradient filter and a watershed segmentation are in turn applied to the raw dataset. In the final step before generating the mesh, constrained elastic nets are applied. The mesh is then compared to the surface model using some existing variance comparison tool. We combine the entire variance comparison process into a single, interactive system.

Table 2 shows the performance of our system, in comparison to the robust surface detection pipeline [6]. The bottle neck in earlier methods has been the surface extraction process. The surface extraction took very long compared to the comparison process because of parameter tweaking. Our method is more automated and requires much less user interaction. A screen shot of our system is shown in figure 13.

Table 2: Comparison of the performance of our system.

	Test-part-1	Test-part-4
Distance analysis (our method)	0.051 sec	0.033 sec
Robust surface detection pipeline	10.23 min	4.58 min
Distance analysis (Geomagic)	9.31 sec	8.51 sec

Distance glyphs and the 3D box plots are additional visualizing techniques for showing deviations and uncertainties. We showed our work to two domain experts to calibrate the usefulness of our visualization techniques. Both of the experts have used various variance comparison systems in their professional capacity. They were

provided a short explanation of the functionality of the distance glyphs and 3D box plots and then they experienced our proposed visualization techniques on a known dataset.

They were both very interested in using distance glyphs to visualize deviations as compared to color coding alone. They acknowledged that they acquired a lot more valuable information about the surface (surface normal), the measurement process (base of the double cone) and the deviations using distance glyphs. One expert liked the idea of showing glyphs in a user specified area whereas the other expert had no specific preferences.

The experts also appreciated the idea of showing the uncertainty of the measurement process along with the distance analysis. They were able to retrieve all the intended information from the 3D box plots but one expert proposed to somewhat reduce the amount of information in the 3D box plots. Both of them however were interested in using 3D box plots in their present form due to the combined visualization of uncertainty and distance analysis.

All the information shown by distance glyphs is included in the 3D box plots as well. Nevertheless, both experts were equally interested in using both the techniques in their professional work. One of the experts preferred 3D box plots over distance glyphs.

5 Conclusion and Future Work

We have presented a variance comparison system that compares a surface model directly to the industrial CT scan of specimens, especially in the preproduction phase and for first part inspection of new industrial products. We avoid intermediate steps for data enhancement and

surface extraction. Two sets of tools, namely geometry-driven and visual-driven techniques provide comprehensive comparison opportunities. The system is implemented on graphics hardware and all the proposed methods work in real time.

In the future we intend to expand the toolset by plugging in more volume visualization and exploration techniques. For instance, feature extraction methods like region growing and 2D histograms can help inspect the concealed structures within the CT data [7]. We use a semi-automatic algorithm for registration. We achieve high accuracy but more automatic registration techniques can be further investigated. Feature-based registration algorithms can be modified to register a surface model to a volumetric dataset.

References

- [1] N. Aspert, D. Santa-Cruz, and T. Ebrahimi. Mesh: Measuring errors between surfaces using the Hausdorff distance. In *Proceedings of the IEEE International Conference on Multimedia and Expo*, volume I, pages 705–708, 2002.
- [2] P. J. Besl and N. D. McKay. A method for registration of 3d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [3] P. Cignoni, C. Montani, and R. Scopigno. A comparison of mesh simplification algorithms. *Computers and Graphics*, 22(1):37–54, 1998.
- [4] P. Cignoni, C. Rocchini, and R. Scopigno. Metro: Measuring error on simplified surfaces. *Computer Graphics Forum*, 17(2):167–174, 1998.
- [5] C. Heinzl, J. Kastner, and E. Gröller. Surface extraction from multi-material components for metrology using dual energy CT. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1520–1527, 2007.
- [6] C. Heinzl, R. Klingsberger, J. Kastner, and M. E. Gröller. Robust surface detection for variance comparison and dimensional measurement. In *Proceedings of Eurographics / IEEE VGTC Symposium on Visualization*, pages 75–82, 2006.
- [7] R. Huang, K.-L. Ma, P. McCormick, and W. Ward. Visualizing industrial CT volume data for nondestructive testing applications. In *Proceedings of IEEE Visualization 2003 (VIS'03)*, pages 547–554, 2003.

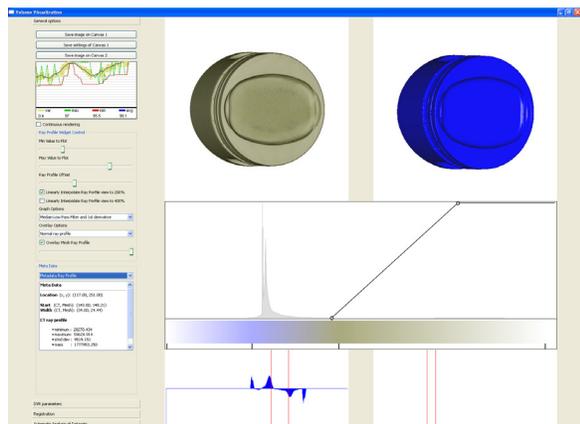


Figure 13: Variance comparison system.

-
- [8] L. P. Kobbelt, M. Botsch, U. Schwanecke, and H.-P. Seidel. Feature-sensitive surface extraction from volume data. In *SIGGRAPH 2001, Computer Graphics Proceedings*, pages 57–66. ACM Press / ACM SIGGRAPH, 2001.
- [9] J. Maintz and M. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
- [10] N. Nikolaidis and I. Pitas. *3-D Image Processing Algorithms*. Wiley-Interscience, 2000.
- [11] E. Pichon, D. Nain, and M. Niethammer. A Laplace equation approach for shape comparison. In *Proceedings of SPIE Medical Imaging*, 2006.
- [12] Raindrop. Raindrop: The magic of making it simple. March 2007. <http://www.geomagic.com/en/products/qualify/>.
- [13] G. Turk. Re-tiling polygonal surfaces. *ACM SIGGRAPH Computer Graphics*, 26(2):55–64, 1992.
- [14] Volume-Graphics. *VGStudio Max 1.2 - Users Manual*. 2004. <http://www.volumegraphics.com/products/vgstudiomax/index.html>.
- [15] C. Weigle and R. M. Taylor. Visualizing intersecting surfaces with nested-surface techniques. *Visualization, 2005. VIS 05. IEEE*, pages 503–510, 2005.
- [16] L. Zhou and A. Pang. Metrics and visualization tools for surface mesh comparison. In *Proceedings of the SPIE Visual Data Exploration and Analysis VIII*, volume 4302, pages 99–110, 2001.