Data Level Comparison
of Wind Tunnel and Computational Fluid Dynamics Data

Qin Shen, Alex Pang
Computer Science Department, UCSC
Sam Uselton
MRJ Technology Solutions
jeanshen@cse.ucsc.edu, pang@cse.ucsc.edu, uselton@nas.nasa.gov

Abstract

This paper describes the architecture of a data level comparative visualization system and experiences using it to study computational fluid dynamics data and experimental wind tunnel data. We illustrate how the system can be used to compare data sets from different sources, data sets with different resolutions and data sets computed using different mathematical models of fluid flow. Suggested improvements to the system based on users' feedback are also discussed.

1 INTRODUCTION

Most comparative visualization studies use juxtaposition where the user is presented with side-by-side images. This is effective only when the differences are significant and easy to detect. Beyond that, it unduly burdens the viewer with the task of finding then estimating the location and degree of difference. Two other comparative visualization techniques are popular: (a) difference images, and (b) overlays or superimposition. Both of these are more effective than juxtaposition. However, they do not provide any additional investigative or analytical capability.

This paper presents a data level comparative system that builds upon these popular techniques. In particular, the users can compare two data sets using a variety of metrics calculated from raw or derived data sets, and visualize them in a variety of ways.

We report on how this system is used to evaluate four aeronautics data sets using a variety of pairwise comparisons. In particular, we demonstrate how this system shows regions of small differences that would be difficult to discern using image level techniques, and the versatility of this system to compare other gridded data sets.

2 COMPARISON FRAMEWORK

2.1 Definition

We distinguish between three general types of comparisons: image level, data level, and feature level. *Image level comparison* [9] techniques are those that use images as the starting point for comparison, particularly images produced by visualization of data. Examples that fall under this caterogy are: side-by-side, image differencing, most Fourier analysis, and summary image statistics such as root mean square. *Data level comparison* [7] techniques are those that use raw data as the starting point for comparison. The main advantages of this approach are the ability to generate different metrics for comparing raw data, access to intermediate calculations while generating the derived data, and apply visualization techniques to the data produced by the comparison. The main drawback of this approach is that comparison systems become less general and more application specific; the choice of operators and metrics must be made in the context of the application. For example, in certain situations such as when comparing direct volume rendering algorithms [5], a reference model may be necessary. Data level comparison does not preclude the use of images when images are the raw data – e.g. pressure sensitive paint images [8] and images of oil streaks [7]. *Feature level comparison* techniques are those that compare extracted features. Features may be represented using geometry produced by some visualization technique such as isosurfaces (e.g. iso-density to represent front of shockwave), streamlines and ribbons (e.g. location of critical or degenerate points), etc. Feature comparisons are usually limited to higher dimensional analogs of image level comparisons, such as side-by-side (e.g. ribbon comparison [7]) or overlays (e.g. vortex core comparison [4]). More recently, feature level comparisons of streamlines are carried out with glyphs and animation [6]. Feature comparisons at a higher semantic level, incorporating application specific knowledge, are of great interest to users.

This paper focuses on the data level comparison capabilities of our system.

2.2 Architecture

Our data sets are in PLOT3D format. Each data set comes in two separate files – the grid file specifies the physical position of each point in the grid, while the solution file specifies up to five physical parameters associated with each grid point. For data level comparison, we want to compare physical parameter values at the same physical coordinates. With this in mind, we base the comparisons on an intermediate mesh or grid (see Figure 1). The intermediate mesh is any grid used to resample values from both data sets. Either of the original grids can be used as the intermediate mesh. In our implementation, a regularly gridded mesh is also provided as a third choice for an intermediate mesh. This mesh may be interactively resized, refined, and positioned over a volume of interest. This ability allows one to zoom in and examine an area in more detail.

The derived field(s) on the intermediate mesh result from a variety of metrics that the user can select from. Examples include: absolute or signed differences between the two input data sets, minimum or maximum values of the two input data sets, or the average between the two input data sets. In general, the metric for the derived field will be any function of the two input data sets that would make sense in the comparison task. For example, to do data level comparison between two vector fields, the dot product between the vector components measures similarity between the two data sets. Currently, the derived fields are selected from a menu of pre-programmed functions. As part of our planned enhancements, users can interactively enter their functions similar to the calculator module of FAST [1]. These functions may be used to generate new derived fields as well as obtain metrics from the derived fields.

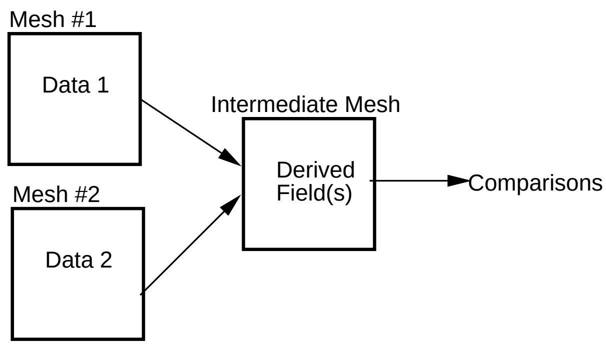

Figure 1: Schematic of data level comparison architecture.

2.3 Implementation

The major task necessary for the comparison framework above is support for point location. That is, given the physical coordinate of each point in the intermediate grid, we need to find the cell that contains that point from each of the two data sets, and then interpolate the value at that point from the values of the surrounding data points. This is one of the main tasks that the Field Encapsulation Library (FEL) [2] supports, providing a uniform interface to a wide variety of grid types. This capability makes it a natural choice for building this tool. In addition, it allows us to investigate the errors introduced by resampling the data to an intermediate grid by experimenting with different interpolation methods and different intermediate grids. Furthermore, FEL supports different data file readers thereby accommodating those wind tunnel experiment data which are not in PLOT3D format. In addition to FEL, our implementation is coded in C++, and uses OpenGL for graphics and xforms [10] for its user interface. Development and testing is carried out on SGI platforms.

Figure 5 (see color plate) shows the user interface for our comparison system. It combines the side-by-side comparison of the data in the left panels, and the data level comparison using the intermediate mesh on the upper right. The lower right panel varies depending on the visualization technique selected, and is used primarily for controlling the comparison window on the upper right panel. Different visualization techniques can be used to visualize the different metrics available to the user. Some of these are demonstrated in Section 4.

Some of the standard “bread and butter” visualization techniques that are supported by this system are cutting planes, isosurfaces, image-based operators and histograms. Cutting planes can be moved, added and deleted. In addition, they can be oriented along any of the three major axes. To calculate isosurfaces, we use a grid independent isosurface extraction routine from the VisTech library – a collection of visualization techniques being built on top of FEL. For small data sets, the user can interactively create new isosurfaces in real time by adjusting the threshold using a slider (see Figure 5 in the color plates). For larger data sets (e.g. the 1.5 and 2.5 million point data sets), however, it takes somewhat longer. Several operators for images or slices are currently available, including signed and unsigned differences, and applying different colormaps to the resulting difference image.

In addition to the standard visualization techniques, an extension was made to make histograms pickable. If the user picks one of the histogram color bars, the range of values for that bin is selected as “active”. Information about global minimum and maximum, active minimum and maximum, and percentage of points in the active range appear on the screen. More importantly, all data points outside of the selected range are made transparent. This is very useful for identifying regions that have data within the selected range of values. Based on user feedback, we have also added the ability to obtain data values directly from the image by clicking on it.

In all cases, the user can apply any one of eight colormaps to both original data sets and the comparison output. The same colormap applies to the input data sets, while a separate colormap applies to the comparison in the intermediate mesh. This separation is essential since they have very different ranges of values.

3 DATA SETS

We use this system to investigate four different data sets. Three of these are from CFD simulations while the fourth is from measurements taken during a wind tunnel experiment. More detailed information about these data sets are available from [3]. These data sets are briefly described below:

1. Experimental. 6,699 grid points (dimension: 11 x 21 x 29). A seven-hole pressure probe was used to measure the velocity and pressure.

2. Computational. 1.5 million grid points (dimension: 115 x 157 x 83) using the Baldwin-Barth (BB) turbulence model.

3. Computational. 1.5 million grid points (dimension: 115 x 157 x 83) using the Spalart-Allmaras (SA) turbulence model. Same grid as previous data set.

4. Computational. 2.5 million grid points (dimension: 115 x 189 x 115) using the SA turbulence model.

In all cases, the computed solutions used the same wing geometry, boundary conditions and initial conditions, and are all steady state solutions. The computational grids are hexahedral curvilinear grids commonly used in high fidelity CFD calculations. The experimental measurements used the same wing geometry and flow conditions. The measurements for each point were taken after sufficient time to allow any transients to die down. The numerical solutions were calculated over a period of several hours each on supercomputers at NASA Ames Research Center. The experimental measurements took several months to collect. The main feature of interest being investigated and validated by these modeling and experimental efforts is the vortex in the flow at the trailing edge of the wingtip.

4 RESULTS

The images in this section are selected to illustrate the different combination of data sets, and the variety of visualization techniques available in our comparative visualization system. In Figure 4 (see color plate), we compare CFD solutions using two different turbulence models over the same computational grid. We use cutting planes to demonstrate how data level comparison is superior to image level comparison.

Note that the differences between the two computed flow fields using two different turbulence models but on the same grid are quite subtle. The differences between the visualizations of the original data are difficult to see, and the histograms also show very little difference. However, the visualization of the data level comparison, in the upper right, shows some interesting variations. The larger magnitude differences near the edges of the cut planes are probably not significant, for two distinct reasons. First, CFD grids are constructed with high sampling density near surfaces, and low sampling density far from the body of interest, so the interpolated values near the edges are less reliable both because of edge effects and the lower sampling density. Secondly, from the application perspective, the reason why sampling is done at lower density far from the surface is that there is little interest in details of the flow there, so even if the larger differences are accurate, they are still less interesting than things near the wing.

The yellow spot just above the wing in the third slice from the right, and the yellow swirling structure near the wing in the two rightmost slices, are quite interesting. They indicate that the two turbulence models do produce different results in the flow velocity precisely in the region of interest at the wingtip.

These features of potential interest are barely visible in the image level comparison. This drawback of image differencing can be attributed to quantization errors. That is, in image differencing, data values are first quantized to color bits before they are compared. On the other hand, in data level comparison, data values are first compared before being quantized to color bits.

In Figure 5 (see color plates), we compare CFD solutions using the SA turbulence model but using two different grid resolutions (1.5 million and 2.5 million points). We show how isosurfaces are used in this comparison. The goal is to see the extent and location of errors when a lower resolution grid is used.

Figure 5 (see color plates) illustrates that while side by side comparison of the isosurfaces on the left panel does not reveal much information, the isosurfaces resulting from the difference between the y component of momentum clearly shows several distinct regions. In particular, the region at the trailing edge of the wingtip show substantial differences as expected. In addition, there are also significant differences near the walls that was not expected and hence more interesting. Note that the isosurfaces on the left are generated directly from the raw data and controlled by a single source data threshold slider, while the isosurfaces in the comparison window are generated from the difference between the two data sets and controlled by a separate comparison result threshold slider.

In Figure 2, we compare the low resolution CFD data using the SA turbulence model to the wind tunnel experiment data set. Large differences can be expected between these two data sets primarily because of the large difference in data resolution. The goal here is to see if gross features can be found in both. We use extensions to histograms to identify these regions. This is achieved by examining the location of different data ranges. Selecting a bin in the histogram by clicking on it will highlight regions in the data volume where those data reside. The middle and right columns illustrate two different bins and their corresponding data extents within the data volume.

Figure 3, illustrates another view that shows how data level comparison is superior to side-by-side comparison. Even with a low resolution intermediate mesh, it is quite easy to see the difference between the two data sets on the single cutting plane. The colormap in the comparison panel is re-adjusted each time a cutting plane is moved, a new plane added, or an existing plane removed. This allows the full range of values in the current comparison set to be visualized. In this image with a single cut plane, it is easy to see where points with maximal values are located. By zooming in to the trailing edge of the wing, by resizing and repositioning the intermediate mesh, one can examine this region further detail.

The electronic version of this paper and additional images are available at www.cse.ucsc.edu/research/avis/cmp.html.

5 LIMITATIONS

We are continuing to incorporate suggestions and improvements based on users' feedback. The most obvious and serious limitation to the current system is the resampling to the intermediate mesh. For this reason, we plan to investigate the errors introduced by different types of interpolation methods employed by the point location algorithm in FEL as a function of grid spacing. Another current limitation is that the software handles comparison only on two data sets at a time. Extending this to multiple data sets is straight forward but will complicate the metric calculations and the user interface. Beyond these, we are also looking at feature level comparison separately. In particular, we are extending previous work on streamline comparisons to other forms of flow visualization techniques. In addition, we are also investigating different ways of comparing isosurfaces.

6 CONCLUSIONS

We have described the features and limitations of our data level comparison system, and have demonstrated its utility and versatility on four different aeronautical data sets. This study shows that this approach is viable and particularly useful when the differences between two data sets become smaller. Hence, as experimental methods become more accurate and as modeling becomes more realistic, we believe our approach will become more important.

ACKNOWLEDGEMENTS

We would like to thank Travis Heppe for help with coding, Han-Wei Shen for his VisTech routines, Patrick Moran and Chris Henze for early versions of FEL2, Jennifer Dacles-Mariani for the computational data sets, and Greg Zilliac for the wind tunnel experiment data set. This project is supported by NASA grant NCC2-5207, NASA contract NAS2-14303, NSF grant IRI-9423881, DARPA grant N66001-97-8900, and ONR grant N00014-92-J-1807.

References

[1] FAST home page. science.nas.nasa.gov/Software/FAST.

[2] Stephen T. Bryson, David Kenwright, and Michael Gerald-Yamasaki. FEL: The Field Encapsulation Library. In R.D. Bergeron and A.E. Kaufman, editors, *Proceedings of Visualization 96*, pages 241–247. ACM, October 1996.

[3] Jennifer Dacles-Mariani, Gregory G. Zilliac, Jim S. Chow, and Peter Bradshaw. Numerical/experimental study of a wingtip vortex in the near field. *AIAA Journal*, 33(9):1561–1568, September 1995.

[4] David Kenwright and Robert Haimes. Vortex identification - applications in aerodynamics: A case study. In *Proceedings of Visualization 97*, pages 413–416, 1997.

[5] Kwansik Kim and Alex Pang. Ray-based data level comparison of direct volume rendering algorithms. Technical Report UCSC-CRL-97-15, UCSC Computer Science Department, 1997.

[6] S. K. Lodha, Alex Pang, Robert E. Sheehan, and Craig M. Wittenbrink. UFLOW: Visualizing uncertainty in fluid flow. In R.D. Bergeron and A.E. Kaufman, editors, *Proceedings of Visualization 96*, pages 249–254. ACM, October 1996.

[7] Hans-Georg Pagendarm and Frits H. Post. Studies in comparative visualization of flow features. In G. Nielson, H. Hagen, and H. Muller, editors, *Scientific Visualization: Overviews, Methodologies, Techniques*, pages 211–227. IEEE Computer Society, 1997.

[8] Samuel P. Uselton. exVis: Developing a wind tunnel data visualization tool. In *Proceedings of Visualization 97*, pages 417–420. IEEE, 1997.

[9] Peter L. Williams and Samuel P. Uselton. Foundations for measuring volume rendering quality. Technical Report NAS-96-021, NASA Numerical Aerospace Simulation, 1996. Also to appear in the Journal of Visualization and Computer Animation.

[10] T.C. Zhao. XFORMS home page. URL: http://bragg.phys.uwm.edu/xforms.

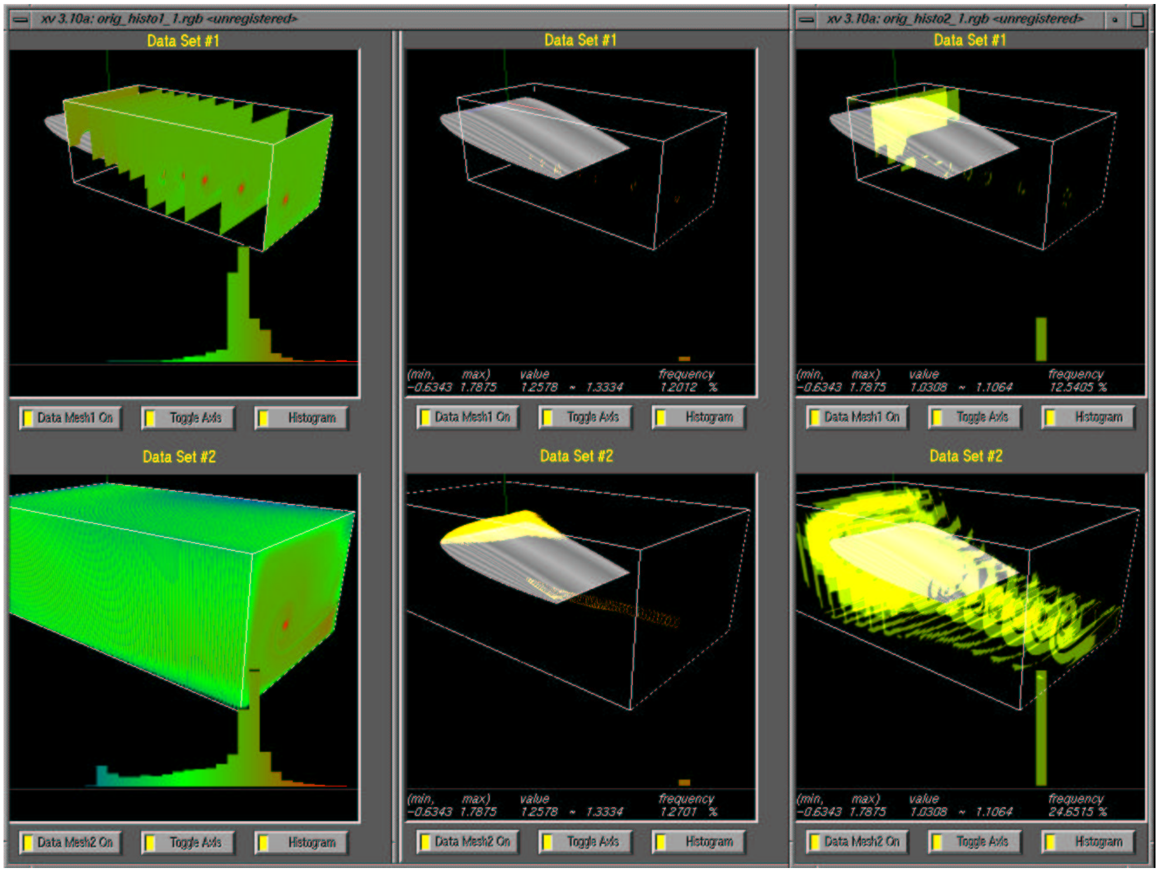

Figure 2: Comparison between the experimental wind tunnel data (top row) and the 1.5 million point SA turbulence model solution (bottom row). The middle and right columns show side by side comparisons of the raw data from these two data sets. In the middle column, only those data with values in the range of (1.2578 - 1.3334) are shown. With the exception of the front of the wing, the tube-like shape coming off the trailing edge of the wingtip are quite similar considering that the CFD data is more than two orders of magnitude denser than the experimental data. The right images show sections of the data corresponding to the values (1.0308 - 1.1064). The histogram bin corresponding to these values are also visible. For this range of values, the two data sets differ significantly.

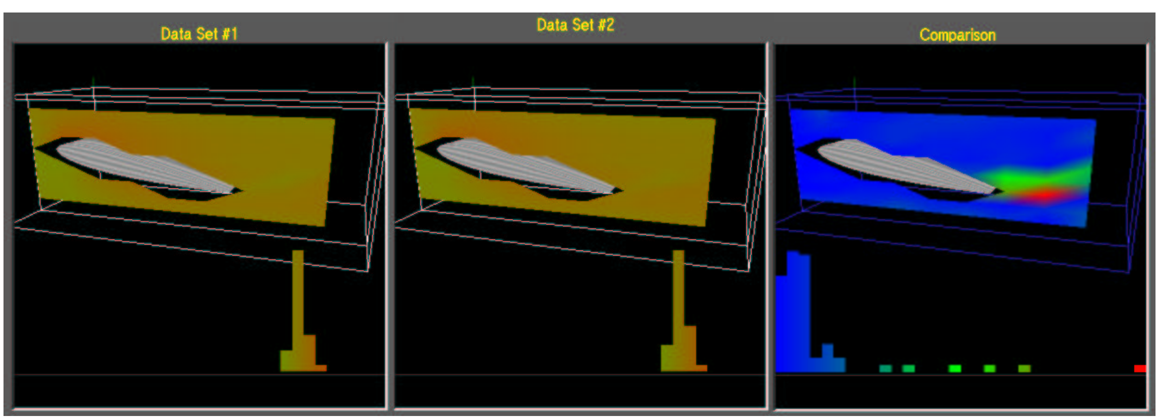

Figure 3: Comparison between a low resolution (1.5 million points) and a high resolution (2.5 million points) solution using the same SA turbulence model. Single cut plane along the flow direction. Even with the low resolution intermediate mesh (10 x 10 x 10), it is easy to see how the two data sets differ at the wake of the flow using data level comparison.

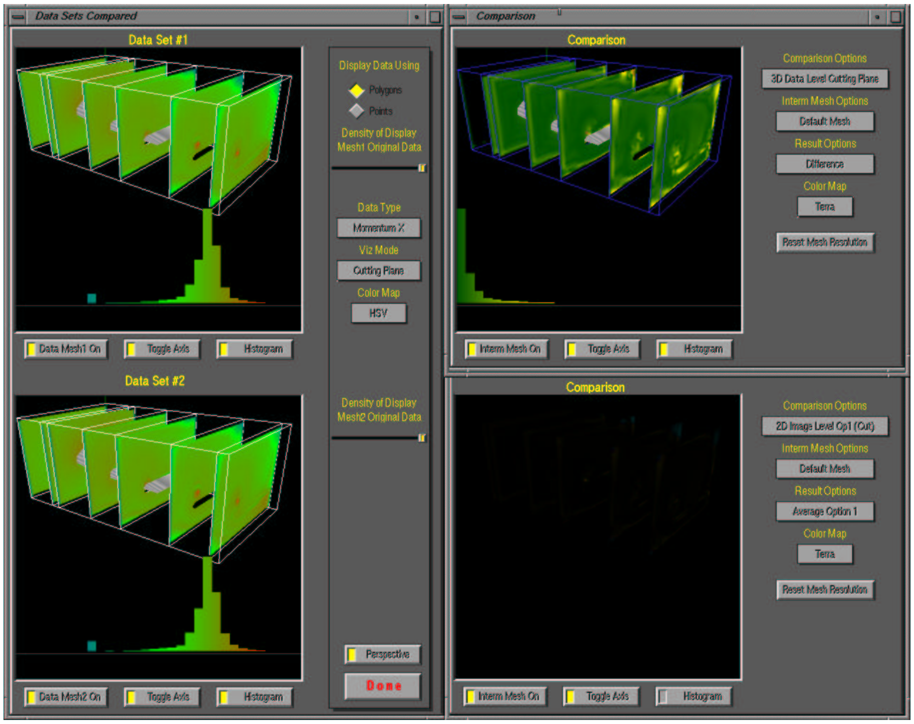

Figure 4: Comparison between two 1.5 million points data sets. The left panel shows selected planes from the CFD data sets. The top left image uses data from the SA turbulence model, while the bottom left uses data from the BB turbulence model. Difference in the solution between these two models are shown on the same cutting planes in the upper right image. The lower right image shows pixel by pixel image differencing between the two images on the left. The data on the left panels have the same range and hence use the same HSV color map, while the data on the right panels have the same range and hence use the same Terra color map.

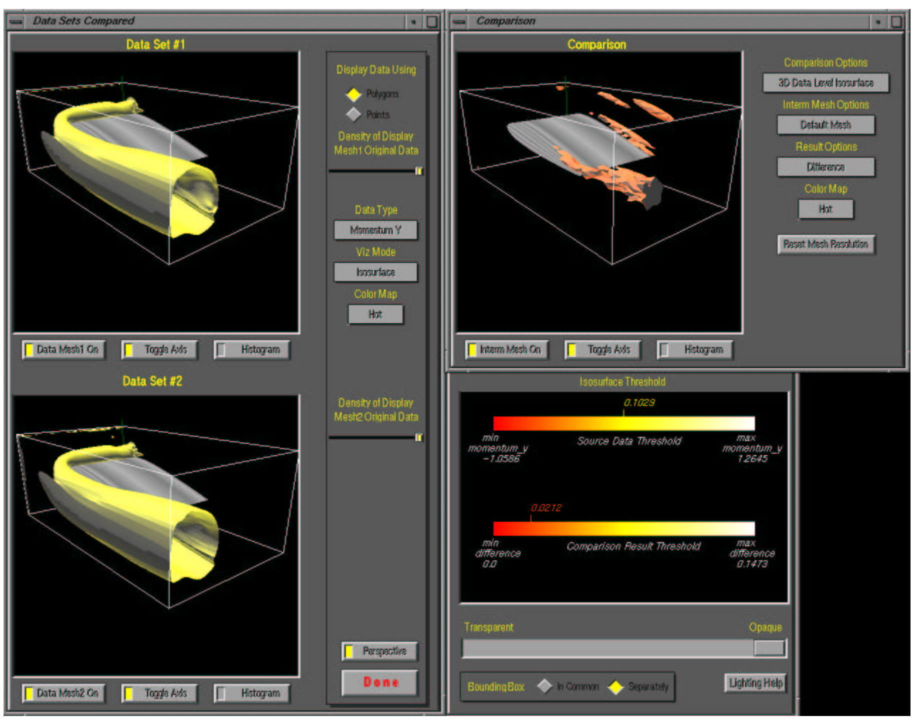

Figure 5: Comparison between a low resolution (1.5 million points) and a high resolution (2.5 million points) solution using the same SA turbulence model. The top left image is the low resolution data, while the lower left image is the high resolution data. Data range and threshold level for the isosurfaces on the left are displayed and controlled by the upper color coded slider on lower right image. We are seeing an isosurface of the y component of momentum on the left panel. On the right, we see an isosurface of the difference between the y momentum values between the two data sets. The threshold level for the differences are controlled separately by the lower slider on the lower right image. Since the data range for the left panels is different than those in the right panel, a different colormap may be used for the isosurfaces on the right.