Interactive Volume Visualization of Complex Flow Semantics

Helwig Hauser and Matej Mlejnek[*]

VRVis Research Center[†]
Vienna, Austria

Figure 1: Volume visualization representing a user's interest in a back-flow region and vortical structures (left side, more details in section 6.1) and in the in-flow of hot fluid in an extended T-junction (right side, more details in section 6.2). Volume rendering (back-flow region, left, and not-hot high pressure areas, and the right), a gradient-dependent iso-surface (316 K, right), and interest-dependent streamlines are used.

ABSTRACT

Comprehending results from 3D CFD simulation is a difficult task. In this paper, we present a semantics-based approach to feature-based volume rendering of 3D flow data. We make use of interactive feature specification to acquire derived data about what is most interesting to the user. This process results in so-called degree-of-interest (DOI) values, associated with the original data items. This information is then used during the visualization mapping to allow for visualization of flow semantics. We present three different approaches in this paper: (a) isosurfacing the degree of interest – the result is a triplet (or quintuplet) of iso-surfaces representing levels of interest; (b) feature volumes – volume rendering is used to depict 3D distributions of DOI values; and (c) interest-based seeding of streamlines, resulting in reduced clutter while focusing on the most interesting parts of the 3D flow. We utilize fast volume rendering (RTVR) for real-time viewing, also providing two-level volume rendering, which allows to seamlessly integrate all of the above-mentioned approaches.

CR Categories: I.6.6 [Simulation, Modelling, and Visualization]: simulation output analysis; I.6.9 [Simulation, Modelling, and Visualization]: visualization—flow visualization, information visualization, multivariate visualization, volume visualization;

Keywords: flow visualization, feature-based visualization, volume visualization

1 INTRODUCTION

The visualization of flow data-sets which result from computational simulation gains increasing importance due to the increased use of

[*]{Hauser,Mlejnek}@VRVis.at
[†]http://www.VRVis.at/

flow simulation in several fields of modern industry – a well-known example is the simulation of air flow around new flight vehicles.

In our case, the visualization users simulate fluid or gas flows within complex 3D geometries, such as combustion processes in the automotive environment.. Also, floods and avalanches (both hot topics in Austria) are investigated by means of computational simulation. Due to the size and the complexity of their data, user-oriented visualization is needed to help answering specific questions about the data during data analysis and exploration.

Also, visualization proves to be very useful when the simulation process itself is investigated – questions such as of whether the grid design was appropriate with respect to the simulated flow features and of whether the boundary conditions actually provided a realistic interface between the "world setting" of the simulation and its inner processes need to be checked after most runs of a simulation.

1.1 Flow data and visualization

The visualization of flow data still is a challenging field of research and applications. This is because usually lots of data-items need to be visualized (up to millions of those), often complex grid structures are involved (such as unstructured grids, for example), and real flows usually are three-dimensional and time-dependent (posing major challenges with 3D visualization such as tremendous resource requirements and handling of occlusion). Results from computational flow simulation also often provide several additional data attributes per data-item, requiring visualization techniques which are fit for multi-dimensional data.

To cope with the above mentioned challenges, visualization techniques are required, which focus on especially interesting sub-sets of the data. If not all of the data is shown simultaneously, occlusion problems are reduced in 3D visualization and resource requirements are moderated, accordingly. Feature-based visualization of

flow data is one very useful approach as well as focus+context visualization of data from flow simulation. While literature [20] already provides a set of very good solutions in feature-based visualization – examples are iconic visualization of flow features [21, 23] or the use of flow topology for visualization [7, 22] –, focus+context visualization of flow data only recently became an active field of research [3]. The work presented here build upon this latter approach.

With respect to the multi-dimensional character of data from computational flow simulation, visualization techniques which extend the usual domain of scientific visualization are needed. In information visualization we find many useful approaches to the visualization of multi-dimensional data [1, 12]. The integration of techniques from information visualization such as scatterplots, histograms, and parallel coordinates [8] already proved to be useful, also in scientific visualization [5].

1.2 Flow data and volume visualization

Volume rendering definitely is one of the most prominent fields of visualization research. Very useful results have been achieved in medical applications. Volume rendering also has been applied to data from flow simulation, however with limited success. The specification of proper transfer functions, for example, which already is a challenging problem in 3D medical visualization, even more is difficult in 3D flow visualization.

Probably the most important reason for this problem is, that medical data and flow data inherently have different data characteristics. Whereas in medical data we usually have tissue boundaries (which are to be shown in medical visualization), in flow data we often miss such (rather sharp) boundaries between sub-sets of the flow. Flow features often are smoothly delimited. Thus tools, which already proved their excellence in medical applications such as iso-surfaces (computed with marching cubes [16]) and alpha-compositing of gradient-weighted voxel data [14], do not produce as meaningful results as in medical visualization.

The respective visualization metaphor of representing data by some kind of thresholding needs to be adapted to data from flow simulation. In this paper, we extend – as one part of our contribution (section 4) – iso-surfacing and alpha-compositing for volume rendering to better match the special situation of flow visualization.

1.3 User questions and visualization

The other part of our contribution in this paper is that we more directly involve the user within the visualization process (more details in section 2). We do so to more specifically respond to the actual user questions. A typical user question would enquire about where in the flow domain there are data-items of certain characteristics – data characteristics which are specified in terms of additional flow dimensions such as pressure or temperature. During the investigation of water-flow through an extended T-junction, the user, for example, is interested in vortical structures and medium-hot flow-regions as the interface between the warm and the hot inflow. Figure 6 shows a volume visualization response to this question.

Therefore, a tool for interactive analysis should, on the one hand, allow the user to easily formulate his or her questions and, on the other hand, should also provide useful visualization answers, which are tailored to the user questions, accordingly. In our software, we use an interactive assessment step in the first place, which is based on multiple, linked views as well as on interactive brushing to determine what the user currently is interested in (see section 2 for more details).


Figure 2: Volume rendering used to represent (a) structures with a significant amount of turbulent-kinetic energy (green: in main flow-direction; red: in opposite direction) and (b) regions of medium-warm flow (in front of the hot water coming in from the front inlet).

The user interest, called degree of interest (DOI), is coded per data-item and ranges from 0 (data-item not interesting) to 1 (maximal interest). If the user is concurrently interested in more than just one question, then several DOI values can be attached to each data-item. During visualization we then use the DOI information (instead of the original data) to visually represent the current data assessment. We have adapted iso-surfacing and direct volume rendering to give meaningful results when applied to DOI data (more details in section 4).

So the main idea beyond this paper is to first assess the data (in terms of what currently is most interesting to the user), and then use this information to visualize the data interpretation instead of the data itself. We therefore talk about this approach as being a *visualization of flow semantics*, because not the data itself is shown but its meaning to the user.

The remainder of this paper is structured as follows: Next, in Sect. 2, we shortly describe our interactive approach to assess the flow data (based on previous work [3, 4]). In Sect. 3 we then describe how we use hardware-accelerated resampling to ease the data access during volume visualization. Sect. 4 then presents three variants of visualization mapping for data representing the data assessment. Afterwards, we shortly describe real-time volume rendering which we use for actual visualization (Sect. 5, also based on previous work [6, 18]). Finally, results are presented in Sect. 6 before the paper is summarized and concluded.

2 INTERACTIVE DATA ASSESSMENT

In this section, we describe how a system of multiple, linked views is used (together with interactive brushing) to determine what the user currently is most interested in. We also describe degree-of-interest (DOI) values which are used to represent the results of the data assessment. Although this part of our approach considerably is based on previous work [3, 4], we nevertheless need to shortly present it here as it is an integral part of our main concept.

2.1 Linking and brushing

In the interactive assessment step, the user investigates the multi-dimensional simulation data by using multiple, linked views, which
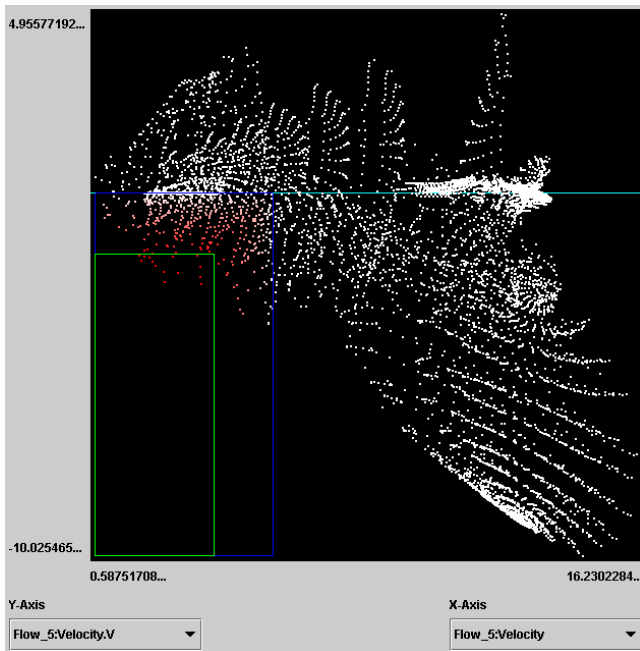
Figure 3: A 2D scatterplot of the catalytic converter data (x: overall velocities, y: vertical flow-components). Smooth brushing was used to select data with a significant "downwards" flow-component.

all show the same data, but usually different data attributes theirof. We provide scatterplots, histograms, and also parallel coordinates for flow data assessment. On demand, the user brings up views as required and adjusts them to show exactly those data dimensions which currently are most related to the user's current interest. In the example of the catalytic converter the user first uses a scatterplot to first visualize a 2D distribution of the flow data with respect to flow velocities (horizontal axis) and the vertical component of the flow vectors (vertical axis) – see figure 3 for this visualization.

After interactive exploration, the user starts to formulate what he or she is currently most interested in: to do so, the user works out the data-items of interest by iteratively restricting the sub-set of interest with respect to certain data dimensions. This is achieved by interactively brushing what actually is shown in the specific views. Continuing with the converter-example, the user would interactively brush the lower left part of the plot to denote that she is currently most interested in flow with a significant downwards-component in flow vectors (figure 3 shows the view directly after this first operation).

In a next step, the user would bring up another scatterplot showing again flow velocities but now versus horizontal flow-components across the flow. With another brush, a second sub-set of interest is described. So, step by step, the user concretizes the description of her current interest. 3D visualization, as described below (section 4), supports this step by interactive spatial feedback (figure 6, for example, shows the response to the above described brushing interactions).

2.2 Degree-of-interest values

As already described, we use degree-of-interest (DOI) values to represent what the user (currently) is most interesting in. We have also already mentioned that in the case of more than one current

user question, several of such DOI values are attached to data-items. Also important to mention is that smooth brushing [4] is used to gain DOI values which non-binarily span the interval between 0 (no interest at all) and 1 (maximal interest). Thereby, the user can work around the case that the smooth distribution of data properties across the flow domain does not allow a sharp and meaningful segmentation into interesting versus not interesting parts of the flow. This also can be interpreted as assigning a fuzzy set membership [17] with respect to the data sub-set of interest.

To enable the user to specifically formulate his or her questions it is necessary to allow complex combinations of simple restrictions. In the interface of our application, logical AND- and OR-combinations are provided through the use of modifier keys while brushing with the mouse. The application explicitly represents a composite brush in terms of a so-called feature definition language [3]; a separate view, similar to a regular tree viewer of a file-system, allows to directly manipulate the settings of each and any brush being used via sliders and numeric input. DOI values are computed from the logical composition tree by hierarchically using a min- or a max-combination where a t-norm (AND) or a t-conorm (OR) is required, respectively.

Although there are interesting alternatives for logically combinations of fuzzy sets [11], we finally chose min/max for our system since this decision enabled fast computations and also the drop-off to zero, when combining two fractional values, is minimal with respect to other alternatives.

3 HW-ACCELERATED RESAMPLING

Flow data-sets from computational simulation usually are laid out on non-Cartesian grids such as curvi-linear grids or unstructured grids. To enable fast volume visualization of 3D flow semantics, we use a variant of a hardware-accelerated resampling approach [24] to provide fast access to the flow data with respect to Cartesian coordinates.

The simulation results, which we are visualizing, are given in the form of cell-based data, i.e., data attributes such as the flow vector, pressure, temperature, etc., are given once per grid cell. For fast volume visualization it is crucial to efficiently interrogate any data attributes of need with respect to (almost) arbitrary spatial locations. Since we usually need access to multiple data attributes during visualization, we decided to not directly resample the all the needed data, attribute by attribute, but to compute a 3D redirection table in the desired resampling resolution (usually we use 512 voxels in the one dimension with greatest extent), and later access the data indirectly through this 3D table.

To compute the redirection table, we first decompose the simulation grid into tetrahedra. Then, for every tetrahedron, we store the cell-index which refers to the respective cell in the unstructured mesh in the R-, G-, and B-component of all the vertices of the tetrahedron. During resampling, then a Cartesian grid of cell-indices is computed which consequently then is used to access all the needed data attributes, available at the original grid cells.

4 VISUALIZATION MAPPING

In this section we describe three adaptations of well-know volume visualization algorithms to better match the situation of visualizing 3D flow semantics. First, we describe two ways of how to adapt iso-surfaces to better reflect their actual sharpness with regard to the underlying data. Then, we discuss volume rendering of DOI

Figure 4: Iso-surface with respect to $316\,\mathrm{K}$ (hot water floating in from the top inlet), whose voxels are more opaque the larger the temperature gradient is at the same voxel. Similarly, the color of the iso-surface varies from red (large gradients) via yellow to greed.


Figure 5: Iso-surface triplet used to represent a user's interest in hot regions of the flow. Red denotes maximal interest (hottest parts), green medium interest, and blue least interest, all according to a DOI-assessment, specified before the visualization interactively.

values. Finally, we present a DOI-based use of streamlines in 3D flow visualization.

4.1 Iso-surfacing the degree of interest

Iso-surfaces are a very useful instrument to visualize scalar data in 3D. This is mostly due to their visually sharp appearance. When shaded, very good depth perception is achieved. The worth of an iso-surface, however, is dependent on how meaningful the respective iso-value is. When visualizing a medical data-set, for example, acquired with a computer tomograph, an iso-value of $+1000\,\mathrm{HU}$ in general is well-suited to separate bones from the rest of the data.

An iso-surface sharply partitions the data-sets into data-items with the respective data attribute being smaller than the iso-value on the one side of the iso-surface and data-items with a larger data attribute of interest on the other side. If the distribution of the investigated data attribute is rather smooth with respect to its spatial locations, then an iso-surface might be even misleading in places where there is not much change in the data attribute of interest – a small variation of the iso-value would cause drastic changes in the geometrical layout of the iso-surface.

Below, we now present two variants of standard iso-surfaces which better reflect their sharpness with respect to the underlying data, thus being better suited for rather smooth data such as flow data from computational simulation. Since all of the here described volume visualization techniques are integrated within a volume rendering setup, the below described iso-surfaces are represented as voxel-surfaces [18].

Gradient-dependent opacity of iso-surfaces

As a first extension, we present iso-surfaces which exhibit a surface opacity which is dependent on the gradient magnitude with respect to the data attribute under investigation. Similar to windowing of intensity values in medical visualization, the opacity $\alpha_{\mathbf{v}}$ of a voxel $\mathbf{v}$ is defined as

$$\alpha_{\mathbf{v}} = \begin{cases} \alpha_{\mathrm{hi}} & \text{if } |\mathbf{g_v}| > g_{\mathrm{hi}} \\ \alpha_{\mathrm{lo}} + \frac{|\mathbf{g_v}| - g_{\mathrm{lo}}}{g_{\mathrm{hi}} - g_{\mathrm{lo}}} \cdot (\alpha_{\mathrm{hi}} - \alpha_{\mathrm{lo}}) & \text{if } g_{\mathrm{lo}} \leq |\mathbf{g_v}| \leq g_{\mathrm{hi}} \\ \alpha_{\mathrm{lo}} & \text{else} \end{cases}$$

where $\mathbf{g_v}$ denotes the gradient of the investigated data attribute, g_{lo} and g_{hi} bound a range of gradient magnitudes in which the surface opacity linearly increases from α_{lo} to α_{hi}. While g_{lo} often is set to 0, setting an upper limit of $g_{\mathrm{hi}} \ll \max_{\mathbf{v}} |\mathbf{g_v}|$ especially is useful

when parts of the iso-surface coincide with the boundary geometry of the flow where usually large gradients occur.

A sample image demonstrating this kind of a iso-surface is shown in figure 4. The more transparent the iso-surface is, the less it actually tells due to small gradients at the respective places. Also, a color transfer function has been used which assigns three different colors with respect to small, medium, and large gradient magnitudes. This color scheme is an additional help with the correct interpretation of the iso-surface.

Iso-surface triplets

Another extension of standard iso-surfaces we present are iso-surface triplets, which also help to correctly read the sharpness of an iso-surface. Instead of one iso-surface, a set of three iso-surfaces is rendered. With respect to one so-called reference iso-surface (iso-value f_{ref} and surface opacity α_{ref}), two additional iso-surfaces are specified. With respect to the settings of iso-values f_{a} and f_{b}, as well as of surface opacities α_{a} and α_{b}, we propose two different specification-modes for iso-surface triplets: *reference-inside* and *intermediate-reference*.

In the *reference-inside* specification-mode of an iso-surface triplet, surface iso-values are sorted $f_{\mathrm{a}} < f_{\mathrm{b}} < f_{\mathrm{ref}}$, i.e., the reference iso-surface is the inner-most, assuming that the data attribute under investigation exhibits increasing values outside-in. Usually an equidistant spacing between the three iso-values is used such that $(f_{\mathrm{ref}} - f_{\mathrm{b}}) = (f_{\mathrm{b}} - f_{\mathrm{a}}) = f_{\mathrm{sep}}$, where the user can adjust f_{sep} in the user interface. For the *reference-inside* type of iso-surface triplets, increasing values of surface opacities proved to be useful: $\alpha_{\mathrm{a}} < \alpha_{\mathrm{b}} < \alpha_{\mathrm{ref}}$ with $\alpha_{\mathrm{ref}} \approx 1$. Again, an equidistant spacing of the three opacities is useful with the distance set by the user. The *reference-inside* iso-surface triplet is useful for DOI visualization, using $f_{\mathrm{ref}} = 1 - \varepsilon$ to represent the 100%-interest sub-sets with the reference iso-surface, and $f_{\mathrm{sep}} = 25\%$, for example, to also denote where interest drops to 75% and 50%, respectively. In figure 5 we demonstrate the use of this kind of DOI visualization.

In the *intermediate-reference* mode of specifying an iso-surface triplet, surface iso-values are sorted $f_{\mathrm{a}} < f_{\mathrm{ref}} < f_{\mathrm{b}}$, i.e., the reference iso-surface is in-between the two others. Two different variants of setting the iso-value spacing proved to be useful: equidistant spacing with $(f_{\mathrm{b}} - f_{\mathrm{ref}}) = (f_{\mathrm{ref}} - f_{\mathrm{a}}) = f_{\mathrm{sep}}$, where the user sets f_{sep}, and proportional spacing with

$$\frac{f_{\mathrm{ref}} - f_{\mathrm{a}}}{f_{\mathrm{ref}} - f_{\mathrm{lo}}} = \frac{f_{\mathrm{b}} - f_{\mathrm{ref}}}{f_{\mathrm{hi}} - f_{\mathrm{ref}}}$$

Figure 6: Volume rendering of all those places in a catalytic converter where flow vectors exhibit a relatively strong vertical flow-component (red: upwards, blue: downwards) – a recirculation bubble becomes apparent in the upper part, before conversion block.

with f_lo and f_hi denoting the value range of the data attribute under investigation. In this case, the iso-values of the additional two iso-surfaces partition the value range below and above the reference iso-value proportionally. Thereby, for example, it is possible to easily adjust the iso-surface triplet in a way such that, regardless which iso-value is used for the reference iso-surface, the outer iso-surface always represents the property boundary half-way to minimum (with respect to the data attribute under investigation), and the inner iso-surface half-way to maximum. For the surface opacity setting, usually $\alpha_\text{a} = \alpha_\text{b} = (\alpha_\text{ref} - \alpha_\text{sep})$ is used with the user setting $\alpha_\text{ref} \gg \alpha_\text{sep} \gg 0$.

In addition to the specific settings of the surface opacities as described above, coloring the three iso-surfaces in red (reference iso-surface), green (inner), and blue (outer) proved to be useful for better visual distinction of the three stacked and semi-transparent iso-surfaces. Future work could clarify whether a more sophisticated modelling of the iso-surface opacity, like the use of curvature-directed strokes [9] could furthermore improve the 3D perception of iso-surface triplets.

4.2 Volume rendering of DOI values

Another useful volume visualization technique for scalar data laid out in three-space is volume rendering by means of alpha-compositing [14]. In our case we use this technique to visualize DOI values instead of original data values. We map DOI values to voxel opacities in a one-to-one manner such that DOI values of 1 appear opaque (or at least with maximal opacity) and DOI values of 0 are not visible at all.

The application of this kind of DOI visualization usually is based on the definition of multiple DOI attributes (in the sense of segmenting the flow data into several interesting sub-sets). We therefore chose to use the hue of voxel colors to visually differentiate between different DOI objects. Usually we vary the color intensity according to Phong illumination [19]. In figure 6 we see two DOI objects (one red and one blue) with varying opacities according to their DOI values.

4.3 Interest-dependent streamlines

Streamlines are a very useful instrument to visualize instantaneous vector fields. Due to their character of visually joining flow locations along virtual paths of massless particles, streamlines give a


Figure 7: Illuminated streamlines representing a user's interest in streamlines whose origins coincide with the hot in-flow.


Figure 8: Illuminated streamlines representing a user's interest in how two flows join in a simple T-junction.

very intuitive image of a flow data-set (however disregarding the actual time up to a certain extent). Evenly-spaced streamlines, which fill up the flow domain in a well-distributed way, are well-accepted for two-dimensional flows [10]. However, in 3D flow visualization, selective streamline placement is needed to avoid problems with occlusion.

In previous work [15] we use the vicinity to topological structures such as characteristic trajectories in the flow data to selectively seed streamlines. In this work, we now use the spatial distribution of DOI values to selectively seed streamlines. To do so, evenly-spaced streamline seeding [10] is extended to 3D. However, streamlines are only seeded in regions where the DOI value is 1. Therefrom the streamlines are integrated forwards and backwards until they either exit the flow domain, violate the minimal distance to other streamlines, or just get too long.

To integrate DOI-based streamlines within our volume visualization setup, we rasterize the streamlines into our voxel grid such that the voxel opacities are set according to the proportional overlap of the streamline and the respective voxel. We use a $5 \times 5 \times 5$ sub-voxels accuracy when computing the overlap of streamlines and voxels, thus gaining anti-aliased streamlines, which is of great importance for useful results. Furthermore, voxe opacities are also dimmed down according to the DOI values exhibited at the respective streamline-points. Thereby, the streamlines fade out according to the degree of interest. For lighting, we use the illuminated streamlines model [25] to furthermore support the 3D perception of streamlines in our visualization. Figure 7 demonstrates the use of DOI-based streamlines on the example of a T-junction.

5 REAL-TIME VOLUME RENDERING

Gradient-dependent iso-surfaces and iso-surface triplet, as well as volume rendering of DOI values and interest-weighted streamlines have all been integrated on the basis of RTVR [18] which is a very fast volume rendering software (based on the shear-warp rendering algorithm [13]) which provides several different modi of volume rendering such as standard alpha-compositing, maximum-intensity projection (MIP), intensity summation, and non-photorealistic contour rendering [2]. Two-level volume rendering [6] is used for data-sets which are composed of several, explicitly separated objects so that different rendering modi can be used for different objects.

In this work, we build up RTVR-objects for every representation of a user's particular interest. Thereby, the user interactively composes a set of volumetric objects (extended iso-surfaces, truly volumetric objects, and streamlines, recall section 4) in dependence on as many data attributes as necessary, in simple or in complex manner (recall section 2). For every object the user can decide in which style the respective object should be rendered.

Iso-surfaces usually are of one color and shaded while their opacity may vary with respect to the data gradient. Iso-surfaces also can be rendered by using the non-photorealistic contour rendering, thereby giving very good context for the entire visualization (see figures 6 and 8 for two examples). For volume rendering, either shaded voxels are used or voxels whose color intensities depend on the DOI attribute. Voxel opacities usually represent the DOI attribute which the voxel object was made dependent on during specification. Streamlines are either illuminated [25] or colored such that the intensity represents the DOI attribute for which the streamlines have been requested. Opacity of streamlines is used for anti-aliasing.

In general, we use hue to discriminate objects from each other. This approach already proved to be very useful in medical applications [18] and also in the context of flow data, color-coding the different 3D objects different from each other significantly eases the perceptual distinction of them.

RTVR provides several useful tools for volume visualization which also in this application helped a lot with comprehending the results of computational flow simulation: interactive change of lighting conditions, the interactive change of an object-wide overall opacity (in addition to the voxel-local opacity transfer function), and, of course, the interactive change of viewing conditions.

6 RESULTS

In the following we shortly describe two application cases in which the previously described visualization technique helped to answer specific user questions.

6.1 Flow through a catalytic converter

One data-set which our visualization users investigate comprises the results of a simulation of gas flow through a catalytic converter. With respect to this case, special interest is put on a recirculation zone, right before the reaction chamber, clearly visualized by the use of selectively placed streamlines in figure 9. Using our setup, the user first describes what he or she is interested in by means of brushing in views from information visualization (sect. 2) – in this example, the specification would encompass (a) a brush on $-x$-velocities (as the positive x-axis is the main flow direction through


Figure 9: Streamlines representing a user's interest in the back-flow region right before the flow enters the reaction chamber of this catalytic converter. Streamlines are started in this particular region and where flow heads backwards – from there streamlines are traced backwards and forwards to put that respective region into the context of the surrounding flow.

the catalytic converter) and (b) the restriction to the region before the reaction chamber, both combined with a logical AND. The thereby defined DOI attribute can then be used to selectively place streamlines as shown in figure 9.

Another approach to this case is to investigate upwards and downwards flow which at the same time does not exhibit a significant flow-component into the main flow-direction. Figure 6 shows a volume rendering with upwards flow in red and downwards flow in blue – note from how deep the upwards flow starts before the reaction chamber as well as the downwards flow component in the part between the inlet-knee and the reaction chamber.

In figure 1, a more complete picture of the here discussed case is given. In red those places within the flow domain are shown which exhibit a relatively large amount of turbulent-kinetic energy, a value which is simulated along with the other flow attributes. In yellow, streamlines are given whose origin coincides with the back-flow region. The back-flow region itself is represented using a greyish volume rendering. The whole catalysator is shown as context using a contour rendering.

6.2 Extended T-junction

The other case presented here is an extended T-junction (with an extended chamber around the junction and an obstacle in front of the secondary inlet) with warm water floating in from the on inlet, and hot water entering the T-junction after the first third of the time-span of the simulation.

In this example, the user primarily is interesting in the mixing behavior of this flow-junction. The obstacle in the middle of the T-junction causes the hot in-flow to split – on part turns directly to the exit whereas the other part traverses an extra-loop before it also develops towards the exit of this structure (see figure 7 for several streamlines demonstrating this behavior). Figures 4 and 5 use the two here described iso-surface extensions to show certain degrees of temperature in response to the hot in-flow from the one side.

In figure 10, several volumetric objects are used to partition parts of the flow in regions which exhibit flow that is mainly in the direction of one of the three main axes (see figure-caption for color scheme). A gradient-dependent iso-surface was used to depict the boundary of the incoming hot flow – this iso-surface is more opaque in regions with larger temperature-gradients. The separation of the

Figure 10: Flow through the extended T-junction (main flow from right to left). Volume rendering was used to encode flow regions with flow which is mostly aligned with one of major axes (blue: $-y$-flow, like from the secondary inlet; yellow: $+y$-flow; red: upwards; green: downwards). A gradient-dependent iso-surface was used to show the extent of medium-hot temperature (hot coming in from the upper inlet).


Figure 11: Flow through the extended T-junction. Volume rendering was used to show flow sub-sets which exhibit relatively high values of turbulent-kinetic energy while also having a strong flow component into(red)/against(red) the main direction of flow. Another volume-rendered sub-sets of the data denotes regions of medium-hot fluid, whereas a gradient-dependent iso-surface was used to work out those parts of the flow where a larger y-component of the flow is given (more details: section 6.2).

`http://www.VRVis.at/vis/research/volflowvis/` for more results and animation sequences.

incoming flow, due to the obstacle, is very well visible in this example, also. Figure 11 comprises the same features than figure 2, just a little later in simulation time. Additionally, a gradient-dependent iso-surface is shown which visualizes flow-regions with a major y-component, i.e., towards the inlet in the front. From this iso-surface, information can be derived about where in the flow vortical structures could be, e.g., before and after the obstacle.

7   SUMMARY AND CONCLUSIONS

In this paper we presented semantics-based visualization of flow data from computational simulation. We described how the user utilizes interactive feature specification to tell the system what he or she currently is interested in. We then proposed (a) gradient-dependent iso-surfaces to visualize the sharpness of an iso-surface in conjunction with its region-separating appearance, (b) iso-surface triplets with specific relations for the three iso-values involved as well as opacities, shortly (c) volume rendering of DOI values, and (d) interest-dependent streamlines, which are based on the evenly-spaced streamline seeding algorithm but only are placed in regions of high user-interest.

Conclusions of this work are (a) that multiple data attributes usually are necessary to really comprehend results from computational flow simulation and that for this reason a flow visualization system should offer easy-to-handle access to all the information which is provided by the simulation, (b) that selective visualization, in our case, focus+context visualization of 3D flow data is (maybe the only) one really good approach to cope with the large amount of information to be communicated – for deciding *what* to show, the user and his/her questions about the data should be considered, and (c) that due to the different requirements of different cases it is very useful to have an integrated solution which provides different means/techniques for the visualization.

The here described framework only minimally supports the visualization of time-dependent flows – single time-steps can be loaded into the visualization through a special interface element whilst the configuration of the visualization is maintained (see also figure 12). Thereby comparisons between different time-steps of unsteady flow can be made. However, a better integration of time-dependent flows still is future work. See the project web page

REFERENCES

[1] Stuart Card, Jock MacKinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, 1998.

[2] Balázs Csébfalvi, Lukas Mroz, Helwig Hauser, Andreas König, and Eduard Gröller. Fast visualization of object contours by non-photorealistic volume rendering. *Computer Graphics Forum*, 20(3), 2001.

[3] Helmut Doleisch, Martin Gasser, and Helwig Hauser. Interactive feature specification for focus+context visualization of complex simulation data. In *Proceedings of the 5th Joint IEEE TCVG - EUROGRAPHICS Symposium on Visualization (VisSym 2003)*, Grenoble, France, May 2003. ACM Press.

[4] Helmut Doleisch and Helwig Hauser. Smooth brushing for focus+context visualization of simulation data in 3D. *Journal of WSCG*, 10(1):147–154, 2002.

[5] Donna Gresh, Bernice Rogowitz, Raimond Winslow, David Scollan, and Christina Yung. WEAVE: A system for visually linking 3-D and statistical visualizations, applied to cardiac simulation and measurement data. In *IEEE Visualization 2000*, pages 489–492, October 2000.

[6] Helwig Hauser, Lukas Mroz, Gian-Italo Bischi, and Eduard Gröller. Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):242–252, 2001.

[7] James Helman and Lambertus Hesselink. Visualization of vector field topology in fluid flows. *IEEE Computer Graphics and Applications*, 11(3):36–46, 1991.

[8] Alfred Inselberg. A survey of parallel coordinates. In Hans-Christian Hege and Konrad Polthier, editors, *Mathematical Visualization*, pages 167–179. Springer Verlag, 1998.

Figure 12: Comparisons of flow features over time (compare to figure 2 and section 6.2 for more about the flow features).

[9] Victoria Interrante, Henry Fuchs, and Stephen Pizer. Illustrating transparent surfaces with curvature-directed strokes. In *IEEE Visualization 1996*, pages 211–218, Oct./Nov. 1996.

[10] Bruno Jobard and Wilfrid Lefer. Creating evenly-spaced streamlines of arbitrary density. In *Visualization in Scientific Computing '97*, pages 43–56. Springer-Verlag, 1997.

[11] Erich Klement, Radko Mesiar, and Endre Pap. *Triangular Norms*, volume 8 of *Trends in Logic*. Kluwer Academic Publishers, Dordrecht, 2000.

[12] Robert Kosara, Helwig Hauser, and Donna Gresh. An interaction view on information visualization. In *Eurographics 2003 State-of-the-Art Reports*, Granada, Spain, September 2003.

[13] Philippe Lacroute and Marc Levoy. Fast volume rendering using a shear-warp factorisation of the viewing transform. In *ACM SIGGRAPH 1994*, pages 451–459, 1994.

[14] Mark Levoy. Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, 8(5):29–37, 1988.

[15] Helwig Löffelmann and Eduard Gröller. Enhancing the visualization of characteristic structures in dynamical systems. In *Visualization in Sientific Computing '98*, pages 59–68. Springer-Verlag, 1998.

[16] William Lorensen and Harvey Cline. Marching cubes: a high resolution 3D surface construction algorithm. In *ACM SIGGRAPH 1987*, pages 163–189, 1987.

[17] Robert Lowen. *Fuzzy set theory: basic concepts, techniques and bibliography*. Kluwer Academic Publishers, 1996.

[18] Lukas Mroz and Helwig Hauser. RTVR - a flexible java library for interactive volume rendering. In *IEEE Visualization 2001*, pages 279–286, October 2001.

[19] Bui-Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.

[20] Frits Post, Benjamin Vrolijk, Helwig Hauser, Robert Laramee, and Helmut Doleisch. Feature extraction and visualization of flow fields. In *Eurographics 2002 State-of-the-Art Reports*, pages 69–100, September 2002.

[21] Freek Reinders. *Feature-based Visualization of Time-dependent Data*. PhD thesis, Delft University of Technology, The Netherlands, 2001.

[22] Gerik Scheuermann, Heinz Krüger, Martin Menzel, and Alyn Rockwood. Visualizing nonlinear vector field topology. *IEEE Transactions on Visualization and Computer Graphics*, 4(2), April – June 1998. ISSN 1077-2626.

[23] Theo van Walsum, Frits H. Post, Deborah Silver, and Frank J. Post. Feature Extraction and Iconic Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2(2):111–119, June 1996.

[24] Manfred Weiler and Thomas Ertl. Hardware-software-balanced resampling for the interactive visualization of unstructured grids. In *IEEE Visualization 2001*, pages 199–206, October 2001.

[25] Malte Zöckler, Detlev Stalling, and Hans-Christian Hege. Interactive visualization of 3D-vector fields using illuminated streamlines. In *IEEE Visualization 1996*, pages 107–113, 1996.