

Paper B

Integrated Multi-aspect Visualization of 3D Fluid Flows

Andrea Brambilla¹, Øyvind Andreassen^{2,3}, Helwig Hauser¹

¹ Department of Informatics, University of Bergen, Norway

² Norwegian Defence Research Establishment, Norway

³ University Graduate Center at Kjeller, Norway

Abstract

The motion of a fluid is affected by several intertwined flow aspects. Analyzing one aspect at a time can only yield partial information about the flow behavior. More details can be revealed by studying their interactions. Our approach enables the investigation of these interactions by simultaneously visualizing meaningful flow aspects, such as swirling motion and shear strain. We adopt the notions of relevance and coherency. Relevance identifies locations where a certain flow aspect is deemed particularly important. The related piece of information is visualized by a specific visual entity, placed at the corresponding location. Coherency instead represents the homogeneity of a flow property in a local neighborhood. It is exploited in order to avoid visual redundancy and to reduce occlusion and cluttering. We have applied our approach to three CFD datasets, obtaining meaningful insights.

This article was published in the proceedings of *VMV 2013: Vision, Modeling & Visualization*, 1–9, 2013.

1 Introduction

Fluid flows are subject to extensive studies in different fields, such as engineering, climatology and medicine. Material properties, external forces and many other factors can heavily affect the motion of a fluid, making it a particularly complex phenomenon to investigate. Computational Fluid Dynamics (CFD) simulations can produce a description of the fluid's motion according to several variables, such as pressure, temperature and boundary geometry. Such a description often takes the form of a flow field, i.e., a function that associates a velocity vector to every spatial location. In order to gain additional insights into the flow behaviour, derived attributes are often computed. This results in multivariate datasets, defined over a 2-, 3- or 4-dimensional domain, and each variable can be either a scalar, a vector or a tensor.

Given the large amount of data, visualization techniques are generally helpful during the exploration and analysis of a flow dataset. However, the substantial complexity of the flow behaviour can hardly be understood by looking at a single attribute alone. In this paper we propose a novel visualization strategy for the simultaneous depiction of multiple flow aspects, such as vortical motion, shearing and stretching. Our approach can effectively display different flow aspects together, so that their interactions can be directly observed. We focus on *flow velocity* (vector), *vorticity* (vector) and *rate of strain* (2nd-order tensor) because of their central role in fluid mechanics. Notice that, while displaying a scalar field is almost straightforward, vectors and tensors require more efforts. In order to guarantee the modularity of our approach, we depict each variable through a particular *visual entity*, in analogy with the *AVO* by Haber and McNabb [HM90].

The core of our technique is a strategy for the placement of visual entities over the spatial domain. Our main challenges are the cluttering and occlusion issues that often arise when visualizing multiple attributes simultaneously. We address these problems by exploiting the concepts of *relevance* and *coherency*. Visual entities are shown only at locations where the associated variable is deemed important according to a *relevance measure*. We evaluate the local homogeneity of an attribute by means of a *coherency measure*. If a region presents highly homogeneous data values, we avoid redundancy by conveying the corresponding information through a single visual entity. Our main contribution is an effective visualization strategy for the simultaneous depiction of multiple flow aspects. Moreover, the appearance and the density of the visualization can be intuitively controlled by tuning the relevance and coherency parameters.

2 Related Work

One of the first approaches that deals with vorticity and rate of strain is the *streampolygon* [SVL91]. Local rotation and strain are conveyed by deforming

a regular polygon. Stream-tubes are then generated by sweeping the polygon along streamlines. Integral curves are used also in the *hyperstreamline* technique [DH93]: 2nd-order tensors (such as the rate of strain) are visualized by integrating a stream-tube along one of the tensor's eigenvectors. The other two eigenvectors determine the shape of the tube's section. Integration of the field lines of the vorticity field leads to the so-called *vorticity lines*. They have been visualized and analyzed for the study of wall-bounded turbulent flows [HRAW07] and in the context of an engineering application [SPS06]. Advection is also used in the work by Schafhitzel et al. [SBV⁺11] for observing the interactions between shear layers and vortices.

Another category of approaches depicts vorticity and rate of strain using specific glyphs. The flow probe by de Leew and van Wijk [dLvW93] simultaneously conveys several quantities, including velocity, rotation and shear. For 2D flows, Kirby et al. [KML99] map every attribute to a different visual entity. By overlapping the various representations, they produce dense visualizations with acceptable cluttering. Even denser visualizations can be achieved [WFK⁺02, UIL⁺04], but handling tensor data becomes then problematic. A survey on multivariate visualization has been recently presented by Fuchs and Hauser [FH09].

Our framework is based on visual overlapping as well, but it presents several advantages over the previous approaches. We can handle both 2D and 3D flow fields and we impose no limitations on the dimensionality of the data attributes. Additionally, we provide control over cluttering through the relevance and coherency measures. The concepts of relevance and coherency are largely adopted in visualization. A relevance measure denotes how important a piece of information is. It is the foundation of most focus+context approaches [Hau03]. It usually takes the form of a scalar value (discrete or continuous) defined over the data samples. It can be specified in several ways, such as querying or brushing the data, and it is normally used to control different rendering aspects. We refer to the tutorial by Viola et al. [VGH⁺05] for an overview of focus+context techniques within illustrative visualization. Coherency instead encodes the homogeneity of a domain region or a set of values. Clustering is a well-known example of a coherency-based technique: a cluster is a set of samples with coherent data values. The simplification of vector fields via clustering has been firstly proposed more than 10 years ago [HWHJ99, TVW99]. Coherency is also a basic concept behind vector field topology [HH91]. The topological description of a vector field is in fact a partitioning of the spatial domain in regions of coherent asymptotic behavior. Recently, the concept of *Shannon's entropy* has been successfully exploited in flow visualization [JBTS08, MCHM10]. As a matter of fact, entropy is inversely proportional to coherency: low entropy corresponds to highly redundant (coherent) data, while high entropy corresponds to highly incoherent data.

Relevance and coherency have been rarely combined. To the best of our knowledge, the only approach that involves both of them has been presented by Bürger

et al. [BKKW08]. They adopt dense data representations, such as deformed glyphs or streamlines, in highly relevant regions. The rest of the spatial domain is instead clustered, and each cluster is depicted by a single arrow glyph. However, this approach does not deal with multivariate data.

3 Physics Fundamentals

In the following, we write vectors and tensors in component form. For example, a vector $\mathbf{u} \in \mathbb{R}^3$ is written u_i . Indices are denoted by i, j and k , and they range from 1 to 3.

The motion of a fluid is described by the *velocity* $u_i(x_j, t)$, which is a vector field defined over a spatial and temporal domain. Useful information can be obtained by computing the spatial derivatives of u_i . Specifically, the *velocity gradient tensor* $\mathbf{U} = \nabla \mathbf{u}$ is a 3×3 matrix with components $U_{ij} = \partial u_i / \partial x_j$. It can be used to express the relative motion near a point as $du_i = U_{ij} dx_j$ (adopting Einstein's summation convention). Moreover, U_{ij} is a square matrix, so it can be decomposed into the sum of a symmetric and an anti-symmetric matrix:

$$U_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) = S_{ij} + \Omega_{ij}.$$

Here

$$S_{ij} = (\partial u_i / \partial x_j + \partial u_j / \partial x_i) / 2$$

is the strain rate tensor, and

$$\Omega_{ij} = (\partial u_i / \partial x_j - \partial u_j / \partial x_i) / 2$$

is the rotation tensor. In other words, the relative motion near a point can be decomposed into straining motion (\mathbf{S}) and rotation ($\mathbf{\Omega}$).

Notice that \mathbf{S} is real and symmetric. Therefore, it has real eigenvalues and orthogonal eigenvectors. This leads to an intuitive interpretation of S_{ij} : the eigenvectors define the principal axis of deformation and the eigenvalues represent the magnitude of the deformation (positive values correspond to expansion, negative values to compression). In contrast, $\mathbf{\Omega}$ is anti-symmetric, so it can be mapped to a vector $\boldsymbol{\omega}$ through the linear relation $\omega_i = -\epsilon_{ijk} \Omega_{jk}$, where ϵ_{ijk} is the permutation tensor. Vector $\boldsymbol{\omega}$ is known as *vorticity* and it represents the local direction and magnitude of rotation.

Vorticity and rate of strain have been thoroughly investigated in the context of fluid mechanics [MK85, VM91]. A notable relationship between these two attributes is given by the *vorticity transport equation* for incompressible flows:

$$\frac{D\omega_i}{Dt} = S_{ij}\omega_j + \nu \frac{\partial^2 \omega_i}{\partial x_k \partial x_k}.$$

D/Dt is the material derivative and ν is the kinematic viscosity. We see that straining acts as a source term to either intensify vorticity (stretching) or to reorient it (tilting and twisting). The relationships between vorticity and strain rate are strongly connected to many flow phenomena. For instance, they play a central role in the evolution of bounded turbulent flows: close to boundaries or obstacles, the fluid undergoes strong shear deformations [NP98]. This can result in the formation of *shear layers*, characterized by high strain and production of vorticity. Shear layers can detach from the wall and roll up into coherent vortices, characterized by strong vorticity [Wil96].

Vortices and shear layers are two types of *coherent structures*, i.e., flow regions identified by specific homogeneous characteristics. Feature detectors have been proposed in order to localize these kinds of structures. Several techniques exist for detecting vortices, e.g., Hunt's Q [HWM88], λ_2 [JH95] and the parallel vector operator [PR99]. In contrast, the literature about the identification of shear layers is substantially narrower: Hunt's Q can be also used to detect strain-dominated areas, while a dedicated shear layers detector has been proposed by Haines and Kenwright [HK99]. Vorticity and rate of strain are in fact the main quantities these feature detectors are based on. More details about Hunt's Q , λ_2 and Haines and Kenwright (HK) detectors are provided in Section 4.3.

The substantial importance of velocity, vorticity and rate of strain, and their continuous interactions, is one of the main motivations behind this paper. Moreover, we take them into account in the context of the related coherent structures, that is, vortices and shear layers.

4 Visualization Strategy

Our first step is to associate a suitable visual entity to every attribute of interest (Sec. 4.1). The visual entities are then distributed over the spatial domain. Our placement strategy (Sec. 4.2) relies on *relevance* (Sec. 4.3) and *coherency* (Sec. 4.4) in order to minimize visibility issues. In the remainder of this section, all the elements of our visualization strategy are described in more detail.

4.1 Design of visual entities

What can be considered a suitable visual entity is often dependent on the task at hand. Specifically, our reference variables are the flow velocity, the vorticity and the rate of strain, so we need appropriate representations for vector and tensor data. Moreover, since we are mainly interested in the interactions between these flow attributes, we require visual entities that effectively convey local information at selected locations in space. For these reasons, we decided to use glyphs as visual entities for all the three flow attributes.




Figure 1: The glyphs we adopted for representing the variables of interest. Arrows depict the flow velocity, a specifically designed glyph is used for the vorticity, while ellipsoids are employed for the rate of strain. In all the three cases, the color encodes the norm of the related variable.

An accurate design of glyphs is of primary importance in many application scenarios [LKH09]. For defining our glyphs (Fig. 1) we followed the Design Guidelines presented by Borgo et al. [BKC⁺13]. Since our placement strategy may lead to locally dense glyph distributions, we adopt simple glyph shapes (D.G. 2 in [BKC⁺13]). Moreover, we opted for glyph shapes which intuitively recall the semantic of the associated attribute (D.G. 10 in [BKC⁺13]):

- The velocity vector \mathbf{u} is mapped to a 3D arrow glyph, oriented according to the direction of the velocity.
- The vorticity $\boldsymbol{\omega}$ conveys rotation instead of linear motion. We mapped it to an ad-hoc glyph that depicts the direction (clockwise or counterclockwise) and the plane of rotation (orthogonal to $\boldsymbol{\omega}$).
- The rate of strain tensor \mathbf{S} is mapped to an ellipsoid glyph. Denoting the eigenvalues and the eigenvectors of \mathbf{S} with λ_i and \mathbf{v}_i respectively, a unit sphere is scaled in the directions \mathbf{v}_i by an amount of e^{λ_i} . In this way the sphere is stretched in the directions of expansion (positive eigenvalues) and squeezed in the directions of compression (negative eigenvalues).

The norm of these quantities is mapped to the color of the glyph. The Euclidean norm is used for vectors, while the norm of the rate of strain tensor is computed as $\|\mathbf{S}\| = \sqrt{S_{ik}S_{ki}}$ (Frobenius norm). We employed different color scale for each quantity. Our color scales, selected from the ColorBrewer tool [HB03], have been chosen in order to help the user tell one glyph type from another. The combination of shape and color makes the glyphs easy to discriminate even in areas of high density (see Figure 6).

The glyph's size is used to encode the local coherency of the data (Sec. 4.4). This is effective only if all the glyphs have the same initial size, but the direc-

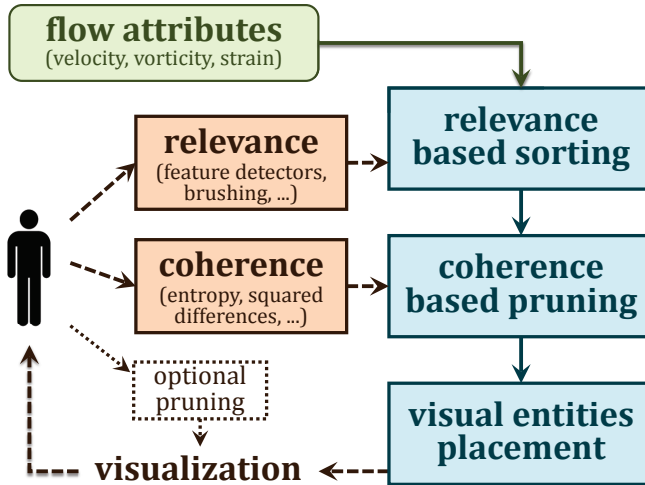


Figure 2: Overview of our multi-aspect visualization system.

tional scaling of the ellipsoids can in fact modify their overall size. Therefore, we uniformly scale every ellipsoid so that its longest axis has unit length. As a matter of fact, no information is lost: the normalized glyph’s shape conveys the relative strain intensities in the principal strain directions, while the overall magnitude of the deformation is conveyed by the color.

4.2 Placement strategy

Our placement algorithm is summarized in Figure 2. For each attribute of interest a we first define a set of locations P_a in the spatial domain. P_a contains all the possible locations where the visual entity associated with a may be shown. We assume the flow data is expressed over a grid, so we initialize the sets P_a using the grid’s vertices.

For each attribute a , the user is required to specify a relevance measure, a coherency measure and a coherency threshold. Relevance leads to a focus+context visualization, where visual resources are primarily assigned to the important portions of the data (the focus). Coherency is instead used to reduce visual redundancy. In practice, every set P_a is sorted in descending order of relevance. Following this ordering, an *area of influence* is computed for every $\mathbf{p} \in P_a$ according to the specified coherency measure and threshold. Then, the sets P_a are reduced: given $\mathbf{p}, \mathbf{q} \in P_a$, if \mathbf{p} lies in the area of influence of \mathbf{q} , and $relevance_a(\mathbf{p}) < relevance_a(\mathbf{q})$, then point \mathbf{p} is discarded. This procedure leads to the reduced sets \bar{P}_a . As a final step, for each attribute a the corresponding visual representation is displayed at the locations in \bar{P}_a .

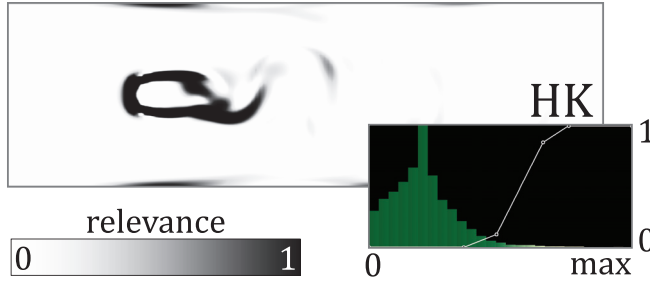


Figure 3: Example of relevance measure specified by brushing over the histogram of Haimes and Kenwright’s shear layers detector [HK99]. The dataset is a $2D$ slice of a flow around a square cylinder (see Section 5).

The resulting visualization can be customized by the user by tuning the relevance and coherency parameters. Cluttering and occlusion can still occur, especially in the case of a $3D$ spatial domain, so we implemented an additional pruning tool (Sec. 5) that eases the inspection of the results.

4.3 Attribute relevance

In accordance with the concept of focus+context visualization, we display a visual entity only when the information it encodes is deemed important. It is a common practice to represent the importance as a scalar attribute defined over the samples in the dataset. However, this approach has a notable limitation: it cannot take into account the fact that different variables can be more or less relevant in different areas of the spatial domain. Therefore we extend this technique by employing multiple relevance attributes, one for each of the variables of interest. Formally, the relevance associated with an attribute a is a function $r_a : P_a \rightarrow [0, 1]$, where 0 denotes the less relevant points and 1 the most relevant ones. The relevance values are directly mapped to the opacity of the related visual entities.

What is more or less relevant depends almost always on the application domain and the task to be accomplished. Therefore we let the user define the various relevance functions. To facilitate this procedure, we provide a simple tool that lets the user specify the relevance values by brushing over the histogram of a support variable. Since we are dealing with flow data, flow feature detectors can serve as effective support variables. Specifically, we integrated in our system the Hunt’s Q , λ_2 and HK detectors. Figure 3 shows an example of a relevance measure; more examples can be found in the additional material.




Figure 4: The left-most image shows a synthetic vector field with a center and a saddle, sampled over a 64×64 grid. All the other images show the discussed coherency measures evaluated over local neighborhoods (radius = 5 cells) of every point in the dataset. Highly coherent areas are white, incoherent areas are black. Arrow glyphs have been placed according to our approach assuming uniform relevance. Streamlines and glyphs are colored according to the velocity magnitude (color scale in Figure 1).

Overview of feature detectors

Hunt's Q [HWM88] is an established method for the detection of vortices. Q is defined as the second invariant of U_{ij} :

$$Q = \frac{1}{2} \left(\left(\frac{\partial u_i}{\partial x_i} \right)^2 - \frac{\partial u_i}{\partial x_j} \frac{\partial u_j}{\partial x_i} \right) = \frac{1}{2} \left(\|\boldsymbol{\omega}\|^2 - \|\mathbf{S}\|^2 \right).$$

Q represents the local balance between strain and vorticity. In fact, $Q < 0$ identifies strain-dominated areas, while $Q > 0$ identifies regions of swirling motion.

The λ_2 method [JH95], introduced by Jeong and Hussain in 1995, is currently one of the most commonly adopted vortex detectors. This method defines a vortex core as a connected region where the second eigenvalue λ_2 of the symmetric tensor $S_{ik}S_{kj} + \Omega_{ik}\Omega_{kj}$ is negative.

Haines and Kenwright [HK99] define a boundary (shear) layer as a region characterized by a strong shear stress. As a matter of fact, they do not present a well-defined shear layer detector, but they suggest to use the following quantity as a measure of shear:

$$HK = \sqrt{((\lambda_1 - \lambda_2)^2 + (\lambda_1 - \lambda_3)^2 + (\lambda_2 - \lambda_3)^2) / 6},$$

where λ_i are the eigenvalues of the rate of strain tensor \mathbf{S} .

4.4 Coherency and areas of influence

Since we are showing multiple attributes simultaneously, an efficient utilization of the visual space is crucial. Thanks to our focus+context approach, we avoid wasting visual resources in areas where no relevant information is present. However, there is no guarantee that highly interesting regions will be free of cluttering and occlusion issues. The situation can be improved by exploiting data redundancy. The basic idea is that data samples which are spatially close and have similar values actually encode the same piece of information. Therefore it is sufficient to show a single visual entity instead of one for each sample.

In order to evaluate the similarity between different samples, we introduce the concept of *coherency measure*. A coherency measure c evaluates a set of data values D and produces a scalar value $c(D)$. Its interpretation varies according to how c is defined, but two main categories can be identified:

- **Coherence without reference:** $c(D)$ is proportional to how close to each other the values in D are; a typical example from statistics is the interquartile range.
- **Coherence with reference \mathbf{v} :** $c_{\mathbf{v}}(D)$ represents how close the values in D are to a reference value \mathbf{v} ; an example is the 2nd moment of a function about a certain value.

In analogy with the aforementioned examples, we assume that $c(D) \geq 0$ and that small values of $c(D)$ represent highly coherent samples. For each attribute of interest a , the user has to specify the desired coherency measure c and a coherency threshold γ . We define the area of influence of a point $\mathbf{p} \in P_a$ as the *largest spherical region* $A_{\mathbf{p}} \subseteq P_a$ around \mathbf{p} that satisfies $c(A_{\mathbf{p}}) < \gamma$. If the chosen coherency measure needs a reference, the value of the attribute in \mathbf{p} is used.

Now we can determine the locations \bar{P}_a where the visual entity for a will be displayed. Constructing \bar{P}_a as explained in Section 4.2 would be highly inefficient. We can exploit the fact that P_a has been ordered in descending order of relevance. Whenever the area of influence $A_{\mathbf{p}}$ of a point \mathbf{p} is computed, \mathbf{p} is added to \bar{P}_a while all the other points in $A_{\mathbf{p}}$ are removed from P_a . In fact, the ordering guarantees that relevance $r_a(\mathbf{p})$ is maximal in $A_{\mathbf{p}}$.

We opted for areas of influence of spherical shape so that they can be easily inferred even without being explicitly shown. We map the radius of every spherical region to the size of the corresponding visual entity. The resulting visualization allows for an intuitive and straightforward interpretation of the data (see Figure 5).

Coherency measures

Taking inspiration from established concepts in statistics and information theory, we designed and integrated in our system four different coherency measures. We developed two measures based on the magnitudes of the data values: The first is a measure without reference and is obtained by directly computing the normalized Shannon's entropy over the magnitudes. The other is a measure with reference and is based on the 2nd statistical moment of the magnitudes:

$$c_mom2_{\mathbf{v}}(D) = \sqrt{\frac{1}{|D|} \sum_{\mathbf{d} \in D} \|\mathbf{d}\| - \|\mathbf{v}\|}.$$

We have also implemented two measures with reference which consider both the orientations and the intensities of the data samples:

$$c_diff_{\mathbf{v}}(D) = \frac{1}{|D|} \sum_{\mathbf{d} \in D} \|\mathbf{d} - \mathbf{v}\|,$$

$$c_dot_{\mathbf{v}}(D) = \frac{1}{|D|} \sum_{\mathbf{d} \in D} \left| 1 - \frac{\langle \mathbf{d}, \mathbf{v} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \right|.$$

Figure 4 shows the effects of all the four measures applied to a synthetic flow dataset. Notice that they can be applied to any kind of attribute (scalars, vectors and tensors).

	Velocity	Vorticity	Rate of Strain
<i>A</i>	high magnitude $c_dot_{\mathbf{v}}$	λ_2 2^{nd} moment	HK $c_diff_{\mathbf{v}}$
<i>B</i>	high magnitude $c_dot_{\mathbf{v}}$	positive Q $c_diff_{\mathbf{v}}$	negative Q 2^{nd} moment
<i>C</i>	see text $c_dot_{\mathbf{v}}$	λ_2 entropy	negative Q $c_dot_{\mathbf{v}}$

Table 1: Relevance (first row) and coherency (second row) measures adopted in our test cases. *A*: flow around a square cylinder. *B*: flow in a box. *C*: exhaust manifold.

5 Demonstration

We analysed three CFD datasets using our approach. Table 1 gives an overview of the adopted measures. The exact parameter settings can be found in the additional material, together with high resolution screenshots of the results. Dataset *A* (Fig. 5) is a $2D$ flow around a confined square cylinder. The fluid flows from left to right. Since the motion is laminar upstream from the cylinder, the velocity vectors are highly coherent, and only two large arrows are displayed. In front of the obstacle, the fluid is characterized by strong shearing, and this leads to the production of vorticity and to a sensible increase of the velocity magnitude. Two shear layers are formed around the cylinder, which, downstream, roll up into distinct vortices. We can see that the sense of rotation of the vortices is always aligned with the high-magnitude velocity vectors. Finally, two other shear layers are clearly distinguishable along the top and bottom walls.

The other two datasets are instead three-dimensional. Since in $3D$ occlusion issues can still arise, we have implemented a simple pruning mechanism that can ease the investigation of particularly dense regions. The user can interactively place a pruning geometry, such as a slice plane or a set of streamlines, in the spatial domain. The simplified visualization is obtained by displaying only the visual entities whose area of influence intersects the pruning geometry.

Dataset *B* (Fig. 6) is a CFD simulation of a flow in a box. The inlet (*i*) is placed in the top-left area, while the outlet (*o*) is on the bottom right, adjacent to the walls. Vortices are generated close to the inlet (v_1), while the region around the outlet is mainly strain-dominated. Where the inflow hits the bottom wall, strong shear is produced (s_1), which leads to the formation of large vortical areas close to the left and bottom walls (v_2). By placing a pruning plane close to the back wall, we can see a thin boundary layer (s_2) detaching from the wall, dragged by the oblique vortex (v_2).

Dataset *C* (Fig. 7) is a simulation of an exhaust manifold, with three inlets (i_1, i_2, i_3) and an outlet (*o*). There is an inflow from i_2 , while the other two inlets are currently inactive. Ideally, the fluid should flow from the active inlet to the outlet only, but in most concrete cases there is also a flow towards the




Figure 5: Dataset A is a $2D$ span-wise slice (192x64 regular grid) extracted from a Direct Numerical Simulation of a flow around a confined square cylinder [BCIS06]. We have used a uniformly resampled version provided by Tino Weinkauff, earlier presented in the work by von Funck et al. [vFWTS08].

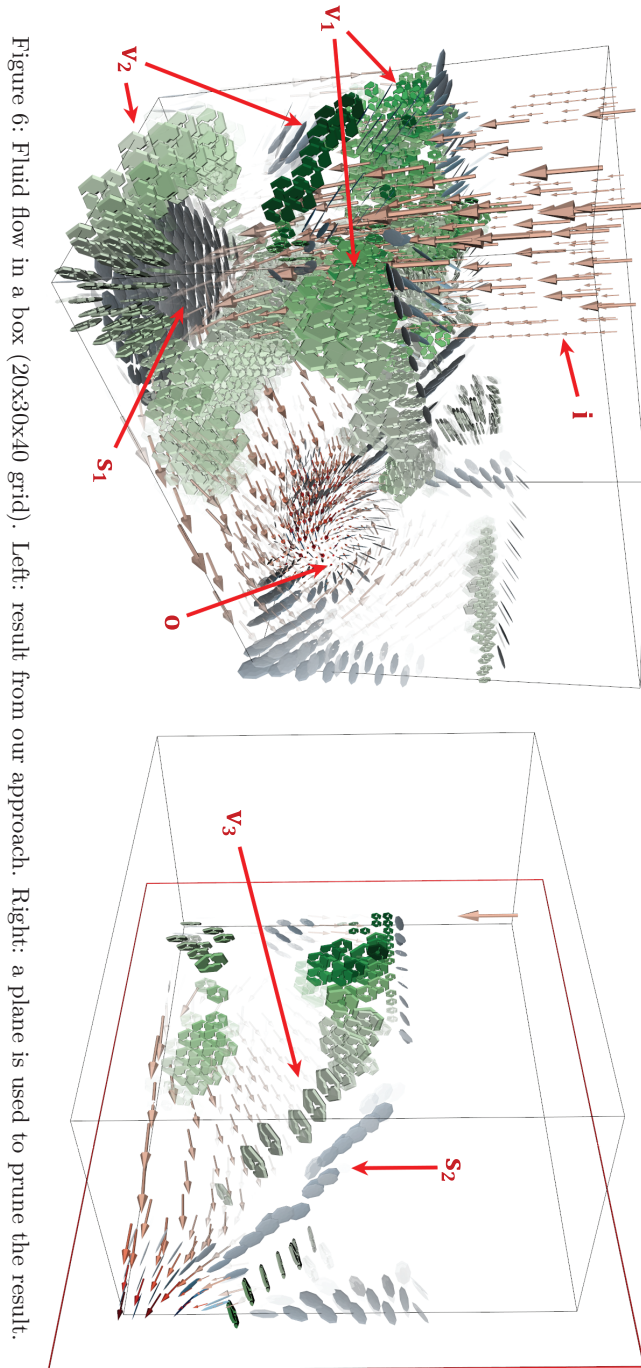


Figure 6: Fluid flow in a box (20x30x40 grid). Left: result from our approach. Right: a plane is used to prune the result.




Figure 7: Simulation of an exhaust manifold (133x81x31 grid). The central inlet (i_2) is currently active. Left: result from our approach. Right: pruning using a set of streamlines seeded from the red circles (full images in the additional material).

inactive inlets. We highlight this undesirable fluid motion by setting as highly relevant all the velocity vectors oriented either upward or rightward (upstream). Our visualization clearly emphasize a significant issue: the curved sections of the manifold are characterized by strong vortical (v_1, v_2, v_3) and shearing (s_1, s_2) motion, which can reduce the overall speed of the exhaust emissions. By pruning the result using streamlines seeded around the red circles, a second issue is exposed. The upstream flow, besides creating back pressure, leads to the creation of vortices (v_4, v_5), which can further slow down the fluid particles. Overall, our visualization suggests that both these issues can be related to the sharp turns present in the manifold's shape.

6 Discussion

Our visualizations can be easily controlled through the relevance and coherency parameters. Specifying the relevance measure is quite intuitive, since it directly corresponds to the user's interest. The coherence measure affects how the visual

entities are distributed. We have explored several parameters settings on different datasets, and none of the coherency measures clearly outperformed the others. The only significant difference is around locations of very low magnitude, such as critical points. These areas are deemed fairly coherent by the measures based on magnitudes only. The other two measures instead identify these regions as incoherent due to changes in directions. Setting the coherency threshold can be sometimes difficult, since the various measures can have different value ranges. In general, the higher the threshold, the sparser the placement of visual entities is. We initialize the threshold to 10% of the coherency value range, then the user can control the density of the visual entities by interactively adjusting it.

The performance of our system is heavily affected by the size of the dataset and the parameters setting. In the presented study cases, the evaluation of a coherency measure is in the order of tenths of second (on a 2.8 GHz CPU). The main bottleneck is instead the generation of the glyphs' geometry, which can take up to a few seconds. This issue could be easily solved by adopting GPU-based glyph rendering techniques, e.g., point sprites or texture atlases.

7 Summary and Future Work

We present an effective visualization strategy for the simultaneous depiction of multiple flow aspects. Each aspect is conveyed through a specific visual entity. Cluttering and occlusion issues are addressed by means of relevance and coherency measures. The overall appearance of the results can be controlled by tuning the relevance and coherency parameters. Moreover, our approach is strongly modular. It can be easily extended with new visual entities, different ways to define relevance, or alternative coherency measures. We are, in fact, planning to include *streamlets* as a new type of representation, and to design a specific coherency measure for tensor data based on tensor invariants.

Currently, visual entities are placed according to the underlying grid structure. In the future we plan to overcome this limitation by replacing the current discrete formulation of the initial sets P_a with a continuous one. We would also like to extend our technique to integral lines and surfaces. Finally, an extension to time-dependant datasets is possible, the main challenge is to guarantee the continuity of the placement locations over the timesteps.

Acknowledgements: We would like to thank Armin Pobitzer for his valuable feedback. The flow in a box and the exhaust manifold are courtesy of AVL List GmbH, Austria.

Paper B

Appendix: Additional Material

Andrea Brambilla¹, Øyvind Andreassen^{2,3}, Helwig Hauser¹

¹ Department of Informatics, University of Bergen, Norway

² Norwegian Defence Research Establishment, Norway

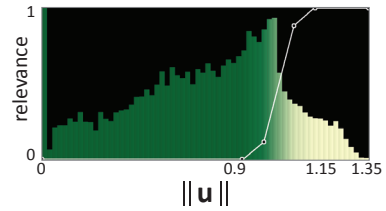
³ University Graduate Center at Kjeller, Norway

This document is a reformatted version of the accompanying material of the paper *Integrated Multi-aspect Visualization of 3D Fluid Flows*. It includes all the parameter settings and the high resolution images of the application scenarios discussed in the demonstration section of the paper.


Flow around a Square Cylinder

	Velocity	Vorticity	Rate of strain
Relevance	high velocity magnitude	positive Q	high Haines & Kenwright
Coherency	c_{dot_v}	2 nd moment of Euclidean norm	c_{diff_v}
Coh. thresh.	0.06	0.1	0.16


Relevance
measure for
Velocity




Relevance
measure for
Vorticity



Relevance
measure for
Rate of Strain





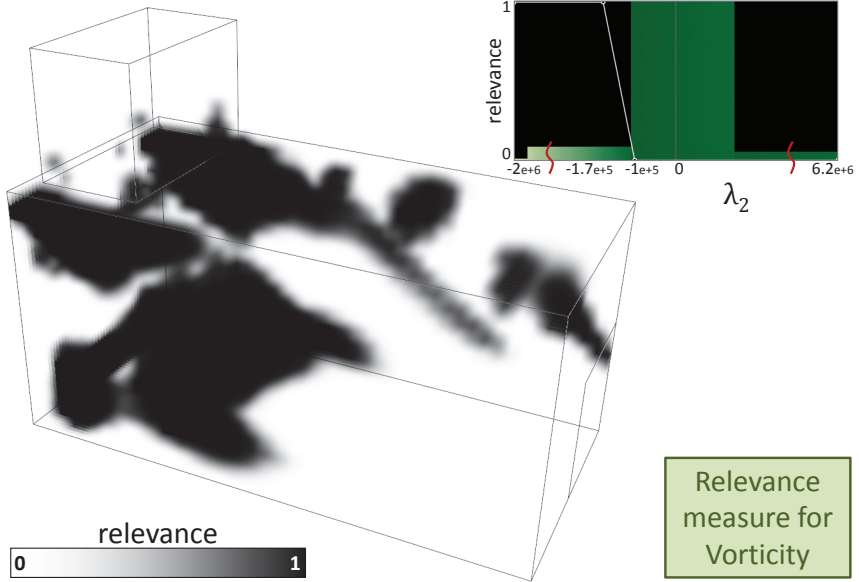



Result

Flow in a Box

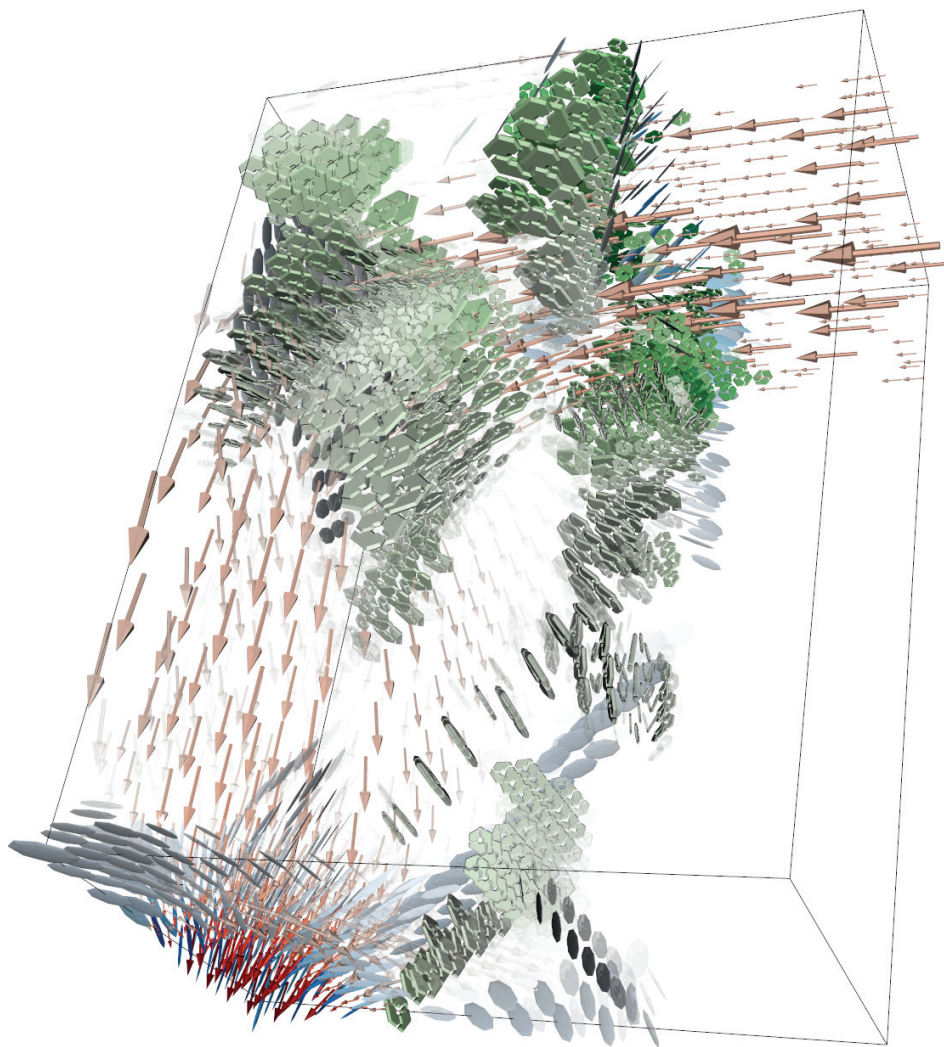
	Velocity	Vorticity	Rate of strain
Relevance	upstream velocity	negative λ_2	negative Q
Coherency	$c_{\dot{v}}$	c_{diff_v}	entropy of Frobenius norm
Coh. thresh.	0.1	0.15	0.3








Result





Exhaust Manifold

	Velocity	Vorticity	Rate of strain
Relevance	upstream velocity	negative λ_2	negative Q
Coherency	c_{dot_v}	entropy of Euclidean norm	c_{dot_v}
Coh. thresh.	0.25	0.4	0.4



Relevance
measure for
Velocity

