Framework and Algorithms for Illustrative Visualizations of Time-Varying Flows on Unstructured Meshes

Alexander S Rattner^(a), Donna Post Guillen^(b), Alark Joshi^(c), Srinivas Garimella^(a) *

^(a) Sustainable Thermal Systems Laboratory
George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332 USA

^(b) Advanced Process and Decision Systems Department
Idaho National Lab
Idaho Falls, ID 83401 USA

^(c) Computer Science Department
College of Arts and Sciences
University of San Francisco
San Francisco, CA 94117 USA

***Corresponding author:**

Email: sgarimella@gatech.edu **Phone:** 1-404-894-7479 **Fax:** 1- 404-894-8496
Address: 801 Ferst Drive, Love Building Room 340, Atlanta, GA 30332, USA

ABSTRACT

Photo- and physically realistic techniques are often insufficient for visualization of fluid flow simulations, especially for 3D and time-varying studies. Substantial research effort has been dedicated to the development of non-photorealistic and illustration-inspired visualization techniques for compact and intuitive presentation of such complex datasets. However, a great deal of work has been reproduced in this field, as many research groups have developed specialized visualization software. Additionally, interoperability between illustrative visualization software is limited due to diverse processing and rendering architectures employed in different studies. In this investigation, a framework for illustrative visualization is proposed, and implemented in MarmotViz, a ParaView plug-in, enabling its use on a variety of computing platforms with various data file formats and mesh geometries. Region-of-interest identification and feature-tracking algorithms incorporated into this tool are described. Implementations of multiple illustrative effect algorithms are also presented to demonstrate the use and flexibility of this framework. By providing an integrated framework for illustrative visualization of CFD data, MarmotViz can serve as a valuable asset for the interpretation of simulations of ever-growing scale.

Keywords

- Flow visualization
- Illustrative visualization
- Feature detection and tracking
- Unstructured meshes

NOMENCLATURE

c	cell
$\bar{c}, \bar{C}$	centroid of a cell or region (m)
d	distance of contour points from a central axis (m)
e	edge
f	feature
g	graph constructed from mesh
h	thickness for feature halos (m)
$\underline{\underline{J}}$	products of inertia tensor for a region (m ⁵)
$\bar{L}$	volumetric angular momentum of a region (m ⁵ s ⁻¹)
n	number of unmatched regions-of-interest
$\hat{n}$	camera projection plane normal
ot	octree used for feature matching
$\bar{p}$	point on a contour (m)
q	queue used for feature matching
r	region-of-interest
$\vec{r}$	position vector (m)
s	strobe silhouette curve
sc	bounding contour around a feature
t	time (s)
T	user-specified threshold
$\vec{u}, \vec{U}$	velocity of a cell, or volume-average velocity of a region (m s ⁻¹)
v, V	volume of a cell or region (m ³)
x, y, z	spatial coordinates (m)

Subscripts

0	initial value
b	bounding contour around a feature

back	rear point on a contour
c	cell count threshold for regions-of-interest
g	gradient threshold for regions-of-interest
i,j,k	counting indices
l	left-most point on a contour
in	inset contour around a feature
min, max	minimum/maximum values along a contour
o	offset contour around a feature
r	right-most point on a contour

Superscripts

*	estimated value
---	-----------------

Greek characters

α	similarity parameter for feature matching (-)
δ	user defined time step (s)
η	similarity criterion for feature matching (-)
θ	relaxation parameter for feature matching (-)
ω	hard lower limit for the feature matching criterion (-)
$\bar{\Omega}$	volumetric average angular velocity of a region (s ⁻¹)

1. INTRODUCTION

1.1 Limitations of photo- and physically realistic visualization techniques

The challenge of presenting intelligible and useful visualizations of data has attracted substantial attention due to the growing prevalence of large-scale computational simulations across most fields in science and engineering [1-3]. Such simulations are tackling increasingly ambitious problems that require unsteady modeling in complex 3D geometries or consider inherently unsteady 3D physics such as *direct numerical simulations* of turbulence in fluid mechanics, and necessitate considerations of multiple physical and temporal scales.

Various *physically accurate* or *photo-realistic* visualization techniques have been employed in the investigation and presentation of such simulation data. Vector plots and streamlines/surfaces can indicate the trajectory and magnitude of vector fields (typically velocity). However, these techniques have limited ability to present 3D vector data, because data are projected to a 2D view plane in most display systems [4], and dense lines/arrows can introduce visual clutter. In isosurface rendering, The 3D dataset is reduced to a set of surfaces that lie along some threshold value of a scalar field. This approach can help observers identify regions of interest but may obscure internal features enclosed by isosurfaces. Semi-transparent isosurfaces and cutaway views can help present such features, but application of such corrections can be a user-intensive process. Volume rendering simulates the passage of light through a 3D dataset that is assigned varying color, absorptivity, and transmissivity properties. These optical material properties are specified through a *transfer function*, which may depend on dataset fields. The flexibility of volume rendering enables users to explore and highlight diverse characteristics of interest, but manual tuning of transfer functions can be tedious.

Time-varying simulations present the additional visualization challenge of indicating histories and development of data in a simulation. Animations offer an intuitive approach to presenting time-varying phenomena, when supported by the presentation platform. However, the value of animations may be limited for complex simulations, because research on human vision indicates that viewers have difficulty simultaneously tracking many moving phenomena [5]. Image series are amenable to more presentation environments than animations, but can take longer to interpret [6]. Additionally, individual image size is often limited in a time series, reducing the clarity of the overall visualization.

1.2 Illustrative visualization techniques

The field of *illustrative visualization* has primarily been inspired by historical work in scientific illustration. Born *et al.* [7] highlighted drawings by Dallman [8] that qualitatively demonstrate key properties of vortical flow using techniques such as half-toning to indicate interior surfaces and dashed lines for hidden sections (Fig. 1a). Similarly, Correa *et al.* [9] discuss work in medical and anatomical illustrations that reveal internal details clearly by deforming bodies, removing extraneous structures, and coloring tissues in vivid fashion (Fig. 1b [10]). Born *et al.* [7] summarized the objectives of illustrative visualization as: “simplification of complex contexts, concentration on relevant features, and neglect of details that obstruct understanding.” Effective illustrative visualizations should present representations that facilitate this *qualitative partitioning and interpretation*. Similarly, analyses of unsteady phenomena often focus on histories and development of individual features. Illustrative visualization tools should therefore have mechanisms for tracking the histories of features of interest

Fig. 1 – Examples of scientific illustrations: **a.** Illustration of vortical flow employing half-toning and hidden sections [8], **b.** Anatomical illustration demonstrating removed extraneous structures and vivid coloring of tissues [10]

1.3 Illustrative visualization approach

Considering the goals and desirable features discussed in Section 1.2, a framework for illustrative visualization of time-varying CFD simulation data on unstructured meshes is presented here. In this framework, source simulations are assumed to record data fields (such as pressure, velocity, etc.) for each point or cell in mesh geometries at discrete *time steps*. Following the formulation of Joshi and Rheingans [11], the illustrative visualization process is divided into three stages for each time step.

1. **Region-of-interest identification** – *Regions-of-interest* (ROI) in the present time step of the simulation are identified and assessed.
2. **Feature matching and tracking** – Identified ROIs in the present time step are matched with corresponding nearby ROIs in other time steps. A *feature* is defined as a time series of such

1
2
3
4 matched ROIs, and could represent a physical object such as a bubble, shock front, vortex,
5
6 etc.
7

- 8
9
10 3. **Illustrative Visualization** – Illustrative techniques or *effects* are applied to specific features
11
12 (or the entire domain), and the complete visualization is rendered and presented to the user.
13
14

15 This illustrative visualization process is distinct from conventional approaches in that it
16 enables application of visualization techniques to specific features, rather than the entire data
17 domain or geometrically defined subdomains. Thus, the resulting visualizations can highlight
18 key elements of simulation data, and better capture the intuitions of subject matter experts.
19
20 Reviews of research conducted in these three tasks for illustrative visualization are presented in
21 the following sections.
22
23
24
25
26
27
28
29

30 31 **1.4 Prior work in ROI identification in simulation datasets** 32

33 Monga *et al.* [12] classified ROI identification algorithms into two groups: those that
34 group “homogeneous” cells and those that identify boundaries around regions of interest. Monga
35 *et al.* found that it was difficult to identify suitable homogeneity conditions for their target
36 applications in medical imaging; therefore, they pursued boundary detection methods. Their
37 study presented formulations for first order (gradient based) and second order (Laplacian based)
38 boundary detection algorithms. They discussed challenges and resolutions in setting threshold
39 gradient values and boundary closing on rectilinear meshes.
40
41
42
43
44
45
46
47
48
49

50
51 Banks and Singer [13] investigated the problem of identifying vortices in flow
52 simulations. Given the specialized problem domain, physical insights could be used as the basis
53 for ROI identification techniques. Banks and Singer developed a seed-and-growth algorithm in
54 which vortex cores are iteratively adjusted to lie along minimum pressure curves, effectively
55
56
57
58
59
60
61
62
63
64
65

being transported by centripetal forces. Such physically inspired approaches to ROI identification can yield accurate and “sensible” results, but are often application specific.

Silver and Wang [14] presented a number of criteria for identifying homogeneous regions, including threshold intervals, vector directions, and neighborhood connectivity. However, these approaches often rely on fine-tuning by a user, and may be geometry specific (for connectivity based approaches). In summary, a variety of ROI identification algorithms have been proposed, from generic techniques (e.g., gradient based) to specialized physics-based approaches.

1.5 Prior work in feature matching and tracking in simulation datasets

The feature-matching or *correspondence* problem has also received significant attention in the computer vision community for image processing applications. Reinders *et al.* [15] investigated pixel- and feature-based approaches. Pixel-based methods measure correspondence of individual cells between nearby ROIs in two dataset time steps. Feature-based approaches operate at a higher level of abstraction, and could range from methods that compare bulk attributes (e.g., volume, mass, and/or products of inertia) to those that compare certain cells in known ROIs. Such methods are less computationally expensive, but yield less specific information about feature history.

Meyer-Spradow *et al.* [4] developed a cell-based matching technique. In their approach, each cell and its surrounding neighbors in one frame are matched to most similarly valued cells in a second frame. Their approach can track the motion of individual cells in a feature, but is somewhat limited as presented, because it only supports datasets with single features on uniform structured rectilinear grids.

1
2
3
4 Kalivas and Sawchuk [16] developed a feature-based matching algorithm that measured
5
6 the difference between 3D regions in a relatively sparse fashion. In their approach, regions are
7
8 projected into a 2D plane and matched based on minimization of best-fit affine transforms.
9

10
11
12 Silver and Wang [17] presented a feature-based algorithm for datasets on regular meshes.
13
14 Their matching algorithm assumes that subsequent time steps are sufficiently close; therefore,
15
16 ROIs corresponding to the same feature at different times overlap in space, significantly reducing
17
18 the computational cost of searches. ROIs from two time steps are matched if the relative volume
19
20 of intersection exceeds some threshold, and pairs with the greatest volume of intersection are
21
22 matched in the case of multiple overlaps. Silver and Wang [18] published a follow-on study in
23
24 which they generalized their algorithm to unstructured static mesh geometries and achieved a
25
26 20× speedup by using refined data structures to compare features.
27
28
29
30
31

32
33 Reinders *et al.* [15] developed a matching algorithm that represented individual ROIs in
34
35 terms of a few bulk attributes including centroid, volume, and mass. In their algorithm, ROIs in
36
37 newly considered time steps are matched by extrapolating attribute values from known feature
38
39 histories. Matching is performed in multiple forward and backward passes with relaxing
40
41 tolerances, yielding rapid matching of closely corresponding ROIs, and subsequent matching of
42
43 less-similar ones.
44
45
46
47

48 This survey of feature-matching algorithms highlights a number of approaches that range
49
50 from individual cell- to bulk-attribute-level comparisons. In general, these approaches trade
51
52 computational cost for additional feature development data and matching accuracy.
53
54
55
56
57
58
59
60
61
62
63
64
65

1.6 Prior work in illustrative visualization effects

A wide variety of illustrative visualization effects has been explored in the literature, ranging from relatively simple techniques that adjust colors to complex approaches that may deform or alter the underlying mesh geometry. This section presents a survey of these effects.

Pagendarm and Walter [19] investigated simplified visualizations for flow simulations. Their goal was to find compact and distinct representations of various phenomena to enable presentation of multiple results in a single visualization. In one figure of hypersonic flow over a wing, they presented vortices (using stream-ribbons), shock-lines, and skin-friction lines.

Post *et al.* [20] presented a visualization approach that performs feature extraction and uses icons as symbolic representations of data. The goal in their study was to replace original complex feature data with clear and compact representations, such as replacing complex features with best-fit ellipsoids.

Silver and Wang [14] proposed a number of illustrative techniques for time-varying data for surface- and volumetric-rendering paradigms. In enhanced surface visualization features are assigned distinct surface colors to assist in tracking over multiple time steps. They also suggested rules for coloring features generated from the coalescence or breakup of other features (*e.g.*, red and blue features could combine into a purple feature). Similarly, in enhanced volume rendering, matched features (perhaps from the same bifurcated source feature) can be assigned specific transfer functions for volumetric rendering. In feature isolation, features of interest can be emphasized by reducing the color saturation or opacity of nearby features.

Ebert and Rheingans [21] developed a number of illustrative visualization techniques for volumetric rendering. Their formulations do not operate on specific features, but instead, are

1
2
3
4 incorporated into complex transfer functions that are applied to the entire volumetric dataset,
5
6 thus reducing visualization flexibility and possibly increasing computational cost. In their
7
8 silhouette enhancement technique, the opacity of cells around the silhouette of a feature is
9
10 increased to assist in distinction between features (also implemented by Stompel *et al.* [22] and
11
12 Born *et al.* [7]). Similarly, in feature halos the region around features, parallel to the view plane,
13
14 is assigned increased opacity. These halos provide depth cues to the viewer, as features partially
15
16 obscured by halos are easily recognized to be in the background. In surface shading, sections of
17
18 feature surfaces facing away from the view direction are assigned desaturated and darkened
19
20 colors. This provides additional cues about the shape and orientation of features.
21
22
23
24
25
26

27 Stompel *et al.* [22] presented a number of illustrative feature enhancement techniques for
28
29 both surface and volume rendering. In their depth enhancement approach, cells closer to the view
30
31 plane are rendered with warmer colors to provide additional depth cues to the viewer. In
32
33 temporal enhancement, cells or surface segments with high temporal derivatives can be rendered
34
35 in warm colors to indicate regions of rapid change. Stompel *et al.* also investigated multivariate
36
37 visualizations using multiple rendering approaches in a single image. For example, in a CFD
38
39 application, the vorticity field could be rendered volumetrically and the velocity field can be
40
41 indicated using sparse strokes.
42
43
44
45
46

47 Joshi and Rheingans presented a number of visualization techniques geared toward time-
48
49 varying datasets – particularly those with moving features. In their speedlines technique,
50
51 superimposed lines and streaks are rendered to indicate feature motion and history. With opacity
52
53 modulation, instances of features from previous time steps are presented at corresponding
54
55 locations with increasing transparency and blurriness. This permits the observer to observe the
56
57 history of a feature in a single image. This technique was adapted by Hsu *et al.* [23]. With stroke
58
59
60
61
62
63
64
65

1
2
3
4 silhouettes [11], the receding sections of feature boundaries are offset along the path of motion
5
6
7 with reduced thickness and detail.
8
9

10 Viola and Gröller [24] surveyed multiple illustrative visualization techniques. Their study
11
12 included importance-driven feature enhancement in which features are assigned “importance”
13
14 values, and low-importance features are rendered transparently or with reduced detail if they
15
16 obscure or lie near more important features. In viewpoint-dependent distortion, geometry is
17
18 distorted to increase the relative size of features of interest, and occluding features may be
19
20 moved aside [25]. In volume splitting and leaflet deformation, outer layers of nested features are
21
22 split and shifted to permit simultaneous visualization of inner and outer regions [26].
23
24
25
26

27 Svakhine *et al.* [27] investigated a number of illustrative visualization techniques geared
28
29 toward the specialized problem of visualizing CFD data. They demonstrated the use of Mixed
30
31 data fields for boundary enhancement. In an example rendering, boundary color was determined
32
33 from temperature data, and opacity from velocity data. Using two-dimensional transfer functions,
34
35 “soft” tangent-lines can be generated automatically in a feature by using a periodic function to
36
37 modulate cell opacity. Svakhine *et al.* also developed computational techniques to simulate
38
39 schlieren and shadowgraph “photographic” techniques.
40
41
42
43
44

45 Correa *et al.* [9] presented a formulation for illustrative visualization motivated by
46
47 illustrations found in anatomy texts. Their algorithms simulate common surgical and dissection
48
49 tools such as peelers, retractors, or pliers. Correa *et al.* also highlighted the value of using feature
50
51 data for generating cut-away views of 3D volumetric data without impinging on key geometry.
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Representative illustrative visualization techniques are summarized and categorized by application in Table 1. Readers are also directed to a review of illustrative visualization by Rautek *et al.* [1] that discusses the history of the field and perspectives on its future.

Table 1 – Summary of illustrative visualization techniques from literature, categorized by application

Feature emphasis	Iconic representation [20]	Silhouette enhancement [7, 21]
	Enhanced surface visualization [14]	Viewpoint dependent distortion [24]
	Feature coloring [14]	Mixed fields for boundary enhancement [27]
	Feature isolation [14]	2D transfer functions [27]
	Enhanced volume rendering [14]	Segment-aligned cuts [9]
Geometry clarification	Iconic representation [20]	Viewpoint dependent distortion [24]
	Silhouette enhancement [7, 21]	Volume splitting and leaf deformation [24]
	Surface shading [21, 22]	Geometry deformation [9]
	Distance color blending [21, 22]	Surface halftones and hatching [7]
Dynamics cues	Temporal enhancement [22]	Opacity modulation [11, 23]
	Speedlines [11]	2D transfer functions [27]
	Strobe silhouettes [11]	Storyboard presentation of data [28]
Other effects	Overlay of multiple fields [19]	Schlieren and shadowgraph techniques [27]
	Mixed rendering techniques [22]	

1.7 Present study

The three primary tasks in the illustrative visualization process (region identification, feature matching, and effect visualization) have received considerable attention in the literature. However, many of these investigations focus on specific geometries (often uniform rectangular grids), and cannot be directly extended to applications on unstructured and dynamic meshes, which are frequently employed in CFD. Additionally, a substantial amount of software development effort has been duplicated by researchers working on proprietary or specialized tools. The objective of the present study is to develop a modular framework for the illustrative visualization of time-varying CFD simulations on general unstructured mesh geometries. It is not feasible to support all illustrative visualization approaches considered in the literature, but by providing a flexible framework, users can harness a variety of ROI identification algorithms,

1
2
3
4 feature matching and tracking algorithms, and illustrative effects, and quickly implement
5
6 additional techniques.
7
8
9

10 The developed framework is implemented as a plug-in filter, MarmotViz, for the
11 ParaView [29] environment. ParaView was selected for this effort because it is an extensible
12 open source visualization tool that supports numerous data formats and mesh geometries, and
13
14 operates on a wide variety of computing platforms. Additionally, ParaView was developed using
15 the Visualization Toolkit (VTK) [30] library, which provides access to useful functionality for
16
17 many illustrative effects – such as gradient filters, computational geometry routines, and
18
19 specialized data structures. Paraview operates in a pipelined fashion, so MarmotViz receives
20
21 generalized input data from a file reader or pre-processing filters, and passes output down the
22
23 pipeline, potentially to other post-processing filters, and eventually to the rendering and display
24
25 system. Thus, the developed plug-in can be used in conjunction with other visualization tools
26
27 available in ParaView.
28
29
30
31
32
33
34
35
36

37 ParaView provides data to the MarmotViz filter at individual time steps in a simulation –
38
39 potentially out-of-order, depending on user input. Thus, at each update call, the MarmotViz filter
40
41 performs the following processes:
42
43
44

- 45 1. Read new user inputs to the MarmotViz GUI (update parameters, create/modify illustrative
46 effects, etc.)
47
48
49
50
- 51 2. Receive dataset from new time step, if the time step has been processed previously, skip to
52 step 3
53
54
55
56
57 a. Identify new ROIs using the user-selected algorithm and parameters. Assess ROIs.
58
59
60
61
62
63
64
65

- b. Match identified ROIs to known features from other time steps using the user-selected algorithm and parameters
3. Construct the initial output data field with cell values corresponding to feature indices
4. Apply illustrative effects to the output dataset
5. Pass the output dataset to the next stage in the ParaView pipeline

A schematic of the MarmotViz program structure is presented in Fig. 2.

The following sections provide detailed descriptions of the subsystems and algorithms developed in this study for ROI identification and assessment (§ 2) and feature matching and tracking (§ 3). Section 4 presents the illustrative effects subsystem and the developed effects, which include: feature coloring and selective visibility, feature smoothing, tube outlines, feature halos, speedlines, and strobe silhouettes.

All implemented ROI identification, feature matching, and illustrative effects algorithms are formulated to support general unstructured meshes.

Fig. 2 – Schematic of the MarmotViz plug-in structure

The algorithms developed here are demonstrated using data from a two-phase flow simulation of a bubble-column – a container of liquid with multiple gas injection ports at its base (Fig. 3). In particular, the fluid-phase field will be considered because it yields intuitive *features* for the following discussion, such as bubbles, droplets, and fluid layers.

2. REGION-OF-INTEREST IDENTIFICATION AND ASSESSMENT

2.1 Subsystem overview

The ROI identification and assessment subsystem receives the input mesh geometry and data fields. It first applies a user-selected ROI identification algorithm to the dataset, which can be controlled with user-specified parameters. A gradient-based ROI identification algorithm was developed in this effort to demonstrate the use and flexibility of this subsystem (Section 2.2). These ROIs are then assessed to evaluate useful attributes (Section 2.3), including: volumes,

Fig. 3 – Schematic of small-scale bubble column simulation employed to demonstrate illustrative visualization effects

centroids, and average velocities. These properties are employed in the feature matching and tracking subsystem and in various illustrative effects. A schematic of the ROI identification and assessment subsystem is presented in Fig. 4.

Fig. 4 – Schematic of the ROI identification and assessment subsystem

Fig. 5 – Transformation from mesh with high gradient-magnitude cells in red (a) to graph with connections between low-gradient magnitude cell face-neighbors (b)

2.2 ROI identification process

The demonstrative ROI identification algorithm developed for the MarmotViz visualization framework employs a boundary detection scheme to partition the geometry, as discussed by Monga *et al.* [12]. This routine begins by evaluating the gradient magnitude of a user-specified scalar field such as phase fraction, pressure, or stream function. Boundary mesh cells are identified as those with gradient magnitude of the selected input field above a user-defined threshold. The mesh geometry is then converted into a graph with each node representing a mesh cell, and edges connecting non-boundary cells that share faces (Fig. 5). This connectivity information can be computationally expensive to evaluate, especially for unstructured meshes, and is therefore cached for reuse after the first evaluation (for static input meshes). Contiguous regions in the graph are identified using a breadth-first search. Following this approach, all contiguous mesh regions separated by high gradient cells are considered to be ROIs, potentially including large background regions with little variation. However, such regions

can easily be excluded from visualizations over multiple time steps once they are classified as persisting features. Additionally, this implementation allows users to specify a minimum ROI cell count to skip small spurious regions. A second, simple threshold value-based ROI identification algorithm is also implemented.

2.3 ROI assessment process

Once ROIs are identified in a time step, they are assessed to yield attributes useful for feature matching and illustrative visualization, including: total volume, centroid, average velocity, products of inertia, and average angular velocity. Also, exterior cells and faces are identified for use in illustrative effect routines, which often operate only on feature boundaries.

Because the MarmotViz utility is developed to support general unstructured meshes, potentially composed of various cell geometries, individual cell properties, such as volume (v_i) and centroid ($\bar{c}_i$), cannot be calculated directly. ROI cells are thus, first, subdivided into tetrahedra using built-in VTK routines. Total region-of-interest (k) volumes (V_k) can be computed as sums of tetrahedral volumes (j) composing cells (i) (Eqn. 1).

$$V_k = \sum_i^{\text{cells}} \sum_j^{\text{tets}} v_{ij} \quad (1)$$

ROI centroids ($\bar{C}_k$) and volume-averaged velocities ($\bar{U}_k$) (using cell velocity data, $\bar{u}_i$) are computed similarly using Eqns. 2-3.

$$\bar{C}_k = \sum_i^{\text{cells}} \bar{c}_i v_i / V_k \quad (2)$$

$$\bar{U}_k = \sum_i^{\text{cells}} \bar{c}_i u_i / V_k \quad (3)$$

Additionally, cell velocity data ($\vec{u}_i$) are used to evaluate ROI volumetric angular momentum ($\vec{L}_k$), products of inertia tensor ($J_{\equiv k}$), and average angular velocity ($\vec{\Omega}$) using derivations by Essén [31] (Eqns. 4-6). Here, coordinates ($\vec{r} = (x, y, z)$) are relative to the ROI centroid.

$$\vec{L}_k = \sum_i^{\text{cells}} v_i (\vec{r}_i \times \vec{u}_i) \quad (4)$$

$$J_{\equiv k} = \sum_i^{\text{cells}} v_i \begin{bmatrix} y_i^2 + z_i^2 & -x_i y_i & -x_i z_i \\ -x_i y_i & x_i^2 + z_i^2 & -y_i z_i \\ -x_i z_i & -y_i z_i & x_i^2 + y_i^2 \end{bmatrix} \quad (5)$$

$$\vec{\Omega}_k = J_{\equiv k}^{-1} \vec{L}_k \quad (6)$$

Exterior cells in a ROI are found by scanning for cells that do not have neighbor cells internal to the region on all faces. The identified exterior cell faces are stored compactly as ordered point lists for later use in illustrative effect algorithms.

3. FEATURE MATCHING AND TRACKING

3.1 Subsystem overview

The feature matching and tracking subsystem is applied once for each newly visited simulation time step. The subsystem receives information about ROIs identified at the present time step and known features from other time steps. It applies a user-selected feature-matching algorithm to assign ROIs to existing features and define new features for unmatched ROIs. A representative adaptive volume based feature matching algorithm is presented in Section 3.2. A schematic of the feature matching and tracking subsystem is presented in Fig. 6.

Fig. 6 – Schematic of the feature matching and tracking subsystem

3.2 Adaptive volume-based feature-matching algorithm

A variety of feature matching strategies ranging from computationally expensive pixel-based- to fast bulk attribute-based-schemes were identified in the literature. To support large simulation datasets, a bulk attribute-based scheme was selected for this effort. The developed matching algorithm begins by seeking feature data from the temporally closest evaluated time step (t') to the current time step (t). All extant features at time step t' are stored in an octree indexed by centroids, which permits efficient comparison of large sets of regions-of-interest. The octree implementation employed in this study was developed by Krafft *et al.* [32].

The actual matching process between ROIs at time t and features at t' proceeds in a cyclic fashion, looping through a queue of unmatched regions. First, the ROI at the head of the queue is dequeued, and its centroid is obtained. An expected centroid at t' is evaluated using the average feature velocity ($\bar{C}_k = \bar{C}_k + \bar{U}_k (t' - t)$). The octree is then searched to identify the feature with the closest centroid at t' . This approach permits matching with greater time differences between data frames than that of Silver and Wang [17], because features are not required to overlap spatially.

Correspondence between the current region at time t and the feature at t' is evaluated using a volume-based heuristic (α) (Eqn. 7).

$$\alpha = 1 - \frac{|V - V'|}{\max(V, V')} \quad (7)$$

This value of the matching parameter, α , ranges from 0 when one of the volumes is zero and the other is not, to unity when they have equal volumes. If α is greater than some similarity criterion η , the region is matched to the feature, and the feature is removed from the octree. Otherwise, the unmatched region is enqueued. Every time the loop completes a pass through all unmatched features, the similarity criterion is relaxed by a user-specified factor ($\eta = \theta \cdot \eta$). This process continues until all regions at t are matched to features at t' (or vice versa), or the similarity parameter becomes smaller than some user-specified hard limit, ω . Typical values of these parameters are: $\eta = 0.9$ (first pass value), $\theta = 0.8$, and $\omega = 0.5$. These values specify that the initial matching criterion requires volume differences of less than 10% ($1 - 0.9$), the criterion relaxes to 80% at each pass, and matching ceases when η falls below 0.5 – thus, ROIs are not matched with feature instances that differ in volume by more than 50%.

One of the key advantages of this algorithm is that it dynamically adjusts the matching criterion during operation, avoiding high sensitivity to a single static parameter as in the approach of Silver and Wang [17]. In future revisions, additional matching parameters (such as products of inertia or additional data fields) and feature data from multiple time steps could be considered, as in the algorithm of Reinders *et al.* [15].

Fig. 7 – Feature coloring illustrative effect applied to bubble-column dataset, assigning unique colors to features assists in interpretation of time series or animations

4. ILLUSTRATIVE VISUALIZATION EFFECTS

4.1 Feature coloring and selective visibility

Feature coloring and selective visibility are presented first, because they are independent of the main illustrative effect subsystem that supports application of general effects (Section 4.2).

After completing feature matching, the MarmotViz filter generates an output field in which mesh cells are colored by feature number. By viewing renderings of this output field, users can easily track distinct features in animations or image series of time varying simulation data. This effect can be observed in the image series of rising bubbles from the example bubble column dataset (Fig. 7). The feature smoothing illustrative effect (Section 4.3) is applied to features in Fig. 7 for aesthetic purposes.

Additionally, the MarmotViz GUI provides user controls for selectively disabling features and removing them from the output dataset. The visualizations achieved from this

Fig. 8 – Application of feature coloring to bubble column dataset: **a.** Original dataset with surface rendering, **b.** Volume rendering reveals internal features, but with limited clarity, **c.** Isosurface rendering can extract certain features of interest, but overlapping features may be difficult to distinguish, **d.** Selective visualization can be used in conjunction with feature coloring to isolate features of interest (background and small features removed) and facilitate feature identification

process yield *importance-driven feature enhancement*, as proposed by Viola and Gröller [24].

Fig. 8 demonstrates the application of this effect to the bubble column test case to conceal obscuring background features.

4.2 Overview of illustrative effects subsystem

The developed framework also supports application of general illustrative visualization effects to features. In the current implementation, illustrative effect objects receive data from a particular feature, and apply illustrative effects by altering output mesh geometry and field data.

A higher-level illustrative effect applicator object manages the creation, modification, and deletion of illustrative effects, and applies activated effects to the output dataset during each update of the MarmotViz filter. A schematic of the illustrative effect subsystem is presented in

Fig. 9.

Fig. 9 – Schematic of the illustrative effects subsystem. Mesh and field data and user controls are received as input, individual effects are applied sequentially, and the modified mesh and field data are returned as output

Formulations for five illustrative effects are presented in the following sections to demonstrate the utility (and limitations) of the framework developed in this study.

4.3 Feature smoothing effect

ROIs identified using the gradient-based approach developed in this effort may appear blocky or jagged depending on the underlying mesh geometry. It is often desirable to present smoothed representations of these feature visualizations – especially in cases like the example dataset used here, where features (bubbles, droplets) are physically smooth. Additionally, rendering costs can be reduced by using a simplified bounding surface instead of a blocky set of cells with more exterior faces.

The smoothing effect developed here begins by extracting all exterior points from a feature at a given time step. These points are passed to the built-in VTK Delaunay meshing filter (`vtkDelaunay3D`) to generate a locally convex representation of the feature comprised of tetrahedra. The exterior surfaces of this representation are then extracted using another built-in VTK filter (`vtkDataSetSurfaceFilter`), and added to the output mesh with the same data

Fig. 10 – Features extracted from the mesh geometry may appear jagged or blocky (a), the feature smoothing effect (b) produces more aesthetically pleasing, and often, more physically realistic visualizations, while reducing the number of exterior surfaces that must be rendered

field values as the original feature. The generated smoothed surfaces are cached for reuse if time steps are revisited. A second, feature smoothing algorithm is also implemented based on the sphere shrink-wrap approach of Doria [25]. Feature smoothing is employed in conjunction with other illustrative effects in Figs. 7, 8d, 12, 14, 16, and 18. A demonstration of the feature smoothing effect is presented in Fig. 10.

4.4 Tube outline effect

The tube outline effect developed in this study is motivated by silhouette enhancement techniques discussed in the literature [21, 22]. These techniques attempt to increase the clarity of bounding edges and contours around features, relative to the view direction – assisting in the identification of and discrimination between features. The effect developed here constructs a tubular contour around features to clarify borders.

Many related efforts in the literature apply silhouette enhancement techniques directly to 2D rendered images of datasets [7, 23, 33]. Due to the pipelined nature of ParaView, effects developed in this study are limited to operating on the 3D field data and meshes. As such, the first step of the tube outline effect is to determine the projection of the selected feature to the

1
2
3
4 view plane. First, the effect retrieves the camera projection plane vector ($\hat{n}$) from ParaView. The
5
6
7 bounding contour around the feature relative to this view direction (silhouette contour) is then
8
9 obtained in a subroutine.

10
11
12 The silhouette contour determination subroutine begins by constructing three uniform
13
14 square grids (100×100 in this implementation): *CoveredCells*, *VisitedCells*, and
15
16 *HeightMap* (Fig. 11a). Next, all exterior face-points and edges on the selected feature are
17
18 retrieved (previously identified during the ROI assessment process, Fig. 11b). For convenience, a
19
20 homogeneous transform is applied to these points so that the projection plane normal ($\hat{n}$) is
21
22 aligned to the z -axis. The algorithm then iterates through all exterior feature edges, and marks
23
24 intersected grid cells in *CoveredCells* (Fig. 11c). During this process, cells in *HeightMap*
25
26 are assigned the highest z -value of intersecting edges (essentially the position of edges closest to
27
28 the view plane). Once this process completes, any weakly connected cells (marked cells with
29
30 only one marked face-neighbor) in *CoveredCells* are unmarked (Fig. 11d). The resulting grid
31
32 yields a connected, closed representation of the projection of the feature into the view plane. An
33
34 exterior marked cell in *CoveredCells* is then selected as a seed point, and the grid is
35
36 traversed in a counter-clockwise fashion between face-neighboring marked cells. Each newly
37
38 visited marked cell in *CoveredCells* is also marked in *VisitedCells*. The contour
39
40 identification process is completed when a previously visited cell is reached – typically the seed
41
42 point, but potentially another cell for cases with unusual features. A closed contour
43
44 approximating the outline of the feature from the view direction is obtained from the marked
45
46 cells in *VisitedCells* and corresponding z -values in *HeightMap* (Fig. 11e). This contour is
47
48 then transformed back into the global coordinate system (Fig. 11f).
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Fig. 11 – The silhouette contour identification process: **a.** Working grids are constructed, **b.** Exterior edges are extracted, **c.** Grid cells that intersect edges are marked (and height map is produced), **d.** Weakly connected marked cells are removed, **e.** The grid is traversed in a counter-clockwise fashion to produce a silhouette contour, **f.** The contour is transformed back to the global coordinate system

The obtained silhouette contour is then smoothed using a discrete cosine transform (DCT) filter to remove a user-specified portion of high frequency components resulting from the mapping of the feature geometry to a rectilinear grid. A built-in VTK filter (`vtkTubeFilter`) is employed to extrude the contour lines into cylindrical tube sections. The tube contour is then added to the unstructured mesh and colored with a user-defined intensity value. User controls are provided for tube thickness, color value/intensity, number of sides per cylinder segment, and the DCT contour smoothing process. A demonstration of this illustrative effect is presented in Fig. 12 for the bubble column test case.

Fig. 12 – In cases where features may be difficult to distinguish (a), the tube outline effect (b) can clarify features of interest

4.5 Feature halos effect

The feature halos illustrative effect developed here is motivated by previous work from Ebert and Rheingans [21]. In this effect, ribbon-like halos are added around selected features to clarify boundaries and provide additional depth cues. The formulation developed here differs from that of Ebert and Rheingans in that halos are constructed by adding 2D surface elements rather than by modulating the opacity of existing cells in a volumetric rendering paradigm.

The feature halos effect begins by generating a silhouette contour around a feature, reusing the routine described in Section 4.2. The contour is then smoothed using a DCT filter, as in the tube outline effect. The silhouette contour (b) is then offset (c) by thickness h_o using averaged cross products between the view-plane normal ($\hat{n}$) and vectors along the segments connected to each point ($\bar{p}_{b,i}$) on the contour (Eqn. 8). This process requires normalization of a number of intermediate vectors, which is indicated with the $\mathbf{N}$ operator in Eqn. 8 for brevity (*i.e.*, $\mathbf{N}(\bar{v}) = \bar{v} / \|\bar{v}\|$).

$$\bar{p}_{o,i} = \bar{p}_{b,i} + h_o \mathbf{N} \left\{ \mathbf{N} \left[\mathbf{N}(\bar{p}_{b,i} - \bar{p}_{b,i-1}) \times \hat{n} \right] + \mathbf{N} \left[\mathbf{N}(\bar{p}_{b,i+1} - \bar{p}_{b,i}) \times \hat{n} \right] \right\} \quad (8)$$

Fig. 13 – Halo effect process: **a.** The feature geometry is retrieved, **b.** A silhouette contour is identified and smoothed, **c.** The halo contour is produced by offsetting each point ($\bar{p}_{b,i}$) along a vector (dark red) to ($\bar{p}_{o,i}$) that is the average of cross products (blue) between silhouette contour segments (black) and the view-plane normal (green), an inset contour (points $\bar{p}_{in,i}$) is similarly produced, **d.** The feature and halo are displayed

A schematic of the contour offsetting process is presented in Fig. 13.

An inset contour is similarly constructed to fill in any gaps between the outer halo and the feature. Triangular surface elements are then generated to fill the space between contours, producing the halo. User controls are provided for halo offset/inset thickness (h_o , h_{in}), color value/intensity, and the DCT contour smoothing process. A demonstration of this illustrative effect is presented in Fig. 14.

Fig. 14 – In cases where feature positions are ambiguous (for example, is the teal feature on the left side of (a) a small feature in the foreground or a large feature in the background?), feature halos (b) can provide additional depth cues

4.6 Speedlines effect

The speedlines effect developed in this study was motivated by prior work from Joshi and Rheingans [11] and Meyer-Spradow *et al.* [4]. This effect adds trailing line-type elements to a feature to indicate its path of motion and history.

As in previously discussed effects, the speedline algorithm begins by generating a silhouette contour for the selected feature, which is smoothed using a Gaussian filter. In general, speedlines are only drawn from the trailing portion of features, so the rear portion of the contour must be extracted. This is achieved by finding the furthest contour points on each side of an axis intersecting the feature centroid ($\bar{C}_k$) parallel to the average feature velocity ($\bar{U}_k$). These points ($\bar{p}_l, \bar{p}_r$) are obtained by scanning the contour points for the two that minimize and maximize the distance function: Eqn. 9.

$$d = (\bar{p}_i - \bar{C}_k) \cdot (\hat{n} \times \bar{U}_k) \quad (9)$$

Fig. 15 – Speedlines effect process: **a.** The feature geometry and motion data are retrieved, **b.** A silhouette contour is identified and smoothed, **c.** The receding portion of the contour is extracted and speedline seed points are evenly distributed, **d.** Speedline cones are constructed on inset seed points using instantaneous velocity and angular velocity data

The trailing portion of the silhouette contour is defined by these two end-points. Speedline seed points are then evenly distributed along this portion of the contour. Local approximate velocities for each of these seed-points ($\bar{u}_i^*$) are evaluated using average feature velocity and angular velocity (Eqn. 10).

$$\bar{u}_i^* = \bar{U}_k + \bar{\Omega}_k \times (\bar{p}_i - \bar{c}_k) \quad (10)$$

Speedlines are then produced using a built-in VTK cone generator (`vtkConeSource`) with base centers inset by the base radius toward the feature centroid, preventing overhang past the feature boundary. Central axes for the cones are oriented in a direction opposite to the estimated velocities ($\bar{u}_i^*$). Cone lengths are set as $\delta |\bar{u}_i^*|$, where δ is a user-defined time difference. User controls are provided for speedline color/intensity, speedline count, cone base

Fig. 16 – Demonstration of the speedlines illustrative effect, which provides an intuitive indication of feature motion

radius, δ , and the number of cone facets. The speedlines algorithm is summarized graphically in Fig. 15, and a demonstration of this effect is presented in Fig. 16.

4.7 Strobe silhouettes effect

The strobe silhouette illustrative effect developed in this study is motivated by earlier work from Joshi and Rheingans [11]. In this effect, the receding portion of a feature silhouette is offset multiple times, indicating the history of feature motion.

The strobe silhouette algorithm begins by identifying and smoothing a silhouette contour around a given feature using a DCT filter. The receding section of the contour is then extracted – as in the speedlines algorithm. The remaining curve (s) is then offset away from the direction of

Fig. 17 – Strobe silhouette effect process: **a.** The feature geometry and velocity data are retrieved, **b.** A silhouette contour is identified and smoothed, **c.** The receding portion of the contour is extracted, **d.** The overlapping portion of the silhouette is excised, **e.** The offset silhouette is produced, **f.** Steps **c-e** are repeated for each silhouette with reduction and simplification stages

motion by the vector $-\delta\vec{U}_k$. For certain features, portions of this offset curve may overlap with the feature, and should be trimmed. First, the rear-most point on the silhouette contour with respect to the direction of motion ($\bar{p}_{back}$) is identified by minimizing the following distance function (Eqn. 11)

$$d = (\bar{p}_i - \bar{C}_k) \cdot \bar{U}_k \quad (11)$$

Then, all points in the offset curve with values of d greater than that of $\bar{p}_{back}$ are removed. The resulting curve is then extruded into cylinders as in the tube outline algorithm, added to the mesh, and colored. Subsequent strobe silhouettes are produced by repeatedly offsetting the curve, removing a fraction of points from the curve ends (shortening the curve), smoothing the resulting curve using a Gaussian filter, and generating tubes with reduced diameters. A summary of the strobe silhouette algorithm is presented graphically in Fig. 17.

Fig. 18 – Demonstration of the strobe silhouettes effect, which indicates the motion history of the receding portion of a feature

Users can control the silhouette color/intensity, initial tube diameter, number of silhouettes, relative reduction in curve length and tube diameter between consecutive silhouettes, δ , number of facets per tube, and the DCT contour smoothing process. A demonstration of this effect is presented in Fig. 18 for the bubble column dataset.

5. DISCUSSION

Illustrative visualization can serve as a valuable asset in the exploration and presentation of CFD simulation data. The ability to autonomously track, selectively visualize, and highlight particular features can aid in the understanding of evolution of phenomena such as bubbly flows and turbulence structures wherein multitudes of similar features are present. Feature halos and similar techniques can also provide depth cues when 3D simulation datasets are rendered in 2D

Fig. 19 – Conventional visualization techniques such as cut-planes (a) and isosurfaces rendering (b) rely on manual user input. Streamlines and vector plots (c) may significantly increase visual clutter. The illustrative visualization approach (d) enables autonomous feature identification, selective visualization, and compact representation of dynamics

media. Illustrative effects, such as speedlines and strobe silhouettes, can provide intuitive representations of dynamics with minimal visual clutter in generated renderings. Comparative presentations of the entire bubble column dataset conventional and the developed illustrative visualization techniques are presented in Fig. 19.

The key feature of MarmotViz is its flexibility in supporting general ROI identification algorithms, feature matching and tracking algorithms, and illustrative visualization effects. MarmotViz supports input data on both structured meshes (*e.g.*, the bubble column test case) and unstructured meshes (shown for bubbly flow in a twisted channel, Fig. 20). This framework permits modification of the underlying mesh geometry in addition to intensities and opacities of output data fields. As such, the current implementation can serve as a basis for experimentation and development of novel algorithms for illustrative visualization. Also, by using the MarmotViz platform, developers can reuse functionality, such as the silhouette contour detection routine employed in the example effects discussed in this study. Furthermore, by operating in the ParaView environment, developers and users can harness built-in filters and visualization tools, and access datasets in a wide variety of file formats and geometries.

Fig. 20 – **a.** MarmotViz applied to simulation of bubbly flow in a twisted channel with an unstructured mesh (detail view), **b.** Resulting visualization of bubbles with speedlines

While MarmotViz includes a variety of illustrative effects, only single ROI identification and feature matching algorithms are provided in the present implementation. In the future, we plan to incorporate additional and more sophisticated algorithms for these subsystems. Furthermore, future implementations of MarmotViz could also benefit from an event detection stage after feature matching is performed, as proposed by Reinders *et al.* [15].

Additionally, while the pipelined nature of ParaView facilitates interaction with other visualization filters and tools, it can be somewhat limiting because the illustrative techniques developed here cannot operate on pre-rendered 2D images. As a further consequence of this architecture, the MarmotViz filter must be manually triggered to update when the view position or orientation changes. Also, certain illustrative techniques that require more flexible display capabilities, such as the storyboard approach presented by Lu and Shen [28], may be infeasible in the current implementation.

6. CONCLUSIONS

In this study, a generalized framework was developed for illustrative visualization of time varying CFD datasets on unstructured meshes. The framework was implemented in MarmotViz, a ParaView plug-in, and supports the use of generalized ROI identification algorithms, feature matching and tracking algorithms, and illustrative effects. This implementation incorporated a gradient-based ROI identification routine, and a novel, adaptive, volume-based feature matching algorithm. A number of illustrative visualization effects were implemented to demonstrate the use and utility of this framework, including: feature coloring, selective visibility, feature smoothing, tube outlines, feature halos, speedlines, and strobe silhouettes.

The current implementation of the MarmotViz plug-in is somewhat limited as it only incorporates basic algorithms for region-of-interest identification and feature matching. Future versions will include additional and more sophisticated algorithms for these subsystems. Additional illustrative techniques will also be incorporated in future releases. An event detection subsystem may improve the functionality of MarmotViz, and enable more sophisticated and expressive illustrative effects. Overall, MarmotViz enables the exploration and presentation of CFD data in an intuitive fashion, and can serve as a valuable resource for interpreting simulations of ever-growing scale.

ACKNOWLEDGEMENTS

The authors wish to acknowledge generous financial support from the U.S. Department of Energy through the Krell Institute (contract DE-FG02-97ER25308) and the DOE Idaho Operations Office (contract DE-AC07-05ID14517).

ROLE OF FUNDING SOURCE

The funding sources: Krell Institute and U.S. Department of Energy Idaho Operations Office did not participate in the design and execution of this study or in the preparation of this manuscript.

REFERENCES

- [1] Rautek P, Bruckner S, Gröller E, Viola I. Illustrative visualization: new technology or useless tautology? SIGGRAPH Computer Graphics. 2008;42: p. 1-8.
- [2] Johnson C, Moorhead R, Munzner T, Pfister H, Rheingans P, Yoo TS. NIH/NSF Visualization Research Challenges. Los Alamitos, CA. 2006.
- [3] Johnson C. Top scientific visualization research problems. IEEE Computer Graphics and Applications. 2004;24: p. 13-7.
- [4] Meyer-Spradow J, Ropinski T, Vahrenhold J, Hinrichs K. Illustrating Dynamics of Time-Varying Volume Datasets in Static Images. of Vision, Modeling and Visualization 2006. p. 333-40.
- [5] Pylyshyn ZW. Seeing and Visualizing: It's Not What You Think: MIT Press; 2003.
- [6] Joshi A, Rheingans P. Evaluation of illustration-inspired techniques for time-varying data visualization. Computer Graphics Forum. 2008;27: p. 999-1006.
- [7] Born S, Wiebel A, Friedrich J, Scheuermann G, Bartz D. Illustrative Stream Surfaces. IEEE Transactions on Visualization and Computer Graphics. 2010;16: p. 1329-38.
- [8] Dallmann U. Topological Structures of Three Dimensional Flow Separations: Deutsches Zentrum für Luft- und Raumfahrt; 1983.
- [9] Correa CD, Silver D, Chen M. Feature aligned volume manipulation for illustration and visualization. IEEE Transactions on Visualization and Computer Graphics. 2006;12: p. 1069-76.
- [10] Gray H, Lewis WH. Anatomy of the human body. Philadelphia, PA: Lea & Febiger; 1918.
- [11] Joshi A, Rheingans P. Illustration-inspired techniques for visualizing time-varying data. IEEE Visualization Conference 2005. p. 679-86.
- [12] Monga O, Deriche R, Rocchisani J-M. 3D edge detection using recursive filtering: Application to scanner images. CVGIP - Image Understanding. 1991;53: p. 76-87.
- [13] Banks DC, Singer BA. A predictor-corrector technique for visualizing unsteady flow. IEEE Transactions on Visualization and Computer Graphics. 1995;1: p. 151-63.
- [14] Silver D, Wang X. Visualizing evolving scalar phenomena. Future Generation Computer Systems. 1999;15: p. 99-108.

- [15] Reinders F, Post FH, Spoelder HJW. Visualization of time-dependent data with feature tracking and event detection. *Visual Computing*. 2001;17: p. 55-71.
- [16] Kalivas DS, Sawchuk AA. A region matching motion estimation algorithm. *CVGIP - Image Understanding*. 1991;54: p. 275-88.
- [17] Silver D, Wang X. Volume tracking. *IEEE Visualization Conference* 1996. p. 157-64.
- [18] Silver D, Wang X. Tracking scalar features in unstructured data sets. *IEEE Visualization Conference* 1998. p. 79-86.
- [19] Pagendarm H-G, Walter B. Feature detection from vector quantities in a numerically simulated hypersonic flow field in combination with experimental flow visualization. *IEEE Visualization Conference* 1994. p. 117-23.
- [20] Post FH, Post FJ, Walsum TV, Silver D. Iconic Techniques for Feature Visualization. *IEEE Visualization Conference* 1995. p. 288-95.
- [21] Ebert D, Rheingans P. Volume illustration: nonphotorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics*. 2001;7: p. 253-64.
- [22] Stompel A, Lum EB, Kwan-Liu M. Visualization of multidimensional, multivariate volume data using hardware-accelerated non-photorealistic rendering techniques. *IEEE Pacific Conference on Computer Graphics and Applications* 2002. p. 394-402.
- [23] Hsu W-H, Mei J, Correa CD, Ma K-L. Depicting Time Evolving Flow with Illustrative Visualization Techniques. In: Huang F, Wang R-C, editors. *Arts and Technology*. Springer Berlin Heidelberg; 2010. p. 136-47.
- [24] Viola I, Gröller ME. Smart Visibility in Visualization. *Computational Aesthetics in Graphics, Visualization and Imaging* 2005. p. 8.
- [25] Doria D. ConvexHullShrinkWrap. <http://www.vtk.org/Wiki/VTK/Examples/Cxx/PolyData/ConvexHullShrinkWrap> 2010 (accessed 12/28/2015).
- [26] McGuffin M, Tancau M, Balakrishnan L. Using deformations for browsing volumetric data. *IEEE Visualization Conference* 2003. p. 401-8.
- [27] Svakhine NA, Jang Y, Ebert D, Gaither K. Illustration and photography inspired visualization of flows and volumes. *IEEE Visualization Conference* 2005. p. 687-94.

- 1
2
3
4 [28] Lu A, Shen H-W. Interactive Storyboard for Overall Time-Varying Data Visualization.
5 IEEE Pacific Visualization Conference 2008. p. 143-50.
6
7
8 [29] Henderson A. The ParaView Guide, A Parallel Visualization Application. 4th ed: Kitware,
9 Inc.; 2012.
10
11
12 [30] The VTK User's Guide. 11 ed: Kitware, Inc.; 2010.
13
14
15 [31] Essén H. Average angular velocity. European Journal of Physics. 1993;14: p. 201-5.
16
17
18 [32] Krafft MF, Harris P, Bougerel S. libkdtree++ 0.7.0. 2008.
19
20
21 [33] Burns M, Klawe J, Rusinkiewicz S, Finkelstein A, DeCarlo D. Line drawings from volume
22 data. ACM Transactions on Graphics. 2005;24: p. 512-8.
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

LIST OF FIGURES

1. Examples of scientific illustrations: **a.** Illustration of vortical flow employing half-toning and hidden sections [8], **b.** Anatomical illustration demonstrating removed extraneous structures and vivid coloring of tissues [33]
2. Schematic of the MarmotViz plug-in structure
3. Schematic of small-scale bubble column simulation employed to demonstrate illustrative visualization effects
4. Schematic of the ROI identification and assessment subsystem
5. Transformation from mesh with high gradient-magnitude cells in red (**a**) to graph with connections between low-gradient magnitude cell face-neighbors (**b**)
6. Schematic of the feature matching and tracking subsystem
7. Feature coloring illustrative effect applied to bubble-column dataset, assigning unique colors to features assists in interpretation of time series or animations
8. Application of feature coloring to bubble column dataset: **a.** Original dataset with surface rendering, **b.** Volume rendering reveals internal features, but with limited clarity, **c.** Isosurface rendering can extract certain features of interest, but overlapping features may be difficult to distinguish, **d.** Selective visualization can be used in conjunction with feature coloring to isolate features of interest (background and small features removed) and facilitate feature identification
9. Schematic of the illustrative effects subsystem. Mesh and field data and user controls are received as input, individual effects are applied sequentially, and the modified mesh and field data are returned as output

- 1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
10. Features extracted from the mesh geometry may appear jagged or blocky (**a**), the feature smoothing effect (**b**) produces more aesthetically pleasing, and often, more physically realistic visualizations, while reducing the number of exterior surfaces that must be rendered
 11. The silhouette contour identification process: **a.** Working grids are constructed, **b.** Exterior edges are extracted, **c.** Grid cells that intersect edges are marked (and height map is produced), **d.** Weakly connected marked cells are removed, **e.** The grid is traversed in a counter-clockwise fashion to produce a silhouette contour, **f.** The contour is transformed back to the global coordinate system
 12. In cases where features may be difficult to distinguish (**a**), the tube outline effect (**b**) can clarify features of interest
 13. Halo effect process: **a.** The feature geometry is retrieved, **b.** A silhouette contour is identified and smoothed, **c.** The halo contour is produced by offsetting each point ($\bar{p}_{b,i}$) along a vector (dark red) to ($\bar{p}_{o,i}$) that is the average of cross products (blue) between silhouette contour segments (black) and the view-plane normal (green), an inset contour (points $\bar{p}_{in,i}$) is similarly produced, **d.** The feature and halo are displayed
 14. In cases where feature positions are ambiguous (for example, is the teal feature on the left side of (**a**) a small feature in the foreground or a large feature in the background?), feature halos (**b**) can provide additional depth cues
 15. Speedlines effect process: **a.** The feature geometry and motion data are retrieved, **b.** A silhouette contour is identified and smoothed, **c.** The receding portion of the contour is extracted and speedline seed points are evenly distributed, **d.** Speedline cones are constructed on inset seed points using instantaneous velocity and angular velocity data

16. Demonstration of the speedlines illustrative effect, which provides an intuitive indication of feature motion
17. Strobe silhouette effect process: **a.** The feature geometry and velocity data are retrieved, **b.** A silhouette contour is identified and smoothed, **c.** The receding portion of the contour is extracted, **d.** The overlapping portion of the silhouette is excised, **e.** The offset silhouette is produced, **f.** Steps **c-e** are repeated for each silhouette with reduction and simplification stages
18. Demonstration of the strobe silhouettes effect, which indicates the motion history of the receding portion of a feature
19. Conventional visualization techniques such as cut-planes (**a**) and isosurfaces rendering (**b**) rely on manual user input. Streamlines and vector plots (**c**) may significantly increase visual clutter. The illustrative visualization approach (**d**) enables autonomous feature identification, selective visualization, and compact representation of dynamics
20. **a.** MarmotViz applied to simulation of bubbly flow in a twisted channel with an unstructured mesh (detail view), **b.** Resulting visualization of bubbles with speedlines

LIST OF TABLES

1. Summary of illustrative visualization techniques from literature, categorized by application

Table 1 – Summary of illustrative visualization techniques from literature, categorized by application

Feature emphasis	Iconic representation [20]	Silhouette enhancement [7, 21]
	Enhanced surface visualization [14]	Viewpoint dependent distortion [24]
	Feature coloring [14]	Mixed fields for boundary enhancement [27]
	Feature isolation [14]	2D transfer functions [27]
	Enhanced volume rendering [14]	Segment-aligned cuts [9]
Geometry clarification	Iconic representation [20]	Viewpoint dependent distortion [24]
	Silhouette enhancement [7, 21]	Volume splitting and leaf deformation [24]
	Surface shading [21, 22]	Geometry deformation [9]
	Distance color blending [21, 22]	Surface halftones and hatching [7]
Dynamics cues	Temporal enhancement [22]	Opacity modulation [11, 23]
	Speedlines [11]	2D transfer functions [27]
	Strobe silhouettes [11]	Storyboard presentation of data [28]
Other effects	Overlay of multiple fields [19]	Schlieren and shadowgraph techniques [27]
	Mixed rendering techniques [22]	