

Visual Analysis of Droplet Dynamics in Large-Scale Multiphase Spray Simulations

Moritz Heinemann · Steffen Frey · Gleb Tkachev · Alexander Straub · Filip Sadlo · Thomas Ertl

Received: date / Accepted: date

Abstract We present a data-driven visual analysis approach for the in-depth exploration of large numbers of droplets. Understanding droplet dynamics in sprays is of interest across many scientific fields for both simulation scientists and engineers. In this paper, we analyze large-scale direct numerical simulation datasets of two-phase flow of non-Newtonian jets. Our interactive visual analysis approach comprises various dedicated exploration modalities that are supplemented by directly linking to ParaView. This hybrid setup supports a detailed investigation of droplets, both in the spatial domain and in terms of physical quantities. Considering a large variety of extracted physical quantities for each droplet enables the investigation of different aspects of interest in our data. To get an overview of different types of characteristic behavior, we cluster massive numbers of droplets to analyze different types of occurring behaviors via domain-specific pre-aggregation, as well as different methods and parameters. Extraordinary temporal patterns are of high interest especially to investigate edge cases and detect potential simulation issues. For this, we use a neural network-based approach to predict the development of these physical quantities, and with this identify irregularly advected droplets.

Keywords Flow Visualization · Time-varying Data · Visualization in Physical Sciences and Engineering.

1 Introduction

Flow visualization has mostly been concerned with the analysis of single-phase flow, i.e., flows where a single type of fluid is involved (e.g., airflow around objects, or liquid flow through machinery). However, in many problems in science and engineering, two or even more phases are involved, such as in the case of water flowing in a domain containing air, or in the dynamics of immiscible liquids. A major difficulty with the analysis of such multiphase flow is, however, its various

M. Heinemann
Allmandring 19, 70569 Stuttgart, Germany
Tel.: +49 (711) 685-60933
E-mail: moritz.heinemann@visus.uni-stuttgart.de

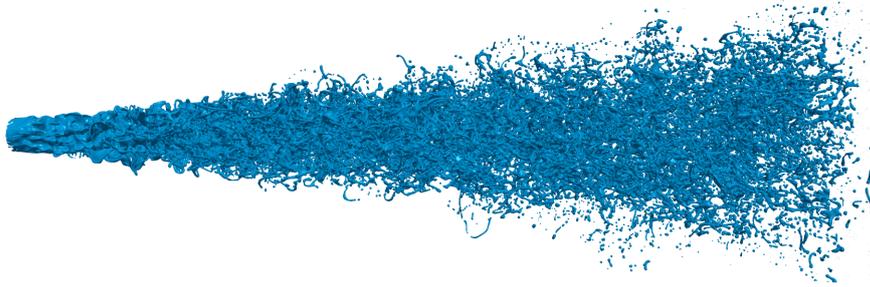


Fig. 1: Interface between liquid and air in the Jet Simulation dataset after the formation of a stable jet. Each time step features a rectilinear grid with a resolution of $1536 \times 512 \times 512$ cells covering a domain of $12 \text{ cm} \times 4 \text{ cm} \times 4 \text{ cm}$; 623 time steps represent a time span of 5.567 ms. The result is a volume of fluid field f and a velocity field \mathbf{u} , both are given in cell-based representation.

degrees of complexity. On the one hand, it inherits all complexity of single-phase flow, whose visualization is subject to ongoing research. On the other hand, the dynamics and physics of the interface between the different phases are closely interrelated with the flow, as well as phenomena from solid mechanics such as collision, deformation, and adhesion. And last but not least, the volume of fluid method (VOF) (Hirt and Nichols, 1981), which is typically used for simulating multiphase flow, further complicates analysis and interpretation. One reason is that the interface between the phases is not tracked, but reconstructed at each time step during simulation, leading to inconsistencies between the flow in the vector field and the motion of the interface.

A phenomenon in two-phase flow with particularly high complexity is the formation of sprays. Sprays play an essential role in a wide range of natural phenomena and production, including precipitation, combustion, food processing, production and application of drugs, and cooling. In technical applications, sprays are typically generated by guiding the liquid through a spray nozzle, which produces an unstable jet that eventually breaks up into droplets. The dynamics of this breakup is highly complex, with primary breakup producing elongated components called ligaments, followed by secondary breakup of ligaments into droplets.

All these processes and complexities make the resulting data extremely hard to navigate and analyze, also due to their spatiotemporal nature. While traditional visualization approaches are applicable to subsets and partial aspects of the data, they cannot provide an effective means for hypothesis forming and hypothesis testing for the entire data, in particular because the importance and interrelations of specific quantities and processes are buried in the discussed degrees of complexity. Another difficult aspect is that in order to be able to resolve the complex dynamics involved at small scales, high temporal and spatial resolution of the data is required.

This results in data sizes quickly reaching terabytes, impeding the direct application of many types of advanced analysis procedures. While interactive exploration is required to analyze the complex, small-scale effects, this is challenging with the large data sizes and numbers of droplets.

In this work, we aim to analyze the droplets in a two-phase DNS flow simulation of the breakup of a liquid jet in air (cf. Figure 1). This simulation was conducted by Ertl (Ertl and Weigand, 2017; Ertl, 2019) using FS3D (Eisenschmidt et al., 2016) and considered water with 0.3% flocculant, leading to non-Newtonian fluid behavior. In this case, the interest of the simulation scientists focused on the fully converged phase of the jet. Accordingly, we mainly consider a respective sub-sequence of 101 time steps covering 0.679 ms from the full simulation (only collision and separation counters cover the full 623 time steps). The total data size is around 7 TB (i.e., ≈ 1 TB for the focused subset of 101 time steps).

Our dedicated visual analysis approach provides different hypothesis forming tools to analyze droplet behavior in the data. Our goal is (i) to allow simulation and domain scientists to analyze what common types of droplets and their behaviors are, and (ii) to identify and study anomalous—i.e. uncommon—cases in-depth. Below, after discussing related work (Section 2), we describe our workflow enabling the interactive visual analysis of large sets of droplets based on the extraction of meaningful quantities and advanced further automated analysis, including clustering and machine learning-based anomaly detection (Section 3). Our visual analytics system then provides different perspectives on abstract and spatial quantities, as well as allowing for detailed flow analysis with the original raw data (Section 4). Finally, we present insights gained from visual analysis (Section 5) before concluding Section 6.

We present the first visual analysis approach for complex time-dependent multiphase flow data. By combing domain specific clustering techniques with an artificial neural network for learning droplet anomaly behavior we developed a highly interactive framework to support flow scientists in the complicated tasks of analyzing terabytes of simulation data.

2 Related Work

Our data stems from a CFD solver for multiphase flow simulation of incompressible fluids. This solver (Eisenschmidt et al., 2016) uses the VOF method (Hirt and Nichols, 1981), in combination with Piecewise Linear Interface Calculation (PLIC) as interface reconstruction algorithm that we also use in this work to analyze droplets (cf. Karch et al. (2013) for a discussion of PLIC in visualization).

The visual analysis of time-dependent flow fields is challenging. Aigner et al. (2007) discuss the consideration of time as an additional dimension in visual analytics. Buerger et al. (2007) integrate local feature detectors in the visual analysis of time-dependent flow simulations. Another line of work concerned itself with the interactive exploration of complex time-dependent flow simulations and real-world data (Doleisch et al., 2003; Helmut et al., 2004; Doleisch et al., 2004). Shi et al. (2007) introduced an approach to visually analyze time-dependent flow fields by means of pathlines in particular. An overview of different feature tracking techniques for flow visualization was presented by Post et al. (2003). Further Theisel and Seidel (2003) show a stream line integration-based method for feature tracking

in instationary vector fields. Garth et al. (2004) discuss an approach for tracking singularities within a vector field. In contrast to these previous works, while we also employ techniques from flow visualization for the detailed investigation of individual droplets, our specific focus on large numbers of droplet data feature quite unique challenges that we aim to address in this work. Further previous work already focused on feature tracking in the context of large-scale datasets (Dutta and Shen, 2016) and calculating tracking information in situ (Biswas et al., 2020).

Combining flow visualization and machine learning, Tzeng and Ma (2005) used neural networks to generate adaptive transfer functions based on user input. Bai et al. (2017) applied linear discriminant analysis to classify experimental images into actuated and unactuated flow. In this paper we also apply neural networks to support flow visualization, but use them to generate visual features that help guiding the analysis. Tkachev et al. (2019) trained neural networks on spatiotemporal volumes to detect irregular behavior. In this work, we use a similar idea for our anomaly detection, but we apply our model to time series of extracted droplet quantities. Machine learning has also been applied in fluid simulations. Artificial neural networks are even used within solvers of the Navier Stokes equation (Tompson et al., 2017) for acceleration of the computation. In visualization in general, machine learning methods are particularly popular in visual analytics approaches (Endert et al., 2017).

This work started in the context of a master thesis (Heinemann, 2018) which was reported in a very early version in a project report (Straub et al., 2019) and within a book chapter, Straub and Ertl (2020). The idea of using ML for prediction on droplet quantity time series is already mentioned there, but due to its overview character, no details of our technique were presented. Further an early version of the surface view of our prototype was depicted there. In this article we introduce our entire frame work and elaborate on new features like clustering, improvements to the 3D surface view and extension to the full system including the quantity and flow view and also an analysis and discussion of the used dataset.

3 Preprocessing: Extraction, Clustering and Anomaly Detection

To make the data (interactively) explorable while still capturing the various degrees of complexity of droplet behavior, we first extract different quantities on the basis of individual droplet instances (Section 3.1, cf. Figure 2). Next, we establish temporal relationships by connecting individual droplet instances between the time steps to time-continuous droplet traces (Section 3.2). To get an overview of the droplet quantities we are using trace-based hierarchical clustering (Section 3.3). In addition, we train a regression model to capture typical temporal patterns in a droplets' quantities, followed by determination of the deviations from this model to guide the researcher to cases that are anomalous in the sense of being uncommon. Akin to previous work (Tkachev et al., 2019), we chose artificial neural networks (ANNs) due to their generality, performance efficiency on large data (compared to, e.g., non-parameteric models), and their successful applications across many diverse tasks (Section 3.4). Below, we particularly focus on the application and adaptation of the different methods involved to enable the interactive exploration of large numbers of droplets (Section 4). [An overview of the order and dependencies of our processing steps can be seen in Figure 2.](#)

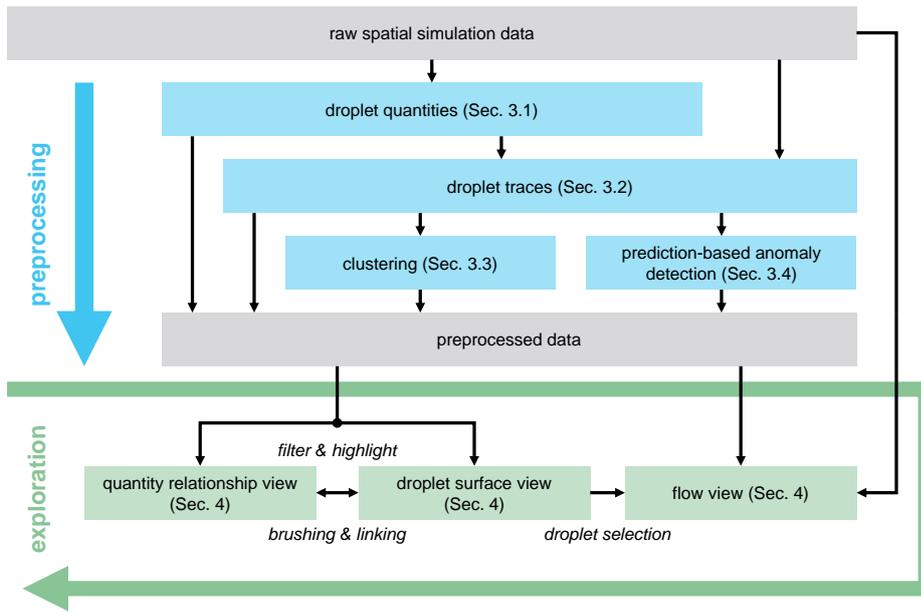


Fig. 2: Overview on the components of our visual analysis approach. This represents the execution order of our processing pipeline and shows data dependency of all steps described within Section 3 and Section 4. Each of the black arrows represents data usage from the previous step.

3.1 Extraction of Quantities

We extract different instantaneous quantities for each droplet instance that capture the various degrees of complexity of droplet behavior. In collaboration with domain experts, the following set of instantaneous quantities has been determined, ranging from purely geometric to purely physical.

Data Representation and Segmentation. The first step consists in the identification of the individual droplets. The employed VOF approach in our two-phase flow simulation maintains a scalar field $f(\mathbf{x}) \in \mathbb{R}$ during simulation. This field stores conceptually for each point $\mathbf{x} \in \mathbb{R}^3$ the volume fraction of the liquid, i.e., $f = 0$ representing only gas, $f = 1$ representing only liquid, and $0 < f < 1$ mixtures of both. Here, $f(\mathbf{x})$ is defined in a cell-based manner, i.e., the field stores this total fraction for each cell i individually, in piecewise constant representation. Thus, f_i represents the value of the f -field for cell i . Similarly, the velocity field $\mathbf{u}(\mathbf{x})$ is given in cell-based representation \mathbf{u}_i . In the domain, a droplet is commonly defined as the face-connected component of cells that exhibit $f_i > 0$. In practice, due to the numerical limitations of the simulation, we use the slightly modified definition $f_i > \tau_f$, with $\tau_f = 10^{-6}$. This value is defined by our domain experts, who run the simulation. Thus, we obtain the droplets by connected component labeling using region growing, resulting in a cell-based label field $l(\mathbf{x})$, where l_i stores the droplet identifier ι of the droplet that cell i belongs to.

A further peculiarity of two-phase flow simulation is the representation of the interface between the two phases via PLIC (Youngs, 1984). During simulation, the

interface is represented in a piecewise planar manner, i.e., within each cell with $0 < f < 1$, a planar patch represents the interface. It is determined by using the negative gradient of the f -field as plane normal and choosing the position such that the volume below the patch is exactly the volume fraction represented by f_i . This piecewise linear representation per cell results in a discontinuous global interface containing gaps (see, e.g., Figure 13). Obviously, using isosurface extraction for interface representation would violate both the connected component definition of droplets as well as interface representation, and is thus dissuaded by domain experts. Therefore, we use the PLIC patches as geometric interface representation for both the computation of derived quantities and rendering, similar to Karch et al. (2013), at the cost of accepting the discontinuities.

Main Instantaneous Quantities. We want the extracted quantities to capture droplet behavior reasonably accurate, but at the same time we are limited to quantities which can be calculated purely from the VOF and velocity field. Further we need a minimum of numeric stability in calculating them. While the overall jet simulation is high-resolution, single droplets within this dataset may only be resolved by few cells. Especially when calculating derivatives on the velocity fields, we have observed instabilities due to the coarse resolution. We therefore focus on derivative-free quantities, and compute a total of eleven scalar quantities for each temporal instance of droplet ι : volume $\bar{\mu}_\iota$, area \bar{A}_ι , area to volume ratio $\bar{\alpha}_\iota$, velocity $\|\bar{\mathbf{u}}_\iota\|$, momentum $\|\bar{\mathbf{p}}_\iota\|$, angular velocity $\|\bar{\boldsymbol{\omega}}_\iota\|$, angular momentum $\|\bar{\mathbf{L}}_\iota\|$, total energy E_ι , kinetic energy $E_\iota^{\mathbf{u}}$, rotational energy $E_\iota^{\boldsymbol{\omega}}$, and residual energy E_ι^{δ} . The formal definition of these quantities is provided within the appendix of this paper (Appendix A). Notice that magnitudes are used in case of vector-valued quantities, for ensuring rotational invariance.

Additional Quantities for Evaluation and Visualization. In addition to the main quantities used for the further preprocessing steps, we use additional quantities for visualization and evaluation. The polygons of PLIC surface are stored to later visualize droplet surfaces. While we deemed derivation-based quantities to be not numerically stable enough to be used in automatic analysis, we still consider them useful when carefully employed in supplementing the analysis. To be able to discern between droplets and ligaments, we compute spherical anisotropy measure from our segmented droplets. We achieve this by computing fractional anisotropy (Rosenberger et al., 2012; Basser and Pierpaoli, 2011) of the covariance matrix of all cell centers comprising a droplet component ι , which we denote c_s , with $c_s = 0$ indicating linear or planar shape, and $c_s = 1$ for a perfectly spherical shape of a droplet. Furthermore, we compute the radius from the smallest surrounding sphere around the center of mass. While our main quantities above are designed to just include droplet local quantities, for analysis purpose we include a location-based quantities to refer to position within the jet. Next to absolute center of mass position, we use the distance of a droplets center of mass to the base axis of the jet and call it radius. In contrast, we do not include different velocity components, i.e. axial and radial velocity of a droplet, because the axial velocity is expected to roughly equal the overall velocity and the radial velocity is assumed to be a relatively small constant only depending of the position within the jet, as described by domain experts. In addition, the number of cells of droplet component ι is provided as a discretized alternative to $\bar{\mu}_\iota$ with two variants: (1)

including the number of cells that are completely filled with the liquid phase and (2) the amount of cells that are at least filled by 50%. In addition, we calculate the vortex core lines of each droplet according to Sujudi and Haines (1995), and count the number of line segments to quantify the presence of a vortex. Furthermore, we save error flags during computation process, i.e., if droplets are too small or we cannot determine the trace of a single droplet. This can later be especially used for filtering.

3.2 Trace Generation

From the segmented **droplet instances**, we now establish a space-time graph depicting their temporal correspondence (i.e., a node is a droplet instance, and each edge a temporal relationship). Initially, the graph consists only of nodes, and is extended by adding edges **if we find correspondence between droplet instances in neighboring time steps. This is achieved by tracing imaginary particles from the center of each cell to the next time step (Karch et al., 2018). Unfortunately, using higher order integration schemes for this, e.g., Runge–Kutta, would require interpolation in space and time, which is very likely to sample data from the gaseous phase and lead to erroneous results as pointed out by Karch et al. (2018). Therefore, we are using a forward Euler step:**

$$\mathbf{x}(t + \Delta t) := \mathbf{x}_i(t) + \Delta t \mathbf{u}_i(t) . \quad (1)$$

An edge is added if $\mathbf{x}(t + \Delta t)$ belongs to a droplet, i.e., $\mathbf{x}(t + \Delta t)$ is located at time $t + \Delta t$ in a cell j with a valid droplet label. We also do a backward Euler step from the center \mathbf{u}_k of each cell k that is part of a droplet at time $t + \Delta t$:

$$\mathbf{x}(t) := \mathbf{x}_k(t + \Delta t) - \Delta t \mathbf{u}_k(t + \Delta t) . \quad (2)$$

If $\mathbf{x}(t)$ belongs to a droplet at time t , i.e., if $\mathbf{x}(t)$ is located in a cell l with a valid droplet label at t , we add the respective edge (if not already present). If a node n_j has $d > 2$ at time t , and only one connected neighbor at time $t - \Delta t$, a breakup event has happened at node n_j at time t . If there is more than one connected neighbor at time $t - \Delta t$, coalescence is involved. Our analysis particularly focuses on the dynamics of single droplets (not considering splits and merges; e.g., cf. Karch et al. (2018) for an analysis of these). Accordingly, we split the graph at nodes where more than two edges meet—i.e. coalescence and breakup events—, and base our approach on the isolated linear subgraphs, each of which represents the *trace* of a single droplet over time. A trace contains droplet quantities as described above, and serves as input to our regression model (Section 3.4).

We did not introduce valency d , is this something well known?

3.3 Clustering

We want to cluster the droplets to reveal their different types of behavior in the simulation. With our distribution of droplet quantities, no meaningful results could be obtained with density-based clustering algorithms such as DBSCAN (Ester et al., 1996) (yielding just one or two major clusters and a large number of outliers). In contrast, hierarchical clustering according to Ward’s method (Johnson, 1967)

has generated more expressive clusters in our experiments. However, hierarchical clustering exhibits quadratic memory complexity. The significant part is a distance matrix requiring $(n^2 - n)/2$ values to store. Handling our 1 000 000 droplet instances ($\approx 10\,000$ droplets over 100 time steps) would require at least 3.6 TB of memory and is therefore not possible on our machine. Instead, we aim to significantly reduce the number of data points. To achieve this, we limit ourselves for this modality to just one aggregate of each trace. Albeit at the cost of omitting temporal information, this is still able to serve the original purpose of identifying different droplet types (cf. Section 5). Note that our learning-based method presented below (Section 3.4) does not require this pre-aggregation.

3.4 Learning-Based Droplet Anomaly Detection

In our droplet anomaly approach, we first train a basic regression model on traces, using the surrogate task of predicting future values from a preceding time window. Assuming that the model does not overfit the training data (which we verify with a hold-out validation), it captures the most common and predictable patterns in droplet behavior. Then we can quickly check droplets against the model to find ones that deviate from the typical behavior. In this work, we employ ANNs for their ability to handle a large number of elements and learn a useful data representation (Bengio et al., 2013).

As surrogate task for training, we define a fixed-sized input window, and slide it along the trace, applying the model at every window position to predict the next value in that trace. That is, given a trace of length n_t , and a fixed window size w , we obtain $n_t - w + 1$ subtraces, and for each subtrace, the model takes the first $w - 1$ time steps as input and predicts the droplet quantities at the last (w -th) time step in the subtrace. This time-window approach has the advantages of allowing for comparably simple models to be used (lowering costs), and especially reducing the risk of overfitting, as shorter subtraces are less unique and harder to “memorize”. We train a separate prediction model for each quantity to affirm that each quantity is given the same importance, avoiding compromises in accuracy as would be the case with multiple output variables. We split 20% of the data into a held-out validation dataset, using the rest for training. Of course, we want the training set so include as much data as possible, but the validation set needs to have a reasonable size. The main purpose is to avoid overfitting, but one detail within this step is, that the validation set needs to be a reasonable sample of the whole dataset. Due to the nature of the dataset of containing a high amount of relatively similar droplets, and a low amount of relatively widely spread outliers, a validation size smaller than 20% may introduce sampling bias. Beforehand, we normalize each quantity to have zero mean and a standard deviation of one, using the mean and standard deviation estimated on the training set. Each model is a fully-connected neural network with two hidden layers of 64 neurons using Rectified Linear Unit (ReLU) (Glorot et al., 2011) activation and single linear unit in the output layer. The models are trained using the Adam optimizer (Kingma and Ba, 2015), with a learning rate 10^{-5} , and batch size of 32 over the course of 100 epochs until convergence of the validation MSE loss (Figure 3). Our setup was chosen empirically, by relying mostly on common values for most ML parameters, as our

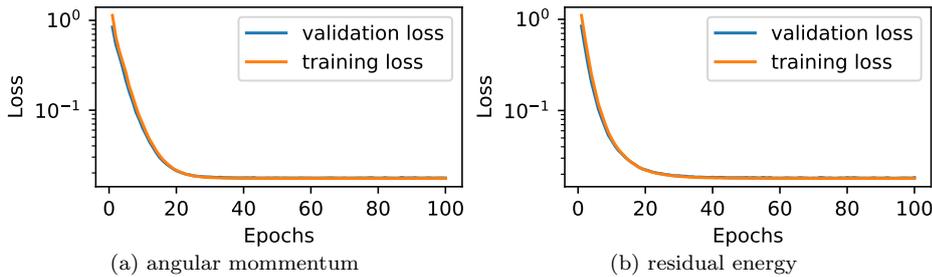


Fig. 3: Example of training and validation loss for the ANNs of angular momentum and residual energy, the curves for the remaining networks have a very similar shape.

focus is more on the overall framework. We think there might be potential for further optimization in future work, starting with a parameter study.

Once the models have been trained, we use them to perform prediction on all available substraces. Note that by design the first $w - 1$ time steps of a trace cannot be predicted, and have to be omitted in the following. We then simply determine the difference between the actual and the predicted value of each quantity. In total, this results in a 11-dimensional vector of deviations, whose Euclidean norm finally yields total error $\bar{\epsilon}$ —our measure of *estimated droplet anomaly*.

4 Visual Analytics System

Our visual analysis system has three different views: the abstract quantity relationship view, the droplet surface view, and the flow view. These components and their interplay enables the detailed analysis of a large droplet collection by supporting various kinds of investigation. The droplet surface view displays a set of droplets in the 3D domain, which allows interactive exploration within spacial and temporal context for selected droplet instances. The quantity relationship view provides an overview on extracted droplet characteristics and supports analyzing the interdependencies between them. Finally, the flow view can be used for the detailed investigation of a single selected droplet using full raw input data. Different views are linked to efficiently supplement each other. An overview of the system can be seen in Figure 4, and the different views are described in detail below.

Droplet Surface View. The 3D droplet surface view allows a user to spatially explore the droplet dataset (Figure 4a (iv)). Droplets of interest can be selected directly via picking (indicated with a crosshair). As showing all droplets would lead to significant occlusion and visual clutter, filtering the data is crucial (Figure 4a (i)). We support filtering w.r.t. arbitrary physical or geometrical quantities, clusters, and prediction errors (and combinations thereof). Color mapping can flexibly depict chosen quantities, with the color-coded anomaly measure $\bar{\epsilon}$ being the default choice (Figure 4a (ii)). A user can navigate within the full trace of a selected droplet and explore its temporal evolution (Figure 4a (v)). Below, similar traces are shown that were identified via feature vector distance, i.e., by Euclidean distance between droplet instances in the 11-dimensional quantity space. This helps, on the one hand, to assess the uniqueness of a droplet evolution, and on the other hand

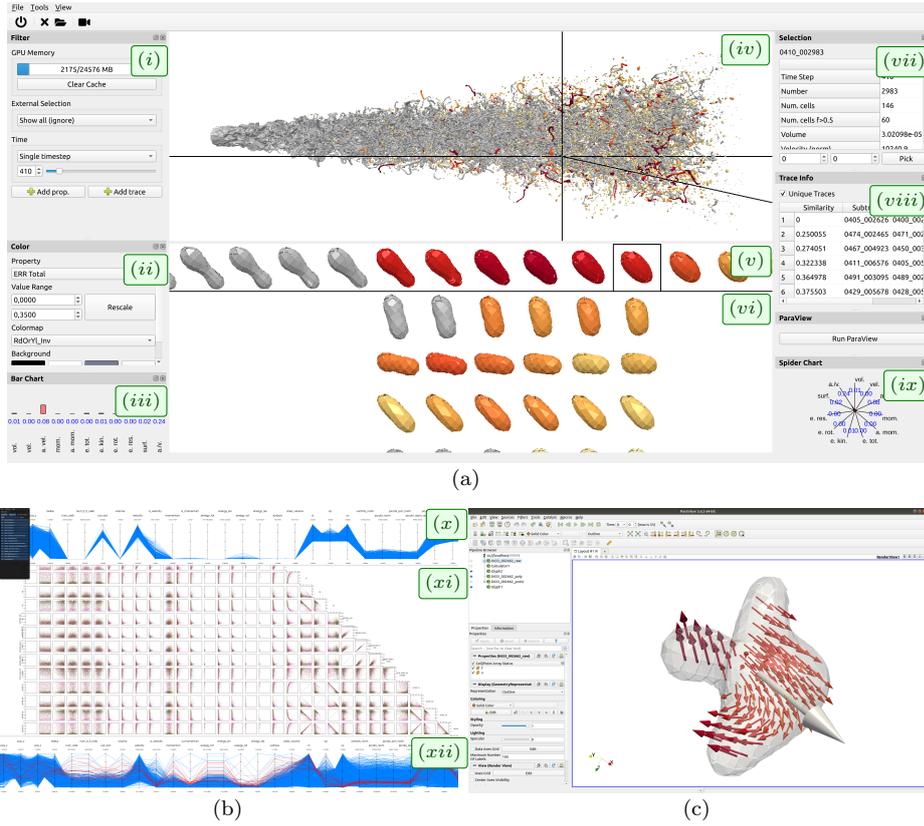


Fig. 4: System overview: (a) 3D surface view for showing droplet surfaces, consisting of filter panel (i), coloring tool (ii), bar chart (iii), main 3D view (iv) (here color mapped to anomaly measure $\bar{\epsilon}$), time view (v) for selected droplet (crosshair in (iv)) with temporal scrolling. The similarity search (vi) provides cases similar to selection based on feature vector distance. In addition, we depict droplet instance information (vii), droplet trace similarities (viii) for (vi), and spider chart as a complement for (iii). (b) Quantity relation view: Input parallel coordinates plot for data filtering (x), scatterplot matrix (xi) and second parallel coordinates plot (xii) allows analysis of data, in addition any quantity can be mapped to color. (c) A ParaView instance is integrated within the system for analysis of raw flow field for a single selected droplet and advanced flow feature extraction. A droplet selected in the 3D surface view can be automatically loaded into ParaView including a useful default filter pipeline as shown within the figure.

the comparison to similar droplets can help to gain further insights into which commonalities or differences have led to certain behavior patterns (Figure 4a (vi)). Additionally, we also provide the values for the selected droplet in our time view with a bar chart (Figure 4a (iii)) and a spider chart (Figure 4a (ix)), which can present its prediction error with respect to different physical quantities. Detailed droplet instance information (Figure 4a (vii)) and droplet trace similarities (Figure 4a (viii)) are provided as well.

Quantity Relationship Views. The quantity relationship view is focused on displaying the different (abstract) droplet quantities in context with each other. Here, we use classical information visualization methods, namely parallel coordinates plots (PCP) and a scatterplot matrix (SPLOM) to present an overview of the physical quantities in the dataset as well as all derived quantities (cluster, prediction error, etc.). In detail, this view consists of three different components as shown in Figure 4b. On top, we see the quantities of all droplets, shown within an interactive PCP, where sliders on each axis can be used to filter the data. Positioned below are a scatterplot matrix and a second PCP, in which we can analyze the filtered data. We provide both SPLOM and PCP to make use of the strengths of both visualization techniques. The PCP is ideal to get an overview of the data and locate single data points with the overall value range context, while the SPLOM can show pairwise relations, and is well-suited to identify patterns and relations within the data. Naturally, highlighting data in one view also will highlight the data in the other view. They are also linked to the droplet surface view (Figure 4a), i.e., brushing within the SPLOM or PCP can be used to filter in the 3D surface view. This can be useful to obtain spatial context regarding location and surface shape to the abstract quantity data points. This component is implemented on top of MegaMol (Gralka et al., 2019), employing OpenGL to render millions of points and lines at highly interactive framerates.

Flow View. For an in-depth exploration of the underlying flow field, e.g., to analyze the reason of a high anomaly, we further incorporate various classic flow visualization techniques by directly integrating ParaView (Ayachit, 2015) into our system (Figure 4c). ParaView is controlled from our application by loading the droplet data of the currently selected droplet within our 3D droplet surface view and automatically setting up the ParaView visualization pipeline. Not only the droplet itself is exported, but its full trace (in a droplet-local coordinate system for convenience), allowing animation. The precomputed surface and all other quantities are provided alongside. As frontend, the user has the classical fully functional ParaView interface.

5 Results

We will now analyze the Jet Simulation (cf. Section 1) in more detail to gain some insights and demonstrate our methods and system in practice. From the dataset we extracted 1 043 168 droplet instances. 273 928 were omitted due to insufficient size (around two third of them are artifacts at the simulation boundary), some more due to not being part of a trace with minimum length six. This results in 575 833 droplet instances which we use in our analysis. They are organized within 23 738 unique traces of sufficient length. These traces were split into 457 143 subtraces for training and validation data.

We used a machine featuring an Intel Core i7-8700K, 64 GB RAM and an NVIDIA RTX Titan. The data needed to be stored on a HDD due to its size; data I/O is the bottleneck for many of the tasks. The computation took ≈ 11 hours for segmentation (Section 3.1, CPU) and ≈ 25 hours for tracking (Section 3.2, CPU). Computation of the droplet quantities took ≈ 9.5 hours (CPU). Trace generation and additional data handling completed in ≈ 2 hours. The machine learning (Section 3.4,

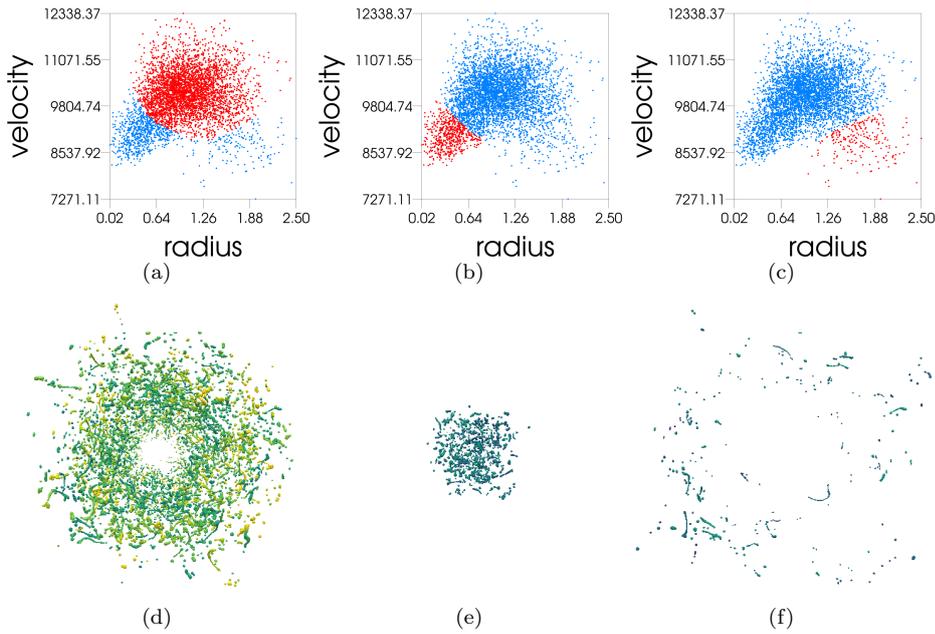


Fig. 5: Analysis of the distance between droplet and jet center (here called radius) to velocity relationship. The top row shows the scatterplot with different selected data points (red). The linked 3D view below provides spatial context for selection.

GPU) took ≈ 4 hours for the eleven ANNs, while clustering was finished in only ≈ 5 minutes (CPU). Our quantity view achieves more than 60 frames per second for a single time step, and drops to around 5 frames per second when looking the whole data at once (GPU).

5.1 Quantity Relationships

We first consider a single time step with the quantity analysis view (Figure 4b). As the jet we are looking at is fully converged, all time steps are quite similar in terms of general structure. The first thing we investigate is the filter PCP at the top to get an overview of the different value ranges. One thing we notice immediately is the huge value range of the droplet quantities due to single outliers, which leads to the majority of droplets being squashed together for some of the quantities, for example, this becomes quite apparent at the volume axis. We use the filter markers on the axis to remove the single droplet with very high volume, the jet base, and also omit droplets with an error flag for being too small—our expert specified a threshold of ≈ 15 cells can show unphysical behavior—or not being part of a trace. With the filtered data also being rescaled on each axis, and we can now see details and structures, which are lost in the first overview PCP (and the SPLOM below).

The first noticeable correlation we see is within the distance between droplet and jet center to velocity scatterplot (Figure 5). In particular, we identified two main clusters which mostly are symmetric to the diagonal. One cluster with high distance and high velocity (Figure 5a) and a second cluster with small distance and

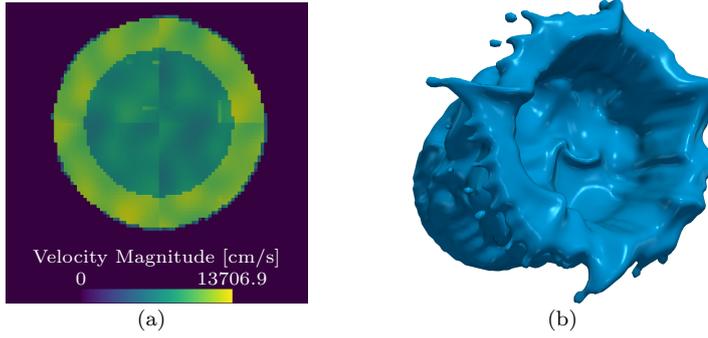


Fig. 6: (a) Inflow velocity magnitude profile at jet nozzle depicting the two-nozzle setup. (b) Isosurface of the jet base in the first simulation time step, reflecting the influence of the two-nozzle setup on the jet.

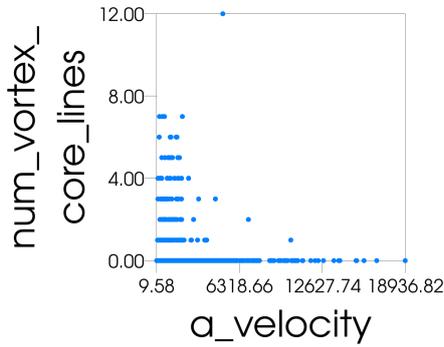


Fig. 7: Comparing angular velocity to the number of vortex coreline segments within a droplet. These two quantities behave inversely, i.e., higher velocity yields a lower number of coreline segments.

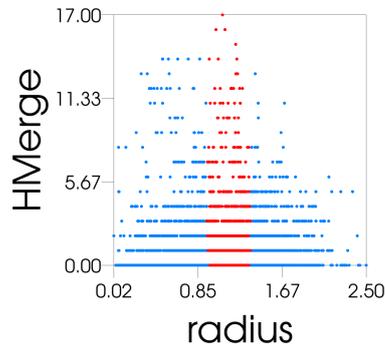


Fig. 8: Comparing radius of the droplet position around the jet center to the number of collision events in the history of a single droplet. Highlighted (red) is the radius range with a peak in the number of collision events.

small velocity (Figure 5b). Furthermore, there are a few more outliers (Figure 5c). With the first two main clusters, we can come to the conclusion, that the droplets in the center of the jet are slower than the droplets away from the center. We compare the scatterplots with the corresponding droplets in the 3D surface view, to provide spatial context (Figure 5d–Figure 5f). In addition the 3D droplets are colored by velocity using the viridis colormap. The reason for this two clusters is that our dataset has a two-nozzle setup where we have an inner and outer zone with different velocity at the boundary of the simulation domain. This can be easily seen by looking at the velocity magnitudes and the surface of the jet base in the very first time step of the simulation (Figure 6). However, it has previously been unclear how exactly this setups affects droplet velocities further within the simulation domain.

Next, we look at the correlation between angular velocity and the number of vortex coreline segments (Figure 7). Originally, our hypothesis was that a high angular velocity in the sense of rigid body rotation could be seen as a vortex,

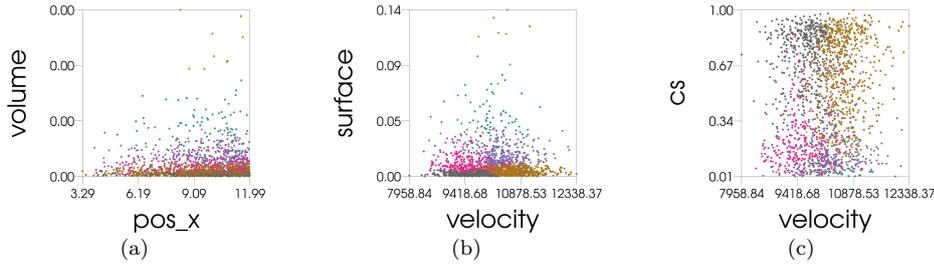


Fig. 9: Data points colored by cluster id from the 8-class hierarchical ward clustering. (a) Position in x -direction (flow direction of the jet) compared to volume. Notice the relation of the clusters for higher volumes. For lower volumes, clusters are not distinguishable. (b) Comparing velocity to surface, we can see separation of clusters in both dimensions. (c) Same holds for comparing velocity to spherical anisotropy (anisotropy is not considered during clustering).

but in fact the plot shows exactly the opposite. We see that only droplets with a relatively small angular velocity have a high number of vortex core line segments, while droplets which have at least one vortex core line segment have a relatively low angular velocity. We consider this to be an interesting finding, but were not yet able to confirm a physical explanation for this phenomenon.

The scatterplot of the distance between droplet and jet center compared to the number of merge events in the history of the jet also looks interesting (Figure 8). It exhibits a triangular shape, except for a few outliers. The most inner and outer droplets seem to have only few merge events, while there is a ring of droplets with average distance, that appears to have a very high number of such occurrences. We attribute this to the dual nozzle injection, where slower droplets in the inner and faster droplets in the outer ring collide in a transition area. The peak in the number of merge events accordingly indicates this contact area.

5.2 Clustering

The clustering results depict different characteristics (cf. Figure 9) We notice that the volume is one of the main factors which separates these clusters. It further shows in the middle view that the velocity is a factor orthogonal to the volume. In Figure 9c, we can see the cluster shape also in reference to the spherical anisotropy. This is especially remarkable because this value was not used for the clustering. Accordingly, we assume that this is an indicator for the correlation of quantities. **A limitation may be that there is no direct physical interpretation of these clusters. While it would be a great result if we had found one, our clustering method is based on the trace average of the extracted quantities. This is a very simplified projection of a droplet, probably not covering all physical laws and may not be fully based on intrinsic physical quantities. Therefore the clusters could only have a phenomenological interpretation by looking at the quantity distribution within the SPLOM view as sketched by Figure 9. Nevertheless, we think these are still structural and dynamical relevant clusters representing groups of similar droplet behavior within the abstract quantity space.**

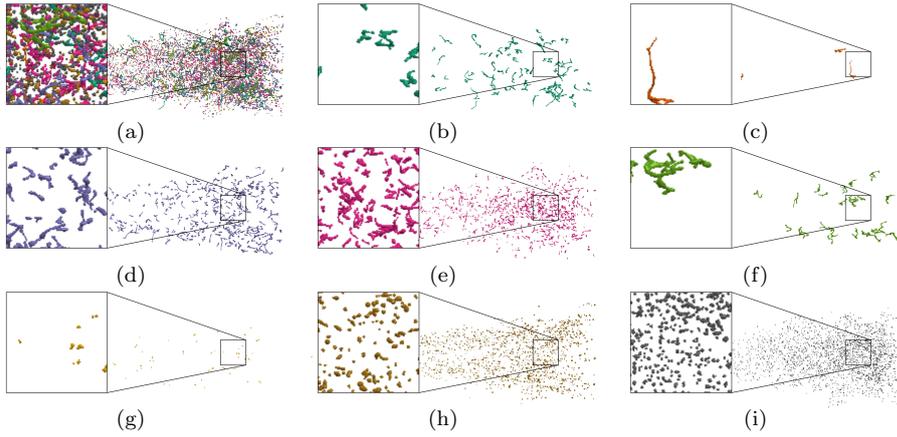


Fig. 10: Same clusters as in Figure 9, but within the 3D surface view. (a) All clusters in one view. (b)–(i) Each cluster on its own. Note the similarity of droplet surface shape and size within clusters.

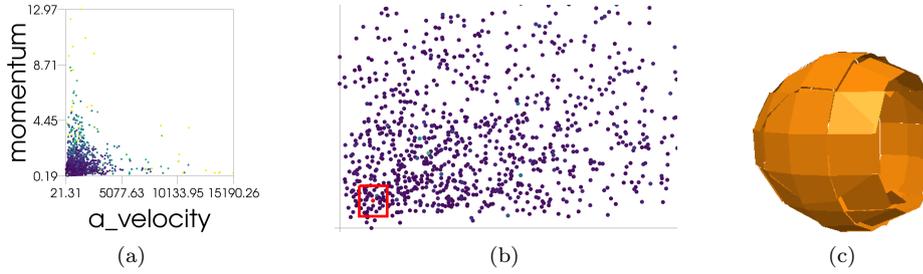


Fig. 11: (a) The comparison of angular velocity to momentum shows a trend of data points aligning along axis. Coloring by prediction anomaly seems to correlate with outliers of this trend. (b) Zoom to bottom left of (a). Selected (red) is the data point with the highest anomaly in this range. (c) 3D surface view of the selected droplet in (b). We can see this droplet is cut by the domain boundary. This implies to be the source for this droplet being an anomaly.

Looking at the 3D surface view, we find that each cluster has similar droplets in reference to its shape and size (Figure 10). Overall this clustering gives us an overview of the different types of droplets within this dataset. It provides additional information to statistical quantity value distribution, for example within the PCP view, by addressing higher order connection in the high dimensional space of the quantity values. However, remember that we neglect temporal effects as we averaged traces prior to clustering (cf. Section 3.3).

5.3 Droplet Anomaly-Guided Analysis

Finally, we analyze the results based on prediction anomaly. In Figure 11a, we see correlation between the angular velocity and the momentum, with most data points being close to the axis. That means a droplet generally exhibit high velocity

or high momentum, but not both. By coloring the data points by the error of the trace prediction, we see that the further droplets are away from this trend along the axis, the higher is the prediction error. Bottom left, we see a dense region with mostly low anomaly droplets. Zooming reveals a few droplets with high anomaly, which attract our attention (Figure 11b). Using brushing on the highest anomaly data point in this region (marked red) and observing it in the linked 3D surface and flow view, we see that this droplet is located at the boundary of the domain. The identification of such problematic cases is important for studying edge cases in the simulation and to discard them from further consideration.

Figure 4a shows the 3D surface view depicting the full time step 410. Color depicts our anomaly estimation $\bar{\varepsilon}$, with higher anomaly indicated by reddish colors and lower anomaly by yellowish ones. Gray indicates droplet instances that are within the $w - 1$ first time steps of a trace, and thus do not provide prediction nor anomaly information $\bar{\varepsilon}$ (this always applies to the jet itself, it is always at the beginning of the main trace due to droplets splitting from it continuously). To reduce occlusion and clutter, a typical first step is to omit those structures, which can be accomplished by requiring $\bar{\varepsilon} \geq 0$, because $\bar{\varepsilon}$ is set to a negative constant in our implementation, if no $\bar{\varepsilon}$ value can be computed (cf. Figure 12a). Here, we observe that ligaments, i.e., the elongated droplets that typically breakup into smaller droplets later on, exhibit very high $\bar{\varepsilon}$, i.e., they show temporal behavior that strongly deviates from the trends of the majority of the droplets. We also note that the structures with the highest $\bar{\varepsilon}$ are typically tiny droplets (Figure 13), which led us to the hypothesis that the quantities that we computed suffered from discretization artifacts for small droplets, even if they are still bigger than the required minimum size originally indicated by our domain scientist. In particular, such droplets tend to exhibit alternating $\bar{\varepsilon}$, switching between high and low values at high temporal frequency, which supports our hypothesis of discretization issues due to insufficient resolution. In fact, we also observed such alternating behavior of $\bar{\varepsilon}$ for very thin ligaments. As a result, we do not consider very small droplets nor ligaments with this analysis component in the following. To accomplish this, we filter by spherical anisotropy c_s , requiring $c_s > 0.4$ in this case. Additionally, we suppress all droplets with low anomaly by adjusting our filter to $\bar{\varepsilon} \geq 0.15$. To obtain a better space-time overview, we now enable the simultaneous display of all time steps with the chosen filtering criteria, leading to Figure 12b. This way, we observe many long traces in the 3D surface view, which provides the point of origin for our further investigations.

We now pick a trace with particularly high anomaly estimation (selection indicated by black cross in Figure 12a). A closer inspection of that trace in the time-view (Figure 4a (v)) reveals that the droplet becomes rounder over time, and that the anomaly indicator $\bar{\varepsilon}$ stays almost constant over time. This is an observation that turned out to be quite rare, since typically the anomaly measure reduces quite quickly, in particular if the respective droplet becomes roundish. A thorough investigation of the respective plots of the original traces, as well as the individual deviations of the predicted quantities from the original ones, did not provide insights on the causes of this behavior. We started to hypothesize that the internal flow within the droplet might provide insights. We thus initiated flow visualization of the liquid phase of the droplet (Figure 14) (left). Please notice that we use linear mapping of velocity magnitude to glyph color, but logarithmic mapping of velocity magnitude to glyph size [within Figure 14 and following](#). Interestingly, we observed

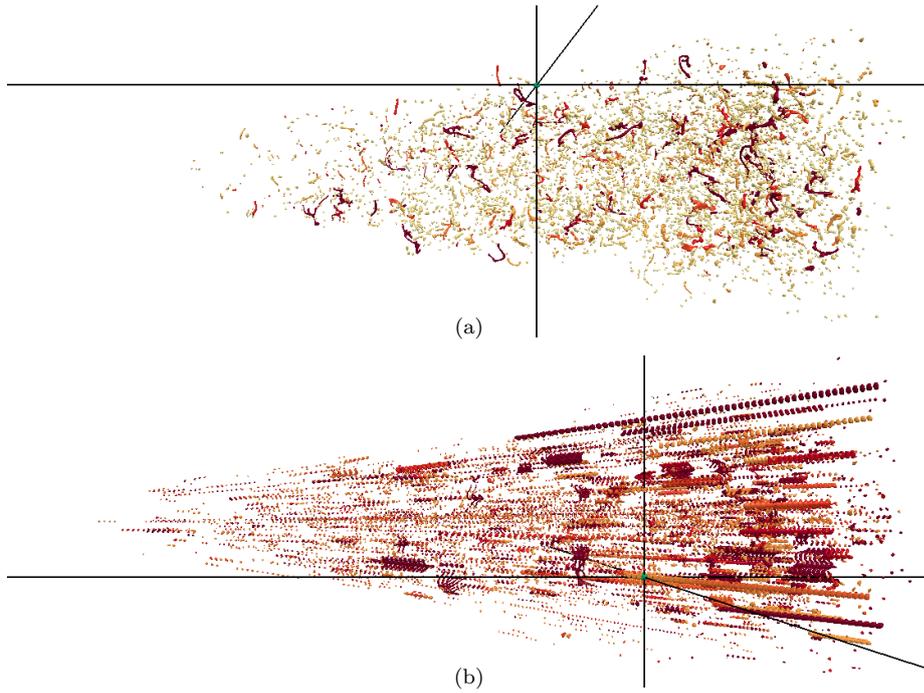


Fig. 12: Exploration of the dataset within the 3D surface view (c.f. Figure 4a (iv)). After loading the dataset the time filter is set to show only a single time step and a second filter omits all droplet instances that lack total error (our anomaly indicator) $\bar{\epsilon}$, leading to the view in (a). The user now selects a droplet of interest for further analysis (black cross), here the droplet shown in Figure 14 (left) is selected. (b) provides a further example. Now the time filter is set to aggregate all timesteps of the dataset, leading to dotted line like structures showing the same droplet over time. To reduce visual clutter further filters are used to only show droplets with anomaly measure $\bar{\epsilon} > 0.15$ and for suppressing ligaments by requiring spherical anisotropy $c_s > 0.4$. Here the droplet in Figure 15 is selected.

a distinguished and strong saddle-type flow pattern in the internal flow, in the frame of reference moving with the velocity of the center of mass of the droplet. We investigated other cases with high anomaly measure $\bar{\epsilon}$, either by direct selection in the 3D view, or by means of our similarity search, as provided in the lower half of our tool. Interestingly, most cases of non-ligament (more or less roundish) droplets with high $\bar{\epsilon}$ turned out to exhibit such saddle-type flow patterns in their interior, including the case shown in Figure 14 (right), found as most similar to that droplet by similarity search.

This motivated us to investigate droplets with different anomaly behavior. While decaying $\bar{\epsilon}$ is quite common, we wanted to have a look at droplets for which $\bar{\epsilon}$ increases over time. The hypothesis behind this reasoning is that anomalous droplet behavior is often caused by a collision of droplets, but while the products of the collision move over time, they tend to “calm down” and become better predictable by our regression approach, and thus $\bar{\epsilon}$ decays. In Figure 12b, we were able to identify such a droplet, selecting it by picking, and investigate its temporal behavior

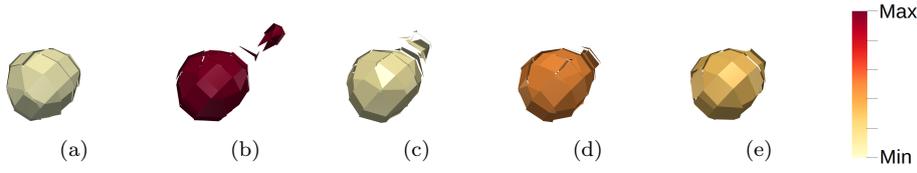


Fig. 13: Droplet exhibiting the highest anomaly. Structures with highest anomaly measure $\bar{\varepsilon}$ are tiny droplets whose quantity computation apparently suffers from discretization issues. In (b) the collision with a smaller droplet artifact can be observed, which was missed by the trace generation (c.f. Section 3.2). Further due to the nature of PLIC, the strongly discontinuous surface reconstruction can be seen in the form of gaps (c.f. Section 3.1). Successive time steps (a)–(e) exhibit alternating behavior of $\bar{\varepsilon}$, supporting the hypothesis of insufficient resolution of the simulation grid.

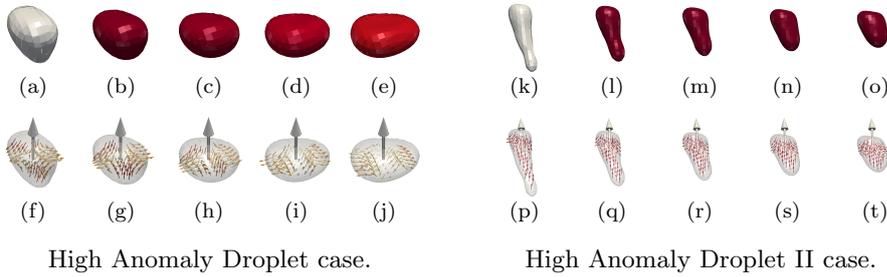


Fig. 14: **Left:** Selected droplet from Figure 12a. We found that droplets with low temporal decay of $\bar{\varepsilon}$ exhibit distinguished saddle-type flow patterns in their interior flow, in frame of reference moving with the center of mass of the droplet. **Right:** This is the droplet found most similar by similarity search in Figure 4a.

in the time view. Detailed flow analysis is shown in Figure 15. Interestingly, this droplet has been the only one we could find with a strong vortex in its interior.

Finally, for comparison and validation of the utility of our anomaly estimation, we investigated droplets with low anomaly measure $\bar{\varepsilon}$. We selected them either by direct picking, or by similarity search. The majority of droplets with low $\bar{\varepsilon}$ exhibit shear flow in its interior, as shown in Figure 16 (left). We were able to find only one case with deviating flow pattern. This droplet (Figure 16) (right) exhibits a “half” saddle-type flow pattern (imagine splitting a 3D saddle along its 2D manifold) but with low flow velocity in the frame of reference moving with the observer.

6 Discussion and Future Work

We developed a visual analysis approach for investigating droplet characteristics and behaviors occurring in large ensembles like sprays. It allows the exploration of what originally are terabytes of simulation data depicting complex small-scale effects with a multi-stage workflow. On the one hand, we reduced the data by extracting droplet surfaces and physical quantities. This allowed to give an overview and investigate large numbers of droplets at once during interactive exploration. The physical

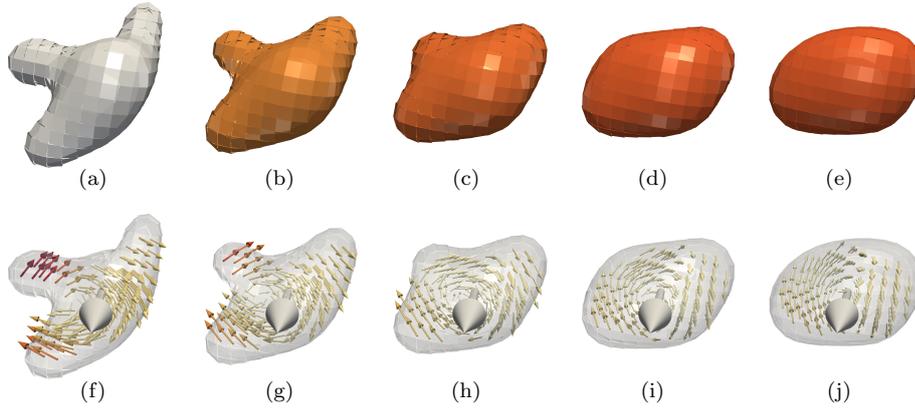


Fig. 15: Increasing Anomaly Droplet case. The very rare case of temporally increasing $\bar{\epsilon}$ brought our attention to this case, which, in frame of reference moving with the center of mass of the droplet, exhibits a strong vortex in its interior flow. This droplet has been the only one we could identify to contain a vortex.

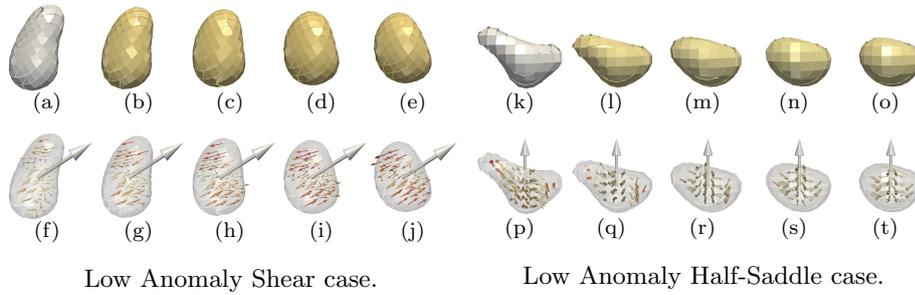


Fig. 16: **Left:** Droplets with low anomaly measure $\bar{\epsilon}$ generally exhibit shear flow patterns in their interior, in frame of reference moving with the center of mass of the droplet. **Right:** Only identified exception to our hypothesis, that low anomaly droplets exhibit shear flow patterns in the interior. This droplet exhibits a weak “half-saddle” pattern, i.e., half of a 3D saddle, similar to a detachment point in its interior, in frame of reference moving with the droplet.

quantities were further used in advanced analysis steps to (1) get an overview on different kinds of droplets via clustering (where we introduced a domain-specific optimization to be able to handle a large number of droplets), and (2) identify droplets with irregular temporal behavior by adapting a recent machine learning-based method. On the other hand, the full original data was still available for the detailed investigation of selected cases. We integrated different analysis components to be able to provide the user with the full breadth of required interactive analysis functionality from investigating relationships of extracted physical quantities, droplet traces and results from clustering and ML-based analysis to a spatial overview, and finally a direct link to ParaView supporting advanced flow analysis for the detailed investigation of selected droplets.

So far, we applied our tool to gain insights from a single large-scale direct numerical simulation datasets of two-phase flow of non-Newtonian jets. However, we believe our approach directly generalizes to different simulations of this kind. We have chosen basic physical quantities which are based on raw simulation data from the VOF method (and accordingly will be available from nearly all VOF based simulations). Depending on the simulation results, the list of properties could be extended, e.g., if pressure is available, or a higher resolution allows for precise calculation of deviation-based quantities.

We would generally expect the ANN-based anomaly detection to work likewise for other data. However, as our prediction is based on traces, we need a certain temporal resolution, especially some highly chaotic datasets might be problematic, where splits and merges happen in nearly every timestep for each droplet. That being said, we generally consider jets to be already among the more difficult real-world examples. However, the trace aggregation approach we use to enable the use of hierarchical clustering might not adequately transfer to different examples. While the representation of traces by means of aggregates proved useful in our analysis scenario, more elaborate, adaptive approaches are potentially required in other contexts.

In future work, we aim to study additional variants of jets, and beyond this we will explore possibilities to generalize our technique for types of simulation data in which large amounts of small entities need to be investigated. We plan to further incorporate split and merge events to better understand the atomization process. While we already collect the number of these events along the full history of a droplet, we additionally aim to integrate a graph view in our tool (enriched with our extracted quantities) and to extend our ML-based anomaly detection to explicitly consider split and merge events. Finally, a main goal would be to allow a user to directly influence our clustering and machine learning component during interactive analysis with respect to specified subsets or partitions of the data.

Acknowledgements We acknowledge the continuous support of our colleagues and domain scientists from the Institute of Aerospace Thermodynamics (ITLR) of the University of Stuttgart and the collaborative research center on droplet dynamics under extreme conditions. We kindly acknowledge the financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project SFB-TRR 75, Project number 84292822.

References

- Aigner W, Bertone A, Miksch S, Tominski C, Schumann H (2007) Towards a conceptual framework for visual analytics of time and time-oriented data. In: Proceedings of 2007 Winter Simulation Conference, pp 721–729
- Ayachit U (2015) The ParaView Guide: A Parallel Visualization Application. Kitware, Inc., Clifton Park, NY, USA
- Bai Z, Brunton SL, Brunton BW, Kutz JN, Kaiser E, Spohn A, Noack BR (2017) Data-driven methods in fluid dynamics: Sparse classification from experimental data. Whither Turbulence and Big Data in the 21st Century? pp 323–342
- Basser PJ, Pierpaoli C (2011) Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance* 213(2):560 – 570, URL <http://www.sciencedirect.com/science/article/pii/S109078071100334X>

- Bengio Y, Courville A, Vincent P (2013) Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(8):1798–1828
- Biswas A, Ahrens J, Dutta S, Musser J, Almgren A, Turton TL (2020) Feature analysis, tracking, and data reduction: An application to multiphase reactor simulation mfix-exa for in-situ use case. *Computing in Science & Engineering* DOI 10.1109/MCSE.2020.3016927
- Buerger R, Muigg P, Ilčík M, Doleisch H, Hauser H (2007) Integrating Local Feature Detectors in the Interactive Visual Analysis of Flow Simulation Data. In: Museth K, Moeller T, Ynnerman A (eds) *Eurographics/ IEEE-VGTC Symposium on Visualization*, The Eurographics Association, DOI 10.2312/VisSym/EuroVis07/171-178
- Doleisch H, Gasser M, Hauser H (2003) Interactive Feature Specification for Focus+Context Visualization of Complex Simulation Data. In: Bonneau GP, Hahmann S, Hansen CD (eds) *Eurographics / IEEE VGTC Symposium on Visualization*, The Eurographics Association, DOI 10.2312/VisSym/VisSym03/239-248
- Doleisch H, Mayer M, Gasser M, Wanker R, Hauser H (2004) Case Study: Visual Analysis of Complex, Time-Dependent Simulation Results of a Diesel Exhaust System. In: Deussen O, Hansen C, Keim D, Saupe D (eds) *Eurographics / IEEE VGTC Symposium on Visualization*, The Eurographics Association, DOI 10.2312/VisSym/VisSym04/091-096
- Dutta S, Shen H (2016) Distribution driven extraction and tracking of features for time-varying data analysis. *IEEE Transactions on Visualization and Computer Graphics* 22(1):837–846, DOI 10.1109/TVCG.2015.2467436
- Eisenschmidt K, Ertl M, Goma H, Kieffer-Roth C, Meister C, Rauschenberger P, Reitzle M, Schlottke K, Weigand B (2016) Direct numerical simulations for multiphase flows: An overview of the multiphase code FS3D. *Applied Mathematics and Computation* 272:508 – 517, URL <http://www.sciencedirect.com/science/article/pii/S0096300315007195>, recent *Advances in Numerical Methods for Hyperbolic Partial Differential Equations*
- Endert A, Ribarsky W, Turkey C, Wong BLW, Nabney I, Blanco ID, Rossi F (2017) The state of the art in integrating machine learning into visual analytics. *Computer Graphics Forum* 36(8):458–486
- Ertl M (2019) Direct numerical investigations of non-newtonian drop oscillations and jet breakup. PhD thesis, University of Stuttgart
- Ertl M, Weigand B (2017) Analysis methods for direct numerical simulations of primary breakup of shear-thinning liquid jets. *Atomization and Sprays* 27(4):303–317
- Ester M, Kriegel HP, Sander J, Xu X (1996) A density-based algorithm for discovering clusters in large spatial databases with noise. AAAI Press, pp 226–231
- Garth C, Tricoche X, Scheuermann G (2004) Tracking of vector field singularities in unstructured 3d time-dependent datasets. In: *IEEE Visualization 2004*, pp 329–336, DOI 10.1109/VISUAL.2004.107
- Glorot X, Bordes A, Bengio Y (2011) Deep sparse rectifier neural networks. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp 315–323, URL <http://proceedings.mlr.press/v15/glorot11a.html>

- Gralka P, Becher M, Braun M, Frieß F, Müller C, Rau T, Schatz K, Schulz C, Krone M, Reina G, Ertl T (2019) MegaMol – A Comprehensive Prototyping Framework for Visualizations. *The European Physical Journal Special Topics* 227(14):1817–1829, DOI 10.1140/epjst/e2019-800167-5, URL <https://doi.org/10.1140/epjst/e2019-800167-5>
- Heinemann M (2018) ML-based visual analysis of droplet behaviour in multiphase flow simulations. Master’s thesis, University of Stuttgart
- Helmut D, Philipp M, Helwig H (2004) Interactive visual analysis of hurricane Isabel with SimVis
- Hirt C, Nichols B (1981) Volume of fluid (VOF) method for the dynamics of free boundaries. *Journal of Computational Physics* 39(1):201 – 225, URL <http://www.sciencedirect.com/science/article/pii/0021999181901455>
- Johnson SC (1967) Hierarchical clustering schemes. *Psychometrika* 32(3):241–254, DOI 10.1007/bf02289588, URL <https://doi.org/10.1007/bf02289588>
- Karch GK, Sadlo F, Meister C, Rauschenberger P, Eisenschmidt K, Weigand B, Ertl T (2013) Visualization of piecewise linear interface calculation. In: *IEEE Pacific Visualization Symposium (PacificVis)*, pp 121–128
- Karch GK, Beck F, Ertl M, Meister C, Schulte K, Weigand B, Ertl T, Sadlo F (2018) Visual analysis of inclusion dynamics in two-phase flow. *IEEE Transactions on Visualization and Computer Graphics* 24(5):1841–1855
- Kingma DP, Ba J (2015) Adam: A method for stochastic optimization. In: *3rd International Conference on Learning Representations, ICLR, Conference Track Proceedings*, URL <http://arxiv.org/abs/1412.6980>, 1412.6980
- Post FH, Vrolijk B, Hauser H, Laramee RS, Doleisch H (2003) The state of the art in flow visualisation: Feature extraction and tracking. *Computer Graphics Forum* 22(4):775–792, DOI <https://doi.org/10.1111/j.1467-8659.2003.00723.x>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2003.00723.x>, <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2003.00723.x>
- Rosenberger G, Nestor PG, Oh JS, Levitt JJ, Kindleman G, Bouix S, Fitzsimmons J, Niznikiewicz M, Westin CF, Kikinis R, McCarley RW, Shenton ME, Kubicki M (2012) Anterior limb of the internal capsule in schizophrenia: a diffusion tensor tractography study. *Brain Imaging and Behavior* 6(3):417–425, URL <https://doi.org/10.1007/s11682-012-9152-9>
- Shi K, Theisel H, Hauser H, Weinkauff T, Matkovic K, Hege HC, Seidel H (2007) Path line attributes - an information visualization approach to analyzing the dynamic behavior of 3D time-dependent flow fields. In: *Topology-Based Methods in Visualization II*, pp 75–88
- Straub A, Ertl T (2020) Visualization techniques for droplet interfaces and multiphase flow. In: Lamanna G, Tonini S, Cossali GE, Weigand B (eds) *Droplet Interactions and Spray Processes*, Springer International Publishing, Cham, pp 203–214
- Straub A, Heinemann M, Ertl T (2019) Visualization and visual analysis for multiphase flow. In: *DIPSI 2019: Droplet Impact Phenomena & Spray Investigations*, Bergamo, Italy, 17th May 2019, Università degli studi di Bergamo, pp 25–27
- Sujudi D, Haimes R (1995) Identification of swirling flow in 3-d vector fields. In: *12th Computational Fluid Dynamics Conference*, p 1715
- Theisel H, Seidel HP (2003) Feature Flow Fields. In: Bonneau GP, Hahmann S, Hansen CD (eds) *Eurographics / IEEE VGTC Symposium on Visualization*,

- The Eurographics Association, DOI 10.2312/VisSym/VisSym03/141-148
- Tkachev G, Frey S, Ertl T (2019) Local Prediction Models for Spatiotemporal Volume Visualization. *IEEE Transactions on Visualization and Computer Graphics* pp 1–1, DOI 10.1109/TVCG.2019.2961893
- Tompson J, Schlachter K, Sprechmann P, Perlin K (2017) Accelerating Eulerian fluid simulation with convolutional networks. In: *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, pp 3424–3433
- Tzeng FY, Ma KL (2005) Intelligent feature extraction and tracking for visualizing large-scale 4D flow simulations. In: *Supercomputing, 2005. Proceedings of the ACM/IEEE SC 2005 Conference*, pp 6–6
- Youngs DL (1984) An interface tracking method for a 3D Eulerian hydrodynamics code. Atomic Weapons Research Establishment (AWRE) Technical Report 44(92):35

A Appendix

A.1 Definition of Physically Motivated Quantities

Droplet Volume. We compute the volume of each droplet by integrating the f -field. That is, the volume $\bar{\mu}_\iota$ of the droplet with identifier ι is obtained by integrating the product of the f -field and the cell volume for all cells i where $l_i = \iota$:

$$\bar{\mu}_\iota := \sum_{i \in \mathcal{C}_\iota} f_i \mu_i , \quad (3)$$

with $\mathcal{C}_\iota := \{j \mid l_j = \iota\}$, and μ_i being the volume of cell i .

Droplet Area to Volume Ratio. As motivated above, we compute the area \bar{A}_ι of droplet ι by integrating the area of the respective PLIC patches, i.e.,

$$\bar{A}_\iota := \sum_{i \in \mathcal{C}_\iota} A_i , \quad (4)$$

with A_i being the area of the PLIC patch in cell i , i.e., $A_i = 0$ if $f_i = 0$ or $f_i = 1$. From that, we compute the area to volume ratio $\bar{\alpha}_\iota$ Karch et al. (2018) of the droplet as

$$\bar{\alpha}_\iota := \frac{\bar{A}_\iota r_s}{\bar{\mu}_\iota \frac{3}{4}\pi} , \quad (5)$$

with $r_s := \sqrt[3]{(3\bar{\mu}_\iota)/(4\pi)}$ being the radius of a sphere with volume $\bar{\mu}_\iota$.

Droplet Velocity and Momentum. The velocity of the center of mass of droplet ι follows to be

$$\bar{\mathbf{u}}_\iota := \frac{1}{\bar{\mu}_\iota} \sum_{i \in \mathcal{C}_\iota} m_i \mathbf{u}_i , \quad (6)$$

with cell-based flow velocity \mathbf{u}_i in cell i , and $m_i := f_i \mu_i$. From that, the droplet's momentum computes

$$\bar{\mathbf{p}}_\iota := \bar{\mu}_\iota \bar{\mathbf{u}}_\iota . \quad (7)$$

Notice that we assume the density of the liquid to be 1, because the density is often not provided explicitly for the liquid phase in two-phase simulations (as in our case), and since density represents only a scaling factor (note that liquids are generally treated as being incompressible). As a consequence, the total mass of droplet ι equals its volume $\bar{\mu}_\iota$.

Auxiliary Measures. A common measure in astrophysics and particle systems is the center of mass. It provides a frame of reference and enables the computation of derived quantities. The center of mass $\bar{\mathbf{x}}_\iota$ of droplet ι computes

$$\bar{\mathbf{x}}_\iota := \frac{1}{\bar{\mu}_\iota} \sum_{i \in \mathcal{C}_\iota} m_i \mathbf{x}_i , \quad (8)$$

with \mathbf{x}_i being the center of cell i .

From that, it is a common step to compute the total angular momentum for a set of particles, resulting in our case in the angular momentum $\bar{\mathbf{L}}_\iota$ of the droplet:

$$\bar{\mathbf{L}}_\iota := \sum_{i \in \mathcal{C}_\iota} \hat{\mathbf{x}}_i \times m_i \mathbf{u}_i , \quad (9)$$

with $\hat{\mathbf{x}}_i := \mathbf{x}_i - \bar{\mathbf{x}}_\iota$ being the center of cell i relative to $\bar{\mathbf{x}}_\iota$. Notice that, for the example of particle systems, this total angular momentum describes the rotational motion of the ‘‘rigid-body aspect’’ of the particle set, or in other words, the rotational motion of the entirety of the particles. Thus, in our context, it represents the ‘‘rigid-body rotation component’’ of the droplet in terms of the flow of the liquid.

In analogy, the droplet's inertia tensor $\bar{\Theta}_\iota$ computes

$$\bar{\Theta}_\iota := \begin{pmatrix} \sum_i m_i (\hat{y}_i^2 + \hat{z}_i^2) & -\sum_i m_i \hat{y}_i \hat{x}_i & -\sum_i m_i \hat{z}_i \hat{x}_i \\ -\sum_i m_i \hat{x}_i \hat{y}_i & \sum_i m_i (\hat{x}_i^2 + \hat{z}_i^2) & -\sum_i m_i \hat{z}_i \hat{y}_i \\ -\sum_i m_i \hat{x}_i \hat{z}_i & -\sum_i m_i \hat{y}_i \hat{z}_i & \sum_i m_i (\hat{x}_i^2 + \hat{y}_i^2) \end{pmatrix}, \quad (10)$$

with $\hat{\mathbf{x}}_i =: (\hat{x}_i, \hat{y}_i, \hat{z}_i)^\top$ and $i \in \mathcal{C}_\iota$. From that, the angular velocity $\bar{\boldsymbol{\omega}}_\iota$ of droplet ι is obtained by

$$\bar{\boldsymbol{\omega}}_\iota = \bar{\Theta}_\iota^{-1} \bar{\mathbf{L}}_\iota. \quad (11)$$

Notice that $\bar{\Theta}_\iota$ can become singular for very small droplets, and this is one of the reasons why we exclude such very small droplets from our analysis.

Total Droplet Energy. The total energy E_ι of droplet ι (notice that the addressed simulations exclude thermodynamical and chemical processes) computes

$$E_\iota := \sum_{i \in \mathcal{C}_\iota} \frac{1}{2} m_i \mathbf{u}_i^2. \quad (12)$$

This total energy can be decomposed into kinetic $E_\iota^{\mathbf{u}}$, rotational $E_\iota^{\boldsymbol{\omega}}$, and residual energy E_ι^δ :

$$E_\iota = E_\iota^{\mathbf{u}} + E_\iota^{\boldsymbol{\omega}} + E_\iota^\delta, \quad (13)$$

as explained next.

Kinetic Droplet Energy. The kinetic energy $E_\iota^{\mathbf{u}}$ of droplet ι with respect to its velocity $\bar{\mathbf{u}}_\iota$ is

$$E_\iota^{\mathbf{u}} := \frac{1}{2} \bar{\mu}_\iota \bar{\mathbf{u}}_\iota^2. \quad (14)$$

This measure represents the kinetic energy of the rigid-body property of the liquid component representing a droplet.

Rotational Droplet Energy. The rotational energy $E_\iota^{\boldsymbol{\omega}}$ of droplet ι computes

$$E_\iota^{\boldsymbol{\omega}} := \frac{1}{2} \bar{\boldsymbol{\omega}}_\iota^\top \bar{\Theta}_\iota \bar{\boldsymbol{\omega}}_\iota. \quad (15)$$

This rotational energy captures the rigid-body part of the liquid dynamics of the droplet with respect to rotational motion.

Residual Droplet Energy. The residual energy E_ι^δ of droplet ι finally computes

$$E_\iota^\delta := E_\iota - E_\iota^{\mathbf{u}} - E_\iota^{\boldsymbol{\omega}}, \quad (16)$$

and includes deformation of the droplet, as well as non-rigid flow in its interior. In that sense, E_ι^δ captures the deviations of a droplet's flow from rigid-body dynamics. Due to limited numerical accuracy, E_ι^δ can turn out to be slightly negative, in which case we clamp it to zero and set $E_\iota^{\boldsymbol{\omega}} := E_\iota - E_\iota^{\mathbf{u}}$.