



**UNIVERSITY
OF LATVIA**

**Summary
of Doctoral Thesis**

Mihails Birjukovs

**INVESTIGATING BUBBLE-
AND PARTICLE-LADEN
LIQUID METAL FLOW AND
ALLOY SOLIDIFICATION USING
NOVEL IMAGE AND DATA
ANALYSIS METHODS**

Riga 2023



UNIVERSITY OF LATVIA

FACULTY OF PHYSICS, MATHEMATICS AND OPTOMETRY
DEPARTMENT OF PHYSICS

Mihails Birjukovs

INVESTIGATING BUBBLE- AND PARTICLE-LADEN LIQUID METAL FLOW AND ALLOY SOLIDIFICATION USING NOVEL IMAGE AND DATA ANALYSIS METHODS

SUMMARY OF THE DOCTORAL THESIS

Submitted as a requirement for the Doctoral Degree in Physics
Specialization: Mechanics of Fluids and Gases

Riga 2023

This doctoral thesis was completed over 2019 to 2023 at the Institute of Numerical Modelling (INM), Department of Physics, Faculty, Physics, Mathematics and Optometry, University of Latvia (UL) with support from Paul Scherrer Institut (PSI, Switzerland) and Helmholtz-Zentrum Dresden-Rossendorf (HZDR, Germany).

The thesis was funded via the ERDF project "Development of numerical modelling approaches to study complex multiphysical interactions in electromagnetic liquid metal technologies" (No. 1.1.1.1/18/A/108), a DAAD Short-Term Grant (2021, 57552336) and the ANR-DFG project FLOTINC (ANR-15-CE08-0040, EC 217/3), as well as directly by PSI, HZDR and UL. In the final year prior to thesis completion, additional funding was provided via the project "Strengthening the capacity of doctoral studies within the framework of the new doctoral model" (No. 8.2.2.0/20/I/006).

This thesis summary contains the introduction, 7 chapters, conclusions and a list of references. Thesis format: a collection of articles published in the field of physics, specializing in fluid and gas mechanics.

Thesis supervisor: *Dr. phys.*, senior researcher **Andris Jakovics**

Reviewers:

1. *Dr. phys.* Tianshu Liu
2. *Dr. phys.* Robert Zboray
3. *Dr. phys.* Guntars Kitenbergs

The thesis will be defended at the public session of the Doctoral Committee of physics, astronomy and mechanics, University of Latvia, at 15:00 on September 29, 2023, Jelgavas iela 3, Riga. The thesis is available at the Library of the University of Latvia, Raiņa blvd. 19.

Chairman of the Doctoral Committee

Secretary of the Doctoral Committee

© University of Latvia, 2023

© Mihails Birjukovs, 2023

ISBN 978-9934-36-071-8

ISBN 978-9934-36-072-5 (PDF)

Contents

I. Introduction	7
A. Relevance of the thesis	7
B. Problems and challenges	8
1. Bubble flow in liquid metal	8
2. Particle-laden liquid metal flow	11
3. Analysis of numerical simulations	13
4. Metal alloy solidification	15
C. Goals and objectives	18
D. Novelty of the presented research	19
E. Theses	21
F. Methodology	22
G. Author's contributions	22
H. Presentation of results	24
1. Relevant publications	24
2. Relevant conferences attended by the thesis author	25
II. Outline of the doctoral thesis	26
A. Neutron imaging of bubble chain flow in liquid metal	26
1. The first round of experiments	26
2. The second round of experiments	34
3. Summary	53
B. X-ray imaging of bubble chain flow in liquid metal	55
1. Image acquisition & properties	55
2. Image processing	56
3. Bubble shape analysis	58
4. Summary	63
C. Bubble trajectory reconstruction	64
1. The concept of MHT-X	64
2. The algorithm	64
3. Main results	71
4. Summary	81
D. Neutron imaging of particle-laden liquid metal flow	82
1. Neutron imaging	82
2. Image processing	83
E. Particle trajectory reconstruction	87
1. Modifications to MHT-X	87
2. Preliminary results	89
3. Further extensions for MHT-X	94
4. Further results	99

5. Summary	113
F. Dynamic mode decomposition of bubble chain flow	114
1. Modelled physical system	114
2. The new DMD algorithm	115
3. Main results	121
4. Summary	156
G. X-ray imaging of metal alloy directional solidification	158
1. Image characterization	158
2. Methodology	160
3. Main results	175
4. Summary	189
III. Conclusions & outlook	190
References	194

Terminology and abbreviations

MF - magnetic field
UDV - ultrasound Doppler velocimetry
MHD - magnetohydrodynamics
Ga - gallium
GaInSn - galinstan
CIFT - contactless inductive flow tomography
UTTT - ultrasound transit time technique
XR - X-ray radiography
XCT - X-ray computed tomography
NR - neutron radiography
SNR - signal-to-noise ratio
FPS - frames per second
CNR - contrast-to-noise ratio
PEPT - positron particle emission tracking
PIV - particle image velocimetry
MHT - multiple hypothesis tracking
DFT - discrete Fourier transform
POD - proper orthogonal decomposition
PCA - principal component analysis
DMD - dynamic mode decomposition
SCTMM - soft color tone map masking
MRIF - multiscale recursive interrogation filter
LFPE - luminance-based false positive elimination
NMM - non-local means masking
SVD - singular value decomposition
MPI - message passing interface
Ar - argon
HMF - horizontal magnetic field
FOV - field of view
SGS - subgrid scale model
LES - large eddy simulation
SST - shear stress transport
VOF - volume of fluid method
SSCF - self-snakes curvature flow
SNIP - statistics-sensitive non-linear iterative peak-clipping
CFD - computational fluid dynamics
TF - transverse field
LF - longitudinal field
 Re - Reynolds number
 Eo - Eötvös number

Fr - Froude number
 N - Stuart number
VMF - vertical magnetic field
DCC - dark current correction
FFC - flat-field correction
CTM - color tone mapping
IW - interrogation window
PM - Perona-Malik
TV - total variation
NM - non-local means
LA - local adaptive
RDP - Ramer-Douglas-Peucker
BCL - bubble camber line
SOI - sphere of influence
PTV - particle tracking velocimetry
 Stk - Stokes number
 Sr - Strouhal number
PDF - probability density function
DFI - divergence-free interpolation
RBF - radial basis function
BC - boundary condition
LIC - line integral convolution
RMS - root-mean-square
SF - solidification front
LZ - liquid zone
SZ - solid zone
BM3D - block-matching 3D
MIL - mean inverse luminance
DGD - dendrite grain decomposition

I. Introduction

A. Relevance of the thesis

Multiphase flow in the presence of external magnetic field (MF) is of great fundamental and practical interest. Bubble flow in liquid metal is eminent in industrial processes such as liquid metal stirring, continuous casting, chemical reactions, etc., where it is or, in principle, could be controlled using applied MF [1–4]. However, intimate understanding of multiphase flows in metallic melts is necessary to enable control and optimization of these processes, and to avoid detrimental flow instabilities. Liquid metal multiphase flow has been extensively studied by means of ultrasound Doppler velocimetry (UDV) [5–8], ultrasound transit time technique [9, 10], X-ray imaging [10] and numerical simulations [11–22], and many of its characteristics and mechanisms are presently sufficiently clear [23–28]. The issue is that many effects related to bubble collective dynamics are not yet properly understood or have not been studied at all [29–32]. This is important because improving effective models for bubble flow (Euler-Euler and Lagrangian) and the above-mentioned industrial processes is not possible without insights into how bubbles interact in magnetohydrodynamic (MHD) flow (or even without applied MF) [33–36].

Bubble interaction with particles is of interest in metal purification [37–41] and froth flotation [42–44] where gas bubbles are injected to remove impurities, which occur in the form of solid particles, from the melt. Purification is accomplished mainly via two mechanisms: firstly, the rising bubbles generate turbulent flow which induces particle agglomeration, increasing the effective particle size and enhancing gravitational separation due to density differences; secondly, direct bubble-particle collisions can occur, trapping the latter at the gas-liquid interface of the bubbles ascending to the free surface. Investigation of such bubble-particle interactions is of utmost importance. The wake flow of bubbles (i.e., the flow pattern that forms behind the bubble) is what primarily determines the trajectories of ascending bubbles if their collective effects are negligible [9, 23, 25, 27]. It has also been demonstrated that free-moving particles are trapped in this wake region, increasing their local concentration and the probability of collision and agglomeration [19–22].

Solidification is a central aspect of many industrial applications, particularly in metallurgy, e.g. production of nickel-based superalloys, lightweight aluminum and magnesium alloys, etc. [45–47]. A well-known and common problem is the risk of defect formation during these processes. Segregation of solute species originates at the microscale, but propagates to and emerges at the macro scale (*macrosegregation*), leading to non-uniformity of the

distribution of the inter-metallic phases in industrial alloys [48]. In addition, during solidification of alloys, partitioning of elements leads to the formation of a solute boundary layer in the vicinity of the liquid–solid interface. In cases where the density of the solute may be lighter than that of the bulk liquid, buoyancy forces in the boundary layer directed back towards the bulk liquid cause the formation of solute plumes that emanate from the solid-liquid interface. Under certain conditions, the escaping solute can form stable channels called *chimneys*. After complete solidification, these disturbances remain as defects in the castings known as *freckles*, which are essentially anisotropic alloy composition inhomogeneities in the form of channels with diameters proportional to a few primary dendrite arm spacings, and lengths varying from millimeters to centimeters [46, 49–54]. It is therefore desirable to control solidification such that defects do not occur. However, control requires understanding the underlying physics, and solidification processes in liquid metal alloys are very complex, with many possible regimes of pattern formation depending on system parameters (e.g. temperature gradient, cooling rate, alloy component mass fractions) [45, 47, 55–57]. There exists an interplay of many physical mechanisms on different length scales: dendrite growth, liquid–solid interface instabilities, natural and forced (in general) liquid mass flow, concentration transport, liquid flow through solidified dendrite structures and remelting, global and local temperature dynamics, etc. [47, 55, 58–62]. One way to control such complex systems is by applying MF to the domain where solidification occurs – however, one then has to consider additional physics, e.g., liquid flow damping or forcing by the Lorentz force, including a thermoelectric contribution. This and other factors introduced by MF application significantly alter solidified microstructures [46, 63–65].

B. Problems and challenges

1. Bubble flow in liquid metal

While of interest, liquid metal systems are notoriously difficult to probe non-invasively due to their opacity, high temperatures and additional complications due to strong MF [5]. Consequently, visual data regarding how applied MF alters bubble shape and bubble collective dynamics are scarce, and experiments with particle visualization in liquid metals are rarer still. The issue with high temperatures and large scales of industrial hardware can be sidestepped by using smaller volumes of low melting point model liquids, such as gallium (Ga) or galinstan (GaInSn) or gallium-tin (GaSn), provided the actual system’s dimensionless parameters are sufficiently close

to model system's [5]. However, to extract physical information from experiments, appropriate experimental and data processing methods must enable tracking bubbles and particles, resolving their collisions, coalescence and/or breakup. In the case of bubbles, resolving their boundaries and capturing dynamics thereof is also necessary.

Regarding bubble flow, several velocimetry methods of varied reliability for liquid metal systems exist, such as contactless inductive flow tomography (CIFT) [66, 67] and UDV [5, 6, 68], but these only yield liquid metal velocity field and simply reveal flow perturbations wherefrom bubble presence may be inferred. With optical methods off the table due to liquid metal opacity to visible light, the ultrasound transit time technique (UTTT), X-ray radiography (XR), X-ray computed tomography (XCT), and neutron radiography (NR) remain among the viable methods. UTTT is very limited in that it only detects the general location and characteristic size of a bubble within metal, and yields no information about its shape, meaning that surface perturbations, being unresolved, may also introduce errors in position and velocity estimates [6–8, 10]. XR and XCT methods allow for direct observation of bubble shapes. However, X-ray based techniques are, in general, very constrained by relatively small liquid metal thickness in the beam direction due to intense X-ray attenuation by liquid metals [29, 30, 69–71]. At the same time, if the studied model system is too thin, wall effects become very significant and the observed bubble flow is much less representative to industrially relevant flow conditions. XCT, while offering very high temporal resolution and sufficient phase boundary detection precision, also entails experimental systems that are rather susceptible to applied MF, rendering them hardly applicable to studies of MHD bubble flow [72–74].

NR, on the other hand, while similar in principle to XR, allows, for some metals, probing much thicker samples and, as such, is a very promising experimental method that should enable broader coverage of the parameter spaces of various liquid metal systems [1, 75, 76]. While considerable progress has been made outside the liquid metal context and without MF [77–79], only a few advances with some notable exceptions, have been achieved for liquid metal systems using NR, and it has not been used to systematically study the MF effects on bubble flow [1, 75, 76, 80–83]. However, it must also be noted that, ultimately, the ability to image thicker liquid metal systems depends on X-ray and neutron fluxes that one is capable of producing, so industrial-grade X-ray tubes could potentially be used to great effect as well, though examples of this are yet to be found in literature for liquid metals.

Through very recent efforts and the advent of dynamic XR and NR for two-phase liquid metal flow [82–88], fundamental investigation of bubble

chain systems mimicking industrially relevant flow conditions is finally underway [1, 11, 29, 30, 69–71, 84]. In bubble chain flow, bubbles are released into a liquid metal system one-by-one with a uniform time delay between each, at a certain gas flow rate, and ascend to the free surface of liquid metal. Such systems are usually rectangular vessels filled with gallium [1] or an eutectic gallium-indium-tin alloy [29, 30, 69, 70, 84] where bubbles are introduced via horizontal or vertical [1, 29, 30, 69, 70] nozzles at the bottom of the vessel, or top-submerged vertical [84] nozzles. Bubble chains flow systems are the next logical step from single-bubble flow investigations, since single-bubble flow, while very informative of the bubble wake flow dynamics and characteristic trajectories without and with applied MF, is not representative of the actual flow conditions typical for the above-mentioned industrial processes where one has columns and swarms with a high number density of deformable bubbles [37, 41, 89, 90].

Bubble chains are still simple enough to enable experimentation with compact systems [1, 69, 71] and contain computationally manageable numbers of bubbles within the liquid metal volume [11, 71]. Meanwhile, they already exhibit collective dynamics between leading and trailing bubbles [11, 70, 71] and, depending on the system geometry and flow rate, bubble agglomeration, coalescence and breakup can occur [29, 30, 84].

Therefore, these systems are a crucial milestone in a transition from studying single-bubble flow to investigations of many-bubble systems that are very close to their actual industrial counterparts. However, despite the *relative simplicity*, dynamics exhibited by bubble chain flow in liquid metal without or with applied MF are still very complex. Depending on the gas flow rate, bubbles produce unstable elongated wake flow regions where periodic vortex detachment occurs, and turbulent pulsations are generated – shed vortices and turbulent wakes of leading bubbles strongly affect the trailing bubbles, leading to bubble pair coupling across the ascending chain [11, 71, 91–94]. There exists a feedback loop involving combined perturbations of bubble shapes and within the bubble chain, surrounding liquid metal flow, and the influence of the free surface at the top of the metal vessels with instabilities and oscillations in the bubble chain shape [11, 71, 94].

In short, despite relevance and the availability of neutron and X-ray imaging facilities and appropriate hardware, bubble flow has not yet been studied systematically via experiments, i.e. a parameter space in terms of dimensionless numbers is yet to be established for bubble chain systems (exists for single-bubble flow [11, 95]). This is because there are other reasons besides experimental difficulties that prevent further research.

It has become abundantly clear that specialized and rather advanced image processing methods and tools are required to extract physically

meaningful data from data sets acquired via dynamic neutron and/or X-ray imaging [1, 29, 30, 84]. This is mainly due to the low signal-to-noise ratio (SNR) associated with imaging thick ($> 20\text{-}30\text{ mm}$) layers of liquid metal at frame rates $\gtrsim 100$ of frames per second (FPS) and the need to resolve many, often closely packed, interacting objects. High frame rates are a requirement to enable capturing fast bubbles, drops and particles flowing in liquid metal and to avoid motion blur [1]. Even with thinner liquid metal layers one faces problems with data interpretation – if bubble agglomeration, breakup and coalescence are frequent, correct detection and resolution of bubble interaction is complicated since they are separated by thin liquid metal films which are corrupted by multiplicative noise in the images and suffer from low contrast-to-noise ratio (CNR) [29, 30]. Meanwhile, neutron flux that can be used in experiments is limited by both the utilized neutron source and the rapid activation of model liquid metals such as gallium. In addition, studying bubble collective dynamics requires not just detection, but also accurate tracking that accounts for bubbles merging and splitting – robust methods that would enable this are currently sorely lacking. Therefore, thus far, researchers investigating bubble flow in liquid metal have been forced to mainly rely on numerical simulations, which in most cases are not directly validated.

2. Particle-laden liquid metal flow

As for particle-laden liquid metal flow, despite the physical interest, much like in the bubble flow case, there is very little experimental work (in contrast to many available simulations [11–18]) where particles in bubble wakes or direct bubble/particle interactions are directly visualized in liquid metal [96]. Again, one of the main reasons for this is the lack of suitable measurement techniques to perform such measurements in opaque liquids (in this case metals), where optical methods cannot be leveraged. UDV has been applied to bubble wake flow characterization [5, 97], but, despite sufficient temporal resolution, currently the spatial resolution is not enough to reliably identify individual particles. Particle tracking in liquid metal using positron emission particle tracking (PEPT) has also been considered for flow analysis [98–102], but this method provides very low temporal resolution, making it not feasible for turbulent flow (e.g. bubble wake flow is turbulent).

It was also proposed some time ago that NR could also be used to directly observe wake flow of bodies and particle flow within optically opaque systems [103]. The first such benchmark study in the context of liquid metal flow with dispersed particles was recently done by Lappan et

al. where gadolinium oxide particle flow around a cylindrical obstacle in a thin liquid metal channel was imaged dynamically with sufficient temporal resolution using high cold neutron flux [96, 104]. The imaged turbulent particle-laden flow was investigated using particle image velocimetry (PIV) and the wake flow velocity field was measured and visualized.

While particle tracking in liquid metals is an important problem that can be solved using dynamic NR, only a very limited number of papers address image processing required to successfully extract physically meaningful information from the acquired image. Notably, Heitkam et al. performed particle detection and tracking in froth using neutron imaging utilizing a particle-mask correlation approach [105]. An original approach for detecting particles and tracking particle flow in the presence of bubbles was demonstrated by Sommer et al., although not in the context of liquid metal [44]. Another approach that is very promising for particle detection within flow with a high particle number density was developed by Anders et al. for optical measurements, but could potentially be generalized [106, 107]. However, the latter two do not seem to be readily applicable to low-SNR images typically associated with high frame rate neutron imaging, and the former would be hard to generalize due to its reliance on preset particle masks. A very comprehensive overview of particle detection and tracking methods with an objective comparison is provided in [108], but again, the showcased methodologies for particle detection are tested at SNR which, compared to the case in this thesis, is very high. Unlike bubble flow imaging, with particles one also faces the issue of correlated noise that produces "phantom" particles, which further complicates detection [105].

Furthermore, it is advantageous to combine a more noise-resilient image processing approach with a more general method for particle tracking based on detections. While advanced particle detection methods are shown, simple nearest neighbor tracking is used in [106, 107], which generally does not perform so well. Tracking methods used in [105] and [44] each use a set of restrictions that make them problematic to use for particle flows with high number densities (i.e. inter-particle distances and particle sizes are of similar magnitudes) where detection is carried out under adverse imaging conditions. Once particles are detected, tracking in [105] essentially relies on combined nearest neighbor and velocity-based predictions. Tracking in [44] is performed using a version of the shake-the-box algorithm [109], wherein tracks of particles are predicted based on preceding time steps and uncertainties between the predicted and actual positions are corrected by varying the predicted position in space until it matches the actual position. Newly entering particles are triangulated as in [110]. While certainly tried and viable, this approach has many inherent constraints which, again, more often than not become a problem when particle number

density is high, particle motion is highly irregular and the false positive and occlusion (i.e. overlaps of visible particle projections in images) rates are not negligible, as it is in this case. This is also true for the methods outlined in [108]. Even in the cases where the most general and robust framework, multiple hypothesis tracking (MHT), is used, the imposed motion restrictions are quite severe. Specifically, near-constant position and/or velocity are assumed, i.e. no physics- or Kalman filter-based motion models are used. These approaches are therefore inapplicable to cases where particles may experience substantial acceleration and changes in position between consecutive frames due to interactions with bubble wake flow – or any benchmark experiment where particle flow in the wake of a body is studied, as in [96].

In other words, despite the fact that approaches to particle detection in metal and particle-laden flow analysis are not quite as scarce as is the case with bubbles, there is still a lack of a robust approach, which prevent further systematic studies. Such studies would require very accurate particle detection and tracking that could generate sufficiently long tracks that encode the underlying system physics, such as the character of flow turbulence. Consequently, studies like [96] have been limited to PIV for flow analysis, which does not enable one to perform an in-depth analysis of turbulence characteristics. This severely limits the means available for research and forces to rely heavily on numerical simulations, which are often not validated.

3. Analysis of numerical simulations

Unfortunately, the problems do not end with bubble and particle imaging and tracking in liquid metal systems. Experimental challenges aside, there is also the matter of analyzing simulation output for the mentioned bubble chain model systems. Given their complicated physics outlined above and the fact that such bubble flow exhibits very diverse temporal and spatial scales [11, 70, 71, 91–94], interpreting simulation results becomes very challenging, especially if it is necessary to assess the differences due to gradual changes in gas flow rate and MF magnitude/orientation. Therefore, it makes sense to look for more convenient representations of the velocity field dynamics, preferably with reduced dimensionality. While vorticity, the Q-factor [111] and other representations based on velocity gradients [112] are very useful for the assessment of turbulent structures within bubble wakes and vortex detection, they do not explicitly inform one about temporal and spatial scales, and dynamics thereof.

One of the most common methods for time series analysis is the discrete

Fourier transform (DFT). If the data exhibits periodic structure, DFT can be used to represent the system in the frequency space, which is a more natural representation of the system. However, if the dataset associated with the system is very high-dimensional in space, it can be very difficult and inconvenient to analyze its structure because FT does not simplify the system in any way. Another way to decompose time series data is through proper orthogonal decomposition (POD) which is equivalent to the principal component analysis (PCA). POD expresses the time series data in a new orthogonal basis. The basis components are chosen such that they optimally cover the data in terms of energy content captured by successive modes. The time evolution is then analyzed in the new reduced coordinate system. While the new coordinate basis is simpler than the original one, it is not guaranteed to be physically meaningful [113].

Dynamic mode decomposition (DMD) is a dimensionality reduction algorithm for spectral analysis of periodic or quasi-periodic data. DMD is effectively a combination of the PCA and the DFT, as it decomposes a series of system states into a number of dominant spatial structures associated with unique frequencies [114]. At the same time, it is also connected to the perturbation theory for partial differential equations [115]. Given a time series of data where each series element is the state of some dynamic system, instead of generating orthogonal basis vectors, a DMD algorithm constructs modes with unique frequencies and growth/decay rates that are not necessarily orthogonal. Each mode's time evolution is independent of other modes. Unlike DFT, DMD generates a *sparse* set of frequencies paired with corresponding spatial modes that account for dominant system dynamics patterns. This enables a simplified yet more physically meaningful representation of the system [113–115]. DMD was originally developed to analyze fluid dynamics systems [114]. Since its inception, it has been used to analyze flow instabilities and vortex shedding in fluids [115–118]. DMD has also been used to study different nonlinear systems in meteorology [119], for video processing [120, 121], electrocorticography analysis [122, 123], sunspot data analysis [124], etc. DMD is a contemporary and actively growing field of research with applications spanning multiple disciplines.

Therefore, in the context of MHD bubble chain flow, DMD is a prospective method that should enable more meaningful and in-depth data interpretation. While there are several cases where DMD has been applied to MHD flows [125–127] and DMD has been used to study two-phase flows and flows containing bubbles [128–132], there are no cases in the literature (to the authors' knowledge at the time of this publication) where DMD was applied to two-phase flow with explicitly resolved bubbles, including, in particular, a chain of bubbles rising due to the buoyancy force, without or with applied MF. In the MHD bubble flow context, there are currently no

instances of DMD applications to UDV measurements or simulations of the velocity field for bubble flow without and with applied MF in liquid metal, and no instances of DMD analysis of bubble shapes from simulations or data obtained via dynamic XR or NR – all of these cases represent important steps to be made towards the understanding of magnetic control of bubble flow in liquid metal. With DMD analysis, one could separate mean flow from large- and small-scale quasi-periodic transients due to wake vortex detachment, back-flow due to free surface oscillations, turbulent pulsations at different spatial/temporal scales, as well as provide quantitative measures for the coherence of bubble motion within the chain and bubble-bubble distances – and compare the characteristic spatial patterns within flow, their frequencies and mode amplitude dynamics over time for different MF orientations/magnitudes and flow rates. This would be crucial for identification and quantitative explanation of bubble chain instabilities to gain insights into their prevention/control, and, potentially, derivation of useful empirical relations between flow characteristics and dimensionless groups that parametrize metal and bubble flow.

4. Metal alloy solidification

A very prospective and commonly used method for studying solidification dynamics without or with applied MF is by using downscaled model systems – Hele-Shaw cells where binary alloy solidification can be observed at the meso-scale (i.e. dendrite grains with spatially resolved individual dendrites) using *in situ* dynamic X-ray transmission contrast radiography. Even though one obtains only the projections of the solidified microstructures, it has proven to be a very effective means of probing systems with solidification processes for physical insights [46, 53, 54, 59, 60, 64, 133–142]. In addition to the challenges associated with imaging, there is also the matter of extracting valuable information from the acquired images. Ideally, to get the full picture of system dynamics, one has to separate the liquid from the solid, identify the solidification front and any liquid enclosures within the solidified microstructure, and obtain the microstructure skeletons. One could then perform orientation analysis for the skeletons, derive the primary dendrite spacing statistics, determine the local velocity with which the solidification front travels, as well as measure the solute concentration near the solidification front as it moves, since the concentration largely determines the front evolution. In this regard, detecting convective plumes and analyzing their shapes is also of interest, as is velocimetry in the liquid flow regions. In addition, it could also be of interest to detect and separate different grains (if any) within the microstructure seen in the images. Of

course, the problem lies in doing all of the above automatically and reliably, which is relevant given the amount of images usually acquired in X-ray radiography experiments and the amount of information captured within each image.

However, while there exist solutions for some of the above problems, most appear to be limited to segmentation/detection of dendritic structures [143–148]. In [143–145], the focus is the detection of the dendrite cores from images of planes normal to the solidification direction – this, and the fact the algorithm presented in [143, 144] uses template matching as one of its stages, makes it hardly applicable to studying X-ray images with directionally solidifying dendrite "forests" (e.g., as in [46, 53, 54, 59, 142]) where planes parallel to the growth direction are imaged. In addition, at least observing the demonstrated application examples, it seems that these methods should be reliably applicable in the cases when the images are fully filled with dendrites, i.e., in instances where there are both liquid and solid regions, one must first be separated from the other using different methods. Dendrite tip tracking is performed in [146] by segmenting the upper part of the solidified structures growing upwards. The key aspect of the segmentation procedure is to use the difference between two consecutive frames to highlight the newly formed solid, segment the relevant region, and then derive the tip coordinates. The utilized approach also enables tracking the solidification front. However, in cases where the differences between the frames are smaller and significant noise is present, the algorithm could be expected to run into performance issues and a more general approach is desired. A versatile approach using neural networks for automated detection of equiaxed dendrite detection was reported in [147]. Another important example of automated dendrite segmentation using neural networks was shown in [148] where, unlike in [147], a binary mask for the solid structure was predicted instead of detecting separate dendrites. It is, however, worth pointing out that the example images/cases presented in [147, 148] do not exhibit significant noise, which is often present even after some temporal averaging in dynamic X-ray radiography experiments where exposure times are relatively low – it is therefore not clear how well these methods will perform under such conditions.

The lack of a systematic approach to image processing beyond methods for segmentation presents a problem, since it has been clearly demonstrated that the microstructure evolution must be analyzed in conjunction with the other processes in solidifying systems. Currently, the most common tool used for image analysis in the field is *ImageJ* with its many custom plugins developed by the community [149, 150]. While *ImageJ* is open-source, with an impressive arsenal of methods, many of them are not automated, robust, or publically available. In contrast, it would be very convenient

to have an open-source all-in-one solution for X-ray image analysis. It should also be noted that such code could be applied to the output of numerical simulations as well, the difference being that the latter do not have the image noise associated with experimental measurements. Thus, more direct comparisons between simulations and experiments, which seem to be largely lacking, could be possible.

C. Goals and objectives

Given the state of the art outlined above and the current issues severely slowing down or outright blocking further progress in the field, the goals of this thesis are:

- Perform systematic experimental investigation of bubble chain flow in liquid metal, with and without applied MF, using NR and XR, numerical simulations and appropriate image and data analysis methods
- Develop the currently missing image and data processing tools necessary for the analysis of experimentally obtained data and output of numerical simulations

Specifically, the objectives are as follows:

1. Design and build a down-scaled model system that can generate bubble chain flow at different gas flow rates, has sufficient thickness to minimize wall effects on bubbles, and is compatible with available X-ray and neutron radiography labs and setups, as well as allows for the application of MF
2. Design and build MF systems that can be used to apply MF of different strength and orientation to the model system
3. Perform dynamic neutron radiography for a range of flow rates and MF strengths and orientations to map the parameter space of the model system in order to observe various bubble flow instabilities and how they are affected by varying the system parameters
4. Perform dynamic X-ray radiography for a Hele-Shaw version of the model system for a range of gas flow rates and with the same MF configurations to study bubble collective dynamics and how they are affected by applied MF
5. Develop image processing methods capable of accurately detecting positions and shapes moving bubbles from low-SNR neutron images
6. Develop image processing methods that can resolve tightly agglomerated bubbles separated by thin liquid metal films from X-ray images
7. Implement methods that enable bubble tracking based on detections from neutron and X-ray images, and account for bubble coalescence and breakup
8. Develop image processing methods for particle detection in neutron images with low SNR/CNR and in presence of phantom particle artifacts
9. Implement methods that enable particle tracking with sufficient accuracy to enable in-detail analysis of turbulent liquid metal flow

10. Validate the methods developed for bubble and particle detection and tracking
11. Adapt DMD for analysis of bubble chain flow simulation data
12. Develop image processing methods to facilitate the analysis of X-ray images of liquid metal alloy directional solidification
13. Make all the resulting code and methods publicly available via publications and open-source *GitHub* repositories

D. Novelty of the presented research

One of the main results of this thesis are comprehensive neutron and X-ray imaging datasets for bubble chain flow under applied MF of different strengths and configurations, for a broad range of gas flow rates. The utilized model systems with rectangular liquid metal vessels, wherein bubble flow is introduced via horizontal/vertical nozzles, were imaged with different liquid metal layer thickness. Data obtained for the thicker system can be used to study how bubble chain instabilities arise from leading/trailing bubble interactions and how they are affected by system parameters, whereas the Hele-Shaw variant of the model system is meant to study bubble agglomeration and direct interactions – collisions, coalescence and breakup. No such datasets were available in the field before, which now sets the stage for further physics-focused research of MHD bubble flow. Preliminary data analysis revealed the multitude of flow regimes that can be achieved by apply different MF strengths and orientations.

The neutron imaging experiments in particular are the first successful campaigns of this kind, i.e. single-bubble or bubble chain flow visualization through a record 30 *mm* liquid metal thickness. The previous attempts suffered from insufficient frame rate or the inability to extract meaningful information from the obtained images. The significance of the greater than before thickness of the imaged liquid metal cell is that it almost completely avoids wall effects and enables one to approach industrially relevant flow regimes, while at the same time being simple enough that fundamental aspects of bubble interaction within chains and chain interaction with the surrounding flow can be assessed.

The reason why the previous attempts at dynamic neutron radiography were not successful is that using higher frame rate and lower exposure time significantly reduces the image SNR as well as the CNR for bubbles, preventing reliable detection with simpler methods, never mind accurate position determination or shape measurements. This issue is at last solved and measurements with sufficient frame rate are now possible despite the low SNR using the developed image processing methods, which are a com-

bination of both known methods and original approaches developed for this thesis. The latter include the *soft color tone map masking* (SCTMM), *multiscale recursive interrogation filtering* (MRIF) and a new luminance-based false positive elimination (LFPE) method. Importantly, these methods are readily applicable well beyond the neutron imaging of bubble flow – this is also demonstrated and further discussed in this thesis. The implemented methods are available on *GitHub* and have been validated indirectly against numerical simulations, as well as directly using reference experiments.

MRIF, originally developed for neutron images of bubble flow, was also instrumental for solving the bubble detection problem for X-ray imaging, where it, in combination with other methods, enabled to successfully resolve clusters into separate bubbles, which, as seen in literature, was not possible before because the thin metal films between the bubbles had low CNR and were corrupted by multiplicative noise.

X-ray imaging of a 30 mm thick liquid metal model system has also been performed for the first time – while the resulting images have noticeably smaller bubble CNR, the SNR is significantly better than for neutron images when comparison is made at identical frame rates.

The problem of particle detection in low SNR/CNR neutron images with phantom particle artifacts has also been overcome using local filtering by image partition into partially overlapping windows – here a combination of known methods, SCTMM and an original non-local means masking (NMM) method ensured particle detection and LFPE efficiently eliminated phantom particles. This is the first robust method developed for such imaging problems for liquid metals. Again, the methodology is extendable beyond the current application and is available via an open-source *GitHub* repository.

An object tracking algorithm *MHT-X* was developed for particle and bubble trajectory reconstruction. The algorithm core is the most general of the tracking methods, *multiple hypothesis tracking* (MHT), which for the first time was made feasible for many-object scientific problems by combining MHT's offline form with *Algorithm X* and using weak mass and momentum conservation constraints for object motion. MHT-X was validated for bubble tracking based on simple benchmarks and by combining it with the developed image processing tools for neutron images. The problem of particle tracking in liquid metal was solved using a novel approach which combined MHT-X and its physics-based particle motion constraints with PIV-assisted motion prediction, where divergence-free PIV field interpolation was used to dramatically boost tracking performance. MHT-X is currently unique in its class and has been validated in combination with the particle detection methods using a benchmark experiment and comparisons with numerical simulations, theory and other experiments

documented in literature. The significance of this result is that, for the first time, direct measurements of flow turbulence characteristics in liquid metal is feasible based in particle flow imaging. MHT-X is currently open-source and is available on *GitHub*.

DMD was applied to bubble flow simulations for the first time. A noise-resilient and memory-efficient DMD implementation based on higher-order DMD combined with the MOSES streaming singular value decomposition (SVD) was applied to the output of numerical simulations and enabled an in-depth analysis of both large-scale velocity field modes in the studied model systems, and the bubble wake flow modes in the bubble reference frame. Modes produced by DMD successfully explained the observed system dynamics with and without applied MF, and allowed to identify the complex coupling between the large-scale modes in the liquid metal vessel and the bubble wake modes. The DMD code is also available on *GitHub*.

Finally, a set of image processing methods were implemented into an all-in-one code for analyzing X-ray images of directional solidification processes. The code is equipped with an original and robust physics-based model for liquid/solid image region separation, a routine for solid region analysis and segmentation of solidified dendrites, liquid metal channel (chimney) and cavity separation, as well as an original method for decomposing the solid region with dendrites into dendrite grains using orientation-to-color space mapping. The code also combines known image processing tools with some of the methods developed for bubbles and particles, such as SCTMM and NMM, which are shown to significantly bolster the existing state-of-art methods. The code is, as with all of the above cases, open-source and available on *GitHub*, and is currently the only open-source tool in the field that covers all of the above functions.

E. Theses

- Developments in neutron radiography presented in this thesis have enabled to image bubble flow in thick liquid metal systems, without or with applied magnetic field, and physical information regarding bubble flow can be extracted using the implemented image and data analysis methods.
- X-ray radiography has been successfully applied to Hele-Shaw systems (without or with applied magnetic field), and splitting, merging and agglomerated bubbles have been resolved and analyzed via specialized image and data analysis methods.
- Using neutron and X-ray radiography, a comprehensive experimental dataset has been obtained that is currently unique in the field, and

covers bubble flow under different magnetic field configurations and degrees of confinement, as well as at different gas flow rates.

- Dynamic neutron radiography in combination with dedicated image and data analysis tools is suitable for particle flow analysis in liquid metals, and can be used to extract turbulence characteristics.
- Dynamic mode decomposition allows to successfully represent numerically simulated spatially/temporally multiscale bubble flow dynamics in terms of much more interpretable elements.
- Developed image processing methods and code enable to perform a much more in-depth assessment of liquid metal alloy solidification processes using dynamic X-ray imaging, accessing physical information that was previously unavailable for automated quantitative analysis.

F. Methodology

Dynamic NR and XR were used for bubble flow imaging in liquid metal – NR was performed at a spallation source beamline (NEUTRA and ICON at PSI), whereas XR was done using an industrial-grade X-ray tube (at HZDR X-ray lab). MF systems for the experiments were designed, simulated and optimized using *COMSOL Multiphysics*. Numerical simulations of the MHD bubble flow were performed using *OpenFOAM* (hydrodynamics) and *Elmer* (electromagnetics) using *EOF-Library* for MPI coupling. *Salome* and *ANSYS Mesher* were used to generate geometry meshes, and *ParaView* was used for post-processing and visualization. The developed image processing methods for bubble- and particle-laden liquid metal flow and solidification process analysis were implemented in *Wolfram Mathematica*, coupling it with *MATLAB* via *MATLink* when necessary. MHT-X and the utilized version of DMD were both implemented in *Python*. Visualization of the DMD modes was done in *ParaView*. Data analysis as well as plot and figure rendering were done with *Wolfram Mathematica* and *Python*.

G. Author’s contributions

Since the results presented in this thesis are the work of a large international team, with the thesis author leading the effort, it is important to explicitly highlight the author’s contributions. These are as follows:

- Planning, organizing and conducting the neutron radiography experiments. This includes: design and parts procurement for the imaged model system; design, optimization and parts procurement for the MF

systems used to apply MF to the model system; estimation of the necessary amount (to account for sample activation by the neutron flux) and procurement of liquid gallium; performing neutron imaging of bubble flow in liquid metal at NEUTRA and ICON beamlines at PSI; reference experiments at the ICON beamline used for validation of the developed image processing methods

- Planning and performing X-ray imaging of bubble flow at the HZDR X-ray lab
- Development and implementation of the image processing methods for neutron (solo) and X-ray (co-developed with Ansis Ziemelis) imaging of bubble flow in liquid metal, including the shape analysis algorithm (co-developed with Peteris Zvejnieks)
- Development and implementation of the image processing methods and code used for neutron images of particle-laden liquid metal flow
- Development and implementation of the image processing methods and code used for the analysis of X-ray images of metal alloy directional solidification processes
- Co-development of the MHT-X object tracking code (main developer Peteris Zvejnieks): contributed mathematical models for physics-based object motion constraints, statistical functions for object motion and interaction, the concept of using PIV for motion prediction and insights into the mathematics of the utilized methods
- Co-development of the DMD code (main developer Martins Klevs): contributed insights into the mathematics of the utilized methods
- Physical interpretation for the results across all papers associated with the thesis
- Main contributor to the first drafts of all the papers linked to this thesis
- Supervision of three completed BSc theses (by Peteris Zvejnieks and Martins Klevs: MHT-X and DMD code development, respectively; by Ansis Ziemelis, X-ray image processing code co-development) and one in-progress MSc thesis by Aleksandrs Jegorovs relevant to this thesis' fields of research
- Establishment and maintenance of collaborations with HZDR, PSI and University of Greenwich (UG)

H. Presentation of results

1. Relevant publications

1. Mihails Birjukovs, Natalia Shevchenko and Sven Eckert. “An image processing pipeline for in-situ dynamic X-ray imaging of directional solidification of metal alloys in thin cells”, *arXiv*, 2022, <http://arxiv.org/abs/2210.00846>
2. Mihails Birjukovs, Peteris Zvejnieks, Tobias Lappan, Martins Klevs, Sascha Heitkam, Pavel Trtik, David Mannes, Sven Eckert, and Andris Jakovics. “Particle tracking velocimetry and trajectory curvature statistics for particle-laden liquid metal flow in the wake of a cylindrical obstacle”, *arXiv*, 2022, <https://arxiv.org/abs/2206.11033>
3. Mihails Birjukovs, Peteris Zvejnieks, Tobias Lappan, Martins Sarma, Sascha Heitkam, Pavel Trtik, David Mannes, Sven Eckert, and Andris Jakovics. “Particle tracking velocimetry in liquid gallium flow around a cylindrical obstacle”. In: *Experiments in Fluids* **63**, 2022, <http://dx.doi.org/10.1007/s00348-022-03445-2>
4. Peteris Zvejnieks, Mihails Birjukovs, Martins Klevs, Megumi Akashi, Sven Eckert, and Andris Jakovics. “MHT-X: offline multiple hypothesis tracking with algorithm X”. In: *Experiments in Fluids* **63**, 2022, <http://dx.doi.org/10.1007/s00348-022-03399-5>
5. Mihails Birjukovs, Pavel Trtik, Anders Kaestner, Jan Hovind, Martins Klevs, Dariusz Jakub Gawryluk, Knud Thomsen, and Andris Jakovics. “Resolving Gas Bubbles Ascending in Liquid Metal from Low-SNR Neutron Radiography Images”. In: *MDPI Applied Sciences* **11(20)**, 2021, <http://dx.doi.org/10.3390/app11209710>
6. Martins Klevs, Mihails Birjukovs, Peteris Zvejnieks, and Andris Jakovics. “Dynamic mode decomposition of magnetohydrodynamic bubble chain flow in a rectangular vessel”. In: *Physics of Fluids* **33(8)**, 2021, <http://dx.doi.org/10.1063/5.0054831>
7. Mihails Birjukovs, Valters Dzelme, Andris Jakovics, Knud Thomsen, and Pavel Trtik. “Phase boundary dynamics of bubble flow in a thick liquid metal layer under an applied magnetic field”. In: *Physical Review Fluids* **5**, 2020, <http://dx.doi.org/10.1103/PhysRevFluids.5.061601>
8. Mihails Birjukovs, Valters Dzelme, Andris Jakovics, Knud Thomsen, and Pavel Trtik. “Argon bubble flow in liquid gallium in external magnetic field”. In: *International Journal of Applied Electromagnetics and Mechanics* **63**, 2020, <http://dx.doi.org/10.3233/JAE-209116>

2. Relevant conferences attended by the thesis author

1. UIE 2021: XIX International UIE Congress on Evolution and New Trends in Electrothermal Processes, Pilsen, Czech Republic, September 1-3, 2021, "Dynamic neutron imaging of argon bubble flow in liquid gallium in horizontal or vertical magnetic field"
2. EPM 2021: Electromagnetic Processing of Materials, Riga, Latvia, June 13-17, 2021, "Image processing methods for neutron and X-ray radiography of liquid and solidified metals"
3. APET-2020: V International Conference Advanced Problems of Electrotechnology, Yekaterinburg, Russia, October 2020, "Image & data processing methods for X-ray and neutron radiography of bubble flow in liquid metal"
4. ENS-2019: The Workshop on Efficient Neutron Sources, Paul Scherrer Institut (PSI), Villigen, Switzerland, September 2-5, 2019, "Argon bubble flow in liquid gallium in external magnetic field"
5. 11TH PAMIR International Conference on Fundamental and Applied MHD, Reims, France, July 1-5, 2019, "Argon bubble flow in liquid gallium in external magnetic field investigated by means of neutron imaging"
6. HES-19: International Conference on Heating by Electromagnetic Sources: Induction, Dielectric and Microwaves, Conduction & Electromagnetic Processing: Padua, Italy, May 22-24, 2019, "Neutron imaging of argon bubble flow in liquid gallium in external magnetic field"

II. Outline of the doctoral thesis

A. Neutron imaging of bubble chain flow in liquid metal

1. The first round of experiments

The first neutron imaging campaign was intended to learn from the previous attempts at NR documented in the literature and assess the feasibility of systematic NR of bubble flow in liquid metal.

a Setup, imaging and simulations

Two neutron imaging campaigns were carried out. During the first one, experiments were conducted at the thermal neutron imaging beamline *NEUTRA* at the Paul Scherrer Institute PSI [151]. The setup consisted of a thin-walled rectangular $150\text{ mm} \times 90\text{ mm} \times 30\text{ mm}$ glass container (Figure 1a) filled with liquid gallium (Ga), wherein argon (Ar) bubble flow was introduced via a submerged horizontal copper tube (gas flow rate adjusted using a digital mass flow controller) which ejected bubbles such that they ascended without wall interactions. To study the influence of applied horizontal MF (HMF), the container was placed between two arrays of neodymium permanent magnets (Figure 1b) with flux density within the container ranging from $\sim 60\text{ mT}$ to $\sim 500\text{ mT}$, and $\sim 300\text{ mT}$ within the bubble flow region, as illustrated [152]. MF at the inlet was $\vec{B} \sim (100; 17; -3.5)\text{ mT}$ [152, 153].

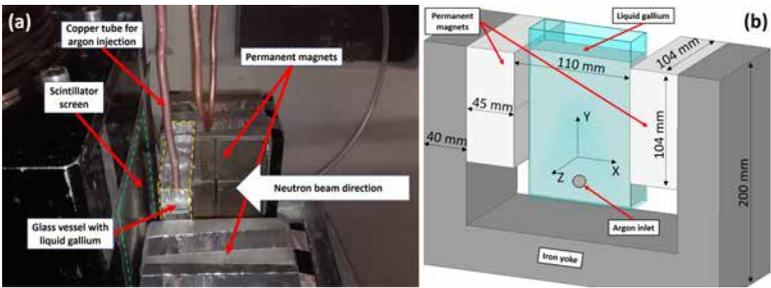


FIG. 1. (a) The Ga vessel used for imaging and (b) a sketch of the vessel inside the MF system.

The experimental setup was imaged such that the field of view (FOV) was $112.8 \times 112.8\text{ mm}$ (enough to capture the entire trajectory of a bubble once it detached from the inlet) with a $55.1\text{ }\mu\text{m}$ pixel size. All images

were acquired at 100 FPS to capture the motion and shape oscillations of ascending bubbles in detail. Bubble trajectories were recorded for 10 to 300 cm^3/min flow rates, with and without applied MF. 30 s sequences were acquired for all flow rates. It is important to note that the 30 mm thickness of the model system is currently the highest in its class. With this, the aim was to push the boundaries of feasible (in terms of NR) thickness of a scaled down system, approaching industrially relevant flow conditions (e.g. minimizing wall effects) [152, 153].

Preliminary analysis and simulations show that the magnetic Reynolds number near the bubbles is $Rm \in [0.01; 0.05]$ for the selected flow rate range, so current induced by metal flow through MF could be important for greater flow rates. However, to accelerate the computations significantly and obtain preliminary results, it was decided to neglect (for now) the MF produced by induced currents. Estimates indicate that $EO \in [2.1; 4.1]$ (6-8 m bubble diameter) and hydrodynamic Reynolds number near bubbles is within $Re \in [10^3; 10^4]$, which corresponds to a flow regime wherein bubbles are of oscillating elliptic shapes. At the same time, $Re \sim 1 - 10$ elsewhere, so an appropriate turbulence model must be used to resolve shape dynamics and trajectories. The k-equation subgrid scale (SGS) large eddy simulation (LES) model was chosen to avoid artificial bubble trajectory stabilization due to overestimated turbulent viscosity, which is introduced by a more standard $k-\omega$ shear stress transport (SST) model. The Froude number is $Fr \in [0.2; 0.3]$, implying that buoyancy is the main driving force, as expected. The Stuart number is $N \in [0.8; 3.5]$ near bubbles, indicating that the Lorentz force should strongly affect metal flow. It is known that, for these values of the Stuart number, MF, horizontal or vertical (or both components, if the field is not homogeneous), radically changes the structure of bubble wake and bubble trajectory characteristics [152, 153].

The experiment was modelled numerically to verify that observed effects are not artifacts due to imperfections in the setup and image processing. *OpenFOAM* (finite volumes) was used to model bubble flow using the volume of fluid (VOF) method (MULES interface scheme), solving the incompressible Navier-Stokes equation (including gravity, Lorentz force, surface tension), the continuity equation and the transport equation for Ga volume fraction, with linear blending functions for material properties. Boundary conditions were: no-slip at the vessel walls; zero relative pressure at the top opening of the vessel to allow gas flow above the free surface; constant mass flow rate at the tube inlet. MF within the Ga container was computed using *Elmer* (finite elements), solving for the magnetic induction equation in terms of magnetic vector potential and the current density continuity equation (including the MF due to permanent magnets and currents due to melt flow) in terms of the scalar electric field potential.

OpenFOAM and *Elmer* were coupled via *Elmer-OpenFOAM* (EOF) library [154]. Simulations were performed on a computational cluster at the University of Latvia (UL) [152, 153].

b Image processing

The first (superseded by the current) version of the image processing pipeline was developed after the first round of experiments. Figure 2 shows the flowchart: neutron images were rid of bright and dark outliers (replacement with the neighborhood median), dark current and flat-field corrected, denoised using the self-snakes curvature flow (SSCF) filter [155, 156], then bubble shapes were extracted via double-Otsu morphological binarization [157] and Shen-Castan edge detection with morphological restoration subroutines for bubble interiors and edges, as well as segment position-based logical filters for outlier removal. An example of what the proposed procedure outputs can be seen in Figure 3. Bubble shape parameters were derived from best-fit ellipses [152, 153].

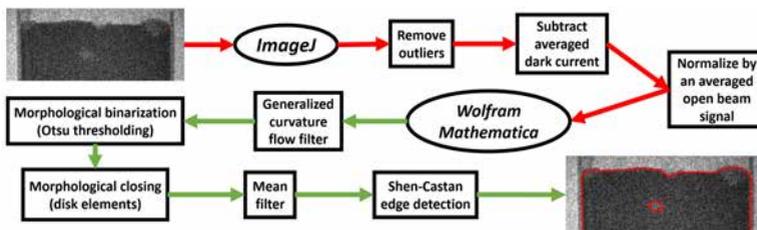


FIG. 2. The image processing pipeline, from raw images to extracted phase boundary shapes.

c Main results

Figure 4 clearly shows the effect of applying MF – bubble trajectory spread is considerably reduced. Stabilization is brought about by electric current due to gallium flow through applied HMF, where flow is induced via fluid displacement by ascending bubbles. The induced current then interacts with MF, resulting in the Lorentz force that acts to reduce the velocity component perpendicular to the applied field. This, in turn, leads to bubble wake laminarization, preventing tail vortex detachment stabilizing bubble trajectory. One can see that simulations and experiments are in very good qualitative agreement. The horizontal displacement of individual bubbles is also greatly reduced, as is evident from both Figure 4 and Figure

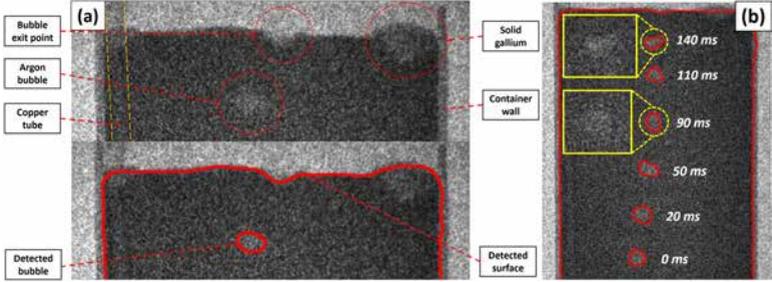


FIG. 3. (a) A sample pre-processed image with highlighted characteristic features and a post-processed image with detected air/Ga and Ar/Ga interfaces, and (b) a composition of several sequential flat-field and dark current-corrected neutron images overlaid with the detected container walls, air/Ga, and Ar/Ga phase boundaries. In (b) close-ups of bubbles and the background noise are highlighted with yellow frames.

5a. While the horizontal velocity component is reduced (Figure 5a), the vertical, conversely, is increased (Figure 5b).

These observations are in qualitative agreement with known experimental and numerical results obtained by other research groups for similar system dimensionless parameters, which indicates that there are no major issues in the experiment, simulations or the image processing algorithm. As expected from preliminary analysis, bubble shapes are slightly oscillating, and these shape perturbations are damped when HMF is applied - however, this requires further quantitative analysis. Previously conducted experiments indicate that slight vertical acceleration is expected at Eo values considered herein - this is observed as well.

However, in addition to bubble position and velocity, shape dynamics are also of interest. Regarding this, main results are presented in Figure 6, wherein bubble tilt angle dynamics averaged over 30 s of bubble flow recordings for $100 \text{ cm}^3/\text{min}$ flow rate are shown with and without applied HMF, as well as respective simulation output averaged over 1 s of flow time. For a 30 second neutron image recording, ~ 2000 -2500 bubble detection events over the entire usable FOV (8-10 cm height depending on the noise in the image sequence) contribute to the statistics. The tilt angle ϕ is defined as the angle between the major semi-axis of the best-fit ellipse for a bubble and the X-axis (the horizontal dimension of the images) — this is illustrated in Figure 6a.

In Figure 6a, notice that the ϕ magnitude in the case with no applied field has a maximum of roughly 40° for both experimental and simulated curves, as opposed to an initial maximum of 20° with applied MF, where

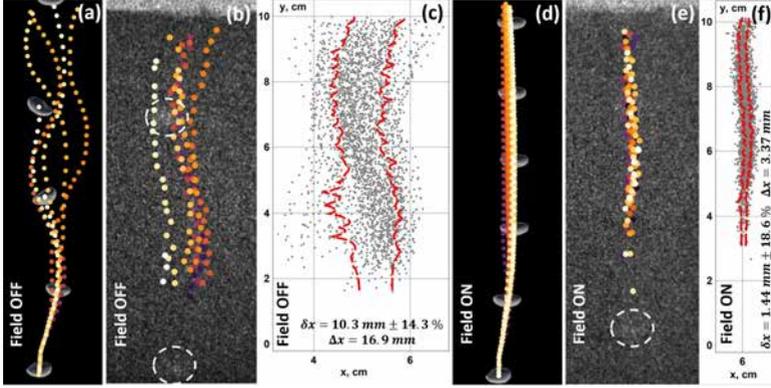


FIG. 4. Several initial bubble trajectories for a $100 \text{ cm}^3/\text{min}$ flow rate derived from simulations (a,d) and experiments (b,e). In cases (a,b) there is no MF, and in (d,e) the field ($\sim 0.3 \text{ T}$) is applied. Simulation frame rate matches that of the experiment. Bubble detection points are color coded by order of appearance, dark purple to white. Inlet and free surface are located right beneath and above vertical boundaries of images, respectively. Bubbles are highlighted in experimental images (b,e) by dashed white circles. In (c,f), entire sets of detected bubbles over all frames are shown, without (c) and with (f) HMF. Dashed red lines indicate bubble set envelopes, derived using the statistics-sensitive nonlinear iterative peak-clipping (SNIP) algorithm. δx and Δx in (c,f) stand for mean bubble set envelope thickness and maximum horizontal bubble spread, respectively. Distance scales in (c,f) are identical.

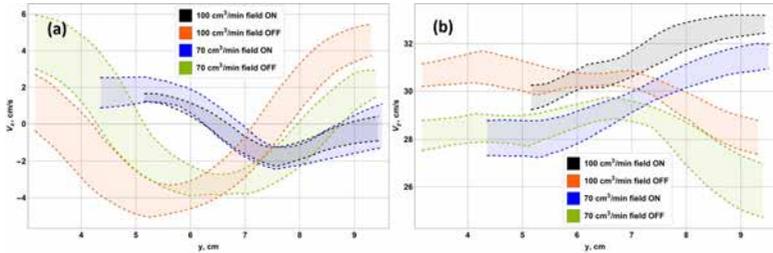


FIG. 5. Experimentally determined averaged (a) horizontal and (b) vertical velocity components of ascending bubbles at different elevations above the inlet, for different gas flow rates, without and with ($\sim 0.3 \text{ T}$) applied MF. Colored bands represent averaged curves plus their local errors.

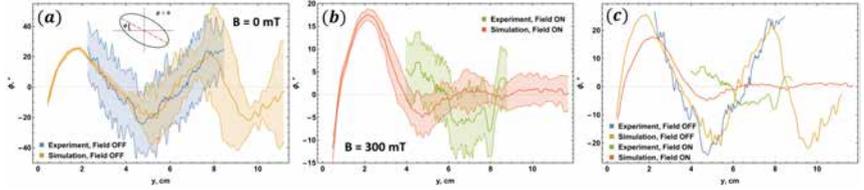


FIG. 6. Experimentally observed and simulated tilt angle and bubble elevation height correlations (a) without and (b) with applied MF, with color bands representing local errors for averaged curves; (c) shows a direct comparison of averaged tilt curves with and without applied field. Flow rate is at $100 \text{ cm}^3/\text{min}$. The positive tilt angle is defined as shown in (a).

one also observes subsequent near zero ϕ for CFD data and greatly reduced amplitude, but with more noise (lower SNR) for experimental data at higher elevations, as seen in Figure 6b – the difference is obvious from Figure 6c.

To interpret these results, consider that without applied MF and for low Re the wake of a solitary rising bubble is axisymmetric and stable [5, 9]. As Re increases, the symmetry is broken and then, as the pair of hairpin-shaped tail vortices on each side of the bubble (with respect to the major semi-axis plane) increases in size further due to do progressively more detached boundary layer, vortex shedding ensues [5, 9, 12, 13, 23]. This occurs even for stable bubble shapes, but is further enhanced by and in turn also affects surface perturbations and asphericity [12, 13, 23]. As a result of vortex shedding, non-zero net lift force is generated beneath the bubble in its wake, forcing it to deflect significantly, which then transitions to periodic lateral motion in alternating directions, as more vortices are shed and higher/lower pressure regions alternate beneath the bubble.

Bubbles with Re and EO ranges considered herein exhibit very unstable and complex turbulent wake structures enhanced by significant shape deformations [24, 95], which is well illustrated in [13][Figure 30]. Bubble paths are expected to be zigzags, with chaotic out-of-plane deviations due to shape instabilities and turbulent pulsations within the wake, accompanied by periodic bubble tilting [9, 13, 23].

This correlates very well with the results obtained without applied MF, as seen in Figure 6, where, in addition, numerical predictions match experimental findings very well within the radiography FOV for which the image processing pipeline reliably resolves phase boundaries. In Figure 6a, note that initially, the error band of the simulated curve is much narrower. During bubble formation at the horizontally directed (and thus asymmetric)

inlet tube [152], the bubble is stretched vertically by the buoyancy forces and laterally due to hydrodynamic gas pressure at the tube opening – then, once the bubble detaches, surface tension forces retract the lagging part of the bubble towards its center of mass, resulting in self-propulsion and rapid initial acceleration responsible for the starting negative ϕ values. This is illustrated in Figure 7.

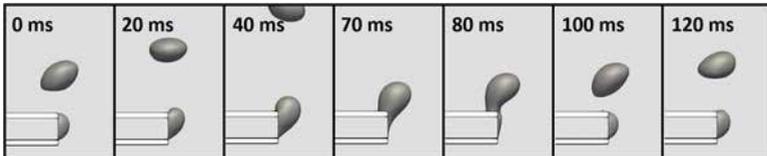


FIG. 7. Simulated bubble (dark gray) detachment at the inlet (white).

Then, violent vortex detachment occurs, resulting in rapid change of sign of ϕ , after which the bubble proceeds to ascend normally. The reason for the initially low dispersion of the ϕ values is that in this initial interval, the effects of self-propulsion and the first vortex detachment dominate over bubble/bubble interactions. For a chain of rising bubbles (here the characteristic bubble/bubble distance is ~ 3 -4 bubble diameters), as the elevation increases, continued vortex shedding and trailing bubble interactions with wakes from leading ones result in chaotic bubble deflections and tilt angle variations, explaining the increased width of the error band, which then coincides with that of the experimental results. However, this initially narrow trajectory corridor is below the data noise threshold.

When HMF is applied, depending on the N value, one observes the recovery of a rectilinear path [6, 10], which in the case of data shown in Figure 6b translates to an initial interval with significant deflection and ϕ , followed by ascension with greatly reduced ϕ . The initially higher ϕ is due to the same mechanisms as in the case of no field. However, note that in Figure 6b the maximum ϕ within the initial region is slightly less when MF is applied. This and the later low ϕ regime are observed because in liquid metal with high electric conductivity, such as Ga, flow about the rising bubble induces electric currents near the bubble and within its wake. These currents interact with the applied MF, resulting in the Lorentz force which opposes and damps fluid flow in the direction depending on the MF orientation (other velocity components are altered indirectly via the continuity constraint), as illustrated in Figure 8 [5]. In this case, magnetic damping is mostly due to the transverse field (TF) alignment (orthogonal to the bubble centroid velocity, Figure 8a), but one must also consider the effects of the partial longitudinal field (LF) alignment (Figure 8b), which

occurs during bubble formation and right after its detachment at the inlet tube (Figure 7).

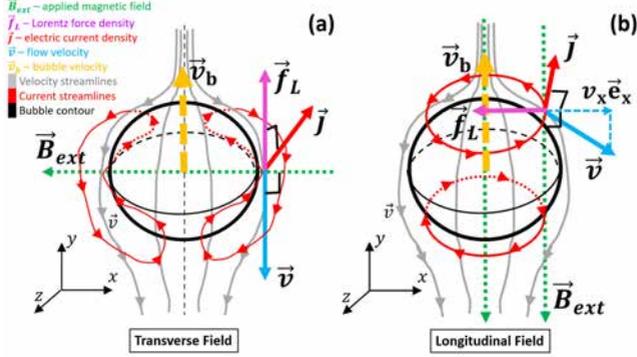


FIG. 8. Idealized representations of the effects of magnetic damping of flow near a bubble for (a) transverse and (b) longitudinal field alignment. Note that in (b) in the lower hemisphere the v_x sign is inverted and so is the direction of the Lorentz force density vector.

TF damping directly affects the velocity components orthogonal to the applied field, which here is the vertical component. Other components are, in turn, affected due to flow velocity field continuity. As seen in [13][Figure 30], applied TF results in several effects. At moderate flux density, the finer wake structure disappears, and vortices are elongated vertically. This can be attributed to flow velocity decrease leading to wake laminarization. Also note the decreased width of the bubble wake stemming from the vortex size diminished by TF damping. However, the wake is still asymmetric. At higher TF flux density, damping is intense enough to reduce the effective bubble Re value below the vortex shedding threshold, resulting in stabilized pairs of counter-rotating vortices and greatly decreased wake length, and thus also a reduced bubble/bubble interaction range. At this point, rectilinear trajectory is recovered, since lift force asymmetry is no longer sufficient to noticeably deflect the bubble. LF damping, on the other hand, directly affects the horizontal velocity component and other components are altered due to the continuity constraint. It has been demonstrated that this results in effects slightly different from the TF case, but also stabilizes bubble trajectory [7, 12].

Actual physics will clearly deviate from this idealized picture, but this simplified treatment may allow for interpretation of results. Qualitatively, it makes sense that during bubble formation and take-off from the inlet, due to its asymmetry, both LF and TF mechanisms are triggered, reducing

lateral deflection and tilt. As the flow around the bubble is damped and laminarized, the trajectory stabilizes and becomes nearly rectilinear (note also that the MF flux increases with elevation, as shown in Figure 1), with the magnetic TF (tilt magnitude is very small) damping maintaining a stable, symmetrized wake until the bubble reaches the surface. Despite minor discrepancies between the experimental and numerical results with applied MF (Figure 6b and 6c), it is evident that in both cases the ϕ magnitude is dramatically decreased and both the experiment and the numerical model indicate ϕ values that agree rather well, especially given the low image SNR.

2. The second round of experiments

The first round of experiments yielded very important insights into what image quality can be expected and what the image processing and bubble detection challenges are. While there were issues, it became abundantly clear that dynamic NR is feasible for bubble flow imaging through liquid metal. Therefore, the (Re, Eo, N) parameter space of a modified model system was sampled and an improved image processing methodology was developed and experimentally validated.

a Setup and imaging

A modified version of the model Ga/Ar system described in [152, 153] was designed for the new experiments. As before, a rectangular $150\text{ mm} \times 90\text{ mm} \times 30\text{ mm}$ glass vessel (boron-pure) filled with liquid gallium up to 130-140 mm mark was imaged. A resistive heating unit with a constant 4.13 W power at the vessel bottom kept Ga above its melting point throughout the experiments. A vertical copper gas inlet tube (1 mm diameter) was inserted 20 mm into the vessel through the bottom glass plate. As before, the gas flow rate was adjusted via a digital mass flow controller [158]. Neutron imaging of bubble flow was performed at the thermal neutron beamline NEUTRA [151] (SINQ, PSI, 20 mm aperture, $10^7\text{ n} \cdot \text{cm}^{-2} \cdot \text{s}^{-1} \cdot \text{mA}^{-1}$ flux) for gas flow rates in the 0-1200 sccm (standard cubic centimeters per minute) range without MF, as well as for {75, 125, 200, 265} mT HMF and {75, 125} mT vertical MF (VMF) in the bubble flow region. Neutron flux was parallel to the 30 mm dimension of the vessel. The distance between the liquid metal layer and the scintillator and was [4; 32] mm depending on the setup (MF system used, if any). A square FOV (123.125 mm sides) above the inlet was imaged at 100 FPS [158]. Static MF was generated by tailored permanent magnet/iron yoke systems assembled at the Institute

of Physics in Salaspils. The MF systems were designed and optimized in *COMSOL Multiphysics* such that, while the overall MF magnitude changes between the systems as specified above, the MF distributions are very similar (Figure 1). For HMF, the field distributions were designed to be close to the HMF counterparts.

Additionally, reference experiments were performed at the cold neutron beamline ICON [159] (SINQ, PSI, 20 mm aperture, $\sim 1.3 \times$ NEUTRA flux) to validate the improved developed image processing methodology. A brass reference body (stationary and moving) with a spherical cavity (5 mm radius) was imaged to reproduce imaging conditions like those for argon bubbles in liquid gallium [158]. Finally, reference neutron imaging with 40- and 80 mm apertures ($4 \times$ and $\sim 11 \times 20$ -mm flux, respectively) was performed to obtain reference images and see how the image processing code performs for different SNR. Also, greater flux was leveraged for FPS up to 600, which in the future should aid the in-depth analysis of bubble shape variations via image-based surface strain rate measurements.

b Image processing

The acquired images are 16-bit 1-channel TIFFs with a 1024×1024 pixel resolution (2×2 average-binned 2048×2048 frames) with a 0.11-0.12 mm pixel size (0.12 mm for the reference experiments). Short exposure times (10 ms) result in strong Poisson (multiplicative) noise from neutrons and converted photons, and salt-and-pepper noise of varying density is present due to overexposed (gamma ray noise) or "dead" camera pixels. The neutron beam flux over the FOV is non-uniform, with a fall-off near the edges of the acquired images.

Figure 9a is an example of an acquired raw image. Here gas flow rate was 120 sccm (standard cubic centimeters per minute) and static HMF of 125 mT was applied to the bubble chain region. Note the bubble regions highlighted in Figures 9a and b – these are the objects of interest that must be segmented and their properties such as centroids, projection areas, tilt angles, aspect ratios, etc. measured. Note that the same color scheme and normalized luminance scale as in Figure 9b are used in all further figures in this section, unless stated otherwise. The image processing challenge with examples of image noise and neutron flux transmission signal for bubbles is further illustrated with examples in [158].

As in the previous neutron imaging campaign, estimates indicate that Eötvös number is $Eu \in [2.1; 4.1]$ and the Reynolds number near the bubbles is within $Re \in [10^3; 10^4]$. This corresponds to a flow regime wherein bubbles have oscillating elliptic shapes [24, 95, 152, 153]. Bubbles have equivalent diameters of $d_b \in [6; 8]$ mm and travel with varying velocities, usually in the

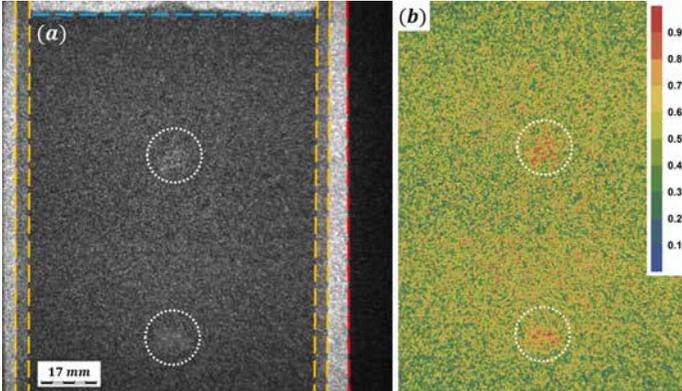


FIG. 9. (a) Original captured FOV after outlier removal and luminance normalization with marked container walls (orange dashed lines), the mean gallium free surface level (light blue), the neutron flux shielding (red, borated polyethylene) and bubble locations within the FOV (white). Note the scale bar in the bottom-left corner. (b) FOV in false color (color bar on the right) after cropping to the container walls and the metal free surface, dark current and flat field corrections, and normalization.

[20; 40] cm/s range [152]. Such velocities and dynamic phase boundaries dictate the high acquisition frame rate required for physical analysis. 100 FPS was considered the optimal trade-off frame rate for the experiments covered herein, in that one already avoids significant motion blur about the bubbles while still maintaining manageable image SNR. Bubble collisions, coalescence or breakup do not occur for image sequences considered here for methodology demonstration, but rather chains with closely packed bubbles manifest for higher flow rates. In the future, the methodology developed herein will also be applied for even higher flow rate cases where bubble collisions do take place. Here, the focus is on capturing the bubble shape variations and positions.

The results obtained via the previous version of the image processing code [152, 153] and shown in Section II A 1 c, as well as the above considerations, determined the objectives that must be met here: improved bubble edge detection stability – more reliable detection, fewer artifacts and false positives; greater bubble shape detection precision; increased bubble detection rates at the bottom of the FOV. The following assumptions regarding the bubbles are in effect: bubbles have perturbed elliptic/circular shapes, i.e. not necessarily convex; bubble phase boundaries do not exhibit local curvature radii less than a pre-defined fraction of their mean curvature,

i.e. smoothness is assumed below a certain length scale due to surface tension; no coalescence or breakup of bubbles is expected at flow rates < 500 *scm* considered herein. Given the image properties/quality, it is not attempted to perform filtering & segmentation such that argon/gallium volume fraction can be recovered, since low SNR will translate to significant errors in the output if both volume fraction field and shape recovery are attempted. Since the 100 FPS frame rate is the lower limit for prevention of considerable motion blur, spatio-temporal noise filtering is not used – only spatial denoising is applied. Furthermore, given the above and to make the methods more general, image sequence frames are treated *separately* at all stages of image processing. As such, it is possible to implement parallelization on the frame level, with simultaneous multi-threading enabled to speed up the processing.

To achieve the above goals, it was decided to separate the bubble segmentation into two stages: *global* and *local* filtering/segmentation. That is, one starts by obtaining first estimates for bubble segments within images (the entire cropped FOV), and then uses a separate routine for local filtering/segmentation about the preliminary segments to improve shape detection precision at local scales (i.e. lower-wavelength corrections to first shape estimates) and to resolve false positives.

The global filtering/segmentation routine, which is essentially a completely overhauled version of the previous image processing approach [152, 153], was designed and adjusted to maximize bubble detection rates at the cost of reduced shape detection precision and higher false positive rates. Further, to improve edge detection stability, *implicit* edge detection was opted for. Global noise filtering is performed in multiple stages, each targeting a certain noise type and/or wavelength range. As for the local filtering/segmentation routine, a recursive multiscale analysis algorithm was implemented that is shown to perform well even for images with especially low SNR and recovers bubble shapes in cases where the first segment estimation outright fails to capture initial bubbles shapes within an acceptable margin of error. The overall structure of the proposed image processing solution is outlined in Algorithm 1.

Image pre-processing, which includes dark current correction (DCC) and flat-field correction (FFC) is performed in *ImageJ* as described in [158]. Images are cropped as indicated in Figures 9a and b. FFC of images results in the spatial dependence of the SNR as a consequence of fewer neutrons detected behind the sample compared to the open beam. After pre-processing, the images are passed to *Wolfram Mathematica* for further processing [158].

First segment estimates are obtained using an algorithm referred to as the *global filter*, which is outlined in Algorithm 2. Segment estimation is

Algorithm 1: Overall structure of the new image processing pipeline

Input: Raw image sequence

- 1 Pre-process images [158]
- 2 Perform *global* filtering for all images (Algorithm 2)
- 3 Process images with the *multiscale recursive interrogation filter* (MRIF, Algorithm 3)
- 4 Apply the luminance map-based false positive filter to the MRIF output (Algorithm 5)
- 5 Apply logical filters

Output: Centroids and shape parameters for bubbles detected in each image

performed in three stages – noise filtering & background removal, implicit edge detection and segment filling & cleanup, all of which are covered in detail in [158]. Noteworthy is the soft color tone map (CTM) masking (SCTMM) which is a non-linear filter designed to clean up the image background by removing large-scale artifacts left over after denoising and to further separate background from bubbles while avoiding excessive erosion of the bubble regions. The large-scale structures in the background were actually one of the sources of the edge detection instability in the previous approach, especially for low-CNR images ($\text{CNR} = |S_b - S_0|/\sigma_n$; S_0 – background signal intensity) with higher bubble number density where bubble detection was often outright impossible due to formed edge artifacts that could not be reliably removed.

Given a normalized original image x , the SCTMM background correction generates a new image y :

$$y = x * \underbrace{\left(x - \underbrace{(1 - \text{CTM}(x, c))}_{\text{Soft thresholding}} \right)}_{\text{Soft background mask}} \quad (1)$$

where $\text{CTM}(x, c)$ is the CTM operation and c is the luminance compression factor. The CTM operation maps the colors (in this case the gray-scale values) of the image using gamma compression with a global compression factor c [167]. The idea behind SCTMM (1) is as follows. A pure $x * x$ product would have the effect of non-uniformly increasing the distances between the nearest luminance values (input x is normalized) and luminance maxima and minima values would be much more distant from the mid-range

Algorithm 2: Global first segment estimator for pre-processed
FOV images

Input: Normalized pre-processed images

Noise filtering & background removal

- 1 Regularized Perona-Malik (PM) filter [160–162]
- 2 Total variation (TV) filtering, Poisson model [163, 164]
- 3 Self-snakes curvature flow (SSCF) filter; *memoize output* [155, 156]
- 4 Soft color tone map masking (SCTMM) (1); *memoize output*

Implicit edge detection

- 5 Compute the luminance using the gradient filter (Gaussian regularization & Bessel derivative kernel)
- 6 2-threshold hysteresis binarization [157]
- 7 (Optional) Morphological erosion [165]
- 8 Thinning transform [157]

Segment filling & cleanup

- 9 Filling transform [157]
- 10 Mean filtering (small radius)
- 11 Otsu binarization [166]
- 12 Remove border components

Output:

- Image mask with first segment estimates
 - Memoized SSCF-filtered image for later use in Algorithm 3
 - Memoized SCTMM-filtered image for later use in Algorithm 5
-

luminance values, which are affected the most. If one masks or lowers the values of certain pixels within one of the images x' , then $x * x'$ would act as a "soft" i.e. weighed mask (as opposed to a "hard" binary mask) that would shift the pixels with reduced values in x' further towards the lower end of the luminance range, ideally making them background. Soft masking is preferred here because masking using binarization and then replacing the removed background using luminance interpolation or other methods will generally produce artificial and potentially very pronounced edges and/or reduce the contrast of actual edges.

Here it is required that x' is such that background and post-filtering artifacts are removed, and bubble contrast is enhanced while bubble features are eroded as little as possible. It was decided to opt for additive masking of the form $x'(x) = x - mask$. An inverted CTM(x, c) was chosen as $mask$ because, if the right c value is set, $1 - CTM(x, c)$ will have high luminance for background and denoising artifacts since CTM(x, c) reduces the difference in their luminance values. This way $x' = x - (1 - CTM(x, c))$ has, conversely, greatly reduced luminance for artifacts and background. The resulting product (1) then has the desired properties and emphasizes

the bubbles while reducing the impact of image artifacts.

The filtering methods utilized in Algorithm 2 are rather aggressive. Testing revealed that, while bubble detection rates are indeed significantly higher than before and bubbles are detected everywhere within the FOV (mostly due to better noise modelling, SCTMM and implicit edge detection), it comes at the cost of decreased shape resolution precision and a higher false positive rate. The former is very important for a more in-depth analysis of the effects of varying flow rate and MF on the behavior of bubble chains. It was decided to fine-tune the global filter such that the bubble detection rates are maximized and good first estimates of bubble shapes/sizes are obtained, and complement it with a routine that would use the first estimates to generate more precise bubble shapes and efficiently filter out false positives. To this end, an original algorithm for iterative segmentation refinement have been developed – the *multiscale recursive interrogation filter* (MRIF) outlined in Algorithm 3 and schematically illustrated in [158].

The key idea is to define *interrogation windows* (IWs) about the initially detected bubbles to exclude irrelevant parts of the image and its intensity histogram from the analysis. This also helps to reduce the influence of any remaining artifacts left over from the global filtering. Especially for images with lower SNR and with closely packed bubbles, it may be the case that the initial segmentation is poor, i.e. two or more bubbles have been segmented as one due to the surrounding artifacts, or a bubble was connected to large-scale artifact structures, forming a large segment that obscures the true object. This means that an appropriate local filtering algorithm must be devised for IWs. However, a single local filtering pass may not be enough for various reasons, e.g. what might appear visually as a poorly segmented bubble at the scale of the current IW might actually turn out to be, after local filtering, multiple bubbles – these would then each require another pass at a finer scale. This means that in general, a series of consecutive interrogation passes could take place. Thus, MRIF performs object filtering at different scales, effectively checking that the segments have been properly resolved by the global filter and/or in the preceding local filtering iterations. A stopping criterion based on the IW size similarity between iterations makes sure that MRIF recognizes that it only makes sense to re-filter an image patch if the object is significantly smaller than the previous IW, since in this case the finer shape features may have been under-resolved.

MRIF consists of several components: an IW generator that centers the IW at the segment location and adjusts its size according to the segment size; a local filter that is responsible for filtering within IWs; a recursive routine that performs a "scale descent" and converges to the "true" segment

Algorithm 3: Multiscale recursive interrogation filter (MRIF)

Input:

- An image mask with first segment estimates (Algorithm 2, Step 12)
 - The global SSCF filter output (Algorithm 2, Step 3)
- 1 Define square (side length L') interrogation windows (IWs) for segments based on their areas and centroids
 - 2 Define the FOV image as an IW with scale L based on FOV dimensions
 - 3 *For every initial segment and updated segments:*
while $L'/L < \varepsilon$, $\varepsilon > 1$ (*user-defined*) **do**
 - Map the global SSCF filter output onto the segment IW
 - Perform local filtering (Algorithm 4)
 - if** *updated segments were found* **then**
 - Define new IWs (side length L') for the updated segments based on their areas and centroids
 - Redefine preceding IW scales as $L' \rightarrow L$
 - else**
 - Break****end**
- 4 *Memoize* converged IWs for all segments
 - 5 Map the centroid coordinates of the resulting segments from converged IWs onto the original image
 - 6 Map the updated segment masks onto the original image and build an updated global mask

Output:

- Updated bubble shape masks for the FOV
 - Converged IWs for later use in Algorithm 5
-

scale starting from the first estimates; a procedure that collects the final, updates segments and maps them onto the original bubble segment mask, substituting the first estimates. The ε criterion is the stopping factor that controls recursion depth, i.e. the lower IW scale threshold. IW generation is explained in detail in [158].

The local filter used for recursive filtering in MRIF (Algorithm 3, Step 3) has elements similar to the global filter, but with important distinctions. It was designed to be less aggressive because MRIF ensures that only the crucial background context from the image is retained within an IW, e.g. the somewhat destructive SCTMM is not required. The local filter operates on an IW as described in Algorithm 4 and its inner workings are explained in [158].

After the MRIF process is finished, the centroids of the detected bubbles are mapped from converged IWs onto original images (full FOV). The

Algorithm 4: Local segment refinement for IWs

Input: Global SSCF filter output (Algorithm 2) mapped onto an IW

- 1 Mean filtering
Implicit edge detection
 - 2 Compute edge halos using the gradient filter (Gaussian regularization + Bessel derivative kernel)
 - 3 Chan-Vese binarization [168]
 - 4 (Optional) Morphological erosion
 - 5 Thinning transform
Segment filling & cleanup
 - 6 Filling transform
 - 7 Mean filtering (small radius)
 - 8 Otsu binarization
 - 9 Remove border components
- Output:** Updated local shape estimate
-

mapping accounts for IWs possibly being out of image bounds. The mapping procedure is described in [158]. MRIF, in addition to its intrinsic false positive filtering capacity stemming from the recursive multiscale analysis, is supplemented by a dedicated luminance-based false positive elimination (LFPE) that processes the segments output by MRIF. The procedure is outlined in Algorithm 5.

Algorithm 5: Luminance map-based false positive filtering for MRIF output

Input:

- Updated bubble shape masks for the FOV (Algorithm 3)
 - Global SCTMM filter output (Algorithm 2, Step 4)
 - Converged IWs (Algorithm 3)
- 1 Map the global SCTMM filter output onto converged IWs
 - 2 Apply the SCTMM correction (1)
 - 3 Multiply by the segment binary mask
 - 4 Normalize the image
 - 5 Compute $\langle I \rangle \cdot \max(I)$ for all segment regions
 - 6 *Thresholding for all segments:*
if $\langle I \rangle \cdot \max(I) < \eta$; $\eta \in [0; 1]$ (*user-defined*) **then**
| Flag the segment as a false positive
else
| Nothing
end
 - 7 Remove the identified false positives from the bubble shape mask
- Output:** FOV bubble shape masks without detected false positives
-

Algorithm 5 exploits the observation that SCTMM applied locally to the global SSCF filter output mapped onto converged IWs should produce very strong intensity maxima and an overall higher intensity in the segment region in the case of a true bubble, while the opposite should hold for false positives. The product of mean and maximum intensity is used so that neither of the two criteria alone are enough to pass the filter, since it might be the case that a region with an otherwise background level intensity might exhibit a tightly localized intensity maximum; similarly, mean intensity filtering alone is not enough since a bubble should have a strong maximum of transmission intensity about its centroid, which in itself is not as strongly correlated to the mean intensity. Another way to interpret this is that, if the maximum and mean thresholding have certain probabilities of accepting a false positive, then the max/mean product thresholding has a false positive acceptance probability at least lower than the greater of the two components.

After filtering the MRIF output using Algorithm 5, all properties of interest are measured for all remaining bubble segments and logical filters can be applied for further false positive elimination. In this case, logical filters check for implausible bubble coordinates, sizes and aspect ratios and remove the outliers from the dataset of measured bubble shapes. Finally, the resulting data can be post-processed and interpreted.

c Main results

Figure 10 shows the effect of subsequent operations that the global filtering routine (Algorithm 2) performs for a pre-processed image. The difference between Figures 10a and 10b is that the noise with low wavelengths (sharply localized luminance maxima and minima left over from pre-processing) have been eliminated. Next, the TV filter (Figure 10c) consolidates the high luminance values within the bubble regions (white dashed circles), increasing the SNR for the bubbles. Also, noise is significantly damped and the wavelength of its features is increased even more. However, the TV filter does not remove the sparse larger-scale luminance maxima still seen in Figure 10c in the background between the bubbles as efficiently as it is desired without degrading the CNR for the bubbles. This function is performed by SSCF (Figure 10c) which specifically diffuses the leftover intensity maxima, increasing bubble SNR (and CNR due to reduced noise in the background) further. A more detailed analysis of how noise filtering stages affect the image is provided in [158]. With SNR increased by the sequence of PM, TV and SSCF filters, SCTMM is now applied to increase the CNR by significantly flattening the background while preserving bubble signal intensity. This enables the gradient filter

to produce bubble edge halos with an even greater CNR, enabling clean segmentation, as seen in Figure 10g.

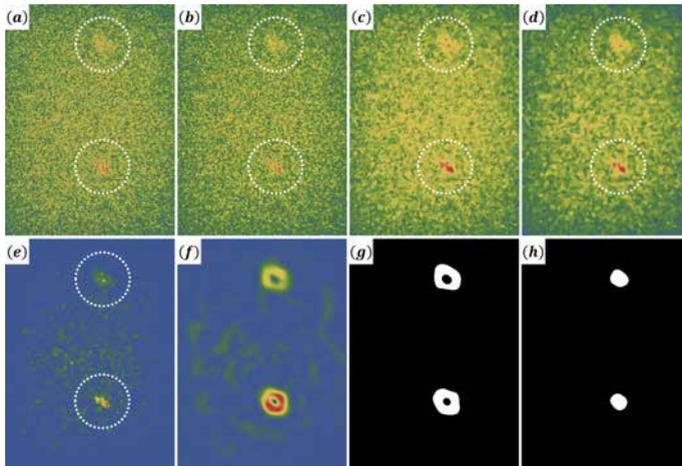


FIG. 10. First segment estimation stages (Algorithm 2): (a) original pre-processed image; (b) PM-filtered image; (c) Poisson TV-filtered image; (d) SSCF-filtered image; (e) SCTMM output; (f) gradient filter output (edge halos); (g) hysteresis-segmented edge halos (prior to erosion); (h) bubble masks obtained after erosion, thinning, filling, small-radius mean filtering, Otsu binarization and border component removal. The color scheme is as in Figure 9.

The first segment estimator could have been enough for bubble shape extraction, when tuned appropriately, if not for the fact that the image considered in Figure 10 is one of the better examples in terms of noise and artifacts present, i.e. a considerable portion of the images captured during the experiments are of a much poorer quality. Not only are the obtained shapes often imprecise or deformed, they can at times be decidedly non-physical – one such example is show in Figure 11 where in (a) a segment is shown that looks like two bubbles that are in the process of merging. Such events are not expected at the flow rate for which this image was acquired, therefore Figure 11a shows an obvious artifact. The bottom-right corner of (b) contains closely packed high-luminance spots, which likely have been combined and merged with the bubble region in the upper-left corner. However, once MRIF targets the segment and the local filter is applied to the SSCF output projected onto the segment IW (c) in stages (d-h), the artifact is no longer present and a single bubble is correctly resolved.

In addition to such artifacts, since the global filter was tuned to maximize

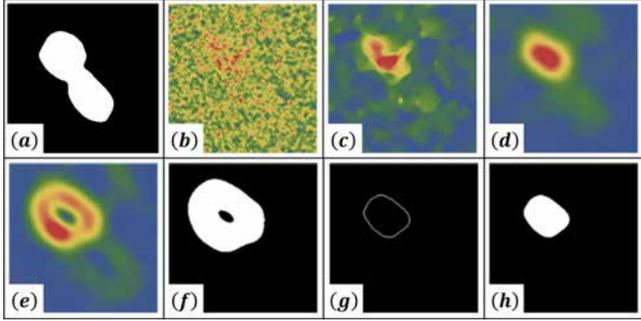


FIG. 11. Updating a first segment estimate using the local filter (Algorithm 4): (a) first estimate in an IW; (b) original pre-processed image projected onto the IW; (c) SSCF output projection onto the IW (Algorithm 2); (d) mean-filtered SSCF output; (e) gradient filter output; (f) Chan-Vese segmentation output; (g) edges extracted via erosion and thinning; (h) updated local segment for the IW. Segment filling and cleanup are similar to what is done in Algorithm 2.

the odds of detecting bubbles in the FOV, there are cases where detected segments are false positives. Two instances of such segments interrogated by MRIF are shown in Figure 12. In (a) one can see that the local filter has revealed that there is indeed no segment contained within the IW. However, interrogating the false positives in an IW might on occasion produce segments yet again, as in (b), where the gradient filter stage (b4) generated a structure that resembles a bubble edge halo. It was then segmented and, through edge cleanup, transformed into a segment that seems eligible – but simply overlaying it over the original image projected onto the IW, one can see that this is not the case. However, MRIF effectively performs a two-factor false-positive check, and in such cases the luminance map-based filter (Algorithm 5) serves as a backup. Once the global SCTMM output is projected onto the IW and SCTMM is applied to the resulting image (b9), one can see that the segment overlay contains only background, and thus this false positive will be eliminated by LFPE since it has $\langle I \rangle \cdot \max(I) < \eta$ (I is the local image luminance).

An example containing several instances of false positives, under-resolved shapes and bubble regions merged with noise patterns is shown in Figure 13. Notice that the image quality, even visually, is much worse than in Figure 10. The artifacts in the upper part of the image stem from lower CNR in (b), whereas one of the bottom artifacts comes from a noise structure in the background that resembles a bubble. However, as seen in (d) and (e), MRIF successfully removes all false-positives and artifacts while improving

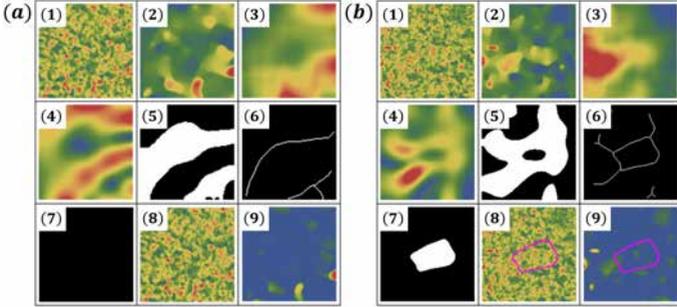


FIG. 12. Instances of false positives revealed by MRIF with (a) no local segments found and (b) an artifact (purple outline) that will be eliminated later by Algorithm 5 based on the segment luminance map. Sub-figures (1-9) in (a) and (b) are: (1) is the original image projected onto an IW; (2-7) are the respective local filtering stages (Figure 11c-h); (8) is (1) with detected edge overlays; (9) is SCTMM applied to (2).

the shape estimates.

To see how MRIF iteratively resolves cases like the one shown above, consider Figure 14 where the updates for the first segment estimates are shown for MRIF iterations. One can see in (e3) that the SNR and CNR are even worse than in the case shown in Figure 13. The largest segment seen in (e1) and (a1) is first resolved into two bubbles (b1) and then each bubble is interrogated once more, obtaining more precise shapes. The bottom-most segment in (e1) requires the most MRIF iterations – the first two, (b3) and (c3), remove portions of the artifact that had obscured the bubble, and the last iteration (d) updates the resolved shape. The resulting bubble shapes are then mapped onto the FOV as indicated in (b-e). Here, it is important to reiterate that the performance of MRIF strongly depends on the user-defined IW scaling factor and the critical IW length scale ratio ε for the latest and the (potential) next recursion iterations. A strategy for determining the optimal parameters is discussed in [158].

Aside from the shown performance examples, it is also of interest how the developed approach performs for entire image sequences in terms of bubble detection density in the FOV and the physicality of obtained results, i.e. bubble trajectory and shape properties. Figures 15-17 demonstrate the differences in performance for the preceding image processing pipeline [152, 153] and the methodology presented here [158].

One can clearly see in Figure 15 that the new version of the image processing code outperforms the previous version by completely avoiding the blind zones in the lower part of the FOV for both image sequences.

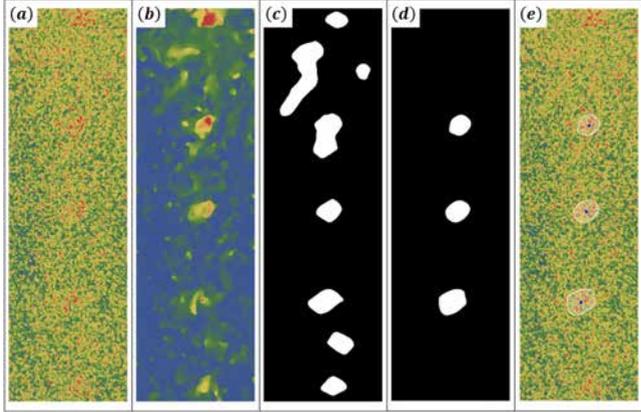


FIG. 13. An example of segmentation improvement by MRIF: (a) horizontally cropped pre-processed image – note the very low CNR in the upper half of the image; (b) global SCTMM filtering result; (c) first segment estimates; (d) updated segments output by MRIF; (e) output segment overlays for (a). Note that the top- and bottom-most segments from (c) are not present in (d) and (e) – this is the correct behavior, since one can visually see in (a) and (b) that the corresponding bubbles are partially outside the FOV, and thus are not eligible for analysis.

Notice also that bubble tracks visible in (b) are much more coherent than in (a). Figure 16, in turn, shows that with the new methods one can now clearly resolve the classical S-shaped mean trajectory cluster formed by zigzag trajectories, as seen in (c) and (d), as opposed to (a) and (b) where a significant portion of the events is missing. The deflection bias in the $x > 0$ direction in (c) and (d) is determined by the horizontal inlet releasing gas in that direction.

Another point of interest are the tilt angle dynamics resolved in [153] versus the current results – this is showcased in Figure 17. First, as a consequence of the blind zone elimination, the new curves extend all the way through the FOV. Second, the average trends yielded by both approaches indicate that the previously used code indeed resolved the dynamics without unacceptable inaccuracy. Finally, the error bands are considerably narrower about the averaged curves for the present results. The latter is especially true for the case with applied MF shown in (b) where the SNR was much lower than in the image sequence corresponding to (a). This indicates that the new approach indeed yields significant improvement not just in bubble detection, but also in shape boundary resolution. The experimentally

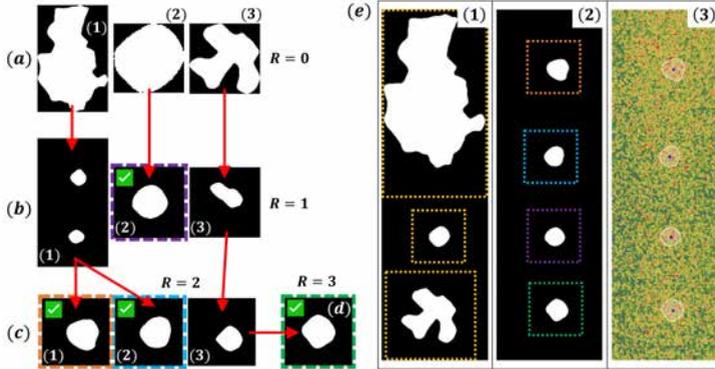


FIG. 14. An illustration of the iterative interrogation process for an image with one of the worst overall SNR values, where R is the MRIF recursion depth. Sub-figures (a-d) show R running 0 through 3 for the segments detected at each depth. Sub-figure (e) displays (1) the initial segments, (2) MRIF output and (3) output overlays for the original image. Converged segments are flagged with green ticked boxes. Colored frames in (e2) correspond to (b2), (c1-2) and (d) via respective colors. Orange frames in (e1) indicate the IWs at $R = 1$: note the high aspect ratio of the topmost IW – its virtual counterpart (see [158]) is significantly out-of-bounds and therefore considerable crop correction is assigned to properly map (c1-2) onto the FOV.

obtained results in [153] were in a rather good agreement with performed simulations, meaning one thus has an indirect yet significant validation of the presented approach.

The developed approach was also applied to the newly acquired data to ensure consistency in the code output across different experimental campaigns – one instance of the new results is shown in Figure 18. Again, the bubbles are resolved over the entire FOV for all three cases shown. An in-depth physical analysis of the bubble dynamics is beyond the scope of this thesis and is reserved for follow-up articles.

While one cannot check how many bubbles the image processing code actually failed to detect (without manual inspection), one can evaluate the amount of detection events that are ruled out as false positives at the various stages of code execution for a sequence of images. The results for the five image sequences considered above (Figures 15 and 18) are presented in Table I. Notice that in the most difficult case of the five (Figure 15d) most of the work is done by the LFPE filter and the object property filter. However, the intrinsic filtering capacity of MRIF is significant because it filters out the detection events that very likely would have passed both of

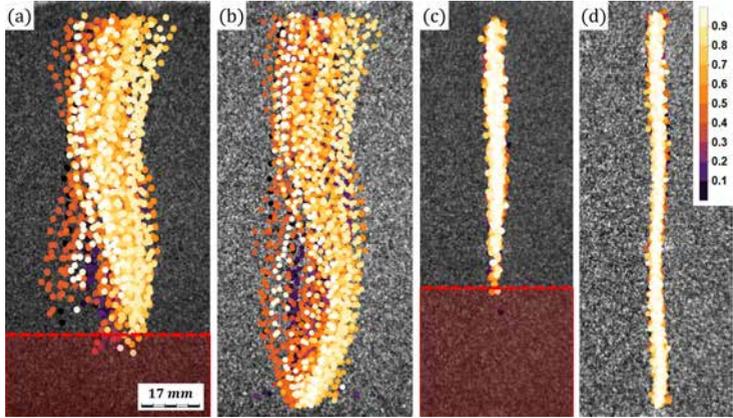


FIG. 15. The locations of bubbles detected in the FOV for a sequence of 3000 images (30 seconds) at a 100 *sccm* flow rate for the model system from [152, 153]: no applied MF, (a) previous and (b) current image processing code; applied ~ 265 *mT* HMF, (c) previous and (d) current image processing code. Bubble locations are marked with dots, color-coded in the chronological order of appearance. Note the color legend in (d): imaging starts at 0 and ends at 1. The red-tinted areas indicate the blind zones of the previously used image processing code. All sub-figures are to scale.

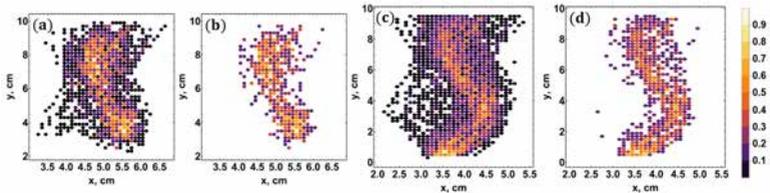


FIG. 16. Normalized bubble detection density histograms with $(dx, dy) = (2, 4)$ *mm* bins color-coded by detection counts: (a-b) the case in Figure 15a with (a) all bins and (b) bins with 3+ detections shown; (c-d) the case in Figure 15b with (c) all bins and (d) bins with 4+ detections shown. Note the color legend to the right of (d).

the two following stages.

As for verifying experimentally that the above methods indeed work correctly, the developed image processing algorithm is first validated by applying it to the images of a stationary reference body. Three imaging cases are considered here: neutron flux transmission through the shorter

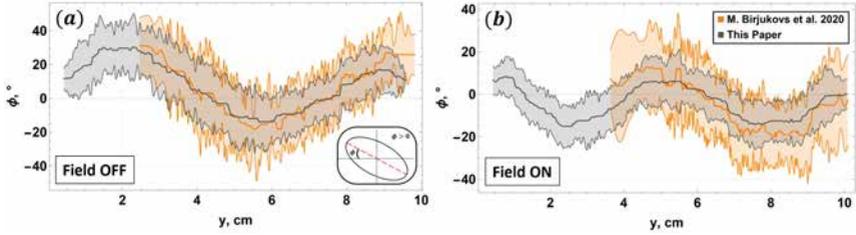


FIG. 17. Bubble tilt angle versus elevation (averaged curves and error bands) over the FOV bottom for (a) the case with no applied MF (Figures 15a and b) and (b) applied horizontal ~ 265 mT MF (Figures 15c and d), both cases at a 100 sccm flow rate. Orange indicates the previous results [153] and the current results are shown in gray. The tilt angle definition is shown in the bottom-right corner of (a).

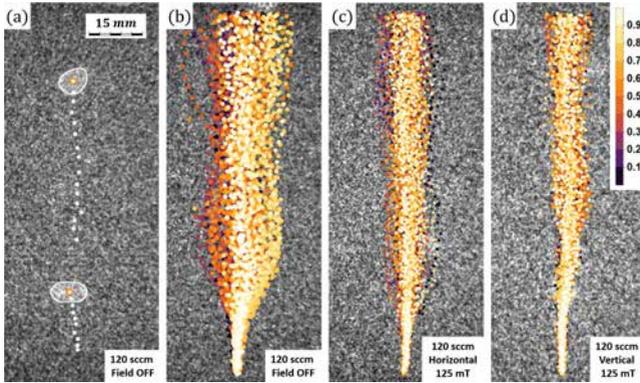


FIG. 18. The locations of bubbles detected in the FOV for a sequence of 3000 images (30 seconds) at a 120 sccm flow rate for the new model system: (a) an example of detected bubbles: white contours are shapes, orange dots are the current positions and white dots are the preceding detections; (b-d) all detected bubble positions with (b) no applied MF, (c) horizontal ~ 125 mT MF and (d) vertical ~ 125 mT MF. Bubble locations in (b-d) are marked as in Figure 15. All sub-figures are to scale.

body axis, the longer axis, and the latter with an extra distance from the body to the scintillator. Thus, the SNR of the neutron-transparent spherical cavity within the body progressively decreases for these cases. This is illustrated in Figure 19.

A neutron radiography image of the reference body (slightly inclined)

Image sequence	MRIF	Algorithm 5	Property filters	Total	Remainder
Figure 15b	1.75	2.63	1.21	5.49	94.5
Figure 15d	1.51	4.63	6.80	12.5	87.5
Figure 18b	1.53	2.05	1.25	4.76	95.2
Figure 18c	0.71	0.51	0.18	1.39	98.6
Figure 18d	1.62	0.89	0.32	2.81	97.2

TABLE I. False positive elimination rates (%) over three stages: MRIF, the luminance map-based filter (Algorithm 5) and the object property filter (with respect to the input that each filter received), and the percentage of detection events input to MRIF that were eliminated in total.

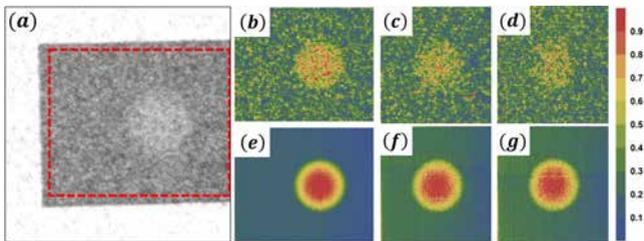


FIG. 19. Static reference body: (a) an example radiograph; (b-c) normalized images cropped as in (a) for neutron flux transmission through the shorter axis, the longer axis, and the latter with an extra distance to the scintillator, respectively; (e-g) normalized mean luminance maps for the respective image sequences. Note the color bar to the right.

is shown in (a) where one can see the rectangular brass frame (darker) and a circular projection of the spherical void (brighter), as well as the surrounding background due to air. Neutron transmission in the case of (a) is along the shortest of the body axes. In all three imaging cases, the images are cropped as indicated in (a). Examples of cropped images of the spherical void within the body with exposure very similar ($\sim 1.3\times$) to that for the bubble images (100 FPS, $\sim 1.3\times$ neutron flux) are shown in (b-d) for the three cases listed above, respectively. The corresponding mean luminance maps shown in (e-g) were obtained by averaging over entire recorded sequences – these averaged images are used to obtain reference shapes. The shapes detected by the image processing code in images like (b-d) are then compared to the reference shapes to compute shape detection error metrics. Note that all images seen in Figure 19 and used for validation are obtained from raw images by pre-processing, as with bubble flow images. Note that the images shown in Figure 19 have the image side length to

sphere diameter ratios that are very similar to what one has for MRIF IWs.

Reference shapes from the images in Figure 19e-g are obtained as explained in [158]. To ensure fair verification, both the global (Algorithm 2) and the local (Algorithm 4) filters are applied to the reference void images with parameters identical to those used for bubble images. The principles behind imaging a moving reference body are as above, except the body is now attached to a pendulum that periodically swings, and thus the body travels back and forth through the FOV, as illustrated in [158].

The motion is mostly horizontal and is initially strongly damped until the pendulum amplitude reaches a state where its oscillations exhibit a very slow decay. This enables us to determine the dynamics of the error metrics and assess the effects of motion blur as the body and the void within it decelerate. Here, before the global and local filters can be applied to the cropped reference body images, one must first segment the body within the FOV, crop the masked image to an IW about the mask centroid, repair the IW images as necessary, and then apply the filters. This procedure is somewhat more involved than the one for the stationary reference body, and is outlined and demonstrated in [158]. Afterward, the resulting repaired regions are used as input for the global and local filters. The settings for the global and local filters are the same as in the cases with stationary reference bodies.

The finer details of experimental verification are beyond the scope and scale of this summary, but are provided in [158], where the error metrics and their statistics, as well as the effects of motion blur and synthetic artifacts (worst-case scenario tests) are explored. To summarize here, it was found that the shape detection errors were within acceptable margins even for high velocities, and the local filtering did in fact provide significant improvement over the first estimates. The observed errors were also consistent with imaging conditions, and detection failure rates for the reference bodies were rather negligible.

3. Summary

An appropriate experimental setup has been developed and neutron imaging of bubble flow in a 30 mm thick model system has been performed for gas flow rates in the 0-1200 *sccm* range without MF, as well as for {75, 125, 200, 265} *mT* HMF and {75, 125} *mT* VMF in the bubble flow region. This covers a large and industrially relevant region of the (Re, Eo, N) parameter space of the studied downscaled model system, and constitutes a unique and comprehensive dataset. Reference neutron imaging with $4\times$ and $\sim 11\times$ the default flux was performed for reference. Also, greater flux was leveraged for up to 600 FPS, which in the future should aid the in-depth analysis of bubble shape variations via image-based surface strain rate measurements, which in turn can be used to validate numerical models.

An image processing methodology has also been developed to analyze the resulting low-SNR images, and provided evidence, including experimental validation, that the current implementation should be safe to apply to the neutron radiography images obtained for the model systems with bubble chain flow, in that physically meaningful results with manageable errors can be expected. Applied to the experimental dataset mentioned above, many new physical insights can be derived in the future. It is expected that the developed image processing pipeline and/or separate elements thereof should be applicable beyond the current application and context, which will be demonstrated in further in this thesis. The image processing code is available on *GitHub*: [Mihails-Birjukovs/Low_C-SNR_Bubble_Detection](#).

The code enabled to observe and quantify the differences between bubble flow patterns for various flow rates and MF configurations. It was shown that, depending on the MF configuration, radically different 1-bubble and bubble chain flow regimes occur. For instance, for a fixed flow rate, bubbles in a chain can exhibit nearly constant, two stage or single stage non-uniform acceleration, each with significantly different trajectories, velocity time series and shape dynamics. This is a clear proof-of-concept for MF control of bubble flow with industrially relevant dimensionless parameters.

Additional reference imaging experiments were also performed at extra high frame rates (up to 800 FPS), to be later used as benchmarks for numerical modelling of bubble shape dynamics, as well as for further development of tools for experimental data analysis.

It is worth mentioning separately that several noteworthy image processing methods were developed that can potentially be very broadly applicable in the scientific community, namely: multiscale recursive interrogation filter (MRIF), soft color tone map masking (SCTMM), and a false positive elimination algorithm based on object luminance and SCTMM.

B. X-ray imaging of bubble chain flow in liquid metal

In addition to the dynamic NR experiments with thicker systems, XR is performed for the Hele-Shaw systems, and the necessary image processing and bubble shape analysis tools are developed. While these results have not yet been published, they are a significant milestone within the context of this thesis.

1. Image acquisition & properties

Previously, bubble coalescence and breakup were studied in systems with 12 mm liquid metal thickness using XR [29, 30]. Here, much more confined systems are examined, with 3- and 6 mm thickness, imaged with gas flow rates 50 - 400 and 50 - 685 cm^3/min , respectively. Imaging with the indicated flow rate ranges was performed for the same MF configurations considered in Section II A 2 a. Injecting large gas volume into flat containers promotes stronger bubble flattening and deformations and intensifies bubble collisions, agglomeration, breakup and coalescence due to partial container obstruction by ascending bubbles. This is clearly seen in Figure 20. Using an XR setup similar to the one in [30], a wide range of different bubble shapes is observed. Their detection and description requires a special approach. Additionally, due to agglomeration, one must accurately resolve thin metal films between bubbles. This is complicated further by rather strong multiplicative noise present within these thin films.

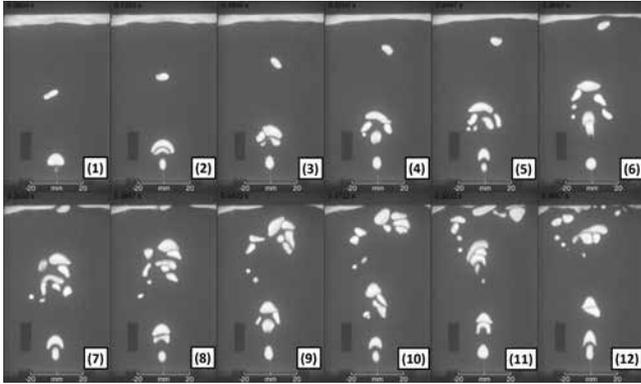


FIG. 20. XR images of argon bubbles in a 3 mm-thick galinstan vessel. Numbers indicate the order of frames.

2. Image processing

The above means that one would be hard-pressed to devise a simple global filter that can correctly segment the bubbles, resolving the thin liquid metal films in between. However, methods previously developed for NR of bubble flow in Section II A 2 b can be leveraged. Specifically, bubble detection is likewise split into two phases: local and global, and use MRIF (Algorithm 3), but restrict the maximum recursion depth to 1, since the SNR in the images seen in Figure 20 is significantly better than in the case of NR (Figure 9). Note that a pre-processing routine identical to the one used for NR images is used here as well. As such, the image SNR also varies over the FOV due to the FFC operation.

The entire image processing procedure is outlined in Algorithm 6. A preliminary threshold is determined for the entire image sequence as an average of thresholds computed for all individual image using the Kapur’s histogram entropy method [169]. One then applies Gaussian filtering to all images in the analyzed sequence and performs segmentation using the global threshold. Then the segments are subjected to size thresholding. This concludes the global filtering stage. Afterward, local filtering is performed for every detected segment. As in Algorithm 3, interrogation windows (IWs) are generated about the segments based on their centroids and areas [158], but with a slight modification. In this case, IWs are not squares, but rather rectangles that adapt to the aspect ratio of the estimated bubble segments.

Once the IWs for a frame are generated, each is treated with a local filter, which operates as follows (illustrated in Figure 21). First, a non-local means (NM) filter with a small kernel is applied [170–172] to increase the SNR of bubbles and the gaps between them, while preserving the shapes of the latter without notable CNR reduction (Figure 21b is the input, and (c) is the output). Local adaptive (LA) binarization [173] is then performed for the filtered IW image – this results in well-segmented bubble interiors, but rather eroded boundaries. Additionally, mask border components are removed, and bubble interior segments are repaired using the filling transform, size thresholding, morphological erosion and opening operations (in this order). However, the next steps address this: first, gradient filtering is performed (Gaussian regularization with the Bessel derivative kernel), greatly boosting the CNR of the bubble boundaries and the liquid metal films in between (Figure 21d); second, the image is re-scaled, and LA binarization is applied, but with a much smaller threshold estimation kernel (1 px by default) – this yields segmented bubble boundary outlines as shown in Figure 21e, but, naturally, the insides of the bubbles are not segmented. This is where the critical step comes in: interior and

Algorithm 6: Bubble segmentation from pre-processed FOV images

Input: Normalized pre-processed images

Global filtering

- 1 Determine the preliminary segmentation threshold for the image sequence
- 2 Gaussian filtering
- 3 Segmentation using the preliminary threshold (Step 1)
- 4 Size thresholding

Local filtering

For every detected segment:

- 5 Generate interrogation windows (IWs), similarly to IWs in Algorithm 3
- 6 Non-local means (NM) filtering
- 7 Local adaptive (LA) binarization
- 8 Segment restoration
- 9 Gradient filtering (Gaussian regularization & Bessel derivative kernel) applied to the output of Step 6
- 10 Image re-scaling & LA binarization applied to the luminance gradient field
- 11 Combine the binary masks generated in Steps 7 & 10

Segment restoration & recombination

- 12 Segment restoration
- 13 Mapping the final bubble segment onto the FOV image

Output: image mask with bubble segments for the FOV

boundary segments masks obtained in Steps 7 and 10, respectively, are added.

Finally, the combined masks are subjected to another segment restoration routine. This routine includes the following. First, filling transform is applied to (e), followed by morphological opening and border component deletion – this eliminates the fine structures seen around the bubble segments in (e) and fills the interiors of the bubble segments, if possible. Then size thresholding removes any leftover artifacts. The second stage of restoration handles any leftover bubble segments that are not filled in, e.g. because the combined interior and boundary masks failed to completely reconstruct some of the bubbles. It starts by decomposing the IW mask resulting from the preceding steps into separate masks for every newly formed segment. Then it is attempted to restore the boundary integrity so that they are closed and the bubble segments can be filled. Closing and filling transforms are applied to the separated segments, and the outputs are Otsu-binarized and recombined into one IW mask. Lastly, the final result seen in (f) is obtained by applying Otsu thresholding to (c), and

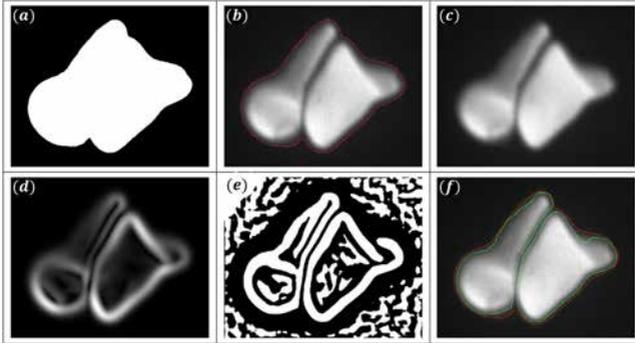


FIG. 21. Local filtering for an IW: (a) bubble segment estimated by the global filter, (b) local IW image with outlines (red) of the estimated segment (a), (c) output of the NM filter, (d) gradient-filtered image, (e) LA binarization applied to (d), and (f) image (b) with the boundaries (green) of updated bubble segments after Steps 12 & 13 of Algorithm 6.

then subtracting the inverted recombined mask. This last operation cuts out the thin films between the bubbles and crops the bubble boundaries that are facing away from the films, resulting in a refined bubble segment estimate. It was found that this procedure reliably separates the initially unresolved bubble clusters (a) into separate entities, while preserving their shapes rather well. It is also capable of resolving inter-bubble spacing down to $\sim 5\text{-}10\text{ px}$ thickness, which is very fine, considering the full FOV image size is $\sim 4\text{K} \times 2\text{K px}$. Note that such resolution was set for the experiments specifically to facilitate individual detection of rising agglomerated bubbles.

More examples of Algorithm 6 are shown in Figures 22 and 23.

In addition to 3- and 6-*mm* thick systems, X-ray imaging of 30 *mm* thick liquid metal layers has also been performed, which was often claimed by other researchers in the field to not be feasible. This was achieved by using industrial-grade X-ray tubes that can output higher-intensity X-ray beams with higher energy, in combination with different scintillation screens. However, these results are not yet published or systematized.

3. Bubble shape analysis

Once the bubbles are segmented and their parameters, such as X-ray transmission projection area, coordinates, etc. are determined, one can already derive much of the dynamics occurring within the FOV. How-

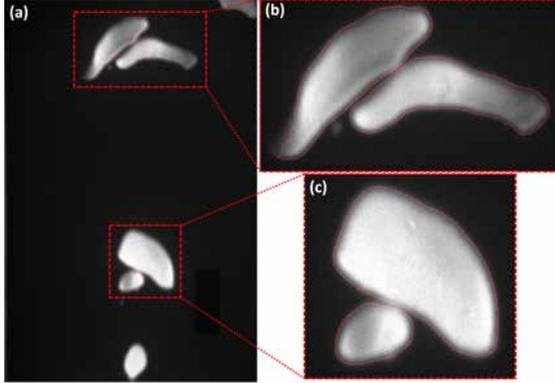


FIG. 22. Output of Algorithm 6 for an image from one of the recorded sequences: (a) the entire FOV (pre-processed) with detected bubbles highlighted with red contours, and (b-c) close-up views of the bubble pairs, with visible inter-bubble thin liquid metal films resolved by the algorithm.

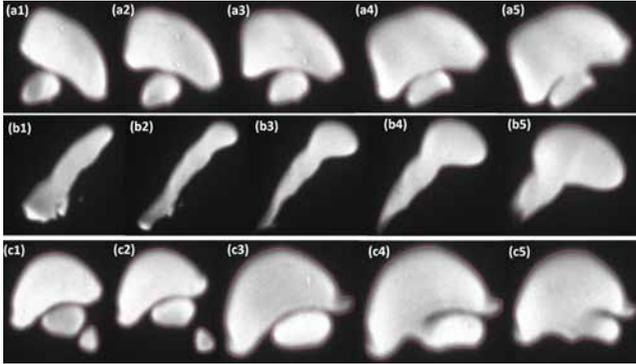


FIG. 23. Examples of resolved merge events and significant bubble deformations. Frames in cases (a-c) are numbered as they occur in time.

ever, unlike the case with NR of bubbles in thicker liquid metal systems, here the bubble exhibit very significant shape variations, generally with no clear symmetries, as well as rapid stretching and contraction during breakup/coalescence events. This means that simple shape parameterization will not suffice, and more advanced methods must be developed. The proposed shape analysis routine is outlined in Algorithm 7, and the key idea is to treat bubbles as strongly deformed and asymmetric (in general)

airfoils, i.e. like aircraft wing profiles analyzed in aerodynamics. A typical way to describe an airfoil is to construct its *camber line*, which is a line that runs through the airfoil shape from its one end to another such that at every point the line cuts the airfoil exactly in half by (differential) area. The endpoints of the camber line are the two points with the greatest curvature along the airfoil boundary.

Algorithm 7: Bubble shape analysis

Input: Bubble segment masks

For every bubble mask

- 1 Get segment centroids & boundary points
- 2 Create the boundary mesh
- 3 Boundary polyline reduction using curvature evolution
- 4 Interpolate the boundary points & upsample
- 5 Compute the boundary curvature profile & find extrema
- 6 Ramer–Douglas–Peucker (RDP) downsampling
- 7 Generate the Voronoi mesh for the boundary points & keep the internal points using the boundary mesh
- 8 Get the minimum spanning tree for the internal Voronoi points
- 9 Find the tree branch endpoints
- 10 Construct the main part of the bubble *camber line* (BCL)
- 11 Connect the main BCL to the boundary

Output: BCLs for the bubble segments

As such, one has to construct the *bubble camber line* (BCL) and correctly snap it to the bubble boundary on both ends. After finding the centroids and boundary points for the segments output by Algorithm 6, a boundary mesh is created to be used later. The boundary can often exhibit noise stemming from segmentation and quantization, both of which may hinder both BCL construction and BCL connection to the boundary. Therefore, given that shape preservation is also a concern, a polyline simplification approach was adopted, meaning that one attempts to reduce the number of points in the boundary mesh such that spurious curvature extrema due to noise are eliminated, but the longer-wavelength features are preserved. It was found that good results are achieved using the curvature evolution approach, particularly its variant proposed in [174]. After reduction, the quantization noise is further reduced by uniformly (by arc length) resampling the boundary with sub-pixel accuracy using closed spline interpolation.

Once the above preparations are done, one can use the now denoised and upsampled boundary to compute the boundary curvature profile over arc length. This is done semi-analytically, i.e. using an analytical expression for line curvature, but in terms of finite differences. Then the curvature extrema are identified. Then, before the BCL construction, the curve

is downsampled to avoid over-fitting, which is done using the Ramer-Douglas-Peucker (RDP) polyline simplification algorithm [175]. The BCL construction is accomplished in several stages. First, it is noted that drawing the BCL through a closed shape is equivalent (at least asymptotically, as the boundary resolution tends to infinity) to computing the Voronoi mesh for the boundary points, and keeping the internal points only. The former is a fairly standard operation, while the latter is done by removing the Voronoi mesh points that are not within the boundary mesh region obtained in Step 2 of Algorithm 7. At this point one has a BCL that exhibits several "tendrils" extending towards higher-curvature areas of the boundary, while not being connected to it, and one must determine the "true" BCL by tracing the right path through the generated Voronoi points. Intuitively, the curvature extrema where the BCL connects to the shape boundary should be the farthest removed along the potential pairs of boundary extrema points, as is the case with airfoil analysis that inspired this approach. However, in general, this may not be the case *exactly*. These problems can be solved by first deriving the minimum spanning tree for the Voronoi points, then identifying the endpoints of its tendrils as graph nodes with vertex degree equal to 1. Afterward, the longest among the paths between all possible endpoints is chosen as the BCL. Finally, the BCL is extended towards the boundary by connecting both its ends to the respective nearest curvature extrema. While the latter is a rather naive solution, the previous steps of the algorithm result in it working in most cases. Examples of the results generated using Algorithm 7 are shown in Figures 24 and 25.

With this, one can compute bubble thickness profile over BCL arc length, BCL curvature, length, etc., accounting for bubble shape irregularity and yielding insight into how bubbles split or merge (e.g., Figures 24d and 24e). One may also construct a principal chord line (PCL) connecting the BCL endpoints – the ratio of the PCL length to that of the BCL is a measure of bubble ellipticity; the distance profiles between PCL and BCL nearest points and integrals thereof over the PCL arc length characterize the degree of bubble bending and irregularity (e.g., Figures 24a, 24b, 24c). PCL orientation, for small instantaneous bubble deformations, informs one of rotational motion. In addition, one may derive thickness profiles for metal films separating agglomerated bubbles to study their interactions. Finally, having obtained the BCL, one can use it for refined bubble centroid estimation. Figure 25 shows how the code treats different bubble shapes and bubble coalescence over a series of frames.

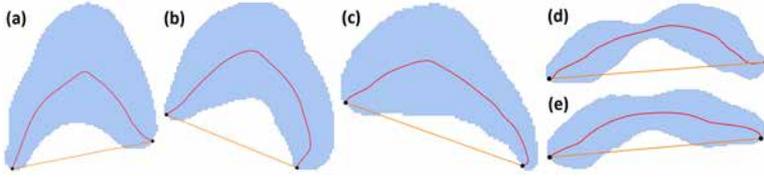


FIG. 24. Examples of bubble segment mesh regions with overlaid BCLs (red) and PCLs (orange).

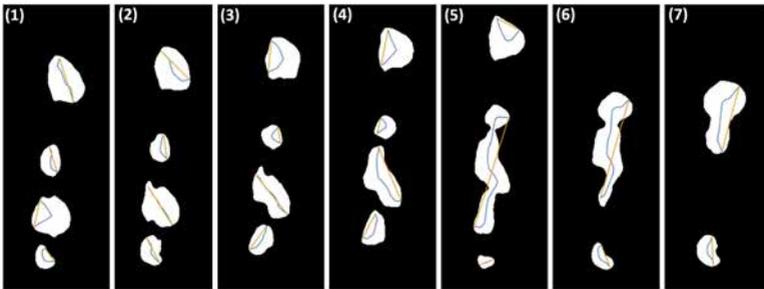


FIG. 25. A demonstration of how the shape analysis algorithm performs for a frame sequence – here frames are numbered in order of appearance in the sequence. Here, the BCLs are blue lines and PCLs are orange.

4. Summary

X-ray imaging of bubble flow in the Hele-Shaw variants of the systems considered in Section II A 2 a has been performed with gas flow rates 50 - 400 and 50 - 685 cm^3/min , for 3- and 6 mm liquid metal thickness, respectively. Imaging with the indicated flow rate ranges was performed for the same MF configurations as in Section II A 2 a. X-ray imaging has also been done for 30 mm thick liquid metal layers successfully for the first time using industrial-grade X-ray tubes with more intense and higher energy X-ray beams, and different scintillation screens.

To analyze the X-ray images with densely agglomerated bubbles separated by thin metal films affected by multiplicative noise, an image processing code has been developed that is capable of reliably resolving these films. It can segment bubble clusters, as well as process bubble coalescence/breakup events, and significant and abrupt bubble shape deformations. In addition, a special bubble shape analysis code has been implemented that should help with an in-depth analysis of bubble interactions. The image processing code is open-source and is available on *GitHub*: [4n515/X-ray-bubble-detection](https://github.com/4n515/X-ray-bubble-detection). The shape analysis code has been published on *GitHub* as well: [Mihails-Birjukovs/Bubble_shape_analysis](https://github.com/Mihails-Birjukovs/Bubble_shape_analysis).

C. Bubble trajectory reconstruction

1. The concept of MHT-X

Multiple hypotheses tracking (MHT) is classically considered to be the most reliable method for finding the optimal solution for data association problems [176–180]. However, it is rarely used so far for real-time (*online*) applications due to its $\mathcal{O}(n) = 2^n$ (brute force) computational complexity when searching for optimal and compatible object trajectories, where n is the number of considered associations. There have been attempts to solve the tracking problem by other means, but this always means trading precision/robustness for decreased complexity [176, 178]. While one could reduce the effective n for MHT, this does not affect complexity scaling, which is a problem in cases of large datasets often encountered in scientific applications. Since its inception, MHT was mainly applied to online tracking problems (military, surveillance, transportation and other related applications), whereas in scientific applications this mostly unnecessary [176]. Rather, reliable solutions are expected for potentially very large scale problems within acceptable time intervals, i.e. not necessarily requiring high performance computing.

Here, a new optimization approach is proposed for *offline* MHT: it is converted to its offline form, which is exploited to significantly reduce the computational complexity. Specifically, the association search problem is formulated as an *exact cover problem*, which is then solved using *Algorithm-X* [181] (hence the algorithm is called MHT-X). This way, one can cover the entire search-space of viable solutions, while entirely omitting contradictory solutions. MHT-X was designed to account for significant shape oscillations and parameter variations for objects that are typical in multiphase hydrodynamic systems (e.g. with bubbles and particles). A directed graph representation for input data and trajectories was chosen because of its intuitive interpretation, which makes the presented tracing framework conceptually easier to implement and customize for the application at hand [177].

2. The algorithm

First comes an initialization step, where the data is used to construct an initial graph. The algorithm then proceeds to scan the graph over time (t) with a time window (i.e. a time interval over which it is attempted to resolve associations between objects) of an iteratively increasing width (Δt). This is done to progressively resolve object associations over different

time scales. Solutions are then stored on directed graph edges, where nodes represent object detections.

Each time a new Δt value is set, unlikely edges are eliminated and a graph sweep over time is executed, where at every time step the following operations are performed:

1. Objects (trajectories) with unresolved endpoints are found
2. Associations are formed for these objects
3. Optimal hypotheses for associations are identified
4. The optimal solution is added to the graph

Initially $\Delta t = 1$ (two neighboring frames are considered) and it is incremented up to a value for which the utilized predictive models of object motion are no longer reliable, i.e. when tracking yields results with nonphysically abrupt trajectory cuts, jumps and/or oscillations. A schematic representation of the framework described above can be seen in Figure 26.

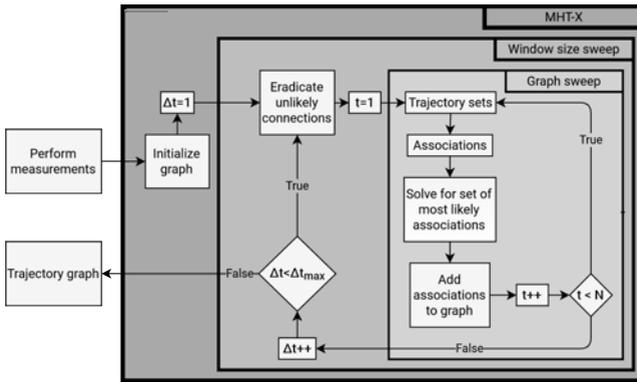


FIG. 26. A flowchart of the developed tracing algorithm MHT-X.

Since the tracking algorithm is offline, a complete set of measurements for object detection events is required as input – in this case these are object positions and projection areas due to optical/x-ray/neutron transmission, as well as the respective time stamps. Generally, however, there are no restrictions on the type and number of properties, but rather this is dictated by the intended application of the present framework.

The time-forward directed graph representation of data is the natural framework for MHT and specifically for the problem at hand, since instead of trajectories are constructed from associations between data points, rather than explicitly. A directed graph is defined by a set of nodes, a set of edges

and an edge function. The node set consists of the detected objects and two *special nodes*. The special nodes are labeled *Entry* and *Exit* which are auxiliary (padding, with no spatio-temporal coordinates) nodes of the graph that represent start and end points of trajectories, respectively. Edges represent associations between the nodes and are assigned likelihood values p to which the edge function maps every edge. Node associations represented by directed edges are aligned with the arrow of time. Edges connecting graph nodes to the special nodes are used to store the p values that the connected trajectory nodes are trajectory endpoints. Multiple edges originating from or leading to a non-special node represent split/merge events, respectively. All edges in such an event have identical p , representing the likelihood of the event. The directed graph formed by this algorithm represents different trajectories encompassed by it, where trajectories are potentially logically connected via split/merge events [182].

A trajectory is a set of sequentially connected objects that were measured over different time frames, i.e. a set of nodes connected by directed edges, bounded by special nodes and/or split or merge events. In cases where one-to-many split or many-to-one merge events occur, the source trajectories are broken, and new ones are formed, generating a set of *logically connected* trajectories – trajectory *families*. Nodes which are only connected to the special nodes are considered false detections – note that this is not a design feature, but rather it is an emergent property of the framework. Trajectory extrapolation is a crucial part of the reconstruction process, so a quantitative model is necessary. For MHD multiphase flow, this version of the code utilizes a rather naive general approach for extrapolation: a spline operator is defined, which maps a trajectory to a piece-wise polynomial function with respect to time. This spline is then used for extrapolation [182].

During the initialization step, the nodes of are defined from the set of measurements. To begin trajectory reconstruction, it is critical to generate boundary conditions (constraints). It is known *a priori* that all nodes corresponding to the first/last frames must be the endpoints of trajectories. As an initial condition for the graph, all nodes in the graph and the special nodes are interconnected, and each of the edges is prescribed a p value based on Boolean statements provided in [182].

The *graph sweep* is the main part of the algorithm – it iteratively performs hypotheses definition and evaluation, optimal association search and graph edge insertion using subroutines described further. For this, a time window of width Δt is defined, i.e. the number of frames covered by the window is $\Delta t + 1$. This window is then translated forward in time frame-by-frame through the graph, inserting new edges within each time window where appropriate.

In the current implementation, this is done repeatedly, constituting a *time window width sweep*. After an initial fully connected graph is formed and edge eradication is invoked (explained further), $\Delta t = 1$ is set defining a 2-frame window with $t \in [t + \Delta t]$. A graph sweep with this window resolves the most obvious associations formed between sequential frames. The resulting graph, after unlikely edge eradication, is used as an initial condition for the next graph sweep with a 3-frame interrogation window ($\Delta t = 2$). This process is repeated with increasing Δt until it exceeds a user defined threshold. This is done so that consecutive graph sweeps resolve associations over greater distance in time, allowing to overcome detection failures. Then, using the p distribution for all the graph edges and a user defined *quantile* parameter $q \in [0, 1]$, a p threshold p_c is computed. Edges with $p \leq p_c$ are eradicated.

At every time step in the graph sweep for a given t and Δt , two unordered sets of trajectories are formed: a set of trajectories which do not have an endpoint within $[t; t + \Delta t]$ and a set of trajectories which do not have an endpoint within $(t; t + \Delta t]$. The different association types – entry, exit, translation, split and merge – are defined in terms of these sets. The latter are also used to set up and test associations (mapped to graph edges), i.e. check that they are plausible and satisfy association constraints. Only plausible associations enter the trajectory association problem. The above is explained in detail in [182].

The following association conditions are applied: self-associations are forbidden; time-forward associations only; limited maximum object displacement per frame; limited association range. With the latter, a primary sphere of influence (SOI) based on the node object’s effective radius is defined about the node [182]. A secondary smaller SOI of a fixed size is also defined, and objects within are always associated. Association constraints, in turn, determine whether an association is plausible. Entry and exit associations are always considered plausible. Translation associations are expected to comply with *weak mass and momentum conservation laws*, which are used to determine if two trajectories are consistent in terms of object motion.

Denote the two trajectory segments within the time window with subscripts 1 and 2 and the connecting edge with subscript k . Translation associations are constrained by the maximum linear acceleration a_c :

$$2 \cdot \frac{\|\vec{v}_k - \vec{v}_1\|}{\Delta t_k + \Delta t_1} < a_c; \quad 2 \cdot \frac{\|\vec{v}_2 - \vec{v}_1\|}{\Delta t_2 + \Delta t_k} < a_c \quad (2)$$

where \vec{v} is velocity at the respective trajectory edges and Δt is the time difference between the nodes of considered edges. The change in direction

of movement is also limited based on the velocity:

$$\begin{aligned} \arccos\left(\frac{\vec{v}_k \cdot \vec{v}_1}{\|\vec{v}_k\| \cdot \|\vec{v}_1\|}\right) &< (\pi + \epsilon) \cdot \exp\left(-\frac{\|\vec{v}_1\|}{\lambda}\right) \\ \arccos\left(\frac{\vec{v}_2 \cdot \vec{v}_k}{\|\vec{v}_2\| \cdot \|\vec{v}_k\|}\right) &< (\pi + \epsilon) \cdot \exp\left(-\frac{\|\vec{v}_k\|}{\lambda}\right) \end{aligned} \quad (3)$$

where ϵ is an arbitrary small constant, and λ controls the maximum tolerable direction deviation with respect to velocity. Higher parameter values mean greater permitted deviations. The direction deviation constraints mimic momentum conservation by expecting objects with greater velocity to be less susceptible to deflection.

Weak mass conservation limits area (e.g. projection area due to transmission contrast imaging) differences between trajectories connected via translation:

$$\frac{|\langle S_1 \rangle - \langle S_2 \rangle|}{\max(\langle S_1 \rangle, \langle S_2 \rangle)} < \varepsilon_t \cdot \left\langle \frac{\sigma_k}{\langle S_k \rangle} \right\rangle \quad (4)$$

where indices 1, 2 refer to two different trajectories, $S_{1,2}$ are the sets of area measurements, $\sigma_{1,2}$ are the standard deviations for $S_{1,2}$, $k = \{1, 2\}$ and ε_t is the area deviation threshold. Greater thresholds permits higher relative deviations of area.

In the case of split/merge events, weak momentum conservation is unreliable due to surface tension effects, therefore only weak mass conservation is used. This essentially checks if the projection areas of objects before and after splits/merges are consistent:

$$\frac{|S_0 - \sum_k \langle S_k \rangle|}{\max(S_0, \sum_k \langle S_k \rangle)} < \varepsilon_s \cdot \left\langle \frac{\sigma_k}{\langle S_k \rangle} \right\rangle \quad (5)$$

where S_k and σ_k , $k \in \mathbb{N}$ correspond to the merge components and/or split products and ε_s is the mass conservation threshold, the lower the value, the more the changes in object areas are expected to comply with mass conservation. Here, a linear relationship is assumed between the mass and the projection area of the bubble. This is *a priori* wrong, since it leads to an apparent mass defect, although the mass is not physically lost. Up to a certain level, this approach provides satisfactory results.

While bubble sizes physically do change the breakup/coalescence frequency, the tracking performance *does not* depend on the bubble sizes, but rather on relative area variations for individual bubbles and the *diversity* of bubble sizes within the field of view (FOV). This is because the

breakup/coalescence events are modeled using Equation 5 and translational motion depends on bubble size measurements via Equation 4. Note that in both cases there is normalization with respect to bubble sizes, meaning that Equations 4 and 5 essentially constrain bubble size variations within trajectories and across trajectories considered for split/merge events. The more diverse bubble sizes are and the more consistent the sizes of individual bubbles are in time, the better the tracking performance. Trade-offs are controlled by ε_s and ε_t parameters.

The Bayesian formulation of MHT is used [177]. Due to the offline nature of the algorithm, the framework is greatly simplified. The formulation for the problem of finding feasible (non-contradictory) sets of trajectories for MHT-X is specified in [182].

The trivial way to solve the offline MHT problem is a brute force search verifying every possible association, which is extremely inefficient due to the $\mathcal{O}(n) = 2^n$ complexity, and is therefore only feasible for a very low number of associations. Reducing the effective n is an option, but that does not solve the scaling problem which becomes critical for very large measurement sets, i.e. measurements with high number density per frame or very long measurement processes, which is the case in many scientific applications. The proposed approach is to recognize this as an exact cover problem [181], since the solutions of the exact cover problem are by definition sets of associations that yield feasible trajectories in this case.

The best known way of solving an exact cover problem is using the *Algorithm X* (Knuth’s algorithm) [181], which reduces the computational complexity significantly down to consistently sub-exponential. Complexity is reduced further by clustering associations into disjoint sets before formulating and solving the exact cover problem [183].

Statistical functions that map likelihoods to graph edges are defined as follows. Let $\mathcal{N}(x, \mu, \sigma)$ be a Gaussian distribution with its mean μ and standard deviation σ . The translation association probability is measured by:

$$f_1 = \alpha \cdot \frac{\mathcal{N}(\delta r, 0, \sigma_{\delta r})}{\mathcal{N}(0, 0, \sigma_{\delta r})} + (1 - \alpha) \cdot \frac{\mathcal{N}(\delta S, 0, \sigma_{\delta S})}{\mathcal{N}(0, 0, \sigma_{\delta S})} \quad (6)$$

where δr is the displacement magnitude, δS is the area difference between node objects and α is the weight adjustment parameter. Both $\sigma_{\delta r}$ and $\sigma_{\delta S}$ are computed for both trajectories considered for connection via a translation edge.

Entry/exit edge probability is:

$$f_2 = \frac{1}{1 + \exp(a(y - b))} \quad (7)$$

where y is the vertical coordinate in a 2- or 3-dimensional image and a , b are control parameters. The value of b determines the location on the predetermined axis (in this case vertical) where there is a 0.5 probability for the existence of an Exit/Entry node. It is recommended that b is set such that this location is near the FOV boundaries where the bubbles are expected to enter/leave. The magnitude of a determines how rapidly the probability for the existence of the Exit/Entry nodes changes over the FOV (from 0 to 1 and vice versa). The value should be chosen such that the false positives can be assigned Exit/Entry anywhere in the FOV, but it should be set high enough to avoid always having Exit/Entry as the likeliest solution. The sign of a determines whether the probability increases or decreases along the axis: $a > 0$ implies a decrease from 1 to 0, and vice versa.

The merge/split event probability is also based on weak mass conservation and is computed as follows:

$$f_3 = \beta \cdot \mathcal{N}(\delta(S_0, S_k), 0, \langle \sigma_{S_k} \rangle) + (1 - \beta) \cdot \frac{M_s}{M} \cdot \mathcal{N}(\langle \vec{r}_k \rangle - \vec{r}_0, 0, \sigma_{\delta r}) \quad (8)$$

where β is the weight adjustment parameter, S_0 is the area of the split source/merge product, S_k are the areas of the merge/split components, \vec{r}_0 is the position of the split source/merge product, \vec{r}_k are the positions of the merge/split components, M and M_s are the number of involved trajectories and the number of trajectories with available motion prediction (i.e. there are enough points in a trajectory), respectively. Therefore, β determines whether the area or the position consistency is emphasized more for the split/merge components.

It is not always desirable to immediately add the feasible associations (edges) to the graph, since new trajectories might be generated or enter the time window in the next iteration of the time sweep that are better solutions to the problem. Therefore, any association considered likely (and the resulting trajectories) must also satisfy extra conditions outlined in [182]. This is the last step of the time sweep. After this, the set of accepted edges is added to the trajectory graph.

3. Main results

To assess the performance of MHT-X for scientific applications, it was applied to three cases of bubble flow in liquid metal where object tracking is necessary and offline tracking is appropriate:

- 2D simulations of argon bubble flow in a rectangular vessel.
- Dynamic X-ray radiography of argon bubble flow in a rectangular vessel filled with GaInSn wherein bubbles are injected via a top submerged lance.
- Dynamic neutron radiography of argon bubble chain flow in a rectangular liquid gallium vessel – bubbles are injected at the vessel bottom via a horizontal/vertical tube.

In all three cases, segmentation is performed prior to tracing and bubble flow regime is such that bubbles are expected to deform considerably while ascending to the free surface of liquid metal at the top of the vessels.

The first test case is the output of a 2-dimensional simulation of bubbles rising through liquid gallium in a rectangular vessel. The model is a volume of fluid (VOF) simulation of two-dimensional bubble flow in a rectangular vessel with a horizontally directed inlet at the bottom. The numerical model is described in [152, 153]. The flow regime has been adjusted such that bubble trajectories are highly irregular and collisions/splits/merges are frequent. The bubbles have been perfectly segmented in that there are no false positives or detection failures in this case. Bubbles with projection areas below a predefined threshold were not considered. Therefore, this is an idealized test of the tracking capabilities of MHT-X in case of a moderate number density of objects of various sizes and variable shapes with frequent interactions. Here, object coordinates and projection area are tracked. Examples of the tracing output are shown in Figures 27-30.

Several frames with overlaid reconstructed bubble trajectories are shown in Figure 27, where trajectories are color coded by bubble ID. The code successfully tracks both large, significantly deforming bubbles, and smaller ones, even in cases of proximity with co-linear motion. Note the two detected bubble splits indicated in Figure 27c with purple arrows. Another important feature of this case is that bubbles move in a variety of patterns: ascension due to buoyancy, slow oscillatory motion due to entrapment in low velocity zones, downward motion due to a large vortex with counter-clockwise mass flow – all of this is representative of realistic flow conditions in two-phase systems.

In Figure 28 one can see several bubbles tracked across consecutive frames. Note that in Figures 28(3)-28(6) two split events in rapid succession are

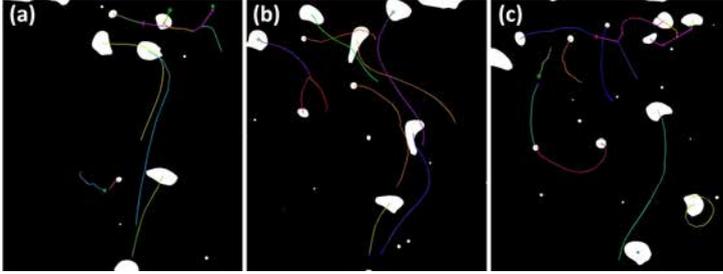


FIG. 27. An example of characteristic trajectory patterns and (c) two splits (purple arrows) in rapid succession resolved in presence of strong deformations and other potentially interfering bubbles. Trajectories are color coded by their IDs.

resolved where bubble shapes before and after breakup are radically different. It is also important to note that MHT-X does not lose track of bubbles despite significant elongation (especially in frames 3-5) and proximity of two more bubbles that, while initially ascending, divert to the left and begin to enter almost co-linear motion (frames 5-7).

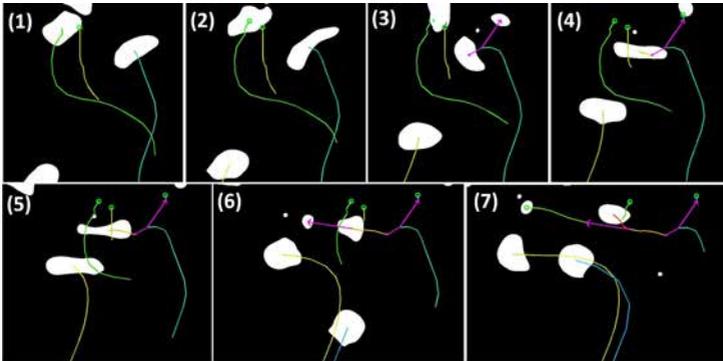


FIG. 28. (1-7) Several nearby bubble trajectories and splitting events resolved over sequential frames from simulation data, and (a,b) examples of characteristic trajectory families.

Figure 29 shows two examples of logically connected trajectory sets (families) derived from the established trajectory graph. Note especially Figure 29a where family members exhibit rather complex trajectories, proximate and even overlapping trajectories. The entire family in Figure 29a originates from a common entry point at the bottom of the field of

view (FOV). The graph allows to directly examine the entire trajectory network and qualitatively assess the intensity of collective dynamics and where the interaction events are localized within the FOV.

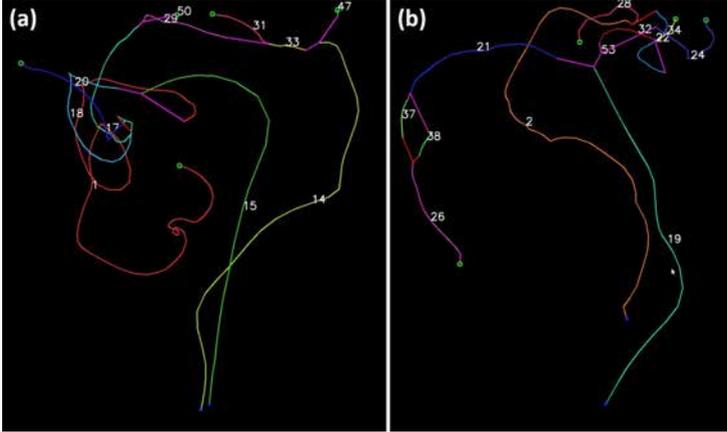


FIG. 29. Two examples of trajectory families recovered from the trajectory graph produced by MHT-X with trajectories color coded by IDs (highlighted). Red arrows represent merging events.

In addition to visual information regarding bubble motion and interactions, families and the corresponding exported timestamped datasets for further processing (velocimetry, trajectory curvature measurements, shape parameter evolution tracking, etc.), it may also be helpful to visualize the constructed trajectory graph itself – the solution graph for this test case is shown in Figure 30.

However, even though this test case demonstrates successful tracing for rather complicated flow patterns and bubble interactions, it is somewhat idealized. Bubble projection area conservation is not violated too strongly, i.e. bubbles do not physically vary in volume (the FOV is above the growth region at the inlet), and the segmented dataset is virtually without error or noise. The following two cases address these conditions.

The second case stems from a dynamic X-ray imaging of bubble flow at 125 frames per second (FPS) in a rectangular vessel with a top submerged lance (inlet) setup [84] where the bubbles are injected within the FOV (data provided by the authors). This means that bubble volume, and therefore also the projection area due to X-ray transmission, are generally very different between a bubble that is being ejected from the inlet tip and already detached bubbles. In addition, the bubbles also exhibit sub-

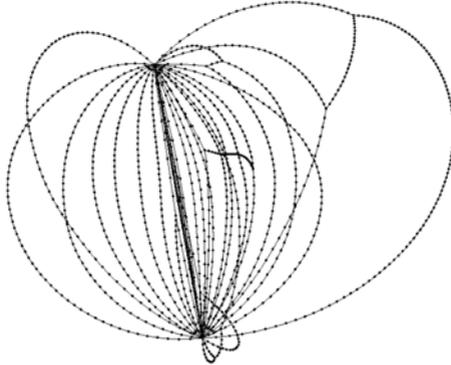


FIG. 30. A solution graph computed by MHT-X for the bubble flow simulation benchmark. This spatially sparse representation of the trajectory graph was rendered in *GePhi*. All trajectories here stem from the entry node (bottom) and converge at the exit node (top). Note that while the coordinates are distorted, the relative vertical (entry to exit direction) positioning is still representative of the actual measurement points (i.e. the relative positions of split/merge events).

stantial deformations, including out-of-plane (with respect to the FOV) motion of the argon/GaInSn interface. Moreover, the bubbles are segmented from X-ray images where noise and potentially artifacts are present, therefore bubble shapes are generally not recovered perfectly. Frequent bubble/bubble interactions are to be expected in this flow regime, even though their number density is less than in the previous test case.

However, in this case the data regarding local volume fraction over the FOV are available from measurements [84] and bubble volume has been derived and supplied to the tracer. A conservation law of the same form as one for the projection area S in (4), (5), (6) and (8) was added for the tracked bubble volume. The developed algorithm performed well under the above conditions, as illustrated in Figures 31-33. As in the previous case, split/merge events proximate in time were resolved (Figure 31) and the many-to-one and one-to-many events, much more frequent in this case, were also correctly identified (Figure 32). More details and examples are shown in [182].

As with the 2D case, trajectory families can be visualized, as shown in Figure 33. The frequent coalescence is clearly seen in Figures 33a and 33b especially, where this results in many logically connected trajectories. Note also that some splits occur at the bottom as well, e.g. in Figure 33a, which occurs because bubbles sometimes detach from a gas pocket that forms at the bottom of the injection lance rather than directly from the lance

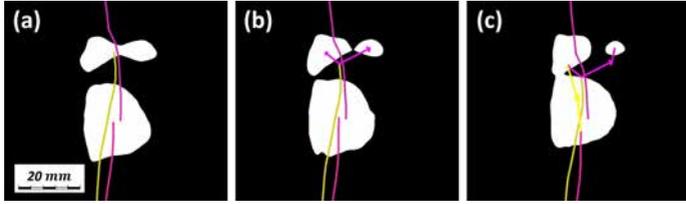


FIG. 31. An instance of a correctly resolved temporally proximate split/merge sequence.

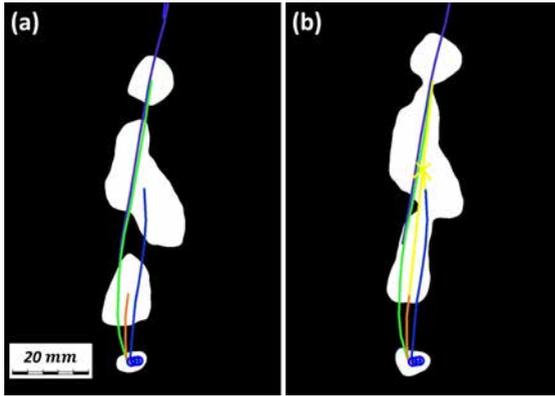


FIG. 32. A three-to-one merge event (yellow arrows) resolved by MHT-X.

nozzle. While this occurs relatively rarely in this test case, some bubbles ascend to the top of the FOV without any interactions whatsoever, e.g. as in Figure 33d.

For some applications it may be of interest to inspect the split/merge event statistics: event locations by type, their spatial density over the FOV, etc. In addition, retrieving angles at which objects, in this case bubbles, merge or split can yield insights into what drives the observed behavior. Figures 34-36 are an example of such an analysis. From Figures 34a and 34b one can see that the merging events mostly occur at the bottom due to trailing bubbles quickly catching up with leading bubbles, whereas splits occur rather often at the top, where bubbles that are often products of mergers are once again split into smaller bubbles, and at the bottom where bubbles detach (technically, split) from the gas pocket that consistently forms at the inlet nozzle.

Figure 35 shows the statistics for merging (a) and splitting (b) angles.

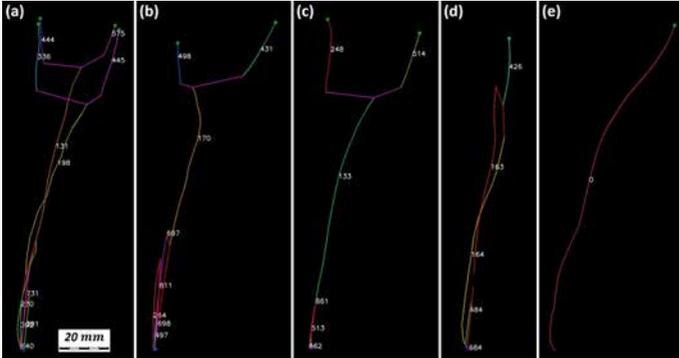


FIG. 33. Representative trajectory families with trajectories labeled and color coded by their IDs. As with the 2D test case, the purple arrows represent split events, while the red arrows stand for merging events.

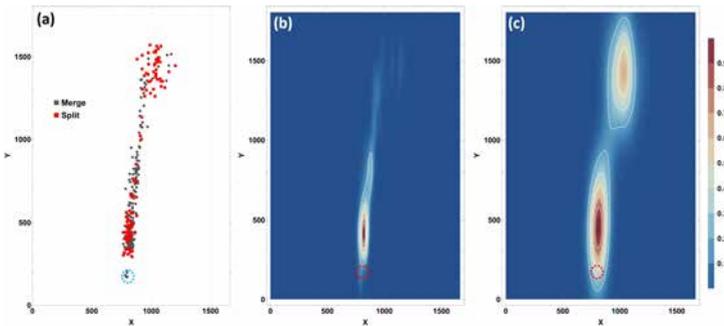


FIG. 34. (a) A map of normalized merge and (b) split event density. (a) and (b) are normalized separately and share the color legend (right). X and Y coordinates are in pixels. The red dashed circles are inlet locations. Density maps are computed using a Gaussian kernel over the count area density function with Silverman’s bandwidth estimation.

Note that merges mostly occur about zero degree angles – this makes sense given mostly vertical collisions between leading and trailing bubbles near and above the inlet nozzle. Splits, however, exhibit a much less ordered distribution with no consistently clear peaks, but there is a pronounced asymmetry in that negative split angles have greater magnitudes. Figure 36 offers further insight, indicating that large angle splits are concentrated in the upper (Figure 36b) left (Figure 36a) area of the FOV and are largely responsible for the upper density maximum zone in Figure 34b. It is also

revealed that high angle merges are distributed mostly in the bottom 2/3 of the FOV.

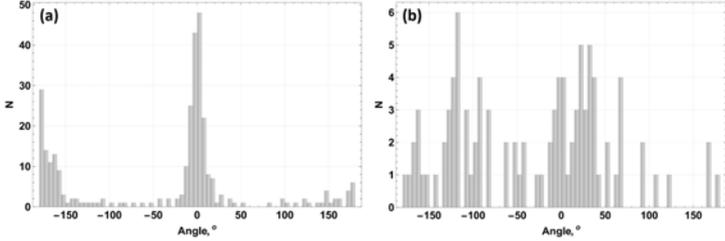


FIG. 35. Direction angle histograms for (a) merge components and (b) split components.

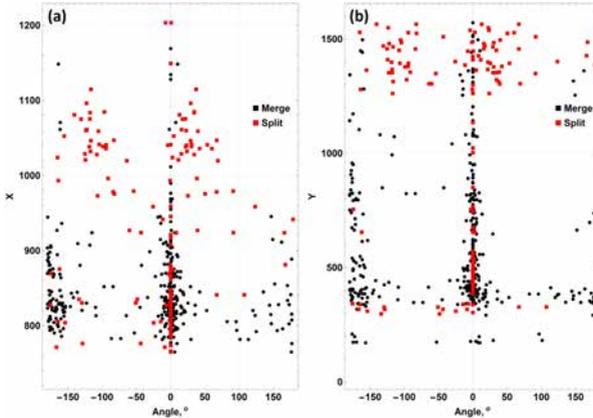


FIG. 36. Merge and split angle distributions over (a) X and (b) Y coordinates (in pixels) in the FOV.

In this test case, the bubble dynamics are a bit more complicated to resolve for MHT-X despite the lower number density of objects compared to the first test case. Nevertheless, the data set is relatively clean because the signal-to-noise ratio (SNR) in the original X-ray transmission images is quite high and false positives are not present. To demonstrate the code's robustness against noisy input and false positives, the results from the third test case are presented in the following section.

The third case uses the data obtained by means of dynamic neutron imaging of argon bubble flow in liquid gallium at 100 FPS in a system

described in detail in [152, 153, 158]. The characteristic feature of neutron radiography images acquired with a high frame rate for thick liquid metal vessels is the very low image SNR. Because of this, even with advanced noise filtering and segmentation, the data provided as input to MHT-X is inevitably noisy in that bubble centroid position uncertainties are considerable and much greater than in the first two test cases shown herein. In addition, occasional false positives and detection failures may occur, further complicating tracing. To isolate these effects, image sequences with no bubble interactions were chosen, i.e. the average bubble spacing is sufficient to avoid collisions. The flow regime expected for this experimental setup [152, 153, 158] and the moderate gas flow rates considered is the bubble chain, where the distance between bubbles is determined by the respective value of the gas flow rate (low flow rates produce quasi single-bubble flow). Figure 37 illustrates characteristic bubble trajectories reconstructed with MHT-X. Here (a) and (b) show zigzagging bubbles and (c) and (d) show mostly rectilinear trajectories due to flow stabilization with applied MF.

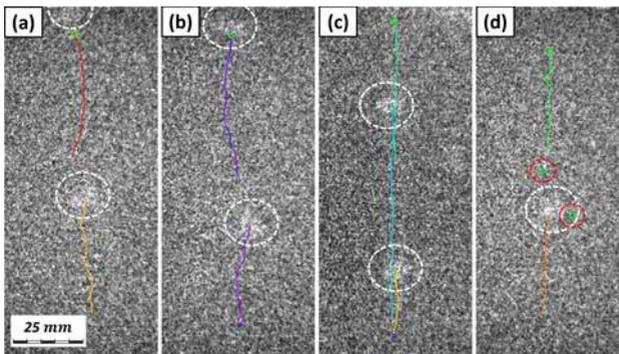


FIG. 37. Neutron radiography images with highlighted trajectories color coded by IDs, bubbles (white dashed circles) and false positives (red dashed circles). Up to a number of latest trajectory edges are shown in (a-d) for visual clarity.

One can clearly see that bubbles (white dashed circles in Figure 37) are largely shrouded by image noise. Despite this and the resulting noise in the object dataset, it is seen that the algorithm performs well and long, consistent trajectories spanning the entire FOV are recovered. Note also that in Figure 37d there are two false positives (red dashed circles) that were overlapping in time with true detections – these were resolved as isolated nodes (1-node trajectories) in that within the solution graph they are only connected to the entry and exit nodes. Further demonstration that MHT-X correctly extrapolates and connects trajectory fragments across a

frame with a detection failure event, as well as shows resilience to noise, is provided in [182].

Since the performance of the entire algorithm mostly depends on that of the Algorithm X, a benchmark is run to test it for the association types that may occur in tracing applications with split/merge events. The benchmark is performed as follows:

1. Generate 2 groups of m and n objects.
2. Association condition fulfillment is set as probabilistic with a likelihood of 0.3.
3. Association constraint fulfillment is set as probabilistic with a likelihood of 0.5.
4. Separate the set of associations into disjoint sets.
5. Log the amount of associations and how much computational time it takes for the Algorithm-X to exhaust the search-space for each disjoint set.

This was repeated 10000 times for all combinations of $6 \leq m, n \leq 10$, restricting the maximum number of many-to-one (and reverse) associations to 3 considering only unique unordered (m, n) pairs. The resulting $\sim 250\text{K}$ data points were separated uniformly into bins by the association count with bin width equal to 1. The resulting bins were then also binned vertically into 100 equally spaced bins from minimum to maximum execution time. Then a discrete normalized density distribution was generated. The resulting computation time distributions versus association count within a disjoint set are shown in Figure 38.

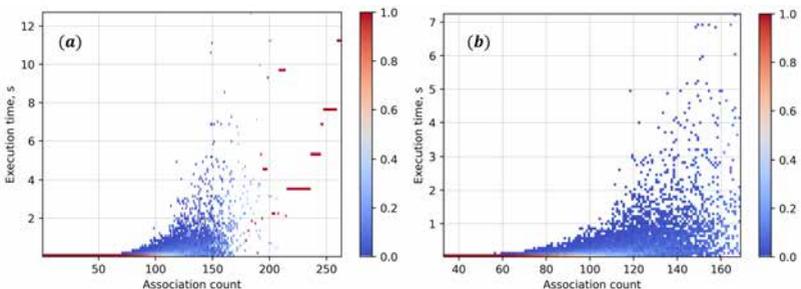


FIG. 38. Computational performance of the Algorithm-X for associations feasible in tracing problems with split/merge events. (a) The entire set of benchmark data points and (b) a subset of the benchmark data with execution time between 10 ms and 8 s which is roughly 15% of the entire data set.

One can see in Figure 38a that the worst case scenario is exponential, as

it should be. The exponential time cases, however, are encountered very rarely and most of the cases are solved significantly faster. It is worth noting that increasing the dimensionality of the problem does not improve or worsen the performance of the algorithm. The limiting factors are object number density and interaction frequency, as well as the mean number of objects in split/merge events. Extra spatial dimensions simply add more coordinates to track.

While this version of the algorithm is already very versatile, as indicated by the above results, there are still potentially many ways it can be improved. The aim is to make it more broadly applicable and enable solving tracing problems with greater object density and more adverse conditions, as seen in the cases considered here. In the related field of research, there are several cases of interest where explicit, precise and robust object tracking is desired: dynamic optical, X-ray and neutron imaging of argon bubble flow in liquid gallium or GaInSn eutectic alloy [29, 30, 84, 152, 153, 158, 184]; neutron imaging of gadolinium oxide particle flow in liquid gallium [85–88]; optical imaging of salt crystals and liquid crystal tracers in water [106, 107]; neutron imaging of gadolinium particles in froth [105]; bubble flow simulations [152, 153, 184]. For bubbles in particular, MHT-X could be used for bubble shape analysis, including shape evolution tracking and phase boundary velocimetry (e.g. by combining it with optical flow velocimetry).

4. Summary

The capabilities of the developed offline Bayesian multiple hypothesis tracker with a directed graph architecture that uses Algorithm X to solve the optimal association problem as an exact cover problem have been demonstrated. The showcased test cases indicate that the current implementation is robust enough to process cases with relatively high object number density, in presence of data noise, false positives and detection failures. The algorithm is capable of resolving one-to-many split and many-to-one merge events for objects with variable shapes and parameters.

In its current state and especially as the outlined improvements are implemented, it is expected that the code will find use in many areas of research. Beyond that, the code is currently in use for the development of a dynamic mode decomposition code for the analysis of output of bubble flow simulations, as well as for bubble shape analysis, including shape evolution tracking and phase boundary velocimetry, particularly in combination with the shape analysis algorithm considered in Section II B 3.

D. Neutron imaging of particle-laden liquid metal flow

The image processing and particle tracking methodology presented below has been developed based on the experimental data obtained by HZDR researchers prior to this thesis' beginning. The objective was to check if particle tracking was feasible for a simple benchmark system with known physics so that the developed methods can then be applied to model systems that would emulate bubble wake flow with and without applied MF. PTV could then be used to study bubble wake flow in liquid metal.

1. Neutron imaging

Gadolinium oxide particle flow in gallium-tin alloy was imaged using dynamic neutron radiography. The experimental setup is shown and described in detail in [96, 185]. Given that neutron transmission imaging yields particle projections, it is desirable to avoid three-dimensional motion, so this experiment is performed for a quasi two-dimensional geometry. The stainless steel cylindrical obstacle with a 5 mm diameter is centered and fixed in the straight section of a closed loop flow channel with a uniform 30 mm × 3 mm rectangular cross-section (flow was imaged through the 3 mm dimension). Continuous liquid metal flow is driven by a disc-type electromagnetic induction pump (permanent magnet-based) [96]. The region of interest where particles must be detected and tracked *explicitly* is the wake flow area highlighted in Figure 39. The particle tracks, which can be visualized using the time projection of minimum luminance (Figure 39b), are strongly affected by the turbulent wake in contrast to the very smooth tracks at the sides of the wake flow zone. The analyzed field of view (FOV) was equal to 408 × 161 pixels (16-bit single-channel images) corresponding to 37.8 mm × 14.9 mm.

Gadolinium oxide particles with a $d_p \in (0.3; 0.5)$ mm diameter have been shown to provide sufficient image contrast for dynamic neutron imaging with a short image exposure time [96]. Note that at the 3 mm thickness in the neutron flux direction, the liquid alloy is rather transparent ($\sim 87\%$ transmission) to neutron radiation. Gadolinium oxide also has a lower paramagnetic susceptibility than its closest alternative, gadolinium, which, conversely, makes particles strongly attracted by the MF of the electromagnetic induction pump [186, 187].

Imaging was carried out with cold neutrons at the ICON beamline [188] of the Swiss spallation neutron source SINQ [189]. The imaging setup was such that the resulting geometrical unsharpness is ~ 0.3 mm which is about the size of the smallest gadolinium oxide particles. The imaging

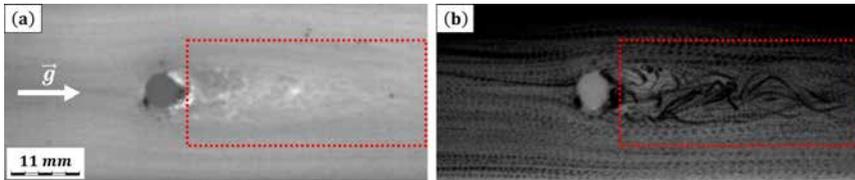


FIG. 39. Pixel-wise (a) standard deviation and (b) minima of luminance values within the imaged flow channel over all captured frames. The region of interest is indicated with a red dashed frame. Note the cylindrical obstacle in both figures. The images shown here were rotated 90 degrees left with respect to the originals – as such, here the originally downward flow is directed from left to right. The white arrow in (a) indicates the gravitational acceleration \vec{g} .

spatial resolution is 10 px/mm . A 10 ms image exposure time was chosen, equivalent to 100 FPS, required to capture individual particles moving in the liquid metal flow.

2. Image processing

It was shown in [105] that the images exhibit correlated noise in the form of grain-like structures with ~ 3 -pixel sizes – this is a considerable fraction of a typical particle size in images, and thus images may contain "phantom" particles (false positives). Particle projections visible due to neutron transmission contrast have strongly varying sizes and SNR/CNR that also change over time as particles travel through the FOV. In addition, the recorded image sequence exhibits a pronounced global luminance non-uniformity, which can be seen in Figure 39b. Given these factors, it was decided to build the image processing procedure around a local filter applied to interrogation windows (IW, not to be confused with PIV terminology) taken from the images. Due to the very high area density of particles in the images, it was decided to scan images entirely with partially overlapping IWs. This approach has three main benefits: first, IW overlaps yield redundancy which should reduce the rate of instances where true particles are excluded; second, local filtering over the entire image means reduced sensitivity to large-wavelength variations in luminance over the image; third, this works well with the false positive exclusion algorithm developed for bubbles (Algorithm 5) and modified for this case (explained further).

The idea of sweeping images with partially overlapping IWs is illustrated in Figure 40. An initial square IW with a side length L' is fitted into the upper-left corner of the image. A set of IWs is then generated from

the initial IW by creating an IW position lattice with horizontal/vertical stepping, where the steps are a fraction of L' . The details regarding how the set of overlapping IWs is generated can be found in [185].

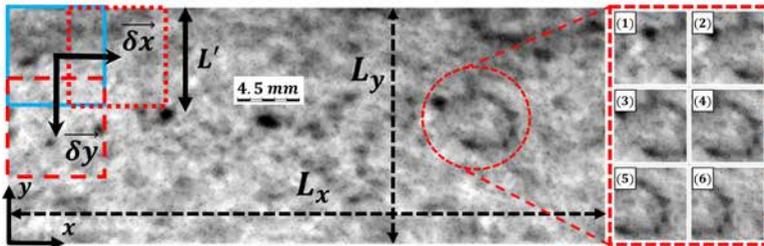


FIG. 40. A schematic representation of the IW sweep for the FOV: the initial IW (light blue) and shifted IWs (red) in the $x > 0$ (dotted) and $y < 0$ (dashed) directions. An example of a horizontal scan in the $x > 0$ direction for an area indicated with a red dashed circle is shown in sub-figures 1-6.

All image processing operations are performed in *Wolfram Mathematica*. The general framework for image processing is as shown in Algorithm 8. L' was set to roughly match the expected scale of particle clusters seen in the FOV images. This is to ensure sufficient redundant detection for particles – it is observed that the selected values are optimal for this case in terms of detection accuracy. However, smaller δx_k make such an approach more computationally expensive and in this case the chosen window step size effectively inflates the amount of data (total image area) by a factor of ~ 16.9 . Therefore, one should take care to optimize the underlying image processing code in terms of memory utilization and parallelize as many of its elements as possible – the way this is done is explained in [185].

The stages of local filtering are shown in Figure 41 and the filter structure is outlined in Algorithm 9 (more details in [185]). The luminance maps are inverted to highlight particles, since by default, due to intense neutron flux absorption by particles, they appear in images as lower luminance zones. This results in an IW image as shown in Figure 41b. Next, non-local means masking (NMM) is performed (Figure 41c) to increase the contrast-to-noise ratio for particles and remove the "haze" (correlated noise due to unsharpness), which is especially important for tightly-packed particle clusters. NMM transforms the original image x into output y as follows:

$$y = 2 * x - w_{\text{nm}} * \text{NM}(x, r_1, r_p) \quad (9)$$

where $\text{NM}(x, r_1, r_p)$ is the non-local means (NM) filter [170], w_{nm} is the NM mask weight, and r_1 and r_p are the filtering neighborhood and neighborhood

Algorithm 8: Image processing framework

Input: A sequence of normalized images (pixel luminance rescaled to $[0; 1]$) with subtracted mean dark current

IW generation

- 1 Generate a $n_x \times n_y$ grid of IW positions [185]
- 2 Disassemble images into their projections onto IWs

Particle detection in IWs

- 3 Normalize the IW images
- 4 Local filtering (Algorithm 9)
- 5 Particle segmentation (Algorithm 10)
- 6 Luminance map-based false positive filtering

Assembly of global detection masks

- 7 Map the false positive-filtered IW particle masks onto the full FOV [185]
- 8 Sum the particle masks over the FOV
- 9 Minimum area thresholding
- 10 Morphological opening
- 11 Compute centroids for the resulting segments

Output: Centroids for particles detected in every FOV image

comparison radii, respectively. This is similar in principle to unsharp masking, but utilizes the NM filter instead of the Gaussian filter. It is observed that here NMM distinctly outperforms simple unsharp masking since the NM filter captures the correlated noise much better. The utilized NM filter computes the normalized neighborhood weights \tilde{w} for averaging as in [171].

Algorithm 9: Local (IW) filtering

Input: A normalized IW image

- 1 Invert the image luminance map
- 2 Non-local means masking (NMM)
- 3 Soft color tone map masking (SCTMM)
- 4 Non-local means (NM) filtering
- 5 Mean filtering

Output: A filtered IW image

Next, SCTMM (1) [158] was applied for background reduction (Figure 41d). Afterward, NM filtering is performed (Figure 41e), followed by the mean filter (Figure 41f). Note that throughout the IW filtering procedure, images are re-normalized after each filtering stage.

Filtered images are then subjected to the segmentation procedure outlined in Algorithm 10. Here local adaptive (LA) binarization (mean- and

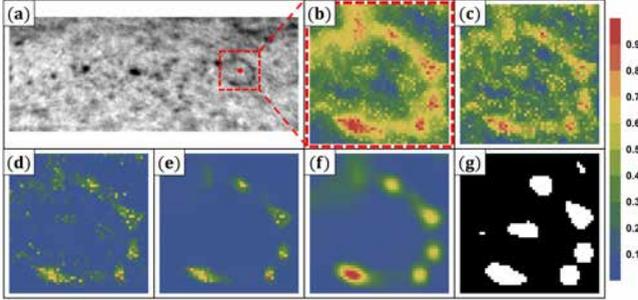


FIG. 41. Local filtering applied to interrogation windows (IWs): (a) original image of the region of interest (Figure 39) with a highlighted IW (red frame), (b) colorized luminance-inverted IW image; IW after sequentially applying (c) non-local means masking (NMM), (d) soft color tone map masking (SCTMM), (e) non-local means (NM) filtering, (f) mean filtering and (g) SCTMM followed by local adaptive binarization.

deviation-based) [173] is used because global thresholding yields very unstable particle detection in filtered IWs due to their dissimilar SNR and CNR, and thus post-filtering luminance distributions. LA binarization, however, is susceptible to the edges of low-luminance particles, many of which are potential false positives. For this reason, the previously covered LFPE filtering procedure was used, with minor modifications explained in [185]. The application of LFPE for this case is illustrated in Figure 42. Finally, image border components are removed to avoid artifacts and artificial particle splitting.

Algorithm 10: Local (IW) segmentation

Input: A normalized filtered IW image (Algorithm 9 and Figure 41f):

- 1 Apply SCTMM
- 2 Local adaptive (LA) binarization
- 3 Remove border components

Output: Particle segment mask for the IW (Figure 41g)



FIG. 42. An illustration of intensity-based false positive filtering: (a) local filter output for an IW (Figure 41f) is multiplied by (b) the IW particle mask (Figure 41g) and $\langle I \rangle \cdot \max(I)$ is computed for (c) the masked luminance I map of each particle, as in [158]. The particles highlighted with red dashed lines were identified as false positives.

Once image filtering, segmentation and luminance-based false positive filtering are complete for the IWs from the original images, the filtered IW particle masks must be assembled into full FOV masks (Algorithm 8). Figure 43 shows the stages of this process. IW particle masks for every FOV image are mapped into the FOV (black background) and summed (Figure 43b). Then segment area thresholding and morphological opening (disk structural elements) [165] are performed (Figure 43c). Finally, particle centroids (Figure 43d) are computed for the remaining particle segments.

It is important to note that persistent artifacts within images may be a problem in that they might introduce systematic errors into trajectories output by a tracing algorithm. Notice that one such artifact is present in Figure 39a – a black spot in the right part of the FOV, which is a particle stuck to the channel wall. In such cases, removing these artifacts with texture synthesis-based inpainting [190] is an effective solution. The artifacts in the considered images are readily segmentable from the mean projection over time for a sequence of images using Otsu binarization [166]. Computational performance analysis and particle detection statistics are provided in [185].

E. Particle trajectory reconstruction

1. Modifications to MHT-X

Once particle centroids were obtained for all images, tracking was carried out with the MHT-X algorithm [182]. Since in this case splitting and merging of particles does not occur in the experiment, only particle translation, entry and exit events must be considered. This makes the association constraints (4,5) redundant and the split/merge statistical functions (6,8) obsolete. The conditions from the original implementation for bubbles [182]

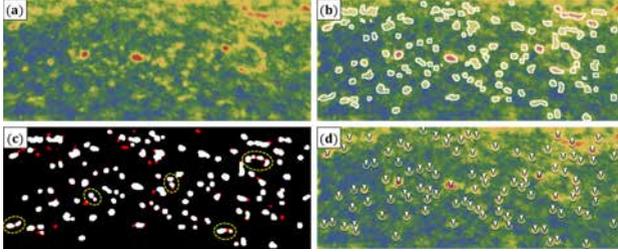


FIG. 43. (a) Colorized luminance-inverted original image, (b) assembled global mask with detected particles overlaid on top of (a), (c) particle segments after area thresholding and small-radius morphological opening using disk elements and (d) particle centroids from (c) indicated with white arrows overlaid on top of (a). Note the red-colored segments in (c) – these are the elements removed from (b) by area thresholding and opening; yellow dashed lines in (c) indicate the segments that were resolved into fragments.

are used with the following changes. First, the association constraints on linear acceleration and deflection (2,3) are used, since constraints on motion are still desirable. However, the parameters are adjusted for the system considered here. Second, the sphere of influence (SOI) used to restrict the association range (only objects with overlapping SOI can be associated) is modified [182]. Instead of defining the SOI about particle locations in frames, a prediction model pinpoints the location \vec{r} of the region that the particle is expected to move to within time Δt and defines the SOI about that point. The prediction consists of the spline extrapolation for particle velocity \vec{v}_s and extrapolation derived from projecting the particle image velocimetry (PIV) field computed in [96] onto particle centroids \vec{v}_{piv} :

$$\vec{r}(t_0 \pm \Delta t) = \vec{r}(t_0) \pm (\alpha \cdot \vec{v}_s + (1 - \alpha) \cdot \vec{v}_{\text{piv}}) \cdot \Delta t \quad (10)$$

where α determines the prediction component weights.

The SOI radius R is based on the velocity magnitude, with higher velocity magnitudes yielding a smaller SOI:

$$R = R_{\text{max}} \cdot \exp\left(-\frac{1}{\lambda_{\text{SOI}}} \cdot \|\alpha \cdot \vec{v}_s + (1 - \alpha) \cdot \vec{v}_{\text{piv}}\|\right) \quad (11)$$

where R_{max} is the upper limit for the SOI radius and λ_{SOI} is a control parameter. If \vec{v}_s is undefined, it and α are set to 0. This effectively assumes that particles with higher velocities are more difficult to deflect and vice

versa, emulating cones of vision for moving particles. If two such cones overlap, an association is formed.

Exit and entry event statistical functions are kept as in 7, except the horizontal x-axis is now the primary one. A model closely resembling the association condition has been implemented for translational motion associations. The translation likelihood estimator consists of three components determined by the location, the linear acceleration and the change in the motion direction.

The location-based likelihood compares the predicted location to the hypothesized location:

$$p_{\text{pos}} = \mathcal{N}(\delta r, 0, \sigma_{\text{pos}} \cdot \Delta t) \quad (12)$$

where $\mathcal{N}(x, \mu, \sigma)$ is a normalized Gaussian distribution with its mean μ and standard deviation σ ; δr is the absolute difference between positions due to the prediction and the hypothesis. The acceleration-based likelihood is calculated as follows:

$$p_{\text{acc}} = \mathcal{N}(a, 0, \sigma_a) \quad (13)$$

The direction-based likelihood component is designed to penalize large changes in the motion direction. The penalty scales with velocity magnitude:

$$p_{\text{dir}} = \mathcal{N}\left(\delta\varphi, 0, \pi \cdot \exp\left(-\frac{\|\vec{v}\|}{\lambda}\right)\right) \quad (14)$$

where $\delta\varphi$ is the change in direction, \vec{v} is velocity and λ is a control parameter. The overall likelihood is computed as a weighted sum of the above contributions.

$$p = \beta_1 \cdot p_{\text{pos}} + (1 - \beta_1) \cdot (\beta_2 \cdot p_{\text{acc}} + (1 - \beta_2) \cdot p_{\text{dir}}) \quad (15)$$

where β_1 and β_2 are weights. Note that while here PIV is used, and in general it is suggested for better tracking, in principle it can be easily disabled in the code if the user does not have the data.

2. Preliminary results

Before proceeding with tracking, the PIV velocity field \vec{v}_{piv} obtained in [96] was interpolated and projected onto the positions of particles detected in each frame (explained in [185]). Particle flow images with \vec{v}_{piv} projections for particles are shown in Figure 44. Note that, according to the \vec{v}_{piv} field,

many particles within the wake often travel in directions opposite (and sometimes normal, as seen in Figure 44d) to the mean flow direction. The obtained \vec{v}_{piv} for particles is used in (10) and (11) for motion prediction.

Figures 45 and 46 present the results of applying MHT-X to the output of image processing. Figure 45 shows some of the reconstructed trajectories within the FOV at four different time stamps. Note that only the last 15 segments of the constructed trajectories are shown. This limitation was introduced for visual clarity, but the trade-off is that the trajectories of slower particles in the wake flow zone are more difficult to show. Despite this, several things can be observed. First, note that trajectories are not broken near the right boundary of the FOV where an image artifact was located before it was removed by texture synthesis inpainting. Second, notice that even with the limitation on the number of segments visible at a time per trajectory, rather long particle tracks can be observed both within and outside the wake flow zone. Third, one can see, especially in Figures 45a and 45d, that densely packed trajectories that cross one another in close temporal proximity are correctly resolved. However, it is also evident that there are quite a few significantly fragmented trajectories, especially within the wake flow zone.

Figure 44 demonstrates the issue of particles being caught within the oscillating wake flow area exhibit both relatively small velocity magnitudes *and* rapid changes in motion direction. This is critical for the current MHT-X implementation since closely packed trajectories with low velocity magnitudes, according to (11) and (14), result in many feasible associations for trajectory connections. With the current spline-based trajectory extrapolation method [182] it is often the case that trajectory fragment mismatch is such that MHT-X opts to assign *Exit* nodes to trajectories prematurely rather than reconstruct longer tracks from fragments. However, MHT-X is still able to resolve quite a number of physically meaningful and long trajectories, examples of which are shown in Figure 46.

Figure 46 shows some of the longer trajectories recovered by MHT-X. In Figures 46a and 46b one can see trajectories of particles that passed by, interacted with and then departed from the wake flow zone. Notice in Figure 46b that as the particle is briefly captured by the wake flow, its velocity becomes lower, as indicated by significantly shorter trajectory segments seen in the middle of the FOV (x direction). This is also seen in Figure 47.

Then, as the particle exits the wake flow, it is again accelerated by the channel flow. Figure 46a, on the other hand, shows that the particle was not entrapped in the wake flow and traversed the FOV much faster. The trajectory in Figure 46c is the longest observed in terms of the number of segments. The underlying particle was first observed and entered the wake

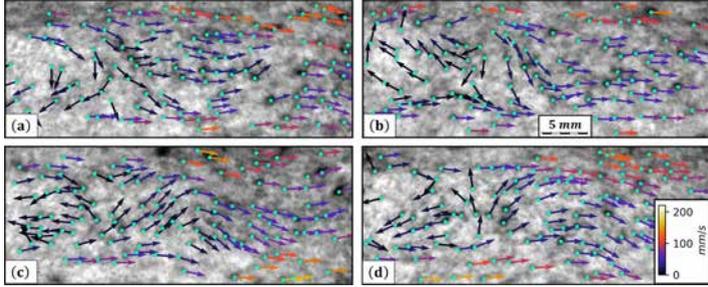


FIG. 44. PIV field projected onto particle centroids at different time stamps. Note the scale bar in (b) and the velocity color bar (normalized for the entire image sequence) in (d).

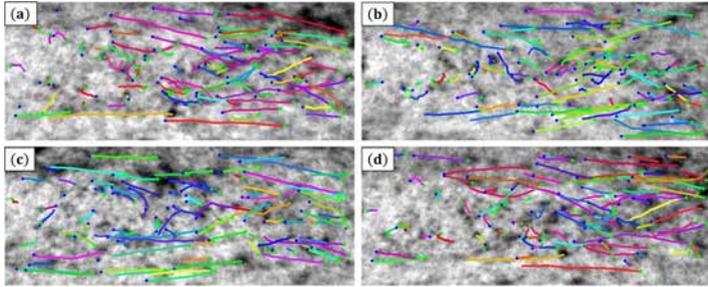


FIG. 45. Snapshots of constructed particle trajectories (indicated with different colors) at different time stamps. Particle *Entry* nodes are indicated with blue dots, while *Exit* nodes are shown as green crosses. Each of the snapshots shows the last 15 segments of reconstructed trajectories. The scale is identical to that shown in Figure 44.

flow zone from the top of the FOV and then had a considerable residence time within the wake before the trajectory was broken off in the left part of the FOV. This particular trajectory is of note for several reasons: first, it clearly shows what is also seen in the PIV projection images – particles captured by the wake behind the obstacle are often diverted towards the center of the wake flow zone and then their direction is reversed such that it is opposite to that of the mean channel flow [44]; second, note the fragment of this trajectory highlighted by a white dashed frame in Figure 46c – one can observe the particle trajectory forming a small loop. It is important that such motion with a low velocity in presence of other potentially interfering particles in the wake flow zone is nonetheless resolved

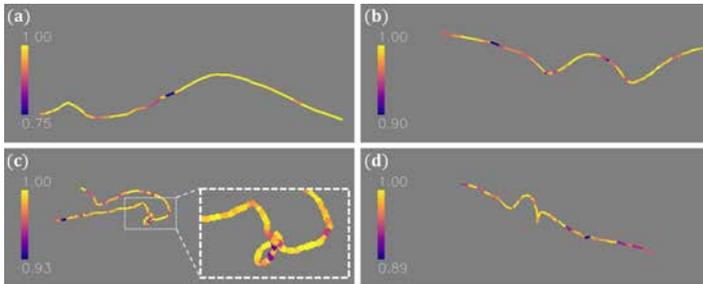


FIG. 46. Representative particle trajectories reconstructed by MHT-X. Trajectory segments are color-coded by their likelihoods. The length scale is identical to that shown in Figure 44.

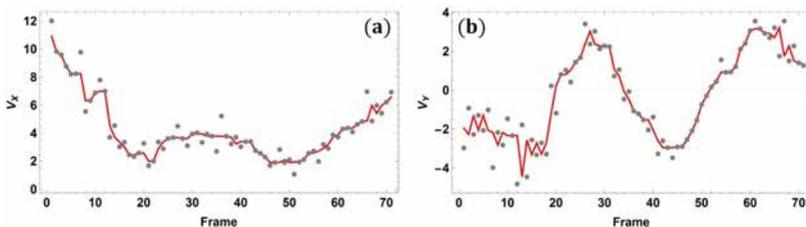


FIG. 47. Velocity time series for the particle with the trajectory shown in Figure 46b: velocity (a) x and (b) y components, pixels per frame. The gray dots are the MHT-X output and the red curves are the median-filtered (1-point radius) velocity components.

with a high degree of confidence – note the minimum segment likelihood is 0.93 (color bar to the left). Figure 46d shows a similar trajectory, except that its motion direction is not reversed during the residence time.

To assess the quality of MHT-X output more quantitatively, several metrics were evaluated for recovered trajectories: segment likelihoods for all segments, mean segment likelihoods and normalized (with respect to mean) dispersion of likelihoods for trajectories, as well as the trajectory size distribution – these are given in Figure 48.

The first thing to note is that most of the segments for trajectories with 4+ nodes (trajectories with < 4 nodes are not usable even for local PTV) have likelihoods mostly in excess of 0.9 with a sharp maximum just below 1 (Figure 48a). This is important since greater likelihoods generally imply less ambiguous trajectory reconstruction and greater confidence that the output is physically accurate. Very high mean segment likelihoods for

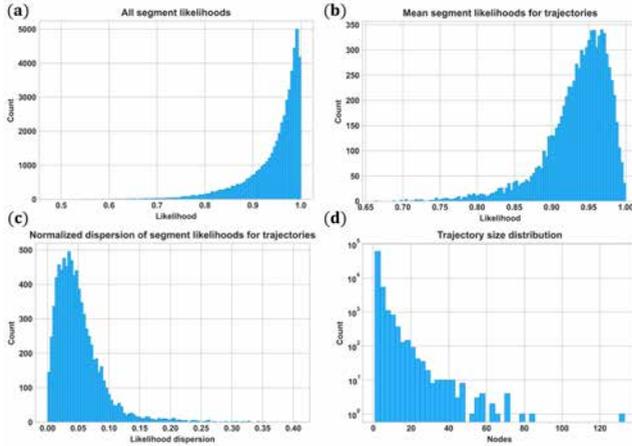


FIG. 48. (a) The likelihoods of all trajectory segments, (b) mean segment likelihoods for all trajectories, (c) standard deviation of segment likelihoods within trajectories, normalized by mean likelihoods, and (d) node count for constructed trajectories. In (a-c) trajectories with 4+ nodes are considered.

constructed trajectories (Figure 48b) and mean likelihood dispersion mostly within 10% of the mean values (Figure 48c) also speak to the quality of the generated results. Finally, Figure 48d indicates that MHT-X produced a few hundred trajectories with ~ 20 nodes and tens of trajectories with 20+ nodes. Note that the trajectory with 130+ nodes seen to the right of the bulk of the distribution is the one shown in Figure 46c. While in-depth physical analysis of particle flow dynamics (residence time within different regions of the FOV, trajectory curvature, etc.) requires more tracks with 30+ nodes (more for slower particles captured by the wake) than are currently produced, shorter tracks can be used for local PTV. The latter could potentially provide insights about the velocity field at smaller length scales than in the case of PIV.

Finally, it was previously assumed in [96] that the utilized gadolinium oxide particles are mostly passive tracers. However, it is worth testing this assumption by estimating the particle Stokes number Stk range. It is found that $Stk \in (0.067; 0.262)$ (the assumptions and derivations can be found in [185]), implying that the particles should be fairly good tracers. Another approach to check this and also qualitatively validate the MHT-X performance is to consider the reconstructed trajectories, perform Fourier analysis of their velocity fluctuations, and check if the fundamental frequency of the cylinder wake oscillation f_0 corresponds to the frequency content extracted

from trajectories f_t . f_0 can be estimated from the Strouhal number St as the inverse of τ_0 – here one has $f_0 \in (3.95; 5.54)$ Hz. On the other hand, taking the first 10 dominant frequencies for 200 longest trajectories and assessing the probability density of the encountered frequencies reveals that there is a distinct peak with a full width at half maximum spanning $f_t \in (2.42; 4.23)$ Hz (Figure 49) which overlaps with f_0 .

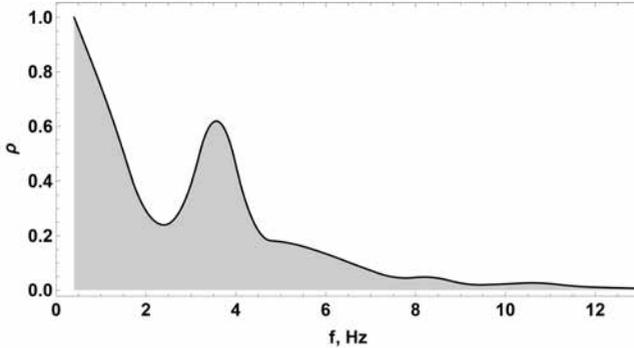


FIG. 49. Smooth normalized density histogram (Scott binning, 2-nd order interpolation) of the first 10 dominant frequencies aggregated for 200 longest (by graph node count) trajectories.

This indicates consistency between the reconstructed dynamics and the expected flow properties. Finally, consider that Stokes’ drag law underestimates the drag force for Re_p that is significantly greater than one, which means that the Stk range derived here is likely to be overestimated.

Therefore, preliminary results indicate that the proposed approach is feasible – for local PTV as is, and for a more in-depth physical analysis of fully reconstructed trajectories, and thus wake flow, after the existing extrapolation scheme [182] is replaced with a better solution. Better interpolation for the PIV field is also desired, as it affects the quality of PIV-based motion prediction.

3. Further extensions for MHT-X

While the extended MHT-X code showed good performance, several aspects of the code needed improvement so that particles could be more reliably tracked throughout the entire imaged FOV [182, 185]. Here the Delaunay triangulation-based cubic interpolation of the PIV field has been replaced with divergence-free interpolation (DFI) which enforces the flow

incompressibility constraint [191–195]. This is necessary for correct PIV-based motion prediction within the obstacle wake where the particles have lower velocity and very oscillatory trajectories, in cases when particles are entering or leaving the wake zone, as well as near the obstacle and channel walls. Improvements have also been made to artifact detection and removal for the neutron images, since artifacts left over from image processing can interrupt trajectories or force incorrect reconstruction thereof. It is shown that the above changes lead to significant increase in particle tracking quality, allowing to reconstruct many trajectories that span the entire imaging FOV; perform PTV and obtain frequency spectra and probability density functions (PDFs) for particle velocity; measure trajectory curvature κ and derive a PDF and a FOV map for κ ; assess particle residence within the FOV – all for quasi stationary flow about a cylindrical obstacle for a range of the cylinder Reynolds number Re_c .

Measuring κ (but also torsion and curvature angle) for particle trajectories and deriving statistics is one of the ways of studying flow in general, and especially turbulent flow, since it yields information about the spatial scales of the flow and how they vary in time. PDFs for κ , torsion and curvature angle are known to exhibit algebraic decay with exponents depending on flow dimensionality, which is something also seen with turbulence energy spectra. There are quite a few instances where this information yields key insights into analyzed dynamical systems [196–211]. In this case the flow is quasi two-dimensional, so torsion calculations are not possible, but one can verify that the observations from particle tracking are adequate by looking at PDFs for κ . If the results are a sufficiently close to the cases where flow is comparable, it means that the presented methods could be applied to geometrical and statistical analysis of flow and turbulence in liquid metals based on particle flow imaging experiments. At the same time, one must point out that such experiment-based data and studies for *liquid metal*, although sought after, to the authors’ knowledge are currently non-existent in literature and one could so far rely only on simulations and theory.

For two-dimensional isotropic and homogeneous turbulence, most relevant studies (experiments, simulations and theory) indicate that for the present system one should expect an algebraic decay with a $k \sim [-2.0; -2.1)$ high- κ exponent, in addition to a $k \sim 0$ low- κ exponent, for the entire Re_c range and persistent over time [196, 201–205]. It was also observed that the κ PDF $k = 0$ region is shifted towards smaller values as flow Re increases and, although not noted explicitly, the $k = -2$ region of the PDF seems to be, conversely, shifted upwards with Re , although the increments seem smaller in comparison to the $k = 0$ case. With this in mind, the above-mentioned flow characteristics are presented for a range of Re_c , compare the results against existing data, and also show how tracking performance of the shown

approach changes with Re_c for a fixed set of parameters for the utilized MHT-X tracking code.

MHT-X performs several scans over the data set (*graph sweep*) with progressively increasing time windows (*window size sweep*) within which trajectory reconstruction is takes place. This enables one to employ motion models which rely on connections between objects made backwards and forwards in time, i.e. make use of *context* due to the associations established between objects during preceding graph sweep iterations. Note that such information is not available *a priori* from the input data. The drawback of such models is that special cases must be defined for when not enough context is available. This results in edges potentially being formed with (arguably) different models during each graph sweep (Figure 50). While this cannot be avoided completely, it can be remedied with trajectory re-evaluation, which is a step taken after every complete graph sweep. It recalculates the likelihood for every edge in the trajectory graph.

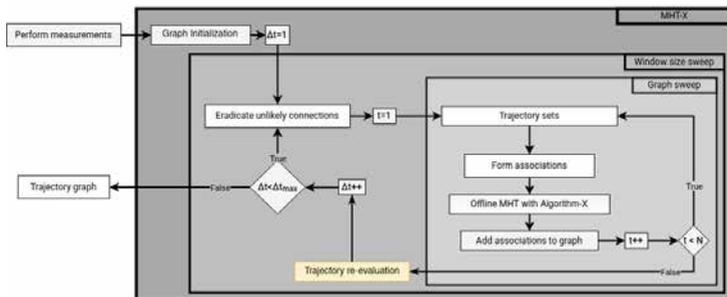


FIG. 50. A flowchart for the updated MHT-X algorithm (versus Figure 1 in [182]) with the new re-evaluation step.

This step is necessary because the connections formed during earlier graph sweeps will have been formed with less context, and might, in retrospect, have overestimated likelihoods. Previously, if an absurd connection for some reason was assigned a high likelihood, it could not be eliminated because a likelihood was assigned only upon forming a connection. Now, after a trajectory is complete, it is checked how well each of its connections fits with the rest. This is done after every complete graph sweep. The procedure is as follows: every trajectory is iteratively split in two at every edge by temporarily deleting the currently tested edge. The likelihood of the temporarily deleted edge is then recalculated based on the two resulting trajectory parts by passing them to the statistical functions (shown in [182] and [185]). Similarly, the trajectory *Entry* and *Exit* connection likelihoods

are also reevaluated.

Previously, when using interpolated PIV fields for particle motion prediction, the PIV vector field components were treated as two independent scalar fields. Interpolation was necessary because PIV yielded vector values on a sparse regular grid and particle positions could be anywhere within the FOV of interest with sub-pixel coordinates. For each of the two PIV components, Delaunay triangulation was performed for the PIV point grid and cubic interpolation was used for particle centroids that were within triangles formed by nearby PIV grid points, while nearest neighbor interpolation is used otherwise. Then the interpolated velocity field was projected onto particle positions. This rather naive approach provided acceptable results. However, here an improved approach using divergence-free interpolation (DFI) of the PIV data is presented. The reason for adopting this method is that the previously used interpolation method is not physics-based and imposes no constraints on the interpolated field. Meanwhile, DFI is specifically designed such that the interpolated field analytically satisfies the flow incompressibility constraint, which was shown to have a significant impact, especially in cases with particle-laden flow [191–195]. In this case, the flow is quasi two-dimensional, so one can expect near-zero divergence.

The div-free interpolant uses matrix-valued radial basis functions (RBFs) and is described in [191]:

$$\Phi(\vec{x}) := (-\nabla^2 I + \nabla \otimes \nabla)\phi_\delta(\vec{x}) \quad (16)$$

where I is the identity matrix. This interpolation method is favorable because it supports arbitrarily sampled data and can be made computationally cheaper with a multilevel approach [193]. The current non-multilevel implementation (to be upgraded to multi-level in the future) utilizes scaled RBFs:

$$\phi_\delta(\vec{x}) := \delta^{-d}\phi_{\nu,k}\left(\frac{\|\vec{x}\|}{\delta}\right) \quad (17)$$

where δ is the support radius, d is the number of motion dimensions of the system and $\phi_{\nu,k}$ is the RBF as defined in [192]. The DFI approach also enables one to define boundary conditions (BCs) within the domain where object tracking occurs. This is important to account for because in addition to the zero-divergence constraint, the flow field must also comply with the no-slip BCs at the channel walls and the perimeter of the obstacle. The no-slip BC was set at discrete, uniformly spaced points for each of the walls and for the obstacle perimeter (black lines in figure 51a). The addition of

this auxiliary information enables us to correctly model boundary layers near the obstacle and walls and eliminate streamline intersections with the obstacle and/or walls. That is, without the BCs for the DFI and based on MHT-X motion constraints and statistical functions alone, the particles could, for instance, "tunnel" through space where the cylindrical obstacle is. Thus, improved tracking performance of MHT-X with DFI for the PIV field can be attributed to several factors, such as the elimination of sources and sinks in the PIV field, greater smoothness and the additional physical constraints from BCs.

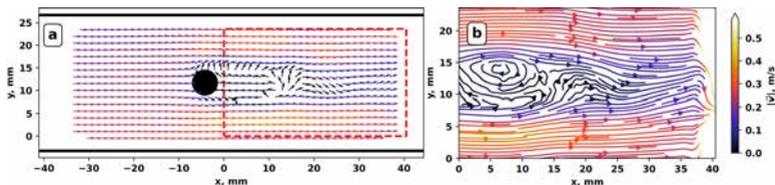


FIG. 51. (a) Vector field acquired with PIV. Black lines are the boundary conditions. The red striped line is the outline of the region of interest, where tracing was performed, and where the streamlines are drawn. (b) Streamlines of resulting interpolation. There is a noticeable artifact on the right side of the interpolation, where the flow abruptly changes direction, and it's velocity rapidly increases (streamlines are not drawn for values too high) due to a lack of PIV data. This error is insignificant as this region is of no significance for the analysis.

It has been previously noted that persistent artifacts within images (e.g. stuck particles) may be a problem in that they might introduce systematic errors into trajectory reconstruction (e.g. trajectory fragmentation). Such artifacts were removed with texture synthesis-based inpainting [190] after segmenting them from the inverse mean projection of the image sequence using the Otsu method [166]. However, as processing more data showed, this segmentation approach is not robust enough (at least for this case). Instead, the method outlined in Algorithm 11 was implemented.

This, incidentally, also allows one to easily detect the obstacle in the flow channel by applying morphological opening [165] and then size thresholding to the output of Algorithm 11. Obstacle segmentation is necessary to prescribe the no-slip BCs for the DFI. The above method is more robust mainly because Step 3 yields a much greater CNR for the artifacts than using just the inverse mean projection, and an optimal manual threshold can be found very quickly based on the initial Otsu method guess (which in some cases can even be enough). In certain instances, though, the CNR of the obstacle can be too low for Algorithm 11, so Step 6 is replaced with K-medoids clustering segmentation [212], after which one can select the

Algorithm 11: Artifact detection

- 1 Compute the inverse of the mean projection for the image sequence
 - 2 Compute the standard deviation projection for the image sequence
 - 3 Divide the output of Step 1 by that of Step 2
 - 4 Perform color tone mapping
 - 5 Apply Gaussian total variation and invert the result
 - 6 Binarize with a manual threshold using the Otsu method output as an initial guess
 - 7 Apply the filling transform and (optionally) size thresholding
 - 8 Perform texture synthesis in the resulting artifact areas for all frames in an image sequence
-

level set containing the cylindrical obstacle and isolate the latter using (in this order) morphological opening, filling transform, morphological dilation [165] and again opening.

4. Further results

To make sure the performance demonstration for the updated methods and code is fair, it was decided to process all image sequences corresponding to the $Re_c \in [988; 4147]$ range using identical settings for both the particle detection code and MHT-X. Image processing is done as described in [185] with parameters identical to what is specified in the paper. For MHT-X, the motion models are as described above and in [182, 185], and their parameters are given in [213].

It is found that the mean particle size seen after image processing is $\langle d_p \rangle = 0.43 \pm 0.13 \text{ mm}$ versus the expected $d_p \in (0.3; 0.5) \text{ mm}$ [96], which is in good agreement. There is a slight bias towards larger particle sizes can be explained by a lower bound on the particle size resolution imposed by the image filter settings used in and adopted from [185]. Given that identical parameters are used in all cases, tracking performance degradation expected as Re_c increases (particle displacements become greater, trajectory fragment connection ambiguities are more frequent) should be clearly observable.

A very strict criterion is defined – only trajectories with 20+ graph nodes (instances over time) are considered *eligible trajectories*. It is strict in that such trajectories will be either the ones that travel past the wake flow zone or interact with it without particle entrapment, and extend throughout the entire FOV length (expected with the chosen frame rate and characteristic particle velocity [185]), or the ones that are caught by the wake and reside inside it over an extended time interval, capturing the small-scale flow

perturbations of the wake stagnation zone and/or its tail which oscillates as vortex shedding occurs. While one can argue that shorter tracks could be used for temporal and geometric statistics acquisition, as well as local PTV, it is best to show how the amount of very long, not just any tracks, changes with Re_c , since it is the longer tracks that encode much more information about the spatial and temporal character of turbulent flow and better capture the lower oscillation frequencies.

As a measure of performance, the ratio of the number of eligible trajectories N_t to the number of frames N_f in an image sequence is used. Figure 52 shows how this ratio changes with increasing Re_c .

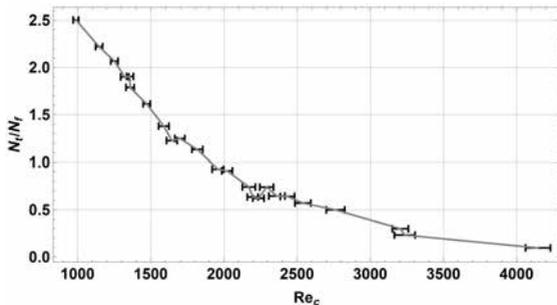


FIG. 52. Tracking performance of MHT-X: the ratio of total eligible trajectories N_t to the number of frames N_f for a range of Re_c .

Note that N_t in this case typically represents 5% to 10% of the data (which is on average several thousand trajectories), but this is because the MHT-X outputs trajectories of all lengths, including 1-node trajectories which in the MHT-X framework should be interpreted as eliminated false positives. The eligible track percentages are much greater if tracks with < 5 nodes are discarded. Observe that the performance degrades at a non-constant rate: it is found that N_t/N_f fits x^n with $n = -2.13 \pm 0.11$ and $R^2 = 0.94$, but does not fit an exponential function ($R^2 \sim 10^{-3}$) which one might expect if the performance was mostly tied to the velocity-dependent exponential scaling factors in the MHT-X motion model equations [185]. To improve tracking for higher Re_c , one could try and circumvent this by making the respective motion models less restrictive, i.e. allow greater angle changes for faster motion. However, this clearly cannot be done indefinitely, since beyond certain Re_c values one will simply need a higher frame rate for imaging.

Figure 53 shows representative eligible trajectories. Note that they capture particle motion in the wake stagnation zone, in the wake tail, fly-by

particle interactions with the wake, as well as the particles that are mostly unaffected. The amount and node count, as well as the physical lengths of eligible tracks in general by far exceeds the best one could previously achieve [185]. While improved artifact detection and cleanup play an important role, most of the improvement can be attributed to a more accurate PIV-assisted motion prediction stemming from a much better interpolated PIV field yielded by DFI.

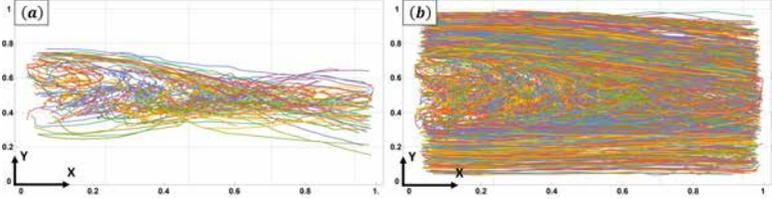


FIG. 53. Examples of eligible trajectories reconstructed within the FOV: (a) the longest 150 and (b) 1500 trajectories. Colors represent different tracks, and axes tick values are normalized to the respective image dimensions.

As before, one can measure particle displacements between frames, but this time there are enough good trajectories to accumulate particle velocity statistics – Figures 54 and 55 show the PDF for streamwise v_x and transverse v_y (axes as in Figure 53) velocity components for all eligible trajectories for each Re_c . Of course, as the tracking performance degrades with Re_c approaching its maximum, the PDFs become less smooth and more noisy, and in the case with $Re_c = 4147$ some PDF bins have much smaller weights than they should have (not shown in Figures 54 and 55). To assess the physical changes in the PDFs as Re_c increases, one could think of the PDFs this way. Particles can be, very roughly, divided into two classes: the ones that are trapped or simply spend a lot of time within the obstacle wake (*captured particles*), and the ones that either pass by unaffected or only weakly interact with the wake (*free particles*), the latter group being faster than the former. One can then expect two things – first, as Re_c increases, more of the would be free particles will become captured and will have much lower velocities with a distribution centered at near zero velocity, with the velocity dispersion increasing as Re_c does; second, the free particles that remain such will have greater velocity overall and also greater velocity dispersion within their class.

This means that one should expect, at very low Re_c , that there is only be one maximum in the v_x PDF, associated with the free particles, since the wake will not have developed sufficiently yet. As Re_c increases, it is expected that the initially solitary PDF maximum will begin to split in

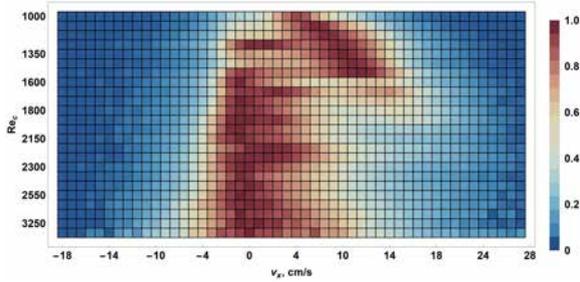


FIG. 54. Separately binned (Freedman-Diaconis method) and normalized PDFs for the streamwise velocity component v_x of eligible trajectories for a range of Re_c .

two as the wake begins to trap particles (the stagnation zone and the tail expand) and the class of captured particles becomes distinguishable. At higher Re_c , the residence time of particles captured by the wake increases greatly, and the corresponding PDF maximum should begin to dominate the free particle class. At the same time, velocity dispersion about both maxima should increase, and the free particle maximum should shift towards greater velocity values. However, since tracking performance degrades with Re_c (Figure 52), one would also expect that the free particle maximum and its neighborhood within the PDF become fainter, since more and more of the faster particles are not properly connected into longer tracks that pass the eligibility criterion. Figure 54 exhibits all of the above trends. The splitting of particles into two classes cannot be seen in greater detail because the Re_c sampling is not very fine in the $Re_c \in \sim [1000; 1200]$ interval, but the transition from free to captured particle domination in the PDFs can be seen fairly clearly in the $Re_c \in \sim [1300; 1800]$ range. The dispersion of the captured particle velocity is asymmetric with a positive v_x bias that becomes more prominent with increasing Re_c since particles enter and leave the wake zone at higher velocities overall.

As expected, Figure 55 shows an increased dispersion for the v_y PDF about zero velocity, which is both due to the more intense velocity field fluctuations in the wake stagnation zones, and the stronger oscillations with higher frequency in the wake tail. The fact that this and the above-mentioned expected trends are observed suggests that particle detection and tracking are performed adequately.

However, this statement still requires further proof and reinforcement. Given the PTV data generated from the eligible trajectories, one can also check if the vortex shedding frequency f_0 of the obstacle wake is

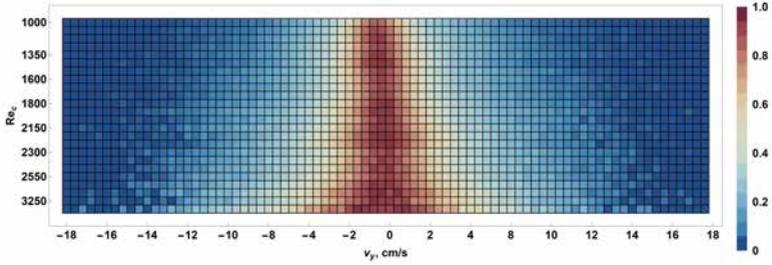


FIG. 55. Separately binned and normalized PDFs for the transverse velocity component v_y of eligible trajectories for a range of Re_c .

captured by MHT-X properly – this is important if one plans to use the presented methods to assess the temporal characteristics of turbulent flow in experiments with liquid metals. Therefore, f spectra are computed for v_x time series (filtered with the median filter with a kernel width of 1 point for outlier removal, i.e. low-pass filtering) of all eligible trajectories for the considered Re_c range. Then, for each Re_c , the f spectra of all trajectories are aggregated by binning all detected f instances using the absolute values of the Fourier coefficients as weights. Doing this for the entire Re_c range and individually normalizing the resulting f PDFs, one obtains Figure 56.

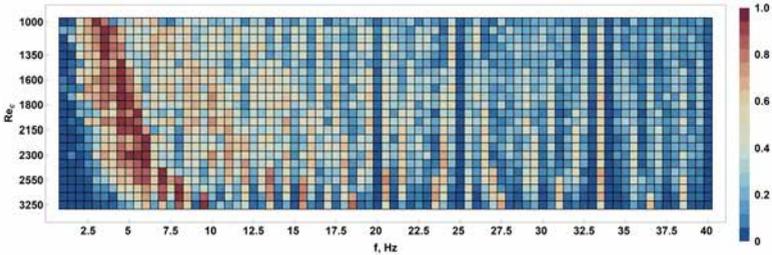


FIG. 56. Frequency f PDFs for v_x over a range of Re_c , aggregated from v_x Fourier spectra of all eligible trajectories in each case. PDF normalization is separate for every Re_c value. Aggregation is done by binning all f instances over the eligible trajectories using Fourier coefficients as weights with a 200-bin resolution.

Notice the lower- f band with a strong maximum that shifts towards greater f values as Re_c is increased – this corresponds to the vortex shedding frequency f_0 . Since the Re_c range was sampled non-uniformly during the imaging experiments, and there are some Re_c instances with rather similar

values, one cannot see a clear diagonal in Figure 56. However, since the observed shifting band produces the most intense maxima, one can take the maximum PDF values for each Re_c , get the corresponding f , and plot the values versus Re_c , comparing the experimentally determined f_0 against what one would expect theoretically. This comparison is shown in Figure 57 where the expected f values were computed from the Re_c range assuming a constant Strouhal number $Sr = 0.197$ [213]. Both visually and upon inspecting inset (a) which shows the relative error histogram for the Re_c range, one can see that the agreement is indeed sound. However, the reader is asked to take the good f match for the maximum Re_c value as an accident, since the corresponding f PDF (not shown in Figure 56) has an extremely low SNR. The f_0 peak for the penultimate Re_c value, however, is very clear despite the objectively worse tracking performance as opposed to the lower Re_c cases. One can also check the quality of f_0 determination with a linear fit of the experimentally derived f_0 . For a constant Sr , one has $f_0 = Re_c \cdot St / \rho d^2$ (derivation shown in [213]). Fitting experimental f_0 values, one gets $Sr = 0.196 \pm 0.009$ which agrees with the expectations.

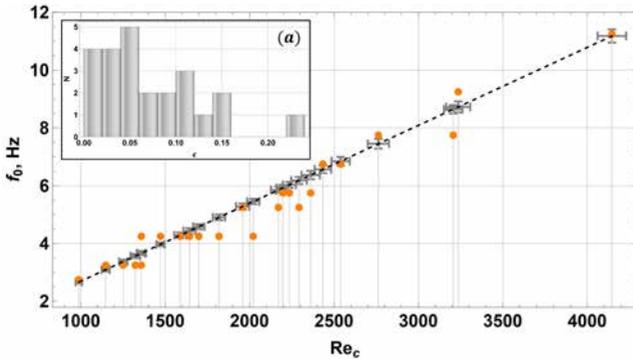


FIG. 57. Vortex shedding frequency f_0 : theoretically expected (gray dots with error bars) and derived from the obtained aggregated frequency PDFs (orange dots) by picking the dominant frequency peak for each Re_c . The black dashed line indicates the expected linear $f(Re_c)$ trend, and the inset (a) shows the histogram for the relative errors (ϵ) between theoretical and experimental f_0 values.

Spatial characterization of turbulence is also of great interest for liquid metals, therefore one must find out if MHT-X constructs trajectories with enough quality, enabling the derivation of trajectory geometry statistics with sufficient accuracy. To this end, curvature κ is computed along every eligible trajectory (procedure explained in detail in [213]). All the curvature values determined for trajectory points are normalized to the inverse mean

particle size $\kappa_0 = 1/\langle d_p \rangle$, and a PDF for κ is computed in the double-log₁₀ domain (Freedman-Diaconis binning). This is done for the entire Re_c range, yielding the PDFs shown in Figure 58.

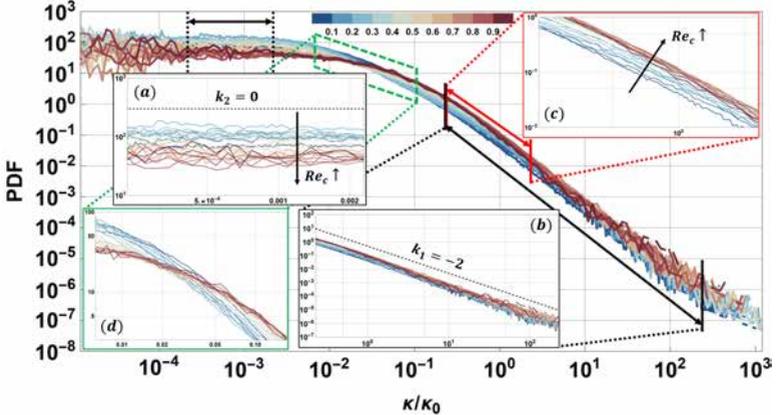


FIG. 58. Trajectory curvature κ PDFs (normalized to the inverse mean particle size κ_0) for all eligible trajectories for different Re_c with the curves color-coded by their respective Re_c values (color bar 0 corresponds to the minimum Re_c and 1 is the maximum), increasing from blue to red. Note the insets (a) and (b) which highlight the constant ($k_2 = 0$ reference dashed line) and algebraic decay ($k_1 = -2$ reference dashed line) intervals of the PDFs, respectively. Inset (a) shows the downwards shift of the PDF curves in the $k_2 = 0$ interval as the Re_c value increases, and an inverse trend can be seen in inset (c) for the $k_1 = -2$ region. Notice also the trend reversal region highlighted with a green dashed frame and shown in greater detail in inset (d).

There are several key features of the PDFs that must be noted. Firstly, one can see very clearly, especially in insets (a) and (b), that the PDFs indeed exhibit a near-constant $k_2 = -0.05 \pm 0.05$ low- κ region (with a much smaller SNR below $\kappa/\kappa_0 \sim 2.5 \cdot 10^{-4}$) and an algebraic decay region with $k_1 = -2.09 \pm 0.04$, both of which are rather close to the expected values. In addition, two trends can be observed: in the low- κ interval, the PDF curves shift downwards as Re_c increases, as shown in inset (a) – this matches what was measured experimentally in [205]; for the high- κ region, an inverse trend is seen in inset (c) – this, while not stated, can still be seen in [205] despite PDFs shown therein being compressed (without modifying $k_{1,2}$) by scaling the curvature values with respect to velocity and acceleration variations. Finally, one can clearly see the Re_c trend inversion in inset (d), also observed in [205]. The fact that these elements of the κ

PDFs and their universal scaling over the Re_c range considered herein were successfully reproduced serves to further validate the applicability of the shown methods to liquid metals.

It is also checked that it is not just the κ PDFs that are computed correctly, but also the κ maps over the FOV. To do this, κ maps are generated from sparse κ fields aggregated from all eligible trajectories for a given Re_c (procedure explained in [213]). An example for one of the analyzed image sequences can be seen in Figure 59 where a color map of $\log_{10}(\kappa/\kappa_0)$ is shown. The higher κ values are concentrated in the obstacle wake zone, as expected, with faint narrower filaments extending over to the right side of the FOV. The filament-like structure seen in Figure 59a can be understood by applying color tone mapping to the $\log_{10}(\kappa/\kappa_0)$ map before colorizing it – compressing the map value histogram reveals clear curves within and outside the wake zone that are nothing other than visualized trajectories, where yellower spots and filaments are higher- κ fragments of trajectories (Figure 59b).

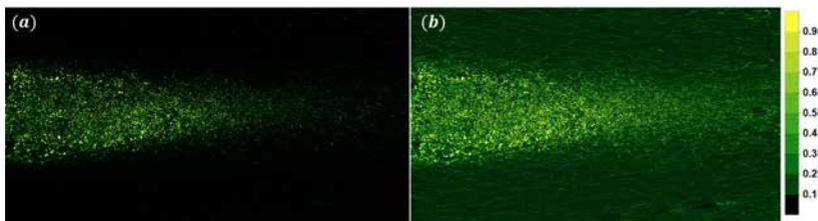


FIG. 59. An example of a trajectory curvature κ map in the imaging FOV obtained from eligible trajectories. (a) Shows the color map of normalized $\log_{10}(\kappa/\kappa_0)$ and (b) is its version with color tone mapping applied.

Interestingly, the κ distribution patterns seen in κ maps for different Re_c are rather similar (shown in [213]). While the X (streamwise) profiles are quite similar, the Y profiles' width slightly scales with Re_c , indicating the obstacle wake expands transversely as the free-stream velocity increases, which makes sense. Figure 60 reveals that the curvature (in the semi- \log_{10} domain) has a plateau with a very small slope that extends from the edge of the cylindrical obstacle ($X = 0$) almost to the middle of the FOV, and then quickly falls off with X . The Y profile is, as expected, quite symmetric.

Of course, another concern one must address is that the κ statistics shown in Figure 58 could be affected by particle collisions. While one would not expect a significant amount of collisions in any given Re_c case, one should provide at least a crude estimate for how much of an effect collisions could have. To do this, a primitive collision estimation model (explained

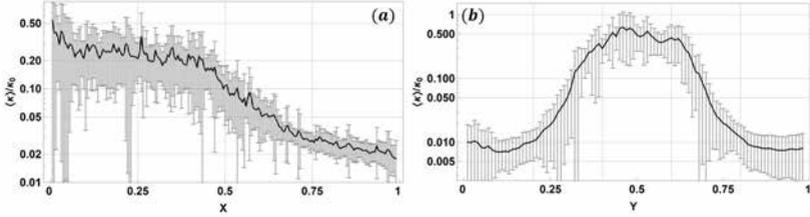


FIG. 60. Mean κ/κ_0 profiles across the Re_c ranges over (a) X and (b) Y directions in the FOV.

in [213]) is used, which considers particles from eligible trajectories which are simultaneously within a given frame. With this primitive model one can roughly estimate how often the collisions might occur, and where – this is reflected in Figure 61 where the mean particle collision count per frame $\langle N_c \rangle$ and its ratio to the mean number of particles per frame $\langle N_p \rangle$ are shown for the considered Re_c range.

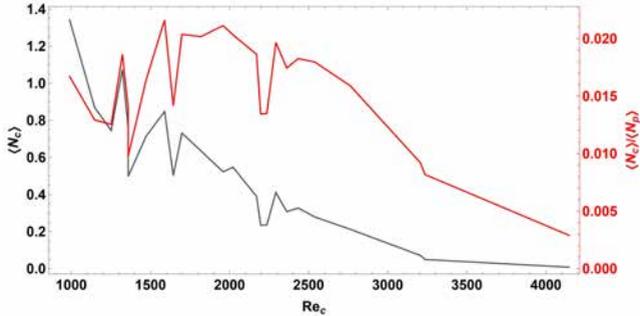


FIG. 61. Mean particle collision count per frame $\langle N_c \rangle$ (gray) and its ratio to the mean number of particles per frame $\langle N_p \rangle$ (red) for the considered Re_c range.

The reader is reminded that one must also consider the degradation of MHT-X tracking performance with increasing Re_c , which explains the rather sharp falloff for $\langle N_c \rangle / \langle N_p \rangle$ past the $Re_c \sim 2500$ mark. However, one can see that even for $Re_c < 2500$ where one should be able to see potential collisions with the above simple model, they are quite rare, especially considering that there are on average ~ 250 particles within the FOV simultaneously. An example of how collisions estimated for an entire image sequence are distributed within the FOV are shown in Figure 62.

It is evident that the overwhelming majority of collisions occur within

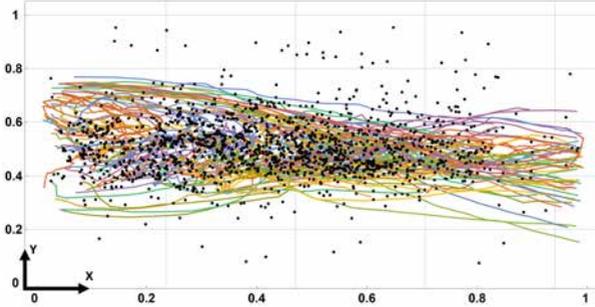


FIG. 62. An example of collision locations (black dots) estimated within the FOV. Trajectories colored by their IDs are shown for context.

the wake, particularly in the stagnation zone, where particles are subject to sharp and rather chaotic perturbations in the velocity field. This is considered in more detail in [213]. The collision number density distributions are quite consistent between the different Re_c values.

Knowing how often and where collisions occur the most, one can estimate their impact on particle motion and trajectories, particularly their κ statistics. One can estimate the collision energy E_c from the relative velocity of colliding particles projected onto the line connecting their centroids. To make it more straightforward yet fair, it is assumed the worst-case scenario and assign the total E_c (particle mass times projected relative velocity) to *both* particles of the colliding pair. Since the particle size deviation σ_d is quite small, one can also assume that all particles have *the same* mass. While E_c itself is not informative, its ratio to the initial kinetic energy E_0 of each particle in a pair should determine to what relative extent the kinetic energy changes and therefore how affected the trajectory shape can be. Computing E_c/E_0 for every particle estimated to participate in a collision, then constructing a PDF in the double- \log_{10} domain (Freedman-Diaconis binning) and doing this for each Re_c case, one has the E_c/E_0 PDFs as seen in Figure 63.

The E_c/E_0 PDFs also exhibit algebraic decay and two distinct regions (high- and low-ratio) – but unlike the κ/κ_0 PDFs in Figure 58, algebraic decay occurs in both of the regions. It is tempting to say that insets (a) and (b) show PDF shifting with Re_c like what is seen for κ/κ_0 in Figure 58, but the SNR here is insufficient for us to state this with confidence anywhere except perhaps the $E_c/E_0 \in \sim (10^{-2}; 10^{-1})$ interval. However, the Re_c PDF shift trend inversion similar to that in Figure 58d is rather clear in inset (c). As for the algebraic decay intervals, here one has a low-ratio

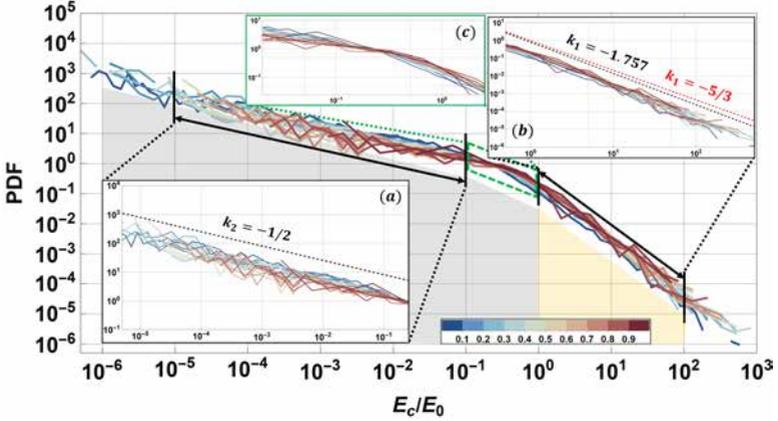


FIG. 63. Collision to initial kinetic energy ratio E_c/E_0 PDFs for all estimated collisions for different Re_c . The curves are color-coded by their respective Re_c values (0 corresponds to the minimum Re_c and 1 is the maximum), increasing from blue to red. Insets (a) and (b) highlight the $k_2 = -1/2$ (reference dashed line) and $k_1 = -1.757/k_1 = -5/3$ (reference black and red dashed lines) intervals of the E_c/E_0 PDFs, respectively. Inset (c) shows the Re_c PDF shift trend inversion like the one seen in Figure 58d. Gray- and yellow-tinted areas represent the cumulative contributions of the low- and high- E_c/E_0 parts of the PDFs, with a boundary at $E_c/E_0 = 1$.

exponent $k_2 = -0.498 \pm 0.009$ (almost $1/2$, although the fit quality is not so good with $R^2 = 0.89$) and a high-ratio exponent $k_1 = -1.757 \pm 0.011$, the latter being very close to $-5/3$, which is a very interesting coincidence. Notice that the transition between the k_2 and k_1 algebraic decay exponents occurs very close to $E_c/E_0 = 1$, starting at $\sim 10^{-1}$ which is, roughly speaking, the threshold beyond which collisions should start noticeably affecting particle trajectories. However, it is abundantly clear from Figure 63 that the dominant contribution to the respective cumulative distribution function stems from low energy ratio collisions. This means that the κ statistics derived from the tracks should indeed encode mostly the effects of turbulent pulsations within the liquid metal flow. Because $\langle N_c \rangle / \langle N_p \rangle$ is very low, and there is a very wide range of E_c/E_0 values, the SNR of the E_c/E_0 profiles for an individual Re_c is also not high enough to determine where the E_c/E_0 maxima are in the FOV within acceptable error margins [213]. To do this, especially for greater Re_c , significantly longer imaging time is required per image sequence. Finally, if one wishes to study collisions in-depth, a rigorous approach is required and ideally collisions should be

built into the MHT-X motion models. One can, however, assess where in the FOV the particles spend more time, which is shown in Figures 64-66.

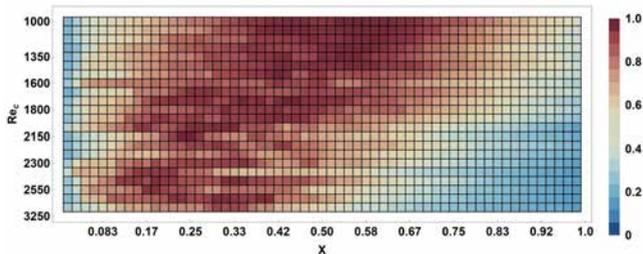


FIG. 64. Mean particle number density ρ_p over X for the range of Re_c . The color map encodes normalized (individually for each Re_c) median filtered (kernel radius is 1 point) ρ_p .

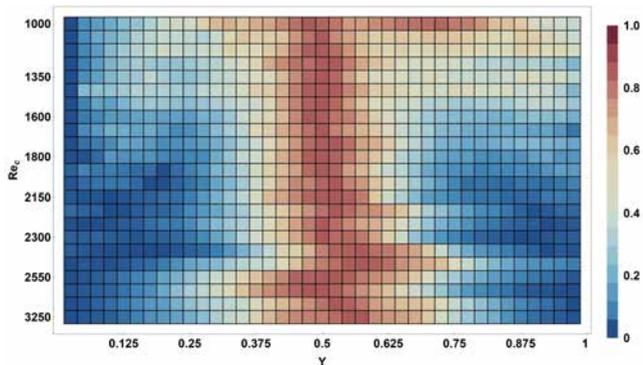


FIG. 65. Mean particle number density ρ_p over Y for the range of Re_c . The color map encodes normalized median filtered ρ_p .

Figure 64 suggests that, as Re_c increases, more and more particles are trapped within the wake, and specifically its stagnation zone, therefore the maximum of the particle density (ρ_p) profile shifts towards smaller X . Averaging over Re_c one also find that the stagnation zone is the most particle-populated area of the FOV. Since the frame rate is constant, ρ_p corresponds to the particle residence time. However, note that due to tracking performance degradation for greater Re_c , some of the particles that travel with faster velocity and only weakly interact with the wake tail will be lost. This is why the ρ_p profiles for higher Re_c do not extend as

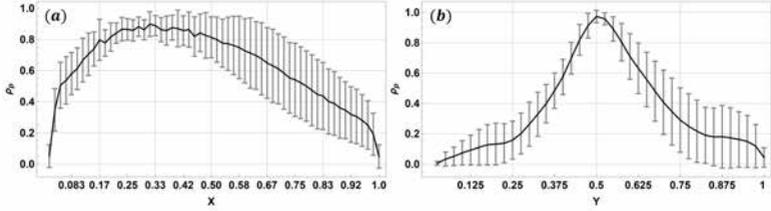


FIG. 66. Mean normalized particle number density ρ_p profiles over (a) X and (b) Y within the FOV for the considered Re_c range.

much beyond the $X = 0.6$ mark. Figure 65, as expected, implies that with increasing Re_c the relative particle residence time, which ρ_p represents, becomes more at the center ($Y = 0.5$) and less away from it, with greater dispersion for higher Re_c due to wake width increase.

Finally, one can use the PTV results from MHT-X to obtain a continuous velocity field, an example of which can be seen in Figure 67, with a spatial resolution superior to that of PIV. This is done using DFI and trajectory resampling, as explained in [213], accounting obstacle and channel wall BCs.

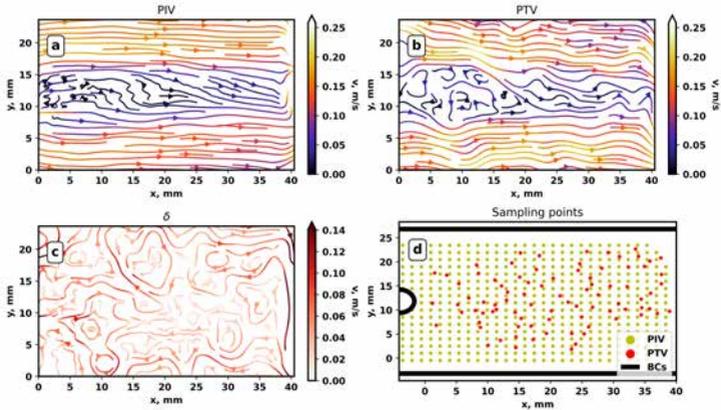


FIG. 67. PTV-based reconstruction of a continuous velocity field in the FOV using DFI: (a) DFI-PIV field used for motion prediction; (b) PTV-based continuous field yielded by DFI; (c) the difference between PIV- and PTV-based fields; (d) sampling points for PIV and PTV.

Indeed, as seen in Figure 67, the resulting PTV-based field (b) indeed has

a greater resolution than PIV (a) and resolves the finer vortices within the wake much better, while still preserving the larger-scale flow features. The difference between the PIV field (which could be considered as the initial guess) and the refined PTV-based DFI field are shown in (c). Note that temporal averaging is not performed – (b) shows an unfiltered instantaneous velocity field reconstructed based only on the points in that frame (d).

To summarize, improved analysis methods have been shown for particle-laden liquid metal flow imaged with neutron radiography, although the shown approach is readily extendable to other modes of imaging as well as physical systems. Specifically, it has been demonstrated how adding a trajectory re-evaluation step to the MHT-X object tracking code, augmenting particle image velocimetry-assisted motion prediction with divergence-free interpolation which enforces boundary conditions and flow incompressibility, as well as upgraded artifact removal methods for raw images results in a significantly improved particle tracking quality. The modified MHT-X has then been validated by showing that it can correctly measure the characteristics of turbulent particle-laden liquid metal flow.

5. Summary

A framework has been developed for particle detection and tracking in liquid metal flow based on neutron radiography images. The solution includes an image processing code leveraging the methods developed for bubbles in Section II A 2 b, as well as the modified MHT-X shown in Section II C 2, which has been extended by implementing DFI-PIV-based motion prediction for particles. The framework has been validated, and it was shown that it is capable of providing tracking performance that enables one to analyze spatio-temporal characteristics of turbulent particle-laden liquid metal flow. The code can reconstruct both particle trajectories, and a continuous velocity field, from rather spatially sparse data.

Of note is the non-local means masking (NMM) method developed for correlated noise removal in images with particles. NMM is a generalized, more efficient version of unsharp masking and is also likely to be broadly applicable outside the present application.

All the developed tools are open-source. The image processing code is available at *GitHub*: [Mihails-Birjukovs/Low_C-SNR_Particle_Detection](#). MHT-X can be found at *GitHub* as well: [Peteris-Zvejnieks/MHT-X](#). There is also a *GitHub* repository for the implementation of divergence-free interpolation [Peteris-Zvejnieks/DivergenceFreeInterpolation](#), which is available as a *Python* PyPi package as well.

F. Dynamic mode decomposition of bubble chain flow

To analyze the simulation results for bubble flow systems, a custom implementation of DMD has been developed. As a test case, it was decided to consider the system that corresponds to the one used in the neutron imaging experiments considered here.

1. Modelled physical system

The physical system of interest is vertical argon bubble chain flow in liquid gallium in a rectangular glass vessel, $150 \times 90 \times 30$ mm [152, 214]. Bubbles are injected vertically/horizontally at the bottom and ascend with acceleration due to buoyancy, exhibiting zigzag trajectories with out-of-plane perturbations. The free surface of gallium is at 130 mm. Static horizontal MF is applied using a system of permanent magnets and an iron yoke wherein the liquid metal vessel is placed [152, 153] and the resulting MF configuration is illustrated in [152, 214]. Four cases are considered to illustrate how DMD can be used to assess the effects of applied MF and varying gas flow rate at the inlet:

- 30 sccm flow rate, no MF
- 100 sccm flow rate, no MF
- 30 sccm flow rate, ~ 265 mT HMF in the bubble flow region
- 100 sccm flow rate, ~ 265 mT HMF in the bubble flow region

where *sccm* stands for *standard cubic centimeters per minute*. Note that 30 sccm flow rate will result in an effectively single-bubble flow regime, which will serve to demonstrate the differences that collective dynamics introduce into the system at 100 *sccm*.

Data for the DMD analysis is generated by simulating the above system under the listed conditions using the previously outlined numerical model [152, 153] in the $Rm \ll 1$ (magnetic Reynolds number) approximation. This is done to exclude the induced contribution to the overall MF – systematic studies of this system including induced MF for higher flow rates is outside the scope of this thesis and is reserved for future papers. The value intervals for other relevant dimensionless groups are provided in [153].

Following [152, 153], bubble chain flow was simulated using *OpenFOAM*, using the *interFlow* solver. As in [152, 153], the k -equation SGS LES model was used for turbulence where the filtering threshold $\Delta = \sqrt[3]{V}$, where V is the grid cell volume. LES was used instead of k - ω shear stress transport (SST) because the SST model greatly overestimates turbulent viscosity and

yields results that do not agree with experiments [152, 153]. Instead of the previously used MULES interface compression method, the *isoAdvector* method was used here to model the volume fraction field [215, 216]. To improve performance and minimize artifacts in the volume fraction field, *isoAdvector* parameters were optimized as stated in [214]. The advantage of *isoAdvector* is that it enables using a rather coarse homogeneous 1 mm hexahedral mesh for the proof-of-concept problems considered here, while avoiding artifacts in bubble shape dynamics. The electromagnetic field was simulated as in [152, 153], with minor modifications explained in [214]. MHD flow was simulated by coupling OpenFOAM and Elmer through *EOF-Library message passing interface* (MPI) [154]. Hydrodynamic and electromagnetic fields were solved for in a segregated fashion by solving one after the other. Here $Rm \ll 1$ is assumed, therefore the MF was only computed once. More details regarding the field coupling are provided in [214].

2. The new DMD algorithm

Consider a dynamic system evolving in time on a manifold M such that, for $x_k \in M$, one has $x_{k+1} = f(x_k)$. Here f is a map from M to itself that governs the time evolution of the system $k \in \mathbb{Z}$. The *Koopman operator* is a linear infinite-dimensional operator K that acts on scalar valued functions on M such that for any scalar-valued function $g : M \rightarrow \mathbb{R}$, K maps \vec{g} to a new function Kg : $Kg(x) = g(f(x))$. Let $\varphi_i : M \rightarrow \mathbb{R}$ denote eigenfunctions and $\lambda_i \in \mathbb{C}$ denote eigenvalues of the Koopman operator: $K\varphi_i(x) = \lambda_i\varphi_i(x)$. For the MHD system described above, DMD treatment of vector fields (velocity, vorticity, etc.) is of interest. Consider a vector-valued observable $\vec{g} : M \rightarrow \mathbb{R}^p$. If each of the \vec{g} components lies within the span of the eigenfunctions φ_i , then one may expand \vec{g} in terms of φ_i as [217]

$$\vec{g}(\vec{x}) = \sum_{i=0}^{\infty} \varphi_i(\vec{x}) \vec{v}_i \quad (18)$$

If the components of \vec{g} *do not* lie within the span of φ_i , one may split K into regular and singular components, and project components of \vec{g} onto the span of the eigenfunctions [217]. The expression (18) is typically viewed as an expansion of $\vec{g}(\vec{x})$ as a linear combination of vectors \vec{v}_i , but it is equivalently an expansion of $\vec{g}(\vec{x})$ as a linear combination of φ_i , where now \vec{v}_i are the vector-valued coefficients of the expansion. Here the

eigenfunctions φ_i will be referred to as Koopman eigenfunctions, and the corresponding \vec{v}_i in (18) are the Koopman modes of the map f for the observable \vec{g} . From the above, one gets the iterates of \vec{x}_0 using

$$\vec{g}(\vec{x}_k) = \sum_{i=0}^{\infty} K^k \varphi_i(\vec{x}_0) \vec{v}_i = \sum_{i=0}^{\infty} \lambda_i^k \varphi_i(\vec{x}_0) \vec{v}_i \quad (19)$$

In its basic implementation, DMD computes an approximate linear operator A that represents the analyzed system (potentially nonlinear) such that A advances the system state one step forward in time. In this framework, the system states are represented by matrices

$$X_{m \times n} = \begin{bmatrix} | & | & & | \\ x_0 & x_1 & \cdots & x_{n-1} \\ | & | & & | \end{bmatrix}, \quad Y_{m \times n} = \begin{bmatrix} | & | & & | \\ x_1 & x_2 & \cdots & x_n \\ | & | & & | \end{bmatrix} \quad (20)$$

where x_i is a column vector that represents the system state for the i -th time step. Y columns are system states advanced by one time step from respective state vectors in X . This framework also assumes the states are sequential and equally spaced in time. With this, the propagation (evolution) equation is given by

$$AX = Y \quad (21)$$

To determine the Koopman mode frequencies, one must compute the complex eigenvalues λ_i and eigenvectors ω_i of A : $A\omega_i = \lambda_i\omega_i$, where ω_i form the span of A . With the Koopman mode decomposition and the system state correlation via A , namely $Ax_i = x_{i+1}$, one can express time stepping in terms of λ_i and ω_i :

$$A^n x_j = x_{j+n} = \sum_k C_k \lambda_k^n \varphi_k = \sum_k C_k e^{a_k t} e^{i\omega_k t} \varphi_k \quad (22)$$

If the system exhibits periodic and/or quasi-periodic processes, ω_i and λ_i represent characteristic time scales within the system. DMD separates dynamics by time scales, and the associated system structures (modes) can then be analyzed separately.

Typically, the matrices X and Y are very large, since they are determined by the (typically fine) resolution of 2D/3D images or measurement point sets from experiments, or simulation grids/meshes. Therefore, solving for A

exactly is not feasible. Usually, most of the system state information stored in X and Y can be represented by orders of magnitudes smaller amount of data, i.e. it is sufficient to encapsulate the main patterns within the system. This stems from the intuition that systems of interest usually have coherent structures that are much larger than a single point of measurement, and therefore much of the fine detail may be discarded without loss of meaningful features. To compute system eigenvalues, *singular value decomposition* (SVD) is used.

The default DMD algorithm is susceptible to noise in the data, in that the noise introduces errors in the generated eigenvalues. It was noticed that for the present system and several benchmarks (other systems), if the magnitudes of the expected eigenvalues are close to unity, i.e. they lie or almost lie on the unit circle in the complex plane, then for even relatively high SNR the noise can somewhat reduce the eigenvalue magnitudes and the associated decay rates become very high, quickly damping the modes to near-zero amplitudes.

The SVD algorithm processes all system state data simultaneously. If the input dataset is very large, then SVD will require large amounts of computer memory to be feasible. This can make it impossible to process datasets from high-resolution numerical simulations and experiments on computers without significant memory resources.

To mitigate this, it is proposed to combine the correlation of multiple adjacent (in time) system states with a streaming (online) SVD algorithm. The system state propagation can be generalized to include more than one preceding system state:

$$x_k = A_1 x_{k-1} + A_2 x_{k-2} + \dots + A_d x_{k-d} \quad (23)$$

This approach correlates successive system states with multiple preceding states, which can be expressed as $x'_{k+1} = K x'_k$, where

$$x'_k = \begin{pmatrix} x_k \\ x_{k-1} \\ \dots \\ x_{k-d+1} \end{pmatrix}, \quad K = \begin{pmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \mathbf{A}_3 & \dots & \mathbf{A}_{d-1} & \mathbf{A}_d \\ \mathbf{I} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \dots & \dots & \dots & \dots & \dots & \dots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{I} & \mathbf{0} \end{pmatrix} \quad (24)$$

and \mathbf{I} and $\mathbf{0}$ are $n \times n$ unit and zero matrices, where n is the dimension of x_i . Essentially, (23) is a linear combination of overlapping system state propagations as in (21). This approach was first proposed to examine systems with a low number of spatial dimensions, but is also applicable to systems with a higher amount of spatial dimensions [218]. The overlap

implies that x_k is given by a "moving average" of the preceding d states, which serves to filter the noise contained within system state snapshots [122, 123].

To introduce this effect, the system state data must be represented ("stacked") appropriately. Let X_i be the following

$$X_i = \begin{bmatrix} | & | & & | \\ x_i & x_{1+i} & \cdots & x_{n+i} \\ | & | & & | \end{bmatrix} \quad (25)$$

Similarly to (20) and (21), the system is expressed as $KM_0 = M_1$, where

$$M_0 = \begin{bmatrix} X_{d-1} \\ X_{d-2} \\ \vdots \\ X_0 \end{bmatrix}, \quad M_1 = \begin{bmatrix} X_d \\ X_{d-1} \\ \vdots \\ X_1 \end{bmatrix} \quad (26)$$

Here, matrices M_0 , M_1 are constructed by vertically stacking time-shifted sequences X_i of system states. This means that a single column of M contains information from d different columns of X . Thus, instead of correlating only neighboring snapshots as in (21), d system states are covered by a moving correlation window. This approach adds extra noise robustness to DMD because a larger "correlation radius" for data effectively averages out the noise, if any. It is observed that this method can significantly increase the accuracy of the DMD eigenvalues for data with low SNR. Equation (23) can be interpreted as a discretized linear differential equation of an order up to d .

Much like SVD was performed for X_0 in the default DMD algorithm, here SVD is performed for M_0 . Then both M_0 and M_1 are projected onto the subspace of U as $\tilde{M}_0 = U^*M_0$ and $\tilde{M}_1 = U^*M_1$. Next, the time-forward and time-backward components of \tilde{K} are computed: $\tilde{K}_+ = \tilde{M}_1\tilde{M}_0^+$ and $\tilde{K}_- = \tilde{M}_0\tilde{M}_1^+$, where \tilde{M}_0^+ and \tilde{M}_1^+ are the pseudo-inverses of \tilde{M}_0 and \tilde{M}_1 . This allows to compute \tilde{K} (similarity-transformed K):

$$\tilde{K} = \sqrt{\tilde{K}_+\tilde{K}_-^{-1}} \quad (27)$$

While one can also simply set $\tilde{K} = \tilde{K}_+$, (27) yields a minor increase in accuracy for next to no computational cost and is therefore worth implementing [219]. Finally, the eigenvalues of \tilde{K} are computed as $\tilde{K}\tilde{\varphi}_i = \lambda_i\tilde{\varphi}_i$, and the eigenvectors in the original basis are recovered via $\varphi_i = U\tilde{\varphi}_i$.

To obtain the modes of X one simply truncates the modes of M_0 to the first m elements (vertically), where m is the length of X columns. While one might argue that much of the system information is lost this way, since M_0, M_1 contain redundant information regarding X due to repeated time-shifted stacking construction, most of the total information can be inferred from the first non-repeating elements. Note also that the partial copies of X evolve in time identically, differing only in phase. To summarize, the proposed approach is as outlined in Algorithm 12.

It is important to note that (27) does not have a unique solution. To address this, it is recommended to choose a solution that is closest to \tilde{K}_+ and \tilde{K}_-^{-1} as the matrix \tilde{K} is expected to be close to these matrices. An alternative approach is to compute \tilde{K} as follows:

$$\tilde{K} = \frac{1}{2} \left(\tilde{K}_+ + \tilde{K}_-^{-1} \right) \quad (36)$$

It is important to note that there exists a closely related algorithm known as HAVOK [220] which uses the same representation as (26) and also uses SVD to extract structures out of the data. HAVOK is closely related to DMD and is capable of extracting the dynamics of highly nonlinear systems.

DMD decomposes the system into complex oscillating modes:

$$f'(t) = \sum_k C_k \varphi_k e^{iz_k t} = \sum_k C_k \varphi_k e^{a_k t} e^{ib_k t} \quad (37)$$

where $C_k \in \mathbb{C}$. In practice, most systems of interest are strictly real-valued:

$$f(t) = \sum_k C'_k \varphi_k e^{a_k t} \cos(b_k t + \phi_k) \quad (38)$$

where $C'_k \in \mathbb{R}$. This means that each real mode will be described with two complex modes with conjugate eigenvalues, which effectively makes half of the generated modes redundant. To address this, one can transform the real-valued input into a complex valued input by adding an imaginary $\pi/2$ phase-shifted version of the original input. This ensures that all the calculated modes are unique. Although this method introduces numerical artifacts at the edges (with respect to time) of system state stacks, it is compensated for by defining cutoff buffers for analyzed datasets. One must be careful, however, since in this case the imaginary mode components do not always "mirror" the real parts. The above is easily achieved using a procedure based on the Hilbert transform, as shown in [214]. With this, an imaginary $\pi/2$ phase-shifted copy is added to each real mode.

Algorithm 12: Custom DMD

- 1 Arrange the data $\begin{bmatrix} | & | & & | \\ x_0 & x_1 & \cdots & x_n \\ | & | & & | \end{bmatrix}$ into matrices

$$X_i = \begin{bmatrix} | & | & & | \\ x_i & x_{1+i} & \cdots & x_{n+i} \\ | & | & & | \end{bmatrix} \quad (28)$$

- 2 Construct matrices M_0, M_1

$$M_0 = \begin{bmatrix} X_{d-1} \\ X_{d-2} \\ \vdots \\ X_0 \end{bmatrix}, \quad M_1 = \begin{bmatrix} X_d \\ X_{d-1} \\ \vdots \\ X_1 \end{bmatrix} \quad (29)$$

- 3 Compute the truncated SVD of M_0

$$M_{m \times n} = U_{m \times r} S_{r \times r} V_{r \times n}^* \quad (30)$$

- 4 Compute matrices $\tilde{M}_0^+, \tilde{M}_1^+$

$$\tilde{M}_0 = U^* M_0, \quad \tilde{M}_1 = U^* M_1 \quad (31)$$

- 5 Compute matrices \tilde{K}_+, \tilde{K}_-

$$\begin{aligned} \tilde{K}_+ &= \tilde{M}_1 \tilde{M}_0^+ \\ \tilde{K}_- &= \tilde{M}_0 \tilde{M}_1^+ \end{aligned} \quad (32)$$

- 6 Compute \tilde{K}

$$\tilde{K} = \sqrt{\tilde{K}_+ \tilde{K}_-^{-1}} \quad (33)$$

- 7 Compute the eigenvalues and eigenvectors of \tilde{K}

$$\tilde{K} \tilde{\varphi}_i = \lambda_i \tilde{\varphi}_i \quad (34)$$

- 8 Compute the full DMD modes of the system

$$\varphi_i = U \tilde{\varphi}_i \quad (35)$$

- 9 Truncate the modes φ_i to the first m elements.

- 10 (Optional) Compute the initial amplitudes of the DMD modes [117]
-

3. Main results

Simulations were run for 20 s of flow time for all 4 cases, which is sufficient for the system to reach a quasi-stationary state starting from initially zero/stationary gas and fluid velocity fields, i.e. initial conditions are as in the previous papers [152, 153].

To provide context for the DMD analysis, one first observes the velocity field and vortex patterns. Characteristic snapshots of velocity fields at different time stamps for the 4 cases considered here are shown in Figures 68-71. Color maps of the midplane velocity magnitude are computed on the vessel walls and are then processed with the *ParaView* plugin *Surface LIC*, which uses the screen space surface line integral convolution (LIC) to transform the color maps into streamline maps with color-coded velocity magnitude [221]. Shaded white-gray overlays indicate bubble interfaces and the free metal surface at the top of the container. The rectangular tube at the bottom of the vessel is the argon inlet.

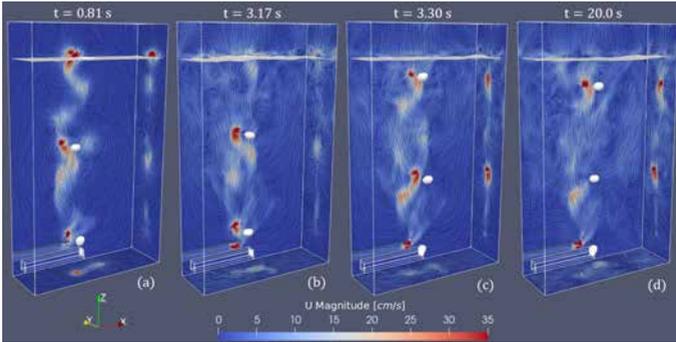


FIG. 68. Velocity LIC plots for bubble flow at 30 *sccm* flow rate without applied MF: (a) $t = 0.81$ s, (b) $t = 3.17$ s (c) $t = 3.30$ s, (d) $t = 20.0$ s.

If no MF is applied, bubble flow exhibits a classic vortex shedding pattern [11, 23, 26, 222] especially evident at earlier time stamps in Figures 68a-68c (*XZ* plane) which later becomes much more disordered, as seen in Figure 68d. Note that in the case with 100 *sccm* without applied MF, larger vortices with greater velocity are shed (Figures 69a-c) and, aside from the obvious global increase in velocity magnitude, one can see that the backflow from the free surface and walls is more pronounced near the bottom of the container, although it is difficult to tell from these plots alone how much further down this mixing layer extends at 30 versus 100 *sccm*. The objective of the DMD analysis will be to determine if there are any special

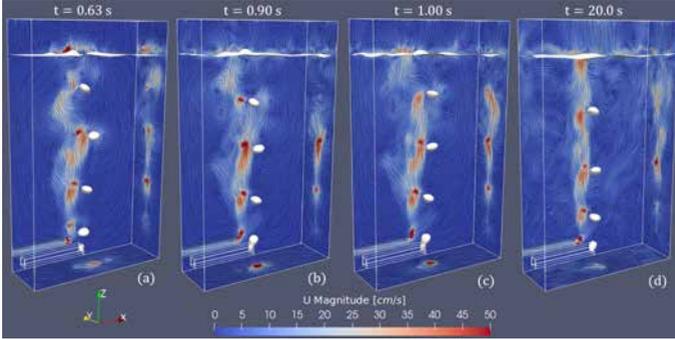


FIG. 69. Velocity LIC plots for bubble flow at 100 *sccm* flow rate without applied MF: (a) $t = 0.63$ s, (b) $t = 0.90$ s (c) $t = 1.00$ s, (d) $t = 20.0$ s.

patterns in both of these cases aside from the trivial mean upward flow within the bubble chain, and how these patterns differ for the two flow rates.

Applying static horizontal MF results in nearly complete flow laminarization for 30 *sccm* (Figure 70) and significant damping for 100 *sccm* (Figure 71), since vortex shedding is suppressed as expected [11, 13, 15, 222]. Note the *YZ* planes in Figure 70 where the LIC plot hints that bubbles ascend via rectilinear trajectories.

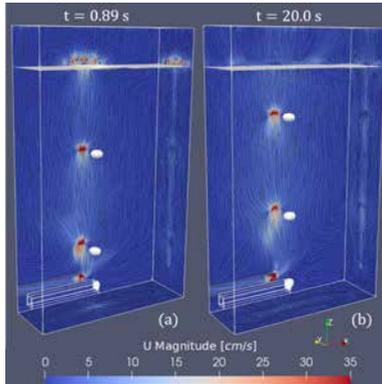


FIG. 70. Velocity LIC plots for bubble flow at 30 *sccm* flow rate with applied MF: (a) $t = 0.89$ s, (b) $t = 20.0$ s.

There is an important distinction in the case of 100 *sccm* – notice that

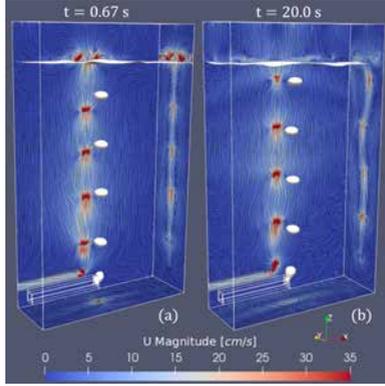


FIG. 71. Velocity LIC plots for bubble flow at 100 *sccm* flow rate with applied MF: (a) $t = 0.67$ s, (b) $t = 20.0$ s.

initially rectilinear ascension, as seen in Figure 71a, later transitions to a different pattern seen in Figure 71b. In the XZ plane, standing velocity magnitude waves form after the first ~ 10 s of flow time while in the XY plane one can see that the trajectory becomes slightly oscillatory in the upper half of the vessel, remaining roughly rectilinear in the lower half. DMD will be used to analyze these patterns in detail.

To assess the flow closer to the bubbles and the bubble chain, it is convenient to use the Q -criterion for vortex detection, since bubble wake flow is what determines the trajectories and collective dynamics. The Q -criterion (Q for brevity) represents both vortex cores ($Q > 0$) and saddle pattern flow zones ($Q < 0$) and is therefore well-suited for assessing the effects of varying the flow rate and MF magnitude [111]. The Q field definition is given in [214], following [111].

Figures 72-75 show the volume rendering of Q for the above treated cases. Q is computed in *ParaView* using the *VTK (Visualization ToolKit) gradient of unstructured data set* filter. For 30 *sccm*, Figure 72 indicates that bubble wakes exhibit a classic configuration with pairs of elongated "hairpin" vortices [11, 23, 26]. One can see that 30 *sccm* indeed corresponds to a quasi single-bubble regime, as trailing bubbles do not run into pronounced vortices behind leading bubbles owing to sufficient spacing that results in vortex decay/relaxation before their collisions with bubbles can take place. This is especially the case when flow stabilizes (Figures 72b-c – note the time stamps: flow stabilization occurs much faster than for 100 *sccm*) and is only violated a few times at the beginning of the simulation, a good example of which can be seen in Figure 72a. For 100 *sccm*, the situation

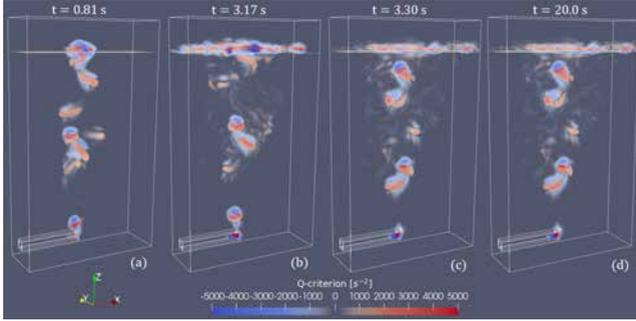


FIG. 72. Q plots for bubble flow at 30 *sccm* flow rate without applied MF: (a) $t = 0.81$ s, (b) $t = 3.17$ s (c) $t = 3.30$ s, (d) $t = 20.0$ s.

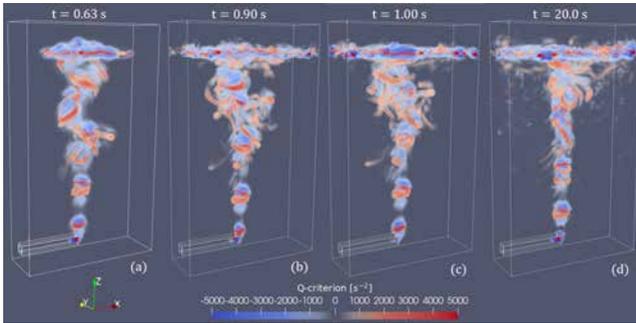


FIG. 73. Q plots for bubble flow at 100 *sccm* flow rate without applied MF: (a) $t = 0.63$ s, (b) $t = 0.90$ s (c) $t = 1.00$ s, (d) $t = 20.0$ s.

differs radically – note how vortex cores in Figures 73 intertwine and collide with trailing bubbles that enter wakes, resulting in a much more disordered flow.

When MF is applied, however, it virtually eliminates vortices outside the near-bubble zones, which is evident from Figures 74 and 75. Notice how, while short vortex core trails are visible in Figures 74 and 75 near the inlet, they completely disappear shortly after bubble detachment from the inlet. Note also that wake vortices are visibly fainter for 30 *sccm*.

To compute the DMD modes for the velocity field in the liquid metal vessel, the system was simulated for 20 *s* of flow time, and the last 600 frames (6 *s*) of the simulation output were used to represent the dynamics of a developed flow. The argon inlet and the free surface at the top is

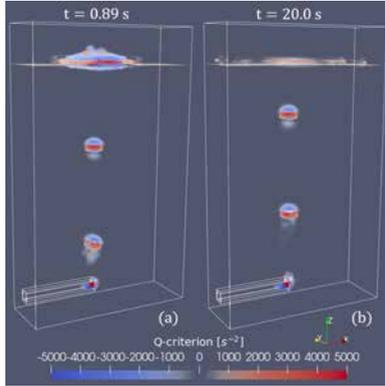


FIG. 74. Q plots for bubble flow at 30 *sccm* flow rate with applied MF: (a) $t = 0.89$ s, (b) $t = 20.0$ s.

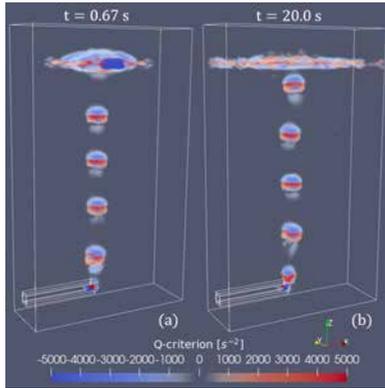


FIG. 75. Q plots for bubble flow at 100 *sccm* flow rate with applied MF: (a) $t = 0.67$ s, (b) $t = 20.0$ s.

cropped out to only observe the liquid metal volume. The cropped 600 frames are then fed to the real-to-complex mapping algorithm [214], after which 50 frames are removed from each end of the 600-frame sampling interval to avoid temporal artifacts caused by the real-to-complex mapping process. DMD is therefore performed for 500 system states. The procedure is identical for all of the above cases. Given the flow patterns seen in Figures 68-75 it is expected that DMD modes for cases with applied MF are going to be simpler/more ordered. Therefore, it makes sense to start

with these as they are easier to interpret and then compare to the cases without applied MF. The dominant modes for 30 and 100 *sccm* when MF is applied are shown in Figures 76-79.

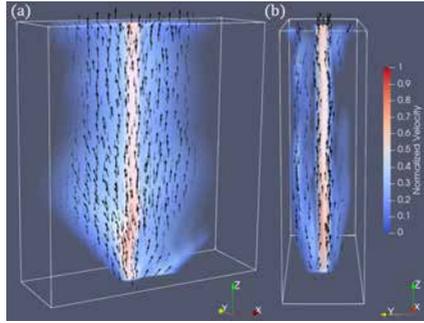


FIG. 76. The zeroth DMD velocity field mode for 30 *sccm* with applied MF: (a) front view, (b) side view.

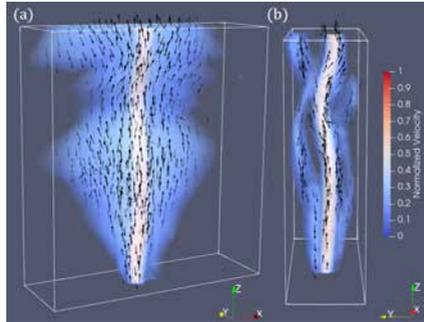


FIG. 77. The zeroth DMD velocity field mode for 100 *sccm* with applied MF: (a) front view, (b) side view.

Starting with 30 *sccm* and ranking the modes in terms of relative amplitudes, the zeroth (strongest) velocity field mode is shown in Figure 76. This mode has a very low frequency ($f_0 \sim 0.15 \text{ mHz}$, $f = \omega/2\pi$) and a negligible growth rate ($a_0 \sim -1.9 \cdot 10^{-3} \text{ s}^{-1}$), and can be considered stationary. Note the pronounced bubble chain flow region in Figure 76a where the maximum normalized velocity is concentrated. Interestingly, as seen in Figure 76b, metal flow in the Y direction is organized in three sheets – one with upwards metal flow about the bubble chain region, extending over the XZ plane (Figure 76a), and two counter-flow sheets to its left and

right in the YZ plane. This mode exhibits a great deal of symmetry with respect to the XZ mid-plane in addition to the YZ plane.

For 100 *sccm*, on the other hand, while a similar pattern appears for the zeroth mode ($f_0 \sim 0.3$ *mHz*, $a_0 \sim 1.7 \cdot 10^{-4}$ s^{-1}), its symmetry in the YZ plane is broken in the upper part of the liquid metal vessel as seen in Figure 77b, and the counter-flow sheets are disordered. Standing waves in the Y component of the mode form in the upper part of the vessel, which is illustrated in Figure 78.

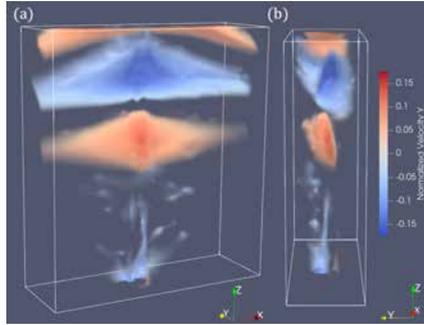


FIG. 78. The Y component of the zeroth DMD velocity field mode for 100 *sccm* with applied MF, normalized with respect to the mode magnitude: (a) front view, (b) side view.

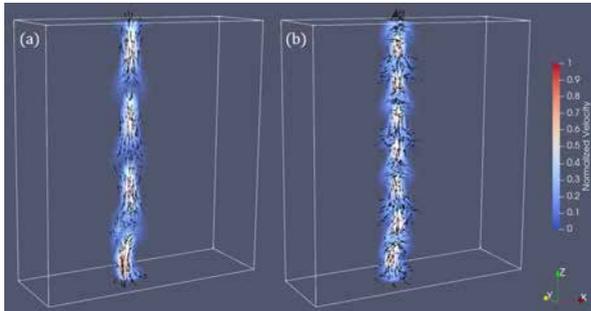


FIG. 79. (a) The 1-st and (b) 2-nd DMD velocity field modes for 30 *sccm* with applied MF.

The 1-st and 2-nd modes for the velocity field for 30 *sccm* with applied MF are shown in Figures 79a and 79b. These modes are, unlike the zeroth modes for 30 and 100 *sccm*, non-stationary and their flow patterns

oscillate at their respective eigenfrequencies ($f_1 \sim 4.2 \text{ Hz}$, $f_2 \sim 8.5 \text{ Hz}$ with $a_1 \sim -6.7 \cdot 10^{-3} \text{ s}^{-1}$ and $a_2 \sim -6.4 \cdot 10^{-3} \text{ s}^{-1}$). These modes, as well as their weaker higher-order spatial harmonics (not shown here), can be interpreted as a measure of coherence of motion within the bubble chain. In this case, strictly periodic vertical patterns can be seen, indicating that applied MF enforces highly ordered bubble chains with stable rectilinear trajectories. For 100 *sccm* these harmonics associated with periodic bubble motion within the chain are very similar but exhibit greater *XY* components. It is then interesting to compare the zeroth, stationary modes for 30 and 100 *sccm* without applied MF against the above cases – this is shown in Figures 80 and 81.

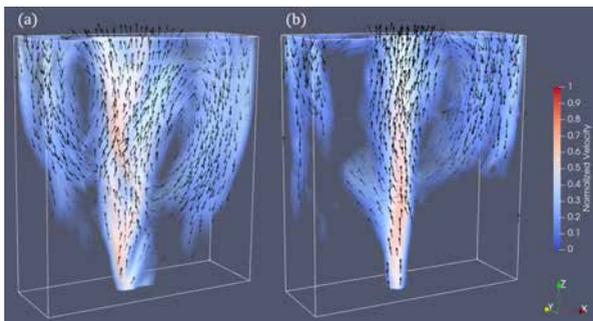


FIG. 80. The zeroth (stationary) velocity field modes for (a) 30 and (b) 100 *sccm* without applied MF. Normalization is separate for each case.

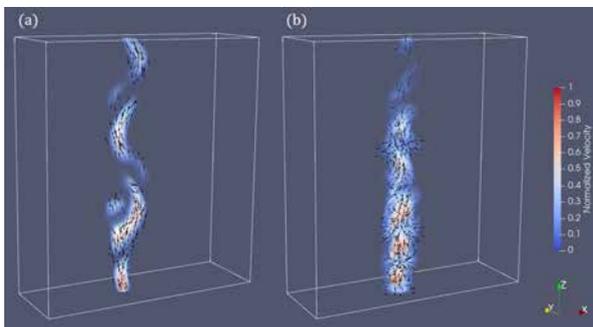


FIG. 81. The (a) 2-nd velocity field mode for 30 *sccm* and (b) the 6-th mode for 100 *sccm* without applied MF. Normalization is separate for each case.

Figures 80a and 80b show the zeroth modes for 30 ($f_0 \sim 3.1 \text{ mHz}$ with

$a_0 \sim -2 \cdot 10^{-3} \text{ s}^{-1}$) and 100 *sccm* ($f_0 \sim 2.1 \text{ mHz}$ with $a_0 \sim 8.3 \cdot 10^{-3} \text{ s}^{-1}$), respectively, without applied MF. The zeroth mode in Figure 80a covers much more space about the bubble chain *core* where its magnitude is relatively very high, whereas in 80b the core zone is much thinner, indicating that some other modes are dominant further away; in other words, metal flow about the ascending bubble chain is much more coherent for 30 *sccm*. The other thing to note is that the 30 *sccm* mode is much more symmetric about the *YZ* mid-plane (though much less than with applied MF). This makes sense given the lower flow rate corresponding to Figure 80a, but what is significant here is that, comparing especially Figures 76 and 80a, there are no longer two symmetry planes. It should therefore be of interest to later study the transition from 1 to 2 symmetry planes with ordered flow sheets for a fixed flow rate as the MF magnitude is swept from zero upwards, as well as how symmetry is disrupted as flow rate is increased (e.g. at a fixed MF magnitude).

The modes presented in Figure 81 for the cases without MF, meanwhile, are noteworthy for several reasons. First, note that they no longer come in first in terms of amplitude, but rather second for 30 (Figure 81a, $f_2 \sim 4.3 \text{ Hz}$, $a_2 \sim 1.6 \cdot 10^{-2} \text{ s}^{-1}$) and sixth for 100 *sccm* (Figure 81b, $f_6 \sim 10.3 \text{ Hz}$, $a_6 \sim 6.0 \cdot 10^{-2} \text{ s}^{-1}$), indicating that, as expected, the overall coherence of bubble motion is much lower than in the cases with applied MF. Moreover, comparing Figure 81a against Figure 81b, one can conclude that, while overall weaker compared to the other flow patterns, Figure 81a shows coherent zigzag motion extending throughout the gallium vessel, whereas in Figure 81b one can see that the mode magnitude falls off dramatically above a certain elevation threshold, indicating a *coherence length*. A criterion could be defined to measure it, which would enable to quantify the effects of varying the flow rate and MF magnitude on flow stability.

One simple way to introduce the coherence length is to fit (with a velocity magnitude threshold) volumes defined by an elliptic cross-section extruded over the *Z* dimension of the vessel to velocity fields of bubble chain modes, such as in Figures 79 and 81, and compute velocity magnitude integrals over elliptic cross-sections along *Z*. As an example, consider the strongest modes associated with bubble motion in the chain for each case – the output of the above procedure is shown in Figure 82. A cutoff threshold of 50% of the velocity value nearest to the inlet is used to define the coherence length – one can see that the coherence lengths of bubble chains for 30 and 100 *sccm* with applied MF and 30 *sccm* without MF extend all the way to the free surface (potentially beyond), while the coherence length for 100 *sccm* without MF is $\sim 7.35 \text{ cm}$ (the inlet is below the 2 *cm* mark). Indeed, Figure 81b suggests that bubble motion within the chain becomes incoherent above the 3/4 of the inlet to free surface distance. This is also

consistent with what is seen in Figure 73, especially near the end of the simulation time interval. Note also that Figure 82 is representative of bubble spatial frequency within the chains for cases with applied MF.

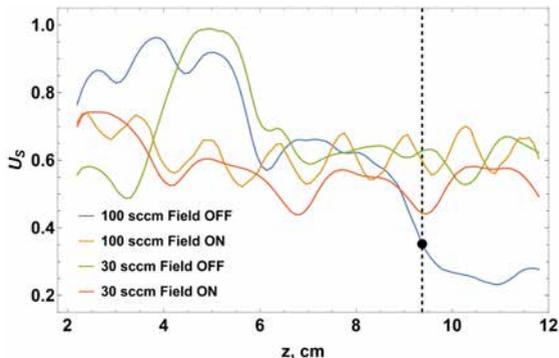


FIG. 82. Normalized velocity magnitude integrals U_S for the strongest velocity field modes associated with bubble motion patterns within the bubble chain. Integrals are computed over elliptic cross-sections at different height (z) containing velocity magnitude values above a threshold that is identical in all cases.

It is also interesting to examine modes for 30 and 100 *sccm* without applied MF that capture flow turbulence, such as the modes shown in Figure 83-85. Consider the 1-st velocity field mode for 30 *sccm* without applied MF ($f_1 \sim 0.18$ Hz, $a_1 \sim 4.6 \cdot 10^{-2}$ s $^{-1}$) seen in Figure 83: Figure 83a indicates, when viewed alongside the zeroth mode in Figure 80a, that mode 1 occupies the space about the central core (with respect to the vessel and the bubble chain) of the zeroth mode. It is also comprised of counter-flowing vertical jets, better seen in Figures 83c and 83d, the latter showing that the counter-flowing regions seem to be delimited (roughly) by the XZ mid-plane, although there is no discernible symmetry like in the zeroth modes with applied MF (Figures 76b and 77b). In Figure 83b one can also see what looks like swirl flow in the upper region of the mode.

In comparison, Figure 84 shows the 1-st velocity field mode for 100 *sccm* without applied MF ($f_1 \sim 0.36$ Hz, $a_1 \sim -1.2 \cdot 10^{-2}$ s $^{-1}$). Here the noteworthy feature is that, unlike its 30 *sccm* counterpart, this mode exhibits two clearly separated regions with counter-flow, evident from Figure 84b. Again, the mode, like in the 30 *sccm* case, occupies the space about the core of the zeroth mode (Figure 80)b, this time with a much clearer symmetry about the XZ mid-plane. Note also that Figure 84c indicates, similarly to Figure 83b, swirl flow in the upper region of the mode, but here it is more ordered. Also, unlike in Figure 83b, one can see

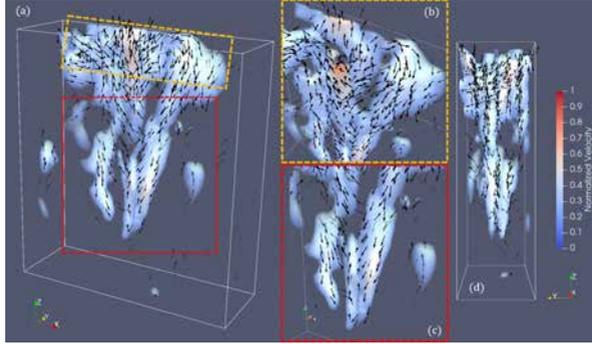


FIG. 83. The first velocity field mode for 30 *sccm* without applied MF: (a) full view, (b) close-up top view, (c) close-up bottom view, (d) side view.

in Figure 84b that the swirl-like flow seems to extend further below the free surface than in the 30 *sccm* case.

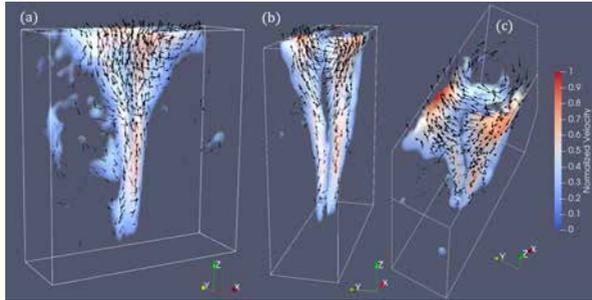


FIG. 84. The first velocity field mode for 100 *sccm* without applied MF: (a) front view, (b) side view, (c) top view.

However, Figures 68 and 69, as well as Figures 72 and 73 suggest that the velocity field is more disordered in the 100 *sccm* case. The reason is that there exists a pronounced 2-nd mode ($f_2 \sim 0.19 \text{ Hz}$, $a_2 \sim -4.9 \cdot 10^{-2} \text{ s}^{-1}$), shown in Figure 85, that spans most of the vessel volume. This mode does not seem to exhibit any discernible symmetry and likely determines the finer disordered velocity field structure seen in Figures 69 and 73. A similar mode for 30 *sccm* has a much smaller amplitude, lower than that of the 2-nd mode in Figure 81, which should explain the observations from the velocity and Q plots.

There are also higher-order modes for the cases without applied MF,

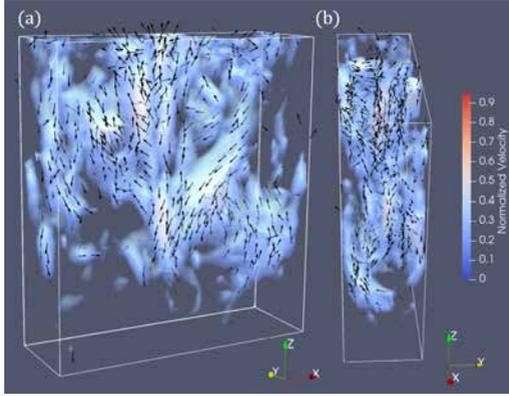


FIG. 85. The 2-nd velocity field mode for 100 *sccm* without applied MF: (a) front view, (b) side view.

but these have lower amplitudes and exhibit patterns similar to the one in Figure 85 with the characteristic scales of their spatial structures decreasing with mode order. They should be taken with a grain of salt for two reasons: first, the 1 *mm* cube mesh should not accurately capture the finer flow structures in the higher-order modes; second, especially for the 100 *sccm* case, longer sampling time is most likely required to capture representative flow field configurations. As such, further modes for the gallium vessel are not shown here.

While many modes may have initially large or conversely very small amplitudes, it is important to consider their growth rates over the DMD sampling time interval (5 seconds). It is also of interest to see what frequencies appear, and what the respective amplitudes and growth rates are in each of the cases considered herein – this is shown in Figures 86-89.

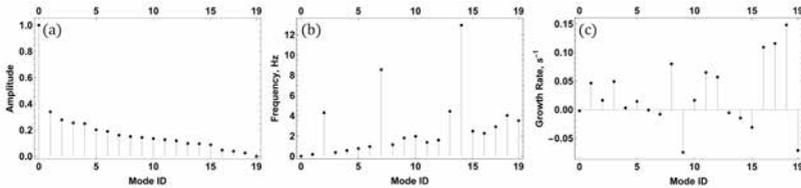


FIG. 86. (a) Initial normalized amplitudes, (b) frequencies and (c) growth rates for modes at 30 *sccm* without applied MF.

In all cases, the zeroth stationary modes have more than double initial

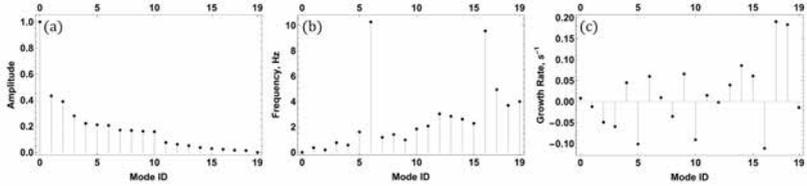


FIG. 87. (a) Initial normalized amplitudes, (b) frequencies and (c) growth rates for modes at 100 *sccm* without applied MF.

amplitudes compared with all other modes. One can note several patterns for the cases without applied MF (Figures 86 and 87). First, modes with higher temporal frequencies generally have lower amplitudes, aside from a few outlier modes that are different for 30 and 100 *sccm* – note that the lower frequency peak seen in Figure 86b is not present in Figure 87b. Another characteristic feature is that the normalized amplitudes decrease with mode order more gradually for 30 than for 100 *sccm* (Figures 86a and 87a) where in the latter case the amplitude slope for modes 1-4 is steeper and there is a slight step-down from mode 10 to 11. It is also of interest that, while the amplitude does not vary too much for higher-order modes despite higher growth coefficients due to their lower initial amplitudes, modes 1-3 and 5 for 100 *sccm* have rather significant negative growth rates (Figure 87c), meaning that their dominance over higher frequency modes will diminish over time. Note also that mode 6 has a notable positive growth rate and the highest frequency (Figure 87b). Meanwhile, for 30 *sccm*, one can see in Figure 86c that the lower frequency modes have either significant positive, or near-zero growth rates and the two modes with the highest frequencies have negative growth rates. This would suggest that the higher flow rate promotes finer flow structures (characteristic for high order modes when no MF is applied) with smaller time scales, and the DMD provides the means of quantifying this. In principle, for simulations with a higher grid resolution, it should be feasible to measure characteristic length scales for DMD mode flow structures via image processing and/or spectral methods for a more in-depth analysis.

When MF is applied (Figures 88 and 89), several key differences appear. First, there is now no clear pattern for frequency versus mode order, as seen in Figures 88b and 89b. Second, Figures 88a and 89a indicate that, unlike the cases without MF, there is a very sharp mode amplitude cutoff after mode 5 for 30 *sccm* and mode 7 for 100 *sccm* beyond which the amplitudes are so insignificant that even the high positive growth rates in the 30 *sccm* case (Figure 88) make no difference to the overall dynamics.

Also, this positive growth for more than half of the modes beyond mode 5 in the 30 *sccm* is in stark contrast with what is seen in Figure 89c for 100 *sccm* where one can see that most of the modes have significant negative growth rate. Figures 88b and 89b also indicate an overall increase in mode temporal frequency – this can be attributed to spatial harmonics associated with the bubble chain like the ones in Figure 79. All of these are quantitative indicators – including the smaller number of modes required to encapsulate the system with applied MF – that can be used so systematically study how MF stabilizes and regularizes the flow field and how this is disrupted at higher flow rates. Note that the mode growth rates in Figure 89c would also explain the gradual transition from rectilinear to *YZ* plane zigzag trajectories for 100 *sccm* with applied MF (Figure 77) – the zeroth mode becomes more pronounced over time versus the other decaying modes.

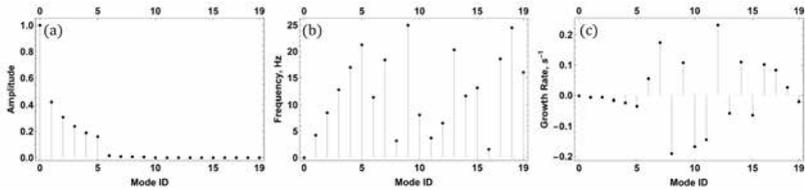


FIG. 88. (a) Initial normalized amplitudes, (b) frequencies and (c) growth rates for modes at 30 *sccm* with applied MF.

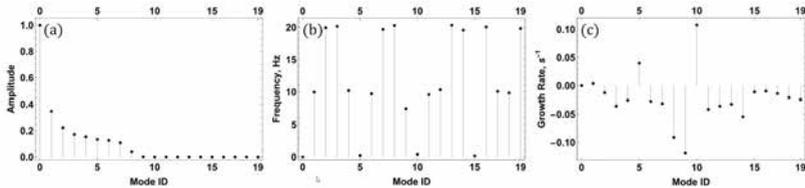


FIG. 89. (a) Initial normalized amplitudes, (b) frequencies and (c) growth rates for modes at 100 *sccm* with applied MF.

Finally, it is important to analyze the degree of spatial correlation between the DMD modes – mode correlation matrices are shown in [214] where one can see that the modes are rather weakly correlated except for several modes in the 100 *sccm* case with applied MF. The reason why some of the mode pairs are less orthogonal than others is that the oscillations of the argon velocity field within bubbles are captured by the DMD. At the current resolution these oscillations are under-resolved and are essentially correlated

noise, hence whatever modes contain these noise patterns are correlated to some degree. Generally, higher-order modes are more susceptible to this effect because argon velocity field noise exhibits fine length scales.

Meanwhile, for the cases without MF, for 30 *sccm* there are a few more pronounced correlations, but overall the values are lower than in the 100 *sccm* where no significant peaks are present, i.e. the off-diagonal values are more diluted. The latter can be explained by the fact that most of the modes in the 100 *sccm* are much more disordered and exhibit finer spatial structures than in the 30 *sccm* case with spatial/temporal timescales often slightly overlapping with the correlated noise, hence the relative homogeneity of the off-diagonal matrix elements. In the former cases where MF is applied flow is largely laminarized and virtually the only fine temporal/spatial structures on scales similar to the noise are found in some of the spatial harmonics associated with the bubble chain, i.e. higher-order modes similar to the ones in Figure 79 – the overlap seems to be stronger for 100 *sccm*.

However, the fact that overall the modes, especially the strongest of the lower-order, are very much orthogonal means that it is quite safe to physically attribute the significant modes to various momentum transfer mechanisms (and spatial harmonics thereof) occurring within the vessel.

To analyze bubble wake flow in greater detail one must first transition to the bubble reference frame by tracking individual bubbles, reconstructing their centroid trajectories and performing velocimetry, wherefrom the relative velocity field can be computed. This is done as outlined in [214]. The representative computed trajectories which can be used to derive the relative velocity fields for bubbles for the 4 cases of interest are shown in [214], as are the respective local Q fields (invariant to the Galilean transformation) about the bubbles. Their interpretation is also given in [214] – in this summary, the focus is the DMD modes stemming from the relative velocity field near the bubbles.

To compute the DMD modes for the wake velocity field in the bubble reference frame, the bubbles with trajectories within the time window of the vessel DMD analysis are used and trajectories shorter than 10 frames are removed. The remaining trajectories are trimmed to include bubbles in $z \in [40 \text{ mm}, 110 \text{ mm}]$. This is done to avoid sharp field transients near the inlet and near the top. The transition to the bubble reference frame is performed as explained in [214], and the velocity field is resampled to a rectangular box that follows bubble centroids. The sampling box is extended in the z^- direction to capture more of the bubble wake velocity field. The box dimensions are: $x^- = x^+ = 15 \text{ mm}$; $y^- = y^+ = 7.5 \text{ mm}$; $z^+ = 5 \text{ mm}$; $z^- = -15 \text{ mm}$ [214]. A spherical mask with a 4.5 mm radius centered at the bubble centroid is applied to the sampled velocity field,

which is done to mask the argon velocity field variations within bubbles. Real-to-complex domain mapping is applied to each field grid point over the time range of a trajectory. The field snapshot stacking strategy is explained in [214], and it is such that DMD yields modes for the entire ensemble of trajectories. Initial amplitudes are calculated for each trajectory separately [117]. As was the case for the vessel, the first 20 modes were computed for the bubble wake DMD.

Starting with 30 *sccm* with applied MF, one can observe that in the case of the modes for the bubble reference frame velocity field in the sampling box – referred to as simply "bubble modes" for brevity – the modes are rather strongly correlated [214]. This means that, unlike the vessel velocity field modes, these must be interpreted jointly and a more in-depth analysis is required. To this end, in addition to mean initial amplitudes (as before, normalized to the zeroth DMD mode) with standard deviations, frequencies and growth rates (Figure 90), one must also carefully examine the dynamics of normalized mode amplitudes (Figure 91), as well as normalized root-mean-square (RMS) mode amplitudes for every analyzed trajectory (Figure 92). Note that in the latter case normalization is performed for modes 1 to 19, for visual purposes – this is because the zeroth mode’s amplitude is roughly an order of magnitude greater than that of the second-highest, in this case mode 12. This pattern for these three types of figures is used for the other three flow cases as well.

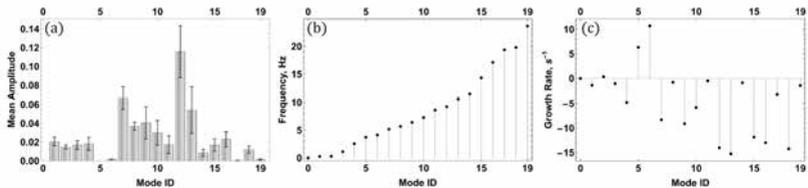


FIG. 90. 30 *sccm* with applied MF: bubble velocity field mode (a) normalized initial amplitudes, (b) frequencies and (c) growth rates. Normalization with respect to the zeroth mode.

The strongest bubble velocity field mode is the zeroth mode, which is essentially a mean flow mode with $f_0 \sim 2.6 \text{ mHz}$ and $a_0 \sim -9.4 \cdot 10^{-3} \text{ s}^{-1}$, shown in Figure 93. Its relative magnitude is consistently about unity for all trajectories, with a $\sim 0.26\%$ standard deviation. This mode aside, as seen in Figures 90 and 91a, the initially dominant modes are 12, 7 and 13 with 9, 8, 10 and 11 having smaller, but still significant amplitudes; other modes of note are 1 to 4, the lower frequency modes, their amplitude dynamics shown in Figure 91b. Note that modes 12, 7, 13, 9 and 10 decline

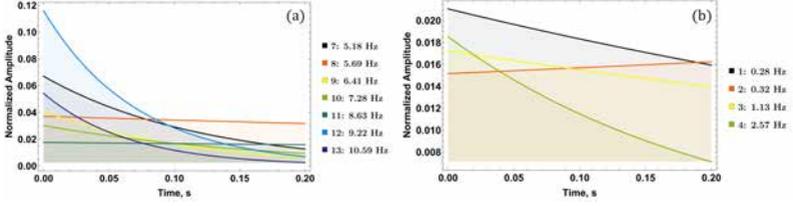


FIG. 91. 30 *sccm* with applied MF: dynamics of normalized amplitudes over the mean trajectory time for significant (a) dominant modes and (b) secondary modes. Normalization with respect to the zeroth mode. Legend: mode IDs and frequencies.

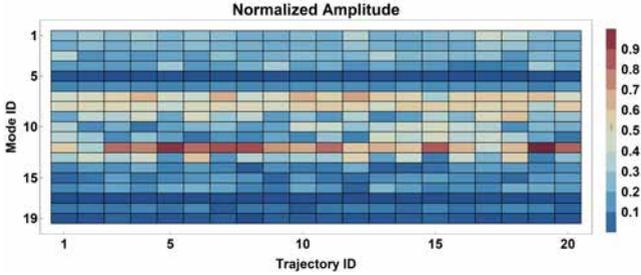


FIG. 92. 30 *sccm* with applied MF: normalized root-mean-square (RMS) amplitudes for modes over all processed trajectories. Normalization performed with modes 1 to 19.

very quickly over trajectory time, while the amplitudes of modes 8 and 11 are almost unchanged. To put this in perspective, consider that initially the sum of amplitudes of the modes shown in Figure 91a amounts to ~ 0.38 of the zeroth mode amplitude, while by the end of the trajectory time interval their contribution is reduced to ~ 0.1 .

To understand the effects that this has on the overall velocity field, one must examine the higher order mode flow patterns – these are shown in Figures 94-98, in order of descending mean initial amplitude. This order of presentation is kept throughout this section. Modes 1 to 4 are not shown here, as modes 1 and 2 are slightly distorted versions of the zeroth mode (and are strongly correlated to it [214]), while modes 3 and 4 are much weaker versions of 12 and 7, respectively. Combined, modes 1 to 4 initially constitute at most $< 8\%$ of the zeroth mode amplitude and considerably decline over trajectory time (Figure 91).

Mode 12 ($f_{12} \sim 9.22 \text{ Hz}$, $a_{12} \sim -14 \text{ s}^{-1}$), as seen in Figure 94, mostly

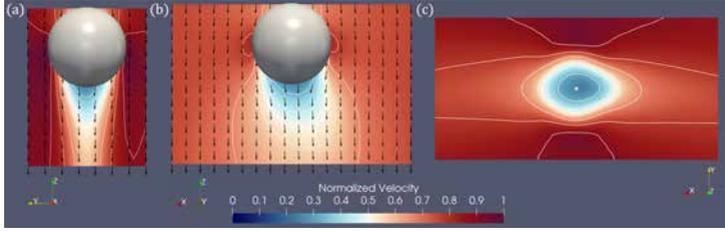


FIG. 93. The zeroth bubble velocity field mode for 30 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane. The gray specular sphere in (a) and (b) is the mask used within the DMD sampling box. Note: the gray dot in (c) in the center of the plane is where the mask intersects the *XY* plane, i.e. the plane is right below the bubble. The *YZ* and *XZ* planes contain the center of the spherical mask. The in-plane velocity field vector lengths in (a-c) are scaled by projecting a grid of equally spaced vectors of equal lengths onto the planes.

has a very pronounced *X* component with a lesser *Y* direction contribution, meaning that this mode represents velocity field oscillations about the bubble mostly in the *XZ* plane with less pronounced *YZ* oscillations. This could be interpreted as the flow pattern responsible for bubble trajectory oscillations in the *XZ* plane in the initial stages of trajectories, as shown in [214] – note the green lines representing trajectory projections onto the *XZ* plane. The exponential decay of mode 12 which decreases in amplitude \sim threefold over the first third of the trajectory time interval and its frequency are consistent with the rapidly damped oscillations observed for the *XZ* projections. Figure 92 suggests this is the case for most of the trajectories.

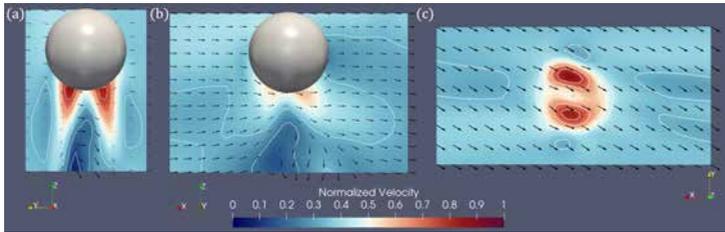


FIG. 94. The 12-th bubble velocity field mode for 30 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

Modes 7 ($f_7 \sim 5.2 \text{ Hz}$, $a_7 \sim -8.3 \text{ s}^{-1}$) and 13 ($f_{13} \sim 11 \text{ Hz}$, $a_{13} \sim -15 \text{ s}^{-1}$), on the other hand, contain velocity field oscillations in the *YZ* plane, *Y* direction with minor out-of-plane perturbations (Figures 95 and

96). While these modes decay rather quickly, modes 8 ($f_8 \sim 5.7 \text{ Hz}$, $a_8 \sim -0.8 \text{ s}^{-1}$) and 11 ($f_{11} \sim 8.6 \text{ Hz}$, $a_{11} \sim -0.5 \text{ s}^{-1}$) which are similar have very small growth rates and thus persist over the mean trajectory time. Mode 9 is not shown here because it is much weaker than and remarkably similar to mode 12.

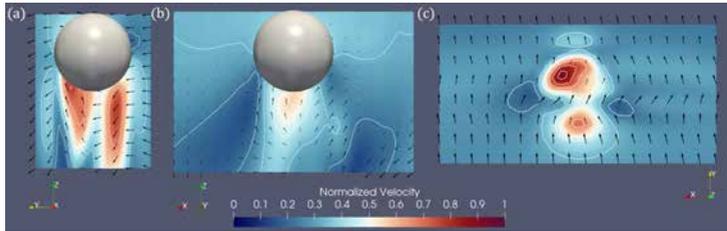


FIG. 95. The 7-th bubble velocity field mode for 30 *sccm* with applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

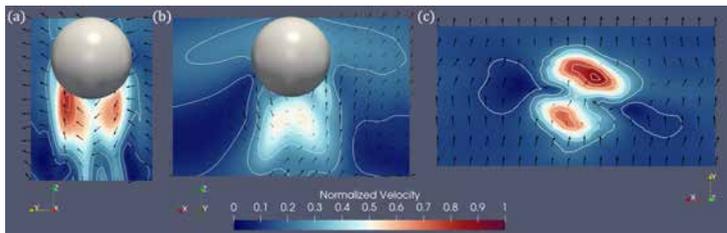


FIG. 96. The 13-th bubble velocity field mode for 30 *sccm* with applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

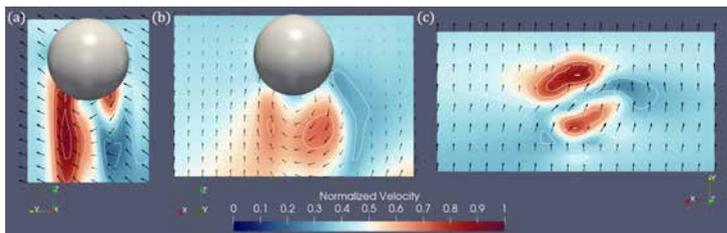


FIG. 97. The 8-th bubble velocity field mode for 30 *sccm* with applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

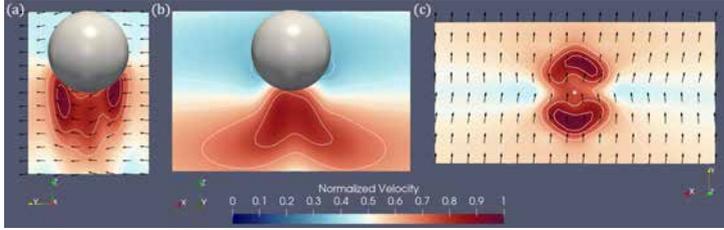


FIG. 98. The 11-th bubble velocity field mode for 30 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

It could be speculated that these modes determine the trajectory oscillations in the *YZ* plane, especially considering that these oscillations have increased amplitudes for later trajectories, which is consistent with somewhat greater RMS amplitudes observed for modes 7, 8 and 11 for later trajectories (Figure 92). Importantly, the mode periods are on the order of or less than the mean trajectory time. Note that this behavior is not necessarily evident from inspections of vessel modes for 30 *sccm* with applied MF (e.g. Figure 76), indicating the potential usefulness of applying DMD to bubble wakes separately. However, longer DMD sampling times are required to be certain.

It should also be noted that considerable correlation is seen between the above modes: $7 \leftrightarrow (8, 12)$, $8 \leftrightarrow 12$, $11 \leftrightarrow 13$ [214]. Note also that the showcased modes exhibit a great deal of symmetry and none of them contain small scale spatial structures, indicating wake flow laminarization by applied MF, which is, again, consistent with observed bubble wakes [214].

Mode analysis results for 100 *sccm* with applied MF are shown in Figures 99-101. As in the 30 *sccm* case, many of the DMD modes are significantly correlated [214]. Figure 99a indicates that there are only two modes with relatively high initial amplitudes but, unlike the 30 *sccm* case, here one can see in Figure 99c that there are quite a few modes with positive growth rates, even though the rate magnitudes are relatively small except mode 16 – this mode, however, has a minimal initial amplitude and despite the highest growth rate has an insignificant RMS amplitude, as seen in Figure 101. It is also interesting to note that the mode growth damping from 30 *sccm* to 100 *sccm* with applied MF observed in Figures 90c and 99c is also observed for vessel DMD modes (Figures 88c and 89c) with the distinction that in the vessel mode case one has mostly negative growth rates for 100 *sccm*, not 30 *sccm* as it is with the bubble modes. Figures 100 and 101 indicate that it makes sense to take a closer look at modes

4, 6, 8-10 and 12 as these are either initially dominant by a considerable margin or persist at or grow to a significant amplitude.

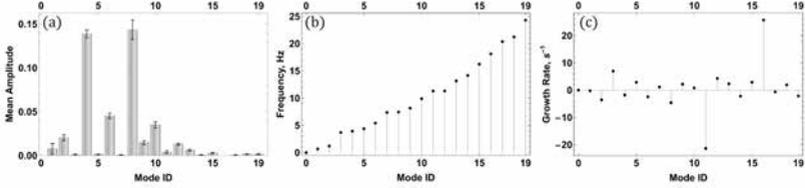


FIG. 99. 100 *sccm* with applied MF: bubble velocity field mode (a) normalized initial amplitudes, (b) frequencies and (c) growth rates.

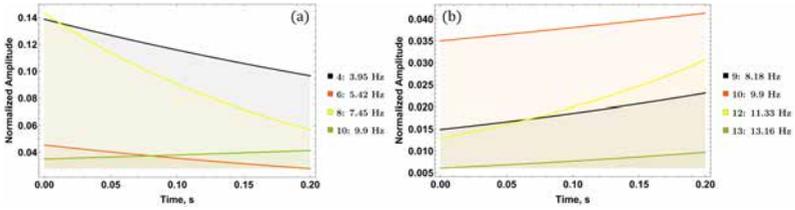


FIG. 100. 100 *sccm* with applied MF: amplitude dynamics over the mean trajectory time for significant modes: (a) dominant modes, (b) secondary modes.

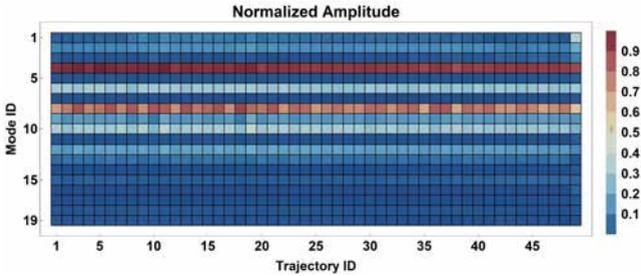


FIG. 101. 100 *sccm* with applied MF: normalized RMS amplitudes for modes over all processed trajectories.

Significant bubble modes for 100 *sccm* with applied MF are presented in Figures 102-108. It is immediately evident that the zeroth mode ($f_0 \sim 2.3$ *mHz*, $a_0 \sim -4.6 \cdot 10^{-3}$ *s*⁻¹, amplitude ~ 1 for all trajectories, with a 0.66% deviation) in Figure 102 is very similar to its 30 *sccm* counterpart

(Figure 93). The difference lies in the more pronounced, larger stagnation zone below the bubble. An important distinction between the modes in this case versus 30 *sccm* is that for 100 *sccm* there are high amplitude modes with much stronger Z components within the bubble wake, i.e. modes 8 ($f_8 \sim 7.5$ Hz, $a_8 \sim -4.6$ s^{-1} , Figure 104), 10 ($f_{10} \sim 9.9$ Hz, $a_{10} \sim 0.83$ s^{-1} , Figure 105), 9 ($f_9 \sim 8.2$ Hz, $a_9 \sim 2.2$ s^{-1} , Figure 107) and especially 12 ($f_{12} \sim 11$ Hz, $a_{12} \sim 4.3$ s^{-1} , Figure 108) versus what is seen in Figures 95a, 96a and 97a.

One difference between the trajectories is that in the 100 *sccm* case, there is no pronounced initial bubble displacement in the XZ plane, unlike for 30 *sccm* [214]. Looking at the modes with significant X components, note that for 30 *sccm* there is mode 12 (Figure 94) that is initially the one with the greatest magnitude after the zeroth mode; for 100 *sccm*, on the other hand, mode 6 ($f_6 \sim 5.4$ Hz, $a_6 \sim -2.4$ s^{-1} , Figure 106) has ~ 3 times lower initial amplitude compared to modes 4 ($f_4 \sim 3.9$ Hz, $a_4 \sim -1.8$ s^{-1} , Figure 103) and 8, and mode 9 has a less pronounced X component and consistently an even lower amplitude throughout trajectory time. The relatively small X component contribution from these modes likely explains the differences in the XZ projections of trajectories.

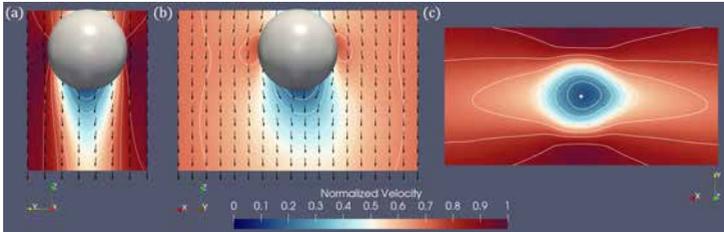


FIG. 102. The zeroth bubble velocity field mode for 100 *sccm* with applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

Consider also the YZ projections of trajectories for 30 *sccm* and 100 *sccm* shown in [214] – the deviations in the YZ plane are overall greater in the 100 *sccm* case, even for initial trajectories. Modes 4 ($f_4 \sim 3.9$ Hz, $a_4 \sim -1.8$ s^{-1} , Figure 103), 8, 10 and 9 all contribute to the bubble wake velocity field oscillations in the XY plane, Z direction, and remain dominant over trajectory time. Moreover, their combined relative contribution is overall significantly greater than that of similar modes in the 30 *sccm* case. While this might explain the trajectory oscillations in the YZ plane – mode frequencies suggest this is plausible – it is not evident from Figure 101 that the above-mentioned modes should cause, over the flow time, the transition from a still roughly rectilinear trajectory to trajectories with

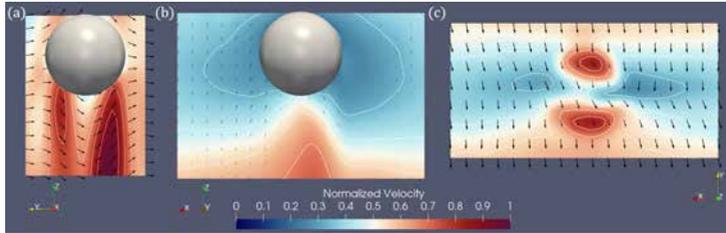


FIG. 103. The 4-th bubble velocity field mode for 100 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

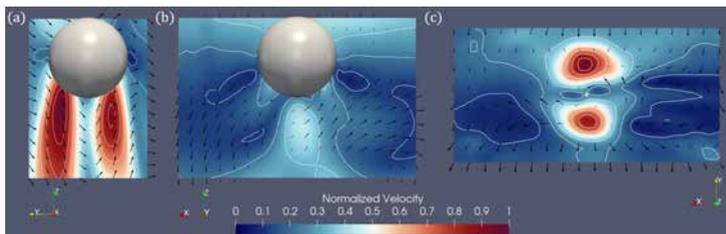


FIG. 104. The 8-th bubble velocity field mode for 100 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

zigzag patterns in the upper half of the metal container.

As mentioned before, this largely stems from the dominant zeroth vessel reference frame velocity field mode (Figure 77) which gradually becomes even more prominent than other quickly decaying modes (Figure 89c). Taking a closer look at the history of RMS amplitudes of the significant bubble modes (Figure 109), one can see that, if anything, modes 4, 8, 10

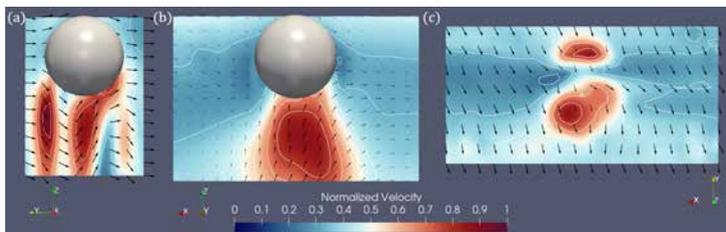


FIG. 105. The 10-th bubble velocity field mode for 100 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

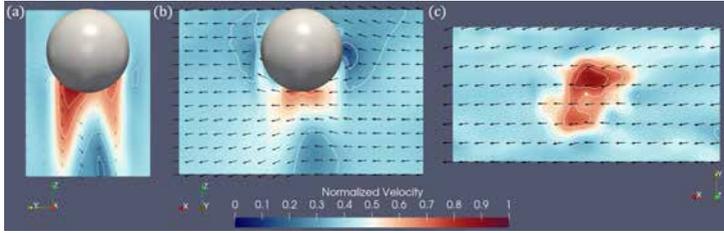


FIG. 106. The 6-th bubble velocity field mode for 100 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

exhibit a slight negative trend. The fact (as far as the presented results indicate) that the trajectory forms and transitions thereof over time in this case are explained by the vessel DMD modes rather than the bubble modes would suggest that larger scale flow structures are responsible, not the bubble wake flow structures. However, it is reasonable to assume that the initial flow asymmetry/perturbations – and therefore the corresponding bubble modes – in the *YZ* plane of the metal vessel are responsible for the onset of the larger scale patterns. The *YZ* plane perturbations themselves might originate from asymmetric bubble detachment from the nozzle – it would make sense that this asymmetry is amplified with increased gas flow rate. It is then of interest to investigate how the *YZ* symmetry is broken as the flow rate varies from 30 *sccm* to 100 *sccm* and up, and to see how that is reflected in the bubble DMD modes. Among the considered modes, the following are significantly correlated: $8 \leftrightarrow (6, 9)$, $10 \leftrightarrow (4, 12)$.

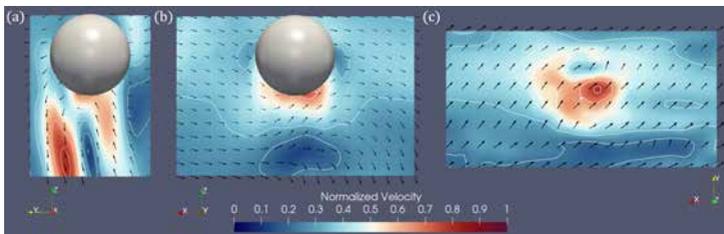


FIG. 107. The 9-th bubble velocity field mode for 100 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

Transitioning to the cases without MF, consider first the 30 *sccm* case: DMD mode analysis results are shown in Figures 110-112. One may notice that the mode frequencies in Figure 110b are overall lower than in the two cases with applied MF (Figures 90b and 99b) – this tendency was

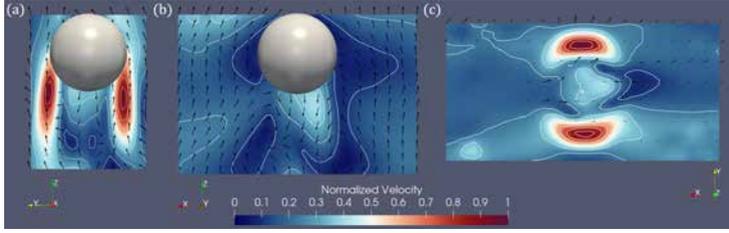


FIG. 108. The 12-th bubble velocity field mode for 100 *sccm* with applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

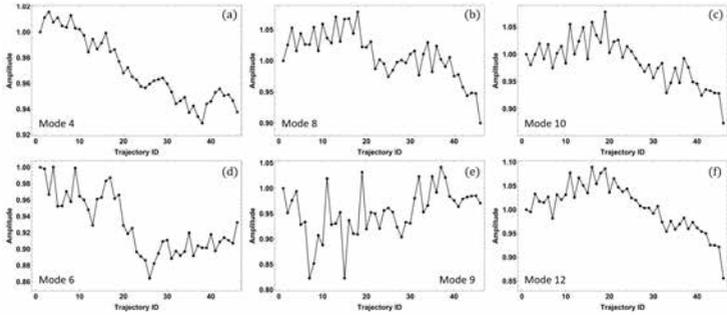


FIG. 109. 100 *sccm* with applied MF: averaged RMS amplitude history over the flow time (trajectory IDs) for significant bubble velocity field modes. The averaging window width is 5 trajectories. The amplitude normalization is post-averaging and separate for each mode: (a) 4th mode, (b) 8th mode, (c) 10th mode, (d) 6th mode, (e) 9th mode, (f) 12th mode.

also observed above for the vessel velocity field modes (Figures 86b-89b). Another difference clearly seen in Figure 111 is that both the first few significant modes that are dominant (a) and the lesser modes (b) constitute a very significant fraction of the zeroth mode's amplitude: ~ 0.48 and ~ 0.29 , respectively.

As seen in Figure 110c, only two modes, the zeroth and 16-th, have slight negative growth rates while every other mode has a positive growth rate leading to an overall increase in the relative importance of non-zeroth modes over the mean trajectory time: amplitude sums for the outlined mode groups become ~ 0.65 and ~ 0.48 , respectively. On this note, it is interesting that the lower frequency modes, 1 to 3, exhibit very small relative amplitude growth. Moreover, Figure 112 indicates that the RMS amplitude of mode 7 (second-strongest mode after mode 0) increases overall

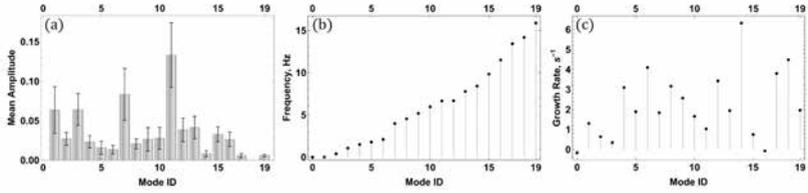


FIG. 110. 30 *sccm* without applied MF: bubble velocity field mode (a) normalized initial amplitudes, (b) frequencies and (c) growth rates.

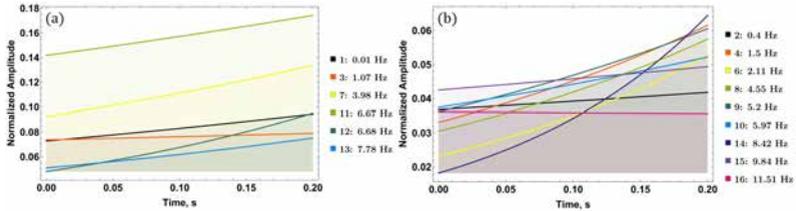


FIG. 111. 30 *sccm* without applied MF: amplitude dynamics over the mean trajectory time for significant modes: (a) dominant modes, (b) secondary modes.

over the flow time, while for mode 11 (the dominant mode second to mode 0) the RMS amplitude shows an overall decline. The most significant modes are shown in Figures 113-119.

The zeroth mode ($f_0 \sim 5.4$ *mHz*, $a_0 \sim -0.16$ s^{-1} , Figure 113) is again with the highest amplitude that is consistently ~ 1 for all trajectories with a $\sim 4.6\%$ deviation. Although its growth rate is higher than in the cases

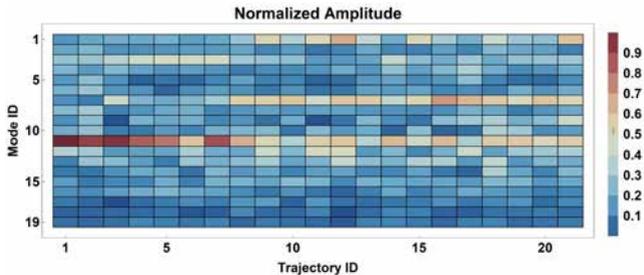


FIG. 112. 30 *sccm* without applied MF: normalized root-mean-square (RMS) amplitudes for modes over all processed trajectories.

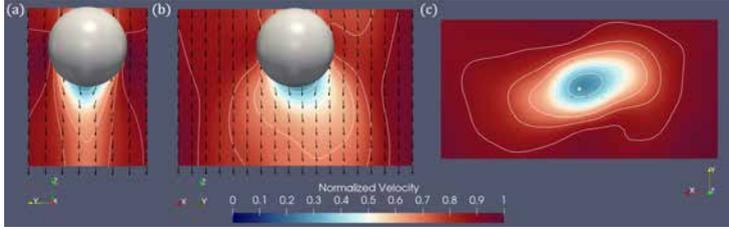


FIG. 113. The zeroth bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

with applied MF roughly by an order of magnitude, its amplitude still only decays to 0.96 on average. It is clear, especially from Figure 113c, that in this case the zeroth mode is very asymmetric in the *XY* plane while mostly retaining symmetry in the *YZ* plane (a). It is interesting to note that the next two dominant modes, 11 ($f_{11} \sim 6.7 \text{ Hz}$, $a_{11} \sim 1.0 \text{ s}^{-1}$, Figure 114) and 7 ($f_7 \sim 4.0 \text{ Hz}$, $a_7 \sim 1.8 \text{ s}^{-1}$, Figure 115), are rather symmetric as well, especially in the *XZ* plane. Note also the very pronounced wake zone beneath the bubble in both cases, and that both modes indicate dominant flow in the *X* direction with a smaller contribution in the *Y* direction. The velocity field patterns seen in Figures 114b and 115b and the mode frequencies consistent with zigzag wavelengths given the mean trajectory time suggest these modes might be responsible for vortex shedding and trajectory zigzagging in the *XZ* plane, since velocity field oscillations in the *X* direction are the strongest in areas where wake vortices form [214]. Mode 12 ($f_{12} \sim 6.7 \text{ Hz}$, $a_{12} \sim 3.4 \text{ s}^{-1}$, Figure 117) is rather similar to 11 and 7, but is less symmetric and has a greater *Y* component. Modes 11 and 12 are also very strongly correlated.

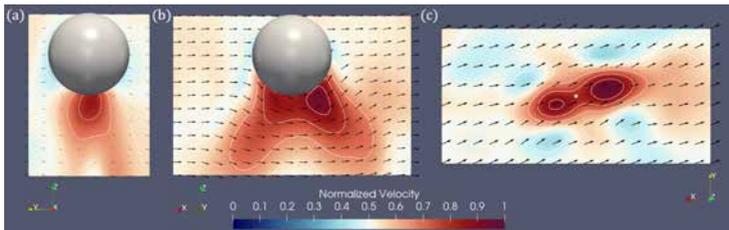


FIG. 114. The 11-th bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

The 1-st mode ($f_1 \sim 12 \text{ mHz}$, $a_1 \sim 1.3 \text{ s}^{-1}$, Figure 116) seems to be

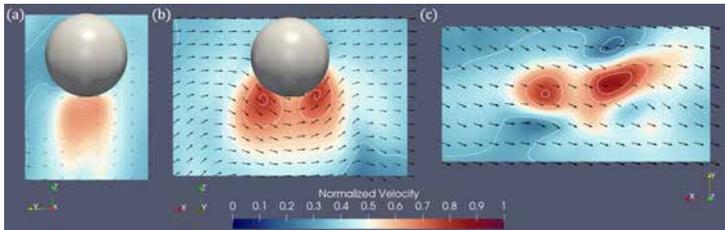


FIG. 115. The 7-th bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

related to the zeroth mode in that it exhibits intense flow in the *Z* direction and the flow pattern in Figure 116c suggests it might be a higher order spatial/temporal harmonic of mode 0 (frequencies differ by a factor of ~ 2.2). Note that modes 0 and 1 are significantly correlated [214].

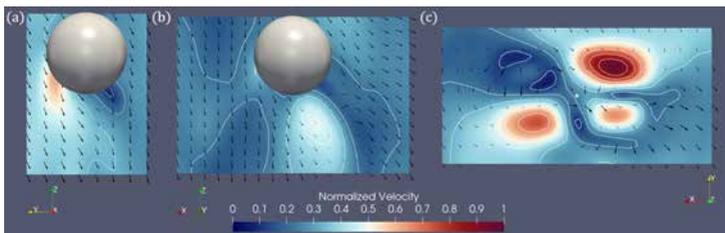


FIG. 116. The 1-st bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

The other two of the dominant modes are 3 ($f_3 \sim 1.1 \text{ Hz}$, $a_3 \sim 0.35 \text{ s}^{-1}$, Figure 118) and 13 ($f_{13} \sim 7.8 \text{ Hz}$, $a_{13} \sim 1.9 \text{ s}^{-1}$, Figure 119) – these exhibit rather complex flow patterns and have no obvious symmetries or clear general velocity field directions. Modes that are grouped in Figure 111b generally contribute velocity fields that are largely distorted versions of modes 11, 7, 1 or exhibit no clear symmetry as in modes 3 and 13.

Modes 15 ($f_{15} \sim 9.8 \text{ Hz}$, $a_{15} \sim 0.75 \text{ s}^{-1}$) and 16 ($f_{16} \sim 12 \text{ Hz}$, $a_{16} \sim -6.9 \cdot 10^{-3} \text{ s}^{-1}$) are interesting, in that they are essentially perturbed variations of mode 12 seen for 100 *sccm* with applied MF (Figure 108) with flow velocity magnitude maxima rotated by $\pi/2$ in the *XY* plane.

Mode 8 ($f_8 \sim 4.6 \text{ Hz}$, $a_8 \sim 3.2 \text{ s}^{-1}$), shown in Figure 120, deserves special attention despite its relatively low amplitude. The velocity field pattern seen in Figure 120b suggests that this mode is likely connected to vortex shedding – not directly like modes 11 and 7, but rather it represents

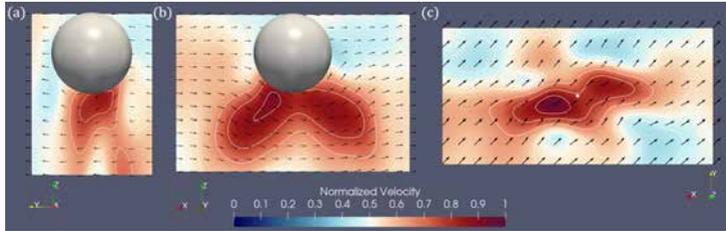


FIG. 117. The 12-th bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

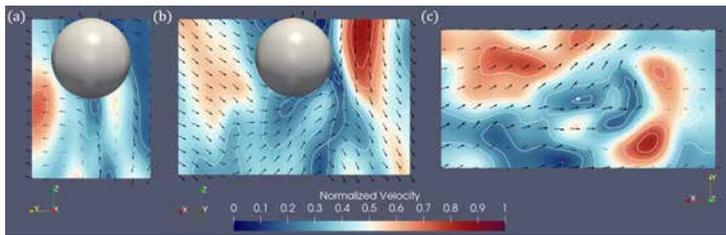


FIG. 118. The 3-rd bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

a common wake pattern right after vortex detachment where there is an upward flow about the bubble at the side (in the *XZ* plane) wherefrom detachment occurred. In this case, or rather looking at the given mode phase, detachment took place to the right of the bubble and a vortex with clockwise rotation induced the upwards flow region. Flow field aside, two more factors suggest this might be the right interpretation. First,

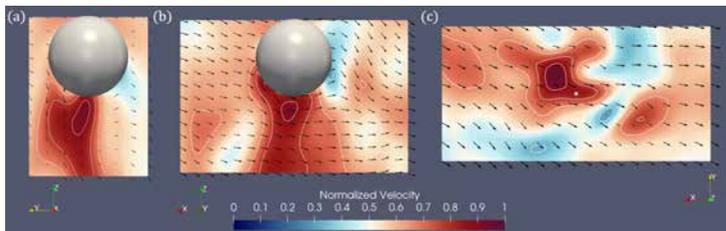


FIG. 119. The 13-th bubble velocity field mode for 30 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

the frequency is roughly adequate, as it corresponds to a period of about the mean trajectory time. The zigzagging motion in the XZ plane of the trajectories fits ~ 1 -1.5 wavelengths into the trajectory, which is slightly off the mode frequency, but is close enough to arouse suspicion. Second, Figure 111b indicates that the 8-th mode's amplitude increases significantly (versus its initial value) over trajectory time, which makes sense given that, as the bubble accelerates during ascension, the detached vortices also exhibit greater velocity and vorticity magnitudes. Another minor argument is that mode 8 is rather strongly correlated with modes 11 and 12 that are also in all likelihood, as noted above, linked to vortex shedding.

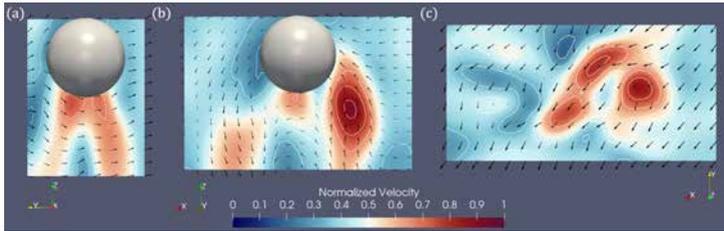


FIG. 120. The 8-th bubble velocity field mode for 30 *sccm* without applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

One may note that trajectory XY projections in the form of distorted spirals and tightly packed self-intersecting lines are due to the relatively much more disordered and asymmetric mode flow patterns seen in this case as opposed to the cases with applied MF. Finally, to reiterate on the correlation between the modes, one has the following dependencies: $0 \leftrightarrow 1$, $13 \leftrightarrow (7, 11)$; correlated triplets $(3, 11, 12)$ and $(8, 11, 12)$ [214].

Given the DMD results for 30 *sccm* above, it is important to see how they are different from 100 *sccm* without applied MF. DMD mode analysis results are shown in Figures 121-123. As with 30 *sccm* without MF, the frequencies for 100 *sccm* are consistently lower than in the cases with applied MF (Figure 121b). Note also that the growth rates for 100 *sccm* (Figure 121c) are, like for 30 *sccm*, positive except for modes 0 and 10 (0 and 16 for 30 *sccm*, as seen in Figure 110c). The growth rate magnitudes are greater overall, though. Another difference is that in this case, higher frequency modes have greater initial amplitudes than with 30 *sccm* (Figure 121a versus Figure 110a). Figures 122a and 122b show the dominant modes and the ones with lesser overall amplitudes, respectively – note that the two overall strongest modes in 122a, modes 14 and 13, have higher frequencies than the two strongest dominant modes in Figure 111a. Since grouping modes into dominant/secondary groups is difficult to do strictly, in this

case the 10-th mode (Figure 122b) was used as a nominal boundary since it changes very little over the mean trajectory time, and its and the 9-th mode's final amplitudes are on average just below that of mode 11 (Figure 111a).

Another major difference is that the dominant modes for 100 *sccm* have lower relative initial amplitudes than in the 30 *sccm* case and have a less sparse value distribution (Figure 122a). In addition, the differences in amplitudes at the start and end of mean trajectory time are greater for most modes at 100 *sccm*. Specifically, amplitudes of dominant dynamic modes initially constitute ~ 0.44 and this sum increases to ~ 0.93 over the mean trajectory time (Figure 122a), while lesser amplitudes initially amount to ~ 0.26 , which becomes ~ 0.48 over trajectory time. The zeroth mode which stands for the mean flow field has an amplitude of ~ 1.08 on average (1 initially) with a 5.2% deviation, meaning that the dominant non-stationary modes practically overshadow the mean flow field with secondary modes also contributing significant perturbations. This leads one to expect much more disturbed wake patterns than in the 30 *sccm* case, as it should be.

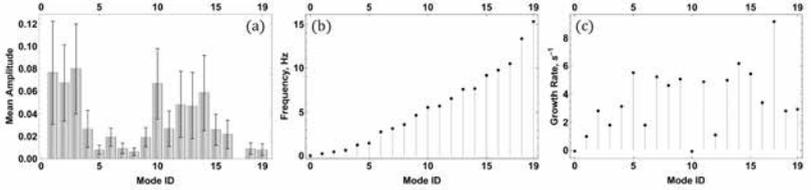


FIG. 121. 100 *sccm* without applied MF: bubble velocity field mode (a) normalized initial amplitudes, (b) frequencies and (c) growth rates.

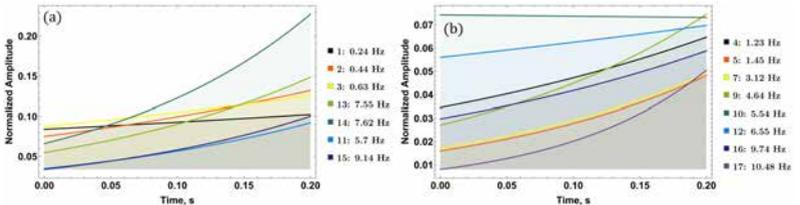


FIG. 122. 100 *sccm* without applied MF: amplitude dynamics over the mean trajectory time for significant modes: (a) dominant modes, (b) secondary modes.

It is also interesting to see that mode RMS amplitude values over consecutive trajectories (Figure 123) seem to exhibit oscillatory trends for five

of the dominant modes as well as mode 12 – close inspection indicates that it is indeed so, as seen in Figure 124. This is interesting as it raises the question of how these trends are connected to the vessel reference frame velocity field modes. Given that the trajectories processed by the DMD algorithm cover ~ 5 s of flow time and the frequencies of significant extrema seen in Figure 124, it seems likely that the considered bubble wake modes are modulated by vessel flow modes with roughly < 1 Hz frequencies – this would correspond to modes 1 to 4 which are the strongest after the zeroth vessel flow mode (Figure 87).

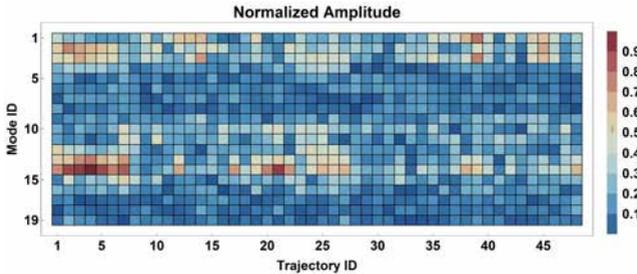


FIG. 123. 100 *sccm* without applied MF: normalized root-mean-square (RMS) amplitudes for modes over all processed trajectories.

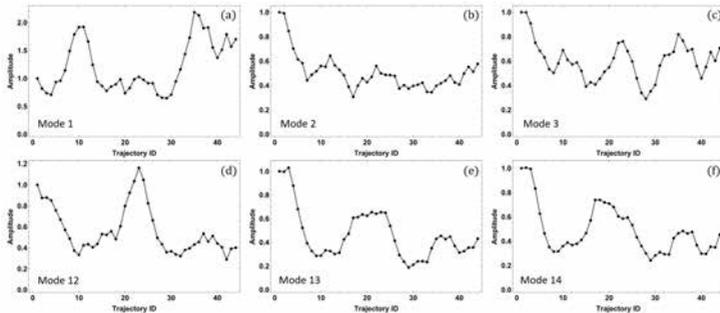


FIG. 124. 100 *sccm* without applied MF: averaged RMS amplitude history over the flow time (trajectory IDs) for significant bubble velocity field modes. The averaging window width is 5 trajectories: (a) 1-st mode, (b) 2-nd mode, (c) 3-rd mode, (d) 12-th mode, (e) 13-th mode, (f) 14-th mode.

The significant bubble wake modes for 100 *sccm* with applied MF are shown in Figures 125-131. The zeroth mode is as in other cases stationary ($f_0 \sim 6.6$ mHz, $a_0 \sim -6.2 \cdot 10^{-2}$ s $^{-1}$, Figure 125) and, rather curiously, is

not quite as asymmetric as its counterpart for 30 *sccm* (Figure 113). It also expectantly has a much more developed bubble wake.

Mode 14 ($f_{14} \sim 7.6 \text{ Hz}$, $a_{14} \sim 6.2 \text{ s}^{-1}$, Figure 126) exhibits a velocity field pattern that is consistent with bubble wake structure during bubble tilting in the XZ plane [214], where a stagnation zone forms on one side of the wake while higher velocity is observed on the other side (XZ plane). Another argument for this conjecture is that the mode frequency is a good match for the mean observed trajectory oscillation frequency. Mode 13 ($f_{13} \sim 7.5 \text{ Hz}$, $a_{13} \sim 5.0 \text{ s}^{-1}$) has the next highest RMS amplitude overall, but is not shown here since its flow pattern is remarkably similar to that of mode 14. In fact, modes 13 and 14 are extremely strongly correlated (~ 0.97 correlation value) and suggests that mode 13 might simply be a slightly phase-shifted version of mode 14, especially given that $f_{14} \approx f_{13}$ [214].

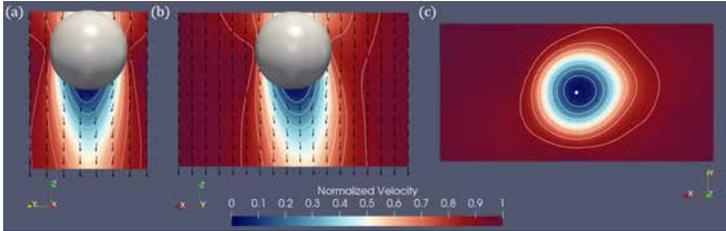


FIG. 125. The zeroth bubble velocity field mode for 100 *sccm* without applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

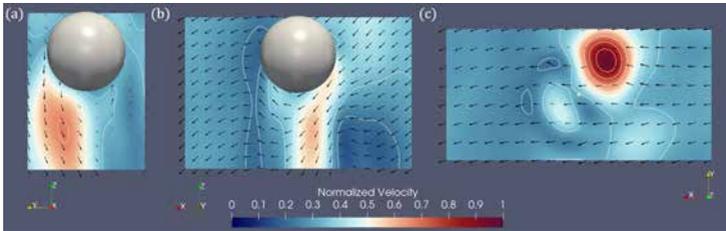


FIG. 126. The 14-th bubble velocity field mode for 100 *sccm* without applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

Mode 2 ($f_2 \sim 0.45 \text{ Hz}$, $a_2 \sim 2.8 \text{ s}^{-1}$, Figure 127) does not seem to be directly linked to a certain momentum transfer mechanism within the wake, but rather exhibits flow oscillations mostly in the XY plane with a very low frequency. Mode 3 ($f_3 \sim 0.63 \text{ Hz}$, $a_3 \sim 1.8 \text{ s}^{-1}$) is likewise

difficult to interpret in terms of clear processes that one would expect in a bubble wake, but rather it shows disordered low frequency oscillations mostly in the Z direction in the form of four "jets" seen in Figure 128c as four velocity field maxima with almost zero XY component – two of these are partially captured in Figures 128a and 128b. It must be noted that modes 2 and 3 are significantly correlated [214].

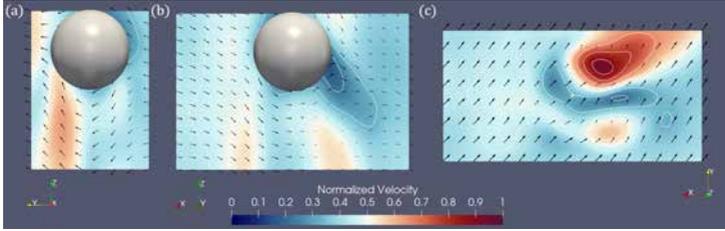


FIG. 127. The 2-nd bubble velocity field mode for 100 *scm* without applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

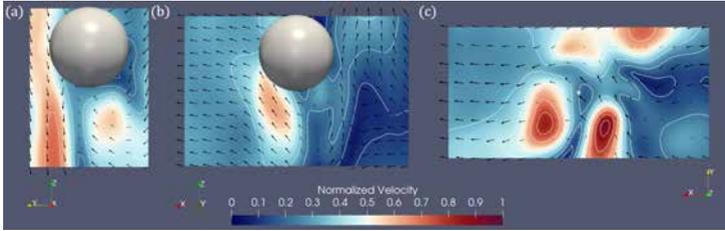


FIG. 128. The 3-rd bubble velocity field mode for 100 *scm* without applied MF: (a) YZ plane, (b) XZ plane and (c) XY plane.

Mode 1 ($f_1 \sim 0.23 \text{ Hz}$, $a_2 \sim 1.0 \text{ s}^{-1}$, Figure 129) contributes velocity field oscillations in the Y direction with a relatively weak Z component and the flow pattern suggests that the mode may be linked to wake oscillations in the YZ plane. It is, however, unclear how, given the low frequency.

Mode 15 ($f_{15} \sim 9.1 \text{ Hz}$, $a_{15} \sim 5.4 \text{ s}^{-1}$, Figure 130) on the other hand has both the timescale and the flow field that fit wake oscillations as the bubble tilts periodically. This mode is also rather symmetric in the YZ plane in terms of velocity magnitude (Figure 130a) and exhibits a swirl flow pattern in the XY plane (Figure 130c). Mode 11 ($f_{11} \sim 5.7 \text{ Hz}$, $a_{11} \sim 4.9 \text{ s}^{-1}$, Figure 131) is interesting, in that it exhibits flow field oscillations almost entirely in the X direction, with an asymmetric velocity maximum located just above (Figure 131b) the zone where the wake vortex

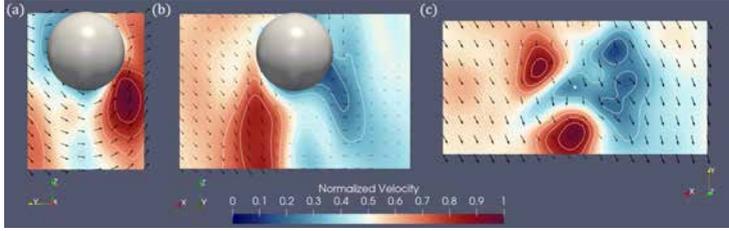


FIG. 129. The 1-st bubble velocity field mode for 100 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

should be expected to form [214]. This mode's frequency is within the range where the mode could be responsible for affecting vortex shedding via velocity pulsations in the *X* direction. Mode 11 is also quite symmetric in the *YZ* plane.

Other notable modes include mode 12 ($f_{12} \sim 5.7 \text{ Hz}$, $a_{12} \sim 4.9 \text{ s}^{-1}$) which is somewhat similar in the flow field structure to mode 8 for 100 *sccm* with applied MF (Figure 104), but with a more developed bubble wake zone. Mode 9 ($f_9 \sim 4.6 \text{ Hz}$, $a_9 \sim 5.1 \text{ s}^{-1}$) has the velocity field resembling mode 7 for 30 *sccm* without applied MF (Figure 115), but with much more pronounced velocity maximum zones beneath the bubble that are elongated in the *Z* direction and an overall greater *Y* component. Mode 10 ($f_{10} \sim 5.5 \text{ Hz}$, $a_{10} \sim -7.5 \cdot 10^{-3} \text{ s}^{-1}$) in turn closely resembles modes 2, 13 and 14 in this case and is somewhat correlated to all three. Mode 16 ($f_{16} \sim 9.7 \text{ Hz}$, $a_{16} \sim 3.4 \text{ s}^{-1}$) is essentially mode 15, but with a smaller RMS amplitude and a slightly greater frequency ($\omega_{15} \sim 9.1 \text{ Hz}$). Note that modes 15 and 16 are very strongly correlated [214].

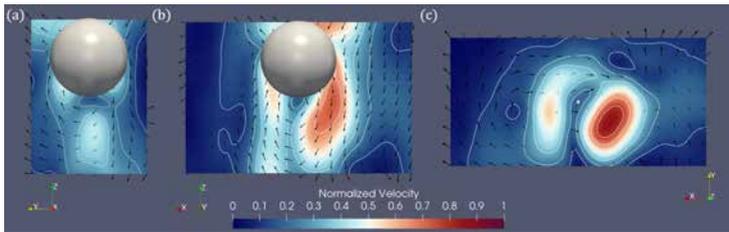


FIG. 130. The 15-th bubble velocity field mode for 100 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

It should be said that, while some of the modes for 100 *sccm* without MF are physically interpretable, at least hypothetically, it is very difficult to

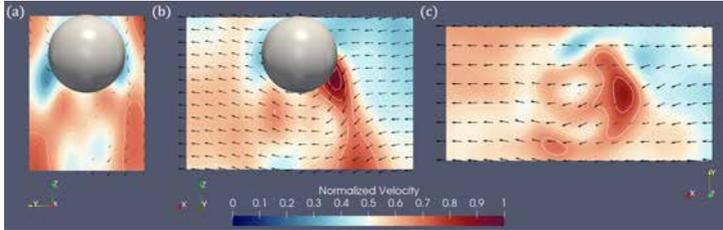


FIG. 131. The 11-th bubble velocity field mode for 100 *sccm* without applied MF: (a) *YZ* plane, (b) *XZ* plane and (c) *XY* plane.

attribute them to the various features of trajectories. While one could argue that the initial few trajectories can be explained by the initial dominance of modes 10, 13 and 14 that would fit the oscillations seen in the *XY* projections of bubble trajectories, later trajectories are highly disordered, some without clearly discernible patterns, making it very difficult to attribute specific modes to the observed behavior given the present analysis [214]. While some of the seemingly random trajectory perturbations could be due to the oscillations in the RMS amplitudes of some of the key modes (Figure 124), it is currently unclear to what extent. To summarize the mode correlations, one has the following dependencies: $14 \leftrightarrow (2, 13)$, $2 \leftrightarrow (3, 14)$, $11 \leftrightarrow (1, 9)$, $15 \leftrightarrow (1, 9, 16)$ [214].

4. Summary

With the above results, the first application of DMD to bubble flow with resolved dynamic liquid/gas boundaries has been demonstrated – specifically, DMD was applied to the output of numerical simulations for a system where chains of bubbles ascend through a rectangular liquid metal vessel, and studied how gas flow rate and applied MF affect dynamic modes of the velocity field. This was done both at the metal volume scale and at the bubble wake scale, and the results of this proof-of-concept study indicate that DMD is indeed a very prospective tool that can yield unique insights into various momentum transfer and bubble interaction mechanisms (including how processes at different spatial and temporal scales interact), and that mode analysis can be used to explain the observed flow patterns. A much more detailed summary of the results and their physical interpretation is provided in [214].

DMD enabled to observe mean flow symmetry changes/breaking in model systems, and how they were caused: it was found that the flow fields are very sensitive to MF asymmetries, and instabilities can arise even at lower gas flow rates, even as MF tends to stabilize the flow overall.

DMD modes corresponding to the bubble chain reveal oscillation frequencies and wake scales, and allow measuring chain coherence length, which is a measure of chain flow stability.

Observing DMD modes of bubble wake in the bubble reference frame enables one to separate the dynamics into different momentum transfer mechanisms. By comparing the statistics and dynamics of relative mode amplitudes and frequencies, one can infer how wake flow affects global flow fields, and vice versa – specifically, observations suggest that flow instabilities and symmetry breaking originates at small scales, possibly within wakes, and propagate to larger scales; with greater spatial resolution, DMD could clearly show how it happens.

Given higher resolution simulations (both spatially and temporally), DMD could also be used to study turbulent flow structures that form as the bubble chain flow evolves. Moreover, with finer meshes, one could readily study how MF affects bubble shape dynamics, specifically which surface perturbation wavelengths are affected by MF the most and how.

The developed DMD code is both noise-resilient and memory-efficient, and available on *GitHub*: [MartinKlevs/PyDMD](#). The newly implemented classes names start with *mosesdmd* and all use implemented MOSES-SVD – its implementation can also be found on *GitHub*: [MartinKlevs/MOSES-SVD](#).

G. X-ray imaging of metal alloy directional solidification

The lack of a convenient and sufficiently general all-in-one solution was the motivation to present the first version of the open-source code developed for automatic analysis of dynamic X-ray radiography images of directional solidification processes studied using Hele-Shaw cells. The current version does not yet have an integrated optical flow component (e.g., like the code used in [142]), but otherwise it meets the above-mentioned analysis functionality requirements. Moreover, it was designed for robustness and is quite resilient to image noise and low image contrast. Utilized approaches to image and data processing combine both well-known state of the art and original methods, particularly for solid structure segmentation and dendrite grain analysis. The performance of the methodology implemented in the code is demonstrated on data from *in situ* X-ray radiography experiments with a gallium-indium alloy [54, 59].

1. Image characterization

All images used in here were acquired at the X-ray lab at Helmholtz-Zentrum Dresden-Rossendorf (HZDR). The Hele-Shaw solidification cell with dimensions $35\text{ mm} \times 25\text{ mm} \times 0.15\text{ mm}$ was imaged at 1 frame per second with a 1 s exposure time. The imaging system utilized the *Phoenix X-ray XS225D-OEM* X-ray tube, and is described in more detail in [46, 53, 54, 59, 64]. For each image sequence recording, dark current signals and X-ray beam profile signals were recorded for subsequent image correction and normalization during pre-processing. For every set of system parameters, repeated recordings were made to ensure the results are reproducible, as well as for redundancy. Images are 16-bit gray-scale TIFFs and the field of view (FOV) typically has a $\sim 760 \times 576\text{ px}$ (pixel) image size with a pixel size [13.7; 37.6] μm (actual image size varies between different image sequences due to boundary cropping). Figure 132 is an example of an acquired image of a dendritic network in the solidifying gallium-indium alloy [54, 59].

To simplify further description and analysis, it is important to at least informally define key image features. In Figure 132, one can see the *solidification front* (SF) outline with a black dashed curve. Here the SF is defined as the envelope of the FOV region with the solidified structures. A more precise *operational* definition in the image processing context will be given below. The region (or multiple) containing solidified structures delimited by the solidification front is the *solid zone* (SZ). This zone may also include liquid *cavities* isolated from the bulk liquid and *channels* that

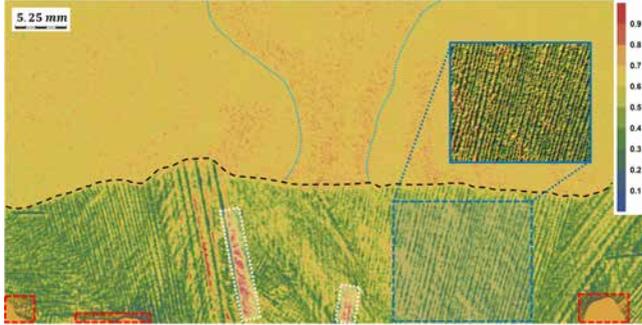


FIG. 132. A typical (false color) FOV in one of the imaging experiments after pre-processing. The black dashed curve is a rough outline of the solidification front. The light-blue dashed line indicates one of the convective plumes within the FOV. Interior liquid metal pools are highlighted with white dashed frames. Note also the image artifacts highlighted with red dashed frames at the bottom of the FOV. The dark blue dashed frame indicates a region that belong to one of the formed dendrite grains – the inset (dark blue frame) shows its relief plot for visual clarity. Both the false-color FOV image and the relief plot share the same color scale (color bar to the right) and scale (top-left corner).

are connected to the bulk liquid (above the SF), e.g. one such channel will later form from the largest of the closed liquid pools highlighted in Figure 132. Thus, the *liquid zone* (LZ) is the difference between the FOV and the SZ, minus the cavities and channels. These definitions will be used throughout the rest of this summary.

The images exhibit Poisson (multiplicative) noise, as well as salt-and-pepper noise due to momentarily overexposed or unresponsive ("dead") camera pixels. The X-ray beam flux over the FOV is nonuniform, with a fall-off near the edges of the acquired images. The contrast-to-noise ratio (CNR) is different for the LZ and the SZ. The convective plume CNR in the LZ is initially rather good, but typically degrades over time as the *solid fill factor* (SFF, the ratio of the SZ area to the FOV area) of the cell increases – this is because the solute is ejected above the SF, the LZ is saturated, and the contrast between the liquid alloy components diminishes. In addition, the CNR in the SZ may also vary over the image since there is solute flow across the solidified structures, potentially occluding them. The signal-to-noise ratio (SNR) is usually adequate for structures in the SZ, but it is rather low for the convective plumes in the LZ. In addition, some of the images may exhibit larger-scale artifacts – for instance, as shown in Figure 132 with red dashed frames. In this case, the artifacts are the spots

where the two parallel walls of the Hele-Shaw cells were fused together.

2. Methodology

a Assumptions and considerations

The developed image processing code must enable in-depth analysis of both the LZ and SZ over time, as well as the dynamics associated with the SF evolution. Therefore, the objectives are as follows: segment the LZ and SZ; derive the SF; identify channels connected to the LZ, and also liquid cavities within the SZ; segment the convective plumes within the LZ for shape analysis; extract skeletons of the structures (in this case dendrites) within the SZ; perform orientation analysis for the structures identified within the LZ; decompose the SZ structures into subdomains (grains) by orientation and connectivity; measure the solute concentration near the SF.

The following assumptions are made regarding the images and the physical system. When treating the SZ, the noise is considered white Gaussian. This is because after pre-processing the image luminance does not vary too much at length scales much greater than the dendrite thickness and inter-dendrite spacing (Figure 132) – this is in contrast with the LZ (compare the luminance distribution within the dark-blue dashed frame against that above the SF). Persistent larger-scale artifacts (i.e. not pixels with outlying luminance values) in the images, if present, are considered stationary over time. Dendrites in the SZ have linear or only slightly curved shapes, and may overlap in the imaging plane, and thus their X-ray radiography projections may cross.

Given the above, the following considerations determine the methods of choice. The pixel size is a significant fraction of the dendrite width, so care must be taken when attempting to remove noise from dendrites. At the same time, the textures in the solid domain are rather fine – the inter-dendrite spacings are roughly of the same order of magnitude as the dendrite width. The SNR is such that methods which are not texture-/morphology-aware cannot achieve satisfactory non-destructive denoising. Liquid flow across the dendrites acts as correlated noise when attempting to derive dendrite morphology. This further complicates the solid structure analysis – treating this issue jointly with the Gaussian noise stemming from under-exposure does not produce good enough results and a separate approach is needed. While the CNR at the SF is rather high in the example seen in Figure 132, in other cases the SF is not as smooth and is less contrast, and segmenting liquid cavities and channels can be even more difficult. Therefore, dedicated filters are required to significantly increase

the CNR of the liquid/solid phase boundary before SZ/LZ segmentation. The segmentation method for SZ/LZ separation must reliably work under potentially varying image quality: the ray produced by the X-ray tube may flicker and has non-uniform intensity; solid fraction increase and solute ejection into the bulk liquid strongly change the image luminance distribution both locally and globally. These effects should be modelled by the utilized segmentation method. Methods used for LZ denoising must not overly deform or smear out the shapes of convective plumes, but denoising here is much less constrained than in the SZ. Image quality varies greatly across different experiments and image sequences, both acquired by the author and those performed by other researchers – it is therefore worthwhile to develop a code that is resilient and can operate under adverse conditions, potentially much worse than what is seen in Figure 132. Such a code with many components and methods will inevitably have a rather large number of parameters – these should either be mostly fixed/general or should be quickly optimizable.

Image processing is organized in stages outlined in Algorithm 13.

b Pre-processing

Image pre-processing is performed in *ImageJ*, and the steps are given in [223]. The procedure involves DCC and FFC, as with neutron and X-ray images of bubble flow. FFC results in the spatial dependence of the SNR. After pre-processing, the images are normalized and saved, then passed to *Wolfram Mathematica* for further processing.

c Liquid/solid zone separation

Sometimes the images will still contain artifacts even after cropping and pre-processing (for example, as in Figure 132). While FFC ensures that such artifacts are no longer strong outliers, these image areas still significantly affect image luminance histograms and may interfere with image filtering and segmentation. A procedure is therefore used that identifies and inpaints these defects, i.e. makes them seamless with respect to the surrounding image textures. The artifact mask for an image sequence is obtained as outlined in Algorithm 14.

Note that dark artifacts can be detected using an identical procedure applied to an inverted image. Morphological dilation using disk structural elements [165] is performed so that the artifact mask has a safety buffer. The latter is necessary for texture synthesis-based inpainting to properly fill the artifact zones using samples from adjacent textures sufficiently far from the artifacts. The parameters are provided in [223].

Algorithm 13: Overall structure of the image processing pipeline

Input: Raw image sequence

- 1 Pre-process images
- 2 Remove image artifacts (Algorithm 14)
- 3 Segment the SZ (Section II G 2 c & Algorithm 15)
- 4 Identify channels connected to the LZ, and also liquid cavities within the SZ (Algorithm 16)
- 5 Derive the SF and segment the LZ (Algorithm 16)
- 6 Extract skeletons of the structures (in this case dendrites) within the SZ (Algorithms 17, 18 & 19)
- 7 Perform orientation analysis for the structures identified within the LZ (Algorithms 20 & 21)
- 8 Decompose the SZ structures into subdomains (grains) by orientation and connectivity (Algorithms 20 & 21)
- 9 Measure the solute concentration near the SF (Algorithm 22)
- 10 Segment the convective plumes within the LZ for shape analysis (Algorithm 23)

Output:

- SF shape, height map and growth rate over time
 - Solute concentration dynamics near the SF
 - Shape dynamics for convective plumes in the LZ
 - Dendrite structure maps with highlighted features
 - Dendrite orientation spectra for the SZ
 - Dendrite orientation spectra and relative areas for grains identified within the SZ
-

Algorithm 14: Artifact removal for pre-processed images

Input:

- Averaged reference images
 - Pre-processed image sequence
- 1 Artifact mask: binarization (possibly multi-pass) with a user-defined threshold or another appropriate method (e.g. Otsu [166])
 - 2 Morphological dilation
 - 3 Artifact inpainting using texture synthesis [190]

Output: Images without artifacts

Prior to segmentation, image filtering is performed to increase the CNR for the LZ/SZ boundaries, including liquid cavities (CNR tends to be especially low) and channels. Here, the filters were applied such that they also eliminate dendrite structures while preserving larger-scale liquid zones and larger spaces between dendrites that are filled with liquid. It was

decided to use block-matching 3D (BM3D) filtering [224–227], since, unlike other tested solutions, it consistently preserved the SZ shape well and also strongly increased the CNR of channels and cavities within the SZ.

Image filtering is performed in two stages. First, BM3D is applied. Its basics and explanation for why a certain version of BM3D is used are outlined in [223], and an in-depth account of the inner workings of BM3D can be found in [224–227]. The latest *MATLAB* implementation by Tampere University of Technology (Python version also available) is used, which is based on [227]. It is integrated into the *Wolfram Mathematica* code using *MATLink* for seamless image processing. The second filtering stage is the non-local means (NM) filter [170, 171] which is used to mitigate any artifacts left over and/or produced by BM3D and further increase the CNR for the SF and cavities and channels within the SZ. NM can be viewed as a simplified/precursor version of BM3D. BM3D and NM filter [172] setup is explained in [223].

Despite significant improvements to SNR and CNR due to filtering (Section II G 2 c), segmentation is still a challenge due to generally imperfect reference-based FFC, spatially varying SNR and CNR, X-ray beam fluctuations and strongly varying SFF and solute concentration in the LZ over time, as well as formation, growth/shrinking and disappearance of cavities and channels. The tests show that (at least in these cases) global segmentation methods, even advanced ones, are not able to detect the SZ stably and accurately over an entire image sequence that usually starts with no solid zone and possibly ends with $SFF \approx 1$. An empirical "physics-aware" model that computes an adaptive binarization threshold for the filtered images has therefore been adopted. The segmentation steps are summarized in Algorithm 15.

Algorithm 15: SZ segmentation

Input:

- Raw image sequence
 - Filtered image sequence (Section II G 2 c)
- 1 (Optional) Apply Gaussian/median filtering to the raw images
 - 2 Compute the mean inverse luminance (MIL) for the raw images
 - 3 (Optional) Filter the MIL time series
 - 4 Compute the MIL-based adaptive threshold time series (39)
 - 5 Segment the SZ from the filtered (i.e. post BM3D and NM) images using the adaptive threshold

Output: SZ masks for the entire image sequence

The initial sequence of *raw* images is used, because it is desirable to capture the beam fluctuations as well. Optionally, small-radius me-

dian/Gaussian filtering may be applied first to mitigate outliers in the images. Afterward, the mean inverse luminance (MIL) $\langle 1 - I(t) \rangle$ is computed for all (normalized) images in a sequence where t is the time/frame index. To avoid over-fitting the adaptive threshold to the MIL time series, it is (optionally) filtered using the Gaussian total variation (TV) filter [228]. The adaptive threshold $\tau(t)$ for images is computed from MIL as follows:

$$\tau(t) = 1 - C_1 \cdot f_1(t) \cdot f_2^p(t); \quad \underbrace{f_1(t) = L_{\text{TV}} \langle 1 - I(t) \rangle}_{\text{SFF correction}}; \quad \underbrace{f_2(t) = (g_1 \circ g_2)(C_2 + 1 - f_1(t))}_{\text{LZ saturation correction}} \quad (39)$$

where $g_1(X) = X/\min(X)$, $g_2(X) = X/\max(X)$, $X(t)$ is time series, $C_1 > 0$, $C_2 \geq 0$, $p \in \mathbb{R}$ and L_{TV} is the TV filtering applied (if necessary) to the MIL time series.

The adaptive part of the threshold consists of two contributions: $f_1(t)$ and $f_2(t)$, where $f_1(t)$ is tied to the SFF of the FOV – the greater the SFF, the greater the MIL is because the SZ attenuates the X-ray significantly more intensely than the LZ; $f_2(t)$ is the correction based on $f_1(t)$ that accounts for the fact that, as the SFF increases, the solute is pushed out of the solid zone and the LZ becomes significantly more saturated. Both $f_1(t)$ and $f_2(t)$ also capture the instantaneous global illumination changes that could result from X-ray beam flickering. Moreover, $f_2(t)$ plays a very important role in cases where flow in the LZ rapidly transitions between natural and forced convection regimes (e.g. under the influence of applied MF), since these transitions are accompanied by significant changes in the LZ luminance. The SZ is segmented by binarizing the inverted filtered images with respective thresholds $\tau(t)$. In the cases considered here, it was decided not to perform Gaussian/median pre-filtering for the raw image before MIL computation. TV filtering is performed for the MIL time series [164, 228] to avoid over-fitting. Parameters C_1 , C_2 and p for $\tau(t)$ can vary rather significantly between cases, so they are provided for each considered example in [223]. There is, however, a general strategy for quickly optimizing these parameters, which is also given in [223]. $f_2(t)$ in (39) is designed such that the threshold $\tau(t)$ is stricter for greater SFF – this is because the difference in luminance between the SZ and LZ is much lower in the initial stages of solidification than it is when the cell with the alloy is almost entirely filled with the solid phase. This is due to the fact that much of the SZ is initially permeated by the solute, some which is later expelled into the LZ.

Next, liquid cavities and channels are identified, the SF is derived and

the LZ is segmented – the steps involved are outlined in Algorithm 16. Steps 1 generates the SZ mask does not contain the areas occupied by artifacts – this is the mask used for the dendrite structure analysis in Section II G 2 d. Then the liquid phase mask is obtained in Step 2. Image padding in Step 3 is done because the cavities can also have a boundary at the bottom and the sides of the FOV. The top boundary is not padded because once the SF has passed the top of the FOV it is impossible to distinguish cavities at the FOV top from channels. This way, the boundary component removal leaves only the cavities in the processed masks. Steps 4 to 9 fill and exclude the detected cavities from the liquid phase mask (Step 4) and, if necessary, one can also remove leftover artifacts (if any) from SZ segmentation (Step 5); the closing transform [165] (Step 6) is performed using disk structuring elements to fill the asperities in the 0- and 1-valued level sets in the mask – the asperities with length scales below a user-defined disk element size are filled conformally to the nearby level set boundary shape; the filling transform [157] (Step 7) completes the channel filling wherever it was imperfect after closing; Steps 8 and 9 detect channels as the differences between the output of Step 7 and the masks after cavity filling while removing small-scale segments that physically do not correspond to channels, i.e. are either too small to be classified as such or simply are artifacts. To segment the LZ, one adds the artifact segments outside the SZ to the masks with filled channels and cavities (Steps 10 to 13). The SF is obtained by applying the Canny edge detection [229] to the LZ mask. The SF is then smoothed with a small-radius Gaussian filter to eliminate noise introduced by the preceding operations, after which Otsu binarization [166] is applied and the thinning transform [157] is performed so that the SF is exactly 1 *px* thick. The physics that can be derived from the sequence of the SF states over time are shown further in Section II G 3. The parameters are specified in [223].

d Solid domain analysis

Once the SZ has been segmented, one can proceed with the analysis of the solidified structures within the SZ. Since the local textures generally vary strongly throughout the SZ, as does their CNR, it was decided to split the FOV image into a number of partitions, determine which ones contain enough of the SZ for significant analysis (*scan regions*), and perform local filtering and feature extraction. This procedure is detailed in Algorithm 17.

The pre-processed images are partitioned into a regular grid of square patches. Their generation is explained in [223]. Once the images are partitioned, the SZ masks are partitioned in the same way, i.e. image and

Algorithm 16: LZ, cavity and channel separation & SF derivation

Input: SZ masks for the entire image sequence (Algorithm 15)

Get the SZ & liquid phase masks with artifact areas excluded

- 1 SZ: multiply the SZ masks by the inverted artifact mask
- 2 Liquid phase: invert the result and multiply by the inverted artifact mask

Segment cavities & channels

- 3 Cavities: apply $1-p \times$ image padding (pixel value 0, all image boundaries except for the top) to the output of Step 2, remove border components, then revert the image padding
- 4 Subtract the cavity masks from the corresponding liquid phase masks (Step 2) and invert the resulting images
- 5 (Optional) Perform segment size thresholding for the resulting masks
- 6 Invert the masks, then apply the closing transform
- 7 Apply the filling transform and invert the result
- 8 From the output of Step 7, subtract the difference between *it* and its segment size-thresholded version
- 9 Channels: find the image difference between the output of Steps 4(5) & 8 and perform segment size thresholding for the result

Derive the SF & segment the LZ

- 10 Apply $1-p \times$ image padding (pixel value 0, all image boundaries except for the bottom) to the artifact mask, remove border components, then reverse the image padding
- 11 Apply the closing transform
- 12 Add the outputs of Steps 7 & 11
- 13 LZ: invert the output of Step 12
- 14 SF: perform edge detection for the output of Step 12

SF correction & smoothing

- 15 Apply small-radius Gaussian filtering to the SF masks
- 16 Normalize the images and apply Otsu thresholding
- 17 Perform morphological thinning

Output:

- SZ with artifacts excluded
 - Segmented LZ with artifacts excluded
 - Cavity masks
 - Channel masks
 - SF states
-

mask patches with identical image position indices (assigned to all patches) correspond pixel-to-pixel. Some patches, especially for the images near the beginning of the sequence, are going to be mostly filled with liquid, and thus are not eligible for solid structure analysis. Identifying which regions to scan for solid structures saves a significant amount of computation time.

Algorithm 17: Image partitioning & scan region identification

Input:

- Pre-processed images
 - SZ masks with artifact areas excluded (Algorithm 16)
- 1 Partition the pre-processed images into grids square patches with side lengths based on image dimensions
 - 2 Partition the SZ masks into corresponding square patches; *memoize output*
 - 3 Assign all patches their position indices; *memoize*
 - 4 Compute the SFF for the SZ mask patches and assign the values to the respective image patches
 - 5 Designate the image patches with $SFF > \varepsilon_{SFF}$ ($\varepsilon_{SFF} > 0$, user-defined) as *scan regions*

Output: Scan regions for further analysis (Algorithm 18)

To determine the patches where significant amounts of the solid phase are present, the SFF is computed for every patch using the SZ mask patches. The patches with the SFF greater than a user-defined threshold ε_{SFF} are designated as *scan regions* and are passed to Algorithm 18 for further analysis. The memoized image position indices for the patches will be used later for the FOV reconstruction.

Algorithm 18: Scan region filtering & dendrite skeleton extraction

Input: scan region images (Algorithm 17)

Prepare filter input

- 1 Re-scale the images and perform color tone mapping (CTM)
- 2 Re-scale the images again and perform FFC

Filter the images

- 3 Apply BM3D filtering
- 4 Perform two iterations of non-local means masking (NMM)
- 5 Apply soft color tone map masking (SCTMM)
- 6 Apply FFC

Extract dendrite skeletons

- 7 Segment dendrites using 2-threshold hysteresis binarization
- 8 Invert the mask
- 9 Apply the thinning transform
- 10 Perform size thresholding
- 11 Multiply the resulting masks with their corresponding SZ mask patches (Algorithm 17, Step 2)
- 12 Perform morphological pruning (optionally in multiple passes)
- 13 Remove border pixels
- 14 Perform size thresholding

Output: Dendrite skeleton masks

Here, the strategy is to use BM3D to denoise the dendrite textures as non-destructively as possible. However, it was found that in general it can be difficult to obtain good results without preparing the images first. This is why the first stage is image normalization, color tone mapping (CTM) [167] and FFC. CTM has the effect of dramatically increasing the CNR of the dendrites. Meanwhile, FFC uses a coarse polynomial fit of the image luminance map to perform background correction (flattening) without reference [230]. In some cases, this helps to reduce the large-wavelength correlated noise due to liquid metal flow across the dendrites.

Next, the scan regions are filtered. One first applies BM3D to restore the dendrite textures in images, then two iterations of non-local means masking (NMM) are applied to mitigate any leftover correlated noise and increase dendrite CNR. NMM could be viewed as a generalized, locally adaptive version of unsharp masking and was previously used for particle detection in neutron radiography images of particle-laden liquid metal flow for a similar purpose, with good results [185, 213]. Afterward, SCTMM (1) is applied for further background reduction and CNR enhancement. This is the same method previously used for bubble and particle detection in neutron images in [158, 185]. Finally, another FFC iteration is performed. The parameters for these stages are given in [223].

Dendrite skeleton extraction is done in eight steps. One starts with double-Otsu hysteresis binarization [157], followed by mask inversion, morphological thinning and size thresholding. Then the resulting skeleton masks are multiplied with their respective SZ mask patches, which crops the skeleton parts that are within the liquid phase areas and thus cannot actually be dendrites. Afterward, the remaining skeleton asperities are removed with multi-pass morphological pruning [157], then border pixels are removed and size thresholding is done again.

After dendrite detection, the skeleton segment orientations (φ , with respect to the image X axis) must be measured. However, it may be the case that the image filters do not properly resolve primary/secondary dendrites or cases with dendrite overlaps. In addition, some segments may not have clearly defined orientations, or could simply be leftover artifacts. To identify and correct such *unresolved* (in the sense of orientation) structures within all IWs, a procedure outlined in Algorithm 19 is performed for all IWs. Algorithm 19 uses the aspect ratio χ as a criterion to determine if the skeleton segments have a resolved (i.e. well-defined) orientation. Since morphological thinning is one of the steps in Algorithm 18 and Algorithm 19 does not add new pixels to masks, most of the skeleton lines should have a 1 px thickness and, provided they are long enough, therefore also high χ . Colorizing the segments by φ will play a key role later during dendrite grain decomposition (Section II G 2 e), but in Step 3 of algorithm 19 it is

used to separate the masks with initially resolved and unresolved skeletons. φ and χ determination is explained in [223].

Algorithm 19: Resolve unoriented skeletons

Input: IWs with dendrite skeletons (Algorithm 18)

Identify unresolved skeletons

- 1 Compute orientation angles (φ , with respect to the image X axis) and aspect ratio (χ) for dendrite segments
- 2 Colorize resolved (oriented) segments ($\chi > \chi_c$, $\chi_c \geq 1$ is user-defined) by their φ
- 3 Separate resolved and unresolved segments into different masks

Resolve the skeletons (for masks with unresolved segments)

- 4 Detect skeleton corner points
- 5 Detect skeleton branch points
- 6 Threshold and filter corner and branch points
- 7 Remove the remaining corner and branch points from skeletons

Reassemble IW skeletons

- 8 Re-evaluate segment orientations for masks with (formerly) unresolved segments (Step 7)
- 9 Colorize resolved (oriented) segments ($\chi > \chi_c$) by their φ
- 10 Add the output of Step 9 to the previously set aside masks with initially resolved segments (Step 3)

Output: IWs with resolved dendrite skeletons colorized by their orientations

Once the masks with unresolved skeletons are separated, corner point detection is performed. The Harris-Stevens method [231, 232] with first-order Gaussian derivatives is used. Then, morphological branch points are detected [233]. Afterward, the detected corner points are filtered by selecting corner pixel clusters with pixel count below a threshold, then combined with the detected branch points, and lastly pixel clusters with pixel counts below a threshold are removed from the resulting mask. Finally, the size-thresholded combined mask is subtracted from the input mask with unresolved skeletons. The parameters for these operations are specified in [223]. With this done, the resolved skeletons are now assigned colors based on φ subject to the χ_c criterion, and the resulting masks are recombined with those containing initially resolved skeletons.

At this point, all the remaining unresolved (white-colored) skeletons within IWs are considered unoriented dendrites and/or artifacts. These are excluded from any subsequent analysis. One can now reassemble the global (FOV) dendrite skeleton image by tiling the IWs according to their position indices from Step 3 of Algorithm 17. It is also now easy to generate maps with color-coded dendrite orientations with highlighted cavities and

liquid/solid boundaries.

e Dendrite grain decomposition

Before proceeding with decomposition of the resulting global dendrite skeleton into grains, a *global* orientation (φ) spectrum must be computed for the assembled skeleton. It is not only of physical interest, but will be used in Algorithm 20 as well. While a global φ spectrum is certainly relevant, it is often desirable to distinguish areas of "coherent" dendrite growth, i.e. *dendrite grains* with their areas and mean dendrite φ . The presented *dendrite grain decomposition* (DGD) method does this by considering φ similarity and proximity of the dendrites detected within the FOV, and it does so by exploiting the color-space representation of φ generated by Algorithm 19. DGD is performed in three stages: a primary scan which detects dominant dendrite grains; a refined scan which checks if the larger grains should be subdivided further and if the smaller grains are eligible; a filtering step which resolves ambiguities and overlaps between the detected dendrite grains. The first step of the DGD procedure is outlined in Algorithm 20.

Algorithm 20: Detecting dominant dendrite grains

Input: Assembled global dendrite skeleton

- 1 Compute the global φ spectrum from the global skeleton
- 2 Filter the φ spectrum & detect dominant peaks
- 3 Find dendrite segments near the φ peaks in the color-space
- 4 Build grain masks that cover the dendrite segments

Output: Separate masks for dominant dendrite grains

The φ spectrum is weighed by dendrite lengths, and computation is explained in [223]. Since there is a mapping between $\varphi \in (-\pi/2; \pi/2]$ and the dendrite skeleton color values (normalized) due to Algorithm 19, one can now find the segments in the global skeleton that correspond to the selected φ peaks. To do this, peak φ values are converted to coordinates r_φ in the *CIELAB* (*CIE76*) color-space and the segments with Euclidean distance within δr_{LAB} from r_φ are selected. The selected segments are further filtered by assigning them weights $w_{\text{LAB}} \in [0; 1]$ based on their distance from r_φ (farthest to closest) and keeping segments with $w_{\text{LAB}} > w_c$. This makes the process more resilient to noise in the φ spectrum and helps to avoid grain mask overlaps later. Sometimes, however, there may be groups of two or more very close peaks r_φ that survive the thresholding by $\rho > \rho_c$. In these cases, such groups of peaks are replaced with mutual Euclidean *CIELAB* distances $< \delta r_\varphi$ by their mean r_φ values. Another issue that might

come up is that $\varphi \in (-\pi/2; \pi/2]$ and the φ spectrum does not have periodic boundary conditions. This means that if there were an *actual* peak of some width for dendrite orientations near $\varphi = \pm\pi/2$, it would be treated as two φ peaks by Algorithm 20. This problem can be solved by checking if the peaks closest to the $(-\pi/2; \pi/2]$ boundaries (*edge peaks*) are close enough to these boundaries and to one another *across* the $\varphi = \pm\pi/2$ boundary. In cases where two or more peaks are detected, the edge peaks φ_1 and φ_2 are subjected to constraints $\min(\pi/2 - |\varphi_{1,2}|) \leq \delta_1$ and $||\varphi_1| - |\varphi_2|| \leq \delta_2$. If both are satisfied, the respective masks with dendrites within the color-space peak ranges are added before color-space proximity thresholding, effectively treating the edge peaks as one. Parameters are provided in [223].

When the dendrite clusters corresponding to each r_φ peak are found, a mask must be created for them that will delimit and separate them as one grain. This is done by applying the closing transform with disk structuring elements to the skeleton clusters, which fills the spaces between the dendrite skeletons while not affecting the outlying parts of the skeletons, i.e. preserving the shapes of the dendrites that protrude from the bulk of the cluster. Note that this may generate more than one grain mask per r_φ peak, since dendrite clusters with very similar orientations may be sufficiently far apart. Thus, grains are identified, accounting for both dendrite orientations and spatial distribution. Afterward, the resulting grain-covering binary masks are thresholded by their area, and the remaining masks are separated for further analysis with a refined scan.

Once dominant grains are identified for each r_φ peak, they are subjected to a secondary scan that is designed to check whether the originally recognized grains need to be further subdivided. This is done to both (implicitly) ensure grain uniqueness, minimize overlaps, and resolve smaller areas within the larger grains that have distinct enough orientations. The scan follows steps similar to those of Algorithm 20, but with the following modifications: steps 1 and 2 are now applied to the dendrites within the grain masks, not the global skeleton; area-adaptive φ spectrum resolution is used.

Prior to φ spectra calculation, the dendrite skeletons belonging to the grains are isolated by multiplying the grain masks by the global dendrite skeleton. The area-adaptive resolution is set up such that, on the one hand, the algorithm can resolve finer differences in orientations within the initial grains and detect the underlying φ peaks, while on the other hand not using exceedingly large resolution for smaller grains with relatively few dendrites. In the latter case, the algorithm would otherwise treat the noise in the spectrum as significant peaks despite the filtering. The adaptive spectrum bin count is calculated as shown in [223]. The other relevant parameters are as in Algorithm 20.

When the initial grains have been scanned again and kept or decomposed further and/or eliminated, the final DGD step is performed – ambiguities and overlaps between the grains are resolved. Two cases must be treated here: the previous stages of DGD have, in separate instances, generated two grains with almost identical (i.e. overlapping) masks and this is indeed one and the same grain; there is partial overlap between the grains, but it is physical since the grains are adjacent and the dendrite orientation changes very slowly from one grain to the other, i.e. there exists a transition zone instead of a sharp boundary. In the former case one of the masks is redundant, and in the latter the overlap zone must be identified and designated as such, since no clear distinction between the two grains can be made in the transition zone. Finally, there is also the matter of potentially leftover dendrites with φ values that are significant outliers with respect to the mean φ for the grains. These issues are addressed by Algorithm 21.

Algorithm 21: Resolving ambiguities & performing grain cleanup

Input: Dendrite grains output by the refined scan

- 1 Compute the overlap masks for dendrite grain pairs
- 2 Compute the *uniqueness factors* $u \in [0; 1]$ for the grain pairs
- 3 Discard redundant grains ($u < u_c$, u_c is user-defined)
- 4 Designate the overlap masks for the pairs of remaining unique grains as *overlap zones* and subtract them from the grain masks
- 5 Perform total area thresholding for the resulting masks, then secondary size thresholding for the underlying segments – the output gives the final grain masks
- 6 Multiply the global skeleton by the final grain masks to isolate the respective dendrites
- 7 Remove dendrites with outlying φ values and perform dendrite size thresholding

Output: final dendrite grain masks & skeletons

The first step is performed by multiplying all possible pairs of dendrite grain masks. Then the uniqueness factors defined as $u = 1 - S_{\cap}/\langle S \rangle$ are computed for all dendrite grain pairs, where S_{\cap} is the overlap mask area and $\langle S \rangle$ is the mean grain area for the pair. Pairs with $u < u_c$ are considered redundant, and only one grain mask from such pairs is kept. The overlap masks for all the other grains are kept as overlap zones and subtracted from the unique masks. The remaining segments are then thresholded both by total and individual areas. Afterward, the resulting masks are multiplied by the local skeleton to isolate the grain dendrites. Finally, the dendrite skeletons are filtered by orientation and length: $\langle \varphi \rangle$ is measured for the grain dendrites and segments outside the $\langle \varphi \rangle \pm 5\sigma$ interval (by default) are eliminated, followed by length thresholding. This concludes the DGD

process. Parameters are specified in [223]

f Liquid domain analysis

Measuring the solute concentration above the SF involves two main considerations. First, filtering the noise in the liquid region above the SF as non-destructively as possible, i.e. not to alter the luminance field too much, as it can later be used with the Beer-Lambert law to assess the concentration of the solute. Second, LZ/SZ segmentation will never be perfect, and dendrite tips could be slightly above the SF, yielding errors in the luminance/concentration measurements – this must be mitigated. Both concerns are addressed by Algorithm 22.

Algorithm 22: Measuring solute concentration above the solidification front

Input:

- Pre-processed FOV images without artifacts (Algorithm 14)
- Solidification front masks (Algorithm 16)

- 1 Apply median filtering to the FOV images
- 2 Apply the bilateral filter
- 3 Perform NM filtering
- 4 Define the SF-conformal buffer zone by shifting the SF contour mask upwards
- 5 Define the concentration sampling zone by extruding the SF-shaped boundary upwards from the buffer zone upper boundary
- 6 Compute the mean luminance for vertical pixel bands within the sampling zone over the FOV width
- 7 (Optional) Use the Beer-Lambert law to convert the luminance to the solute concentration

Output: Mean luminance/concentration above the SF over the FOV width for all time stamps

It is found that the combination of median, bilateral [234] and NM (in this order) filters in Steps 1-3, after some parameter tuning, yields a sufficiently well-filtered luminance/concentration field without offsetting the values too much or significantly affecting the larger-wavelength features. The bilateral filter is chosen in particular because of its luminance value range filter component, since with the right settings it should preserve luminance level sets within the images that are fairly close. Once the images are processed, sampling zones are defined for every frame. To address the above-mentioned issue with dendrite tips possibly being above the SF, a buffer zone is created where no sampling occurs – this is done by shifting the SF mask (curve) upwards by a distance d_{buf} . Starting from

this level, the SF curve is then extruded over a distance d_{samp} to create the SF-conformal sampling zone for the concentration measurements from the filtered images. Here, the mean values are computed for 1 px -wide vertical strips of d_{samp} pixels above the buffer zone over the width of the FOV. Parameters are provided in [223].

With convective plume segmentation, the luminance value preservation is not as serious a concern as long as the shapes are preserved. Here the idea is to filter the images and then decompose the resulting filtered LZ luminance map into N_{level} level sets (clusters) ranked by luminance. One can then, depending on the case, re-assemble N_{plume} level sets with the highest luminance back together to obtain the masks for the convective plumes above the SF. This is done via Algorithm 23.

Algorithm 23: Convective plume segmentation

Input:

- Pre-processed FOV images without artifacts (Algorithm 14)
- LZ masks with filled channels and cavities (Algorithm 16)

- 1 Perform Steps 1-3 from Algorithm 22 (with different parameters)
- 2 Gaussian filtering
- 3 NM filtering
- 4 Create a buffer zone above the LZ boundary (SF) in the LZ mask with filled cavities and channels using morphological dilation
- 5 Multiply the filtered image by the resulting mask
- 6 Decompose the resulting image into level sets
- 7 Re-assemble selected level sets to get convective plume masks

Output: Mean luminance/concentration above the SF over the FOV width for all time stamps

To isolate the liquid region within the FOV, the image is multiplied by the LZ mask. However, before multiplication, a buffer zone is created above the LZ mask by morphological dilation using disk structuring elements with a radius of 15 px . The rationale here is the same as in the case with the concentration measurements above the SF – to eliminate the dendrite tips that may have been imperfectly segmented. If not removed, these would form "parasitic" level sets and reduce the actual level set count within the LZ. The luminance (concentration) level sets for the convective plumes are obtained using the K-medoids method [212] with N_{level} medoids – the level sets are then ranked by their mean luminance (with the excluded solid and buffer region having the lowest, zero luminance after mask multiplication) and the top N_{plume} are re-assembled into the output mask for the convective plumes. Plume segmentation parameters are given in [223].

3. Main results

To demonstrate the resilience of the proposed solution, it was decided to test the developed methodology and code for conditions that are purposefully made worse than one would expect. Specifically, it was decided to use only a *single* frame for FFC during pre-processing for an image sequence. This way the background features, i.e. setup elements caught within the FOV, artifacts, and the X-ray beam profile are still compensated for, but the resulting SNR is significantly lower.

a Image filtering & segmentation

First, consider the stages of FOV processing prior to SZ and LZ segmentation. Figure 133 shows a characteristic example of an input image versus the pre-processed and then filtered output, and is a good example where many of the image features that can realistically be expected are present. Artifacts appear as overexposed corners and areas at the bottom image boundary in (a). A thermocouple is located in the upper-left corner of (a), visibly protruding from the left image boundary). Note also the tape for affixing the thermocouples, visible as rectangular areas with greater opacity at the left and upper boundaries of (a). Liquid channels and cavities are present in the lower part of the FOV, better visible in (b). One can observe that even with single frame FFC, the luminance distribution in (a) about the X-ray beam axis is almost entirely negated in (b), and so are the elements attached to the imaged liquid metal cell. However, as opposed to (a), noise is amplified in (b) and becomes visibly coarser-grained. Afterward, the artifacts are removed, and the result is shown in (c) – the formed artifact areas should no longer affect the NM and BM3D filter patch matching procedures. Finally, BM3D and NM are applied, in that order, resulting in an output seen in (d) with much more contrast LZ/SZ boundaries, as well as cleaner liquid metal cavities and channels. While the liquid metal plume in the upper part of (b) has been dramatically diffused, it is of no concern here, as the result of this image filtering routine is used only for LZ/SZ separation (filters in Algorithm 22 are used for concentration measurements in the LZ instead).

The results of segmentation (Algorithm 15, using (39)) after FOV filtering are presented in Figure 134. The less transparent (greener) structures seen in (b) in the middle of the FOV were formed under externally induced forced convection, and the upper, more transparent and less contrasting part of the dendrite structures developed after abruptly switching forced convection off and leading back to the regime of natural convection (more details in [59]). Here it is demonstrated that the segmentation algorithm

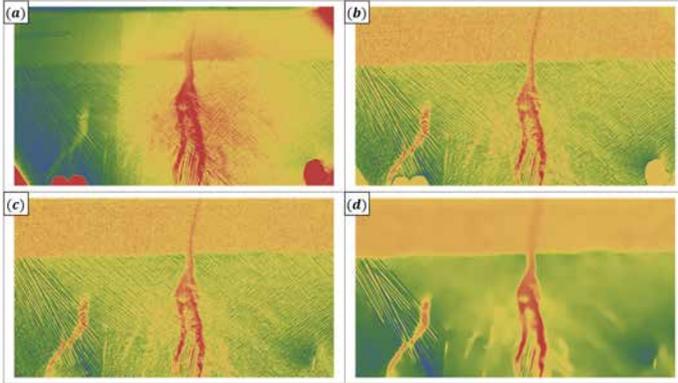


FIG. 133. FOV image processing: (a) raw image, (b) pre-processed image, (c) image with artifacts (note the image corners in (a)) inpainted via Algorithm 14 and (d) final image after BM3D and NM filtering (Section II G 2 c). The color scheme here and further, unless stated otherwise, is identical to Figure 132. Note the enhanced contrast of the SF, liquid channels and cavities.

has no issues detecting the earlier formations in (a) and (b) where forced convection is observed, is not hindered by the rapid onset of convective plumes in (c), and also captures the much less contrast structures seen in the upper part of (d) where the SFF is close to 1. More examples of filtering and segmentation results, as well as a more in-depth analysis, are provided in [223].

After SZ/LZ segmentation is complete, these segments are further differentiated to separate liquid cavities and channels from the SZ, and to determine the shape of the SF. There are also safeguards against artifacts that can potentially be left over after LZ/SZ segmentation, as it is hardly possible to *always* find optimal segmentation parameters immediately. Therefore, it is good if the code has backup options. Figure 135 illustrates the steps involved in separating segment classes.

In both cases, once cavities are identified (Algorithm 16), one readily obtains the *bulk* LZ segment, and can then derive the SF and find the channels extending below the SF. Examples of how the SF is traced along the SZ segment boundaries is shown in Figure 136. The resulting SF edge masks will be used later for concentration measurements above the SF and for convective plume segmentation in the LZ.

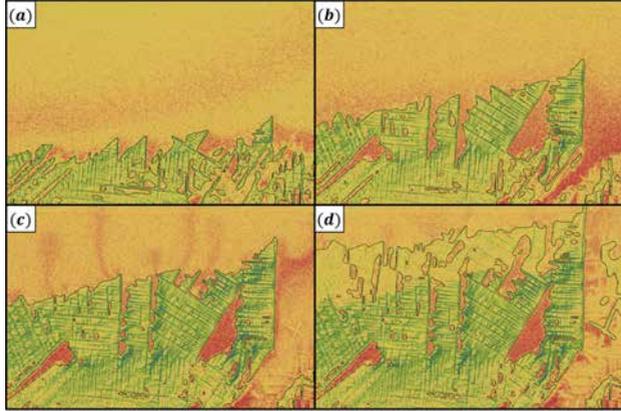


FIG. 134. An example demonstrating the performance of the developed SZ/LZ segmentation method. Boundaries between segmented SZ/LZ (black contours), accounting for areas with removed persistent artifacts, for different time stamps (a-d) in ascending order obtained by segmentation using Algorithm 15 and 39.

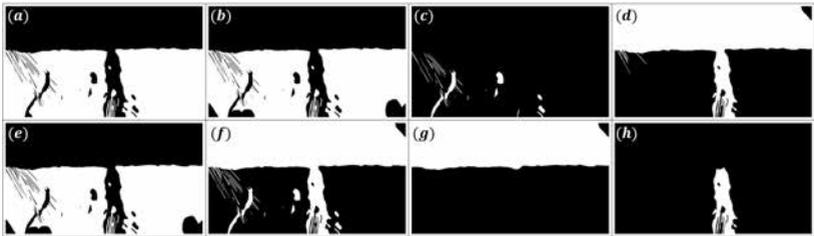


FIG. 135. Stages of segment classification for separated SZ/LZ: (a) segmented SZ, (b) mask (a) with non-border artifacts removed (case none present here), (c) liquid cavity mask, (d) LZ mask without cavities and border or persistent artifacts, (e) mask (b) but without border artifacts (none were present), (f) combined masks (c) and (d), (g) masks separated by the SF, and (h) liquid channel mask. This case corresponds to the image sequence considered in Figure 133.

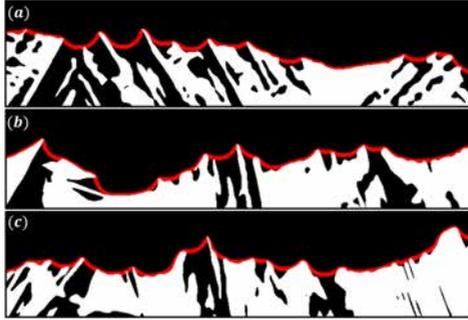


FIG. 136. Examples showing how the SF (red curves) is traced based on the final SZ mask, i.e. how Figure 135g results from Figure 135e/f.

b Analysis of solidified structures

With the SZ identified along with cavities and channels, one can now turn to the analysis of the solidified structures within the SZ. As outlined in illustrated in [223], first the FOV image is partitioned into IWs, and then each partition is processed using Algorithm 18. An example of this is provided in Figure 137. Notice that the noise makes the identification of dendrites in the initial images (a) quite difficult, and Figure 137a additionally exhibits very low CNR for dendrites, mainly due to larger-wavelength correlated noise stemming from liquid flow. Note also that the dendrites have lower X-ray transparency than the surrounding liquid, and therefore it is the liquid that is colored white in (a).

Sub-figures (b) show that CTM and reference-less FFC make the dendrites and the spacings in between much clearer, but the SNR and CNR are unchanged. However, it was found that this stage dramatically boosts the performance of BM3D, which is the next stage, the output of which is seen in (c). While the structures in Figure 137c are already more clearly discernible to the human eye, the CNR is still lower than desired for reliable segmentation. This is because BM3D does not mitigate the correlated noise (uncorrelated Gaussian noise model is used for BM3D). This is where the next stage comes in, with two iterations of NMM correction and SCTMM, after which reference-less FFC is applied again. Here NMM takes care of much of the correlated noise, SCTMM boosts CNR, and FFC acts as post-NMM large-wavelength background cleanup.

Afterward, double-Otsu hysteresis segmentation and image inversion yields the dendrite mask (e), and morphological thinning with size thresholding generates the IW dendrite skeleton (f). Now all that remains is to crop the dendrite skeletons, multiplying their masks by the SZ mask

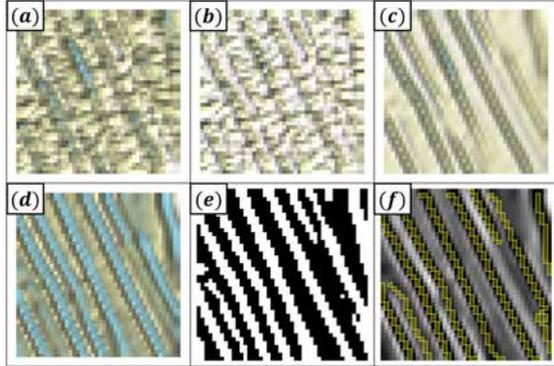


FIG. 137. IW processing via Algorithm 18: (a) a colorized relief plot of an IW (projection of a pre-processed image), (b) IW after re-scaling, CTM and FFC, (c) BM3D output, (d) results after 2 iterations of NMM, one SCTMM iteration and FFC, (e) output after 2-Otsu (hysteresis) binarization and image inversion, and (f) final result after size thresholding and thinning, overlaid on top of the grayscale version of (d). Image luminance increases from color light blue to white.

projected onto IWs – this process is demonstrated in Figure 138. After cropping is done for every IW, one has solid structure skeletons for the entire FOV – an example of IW skeletons output by Algorithm 18 for the case considered in Figures 133 and 135 is shown in Figure 139.

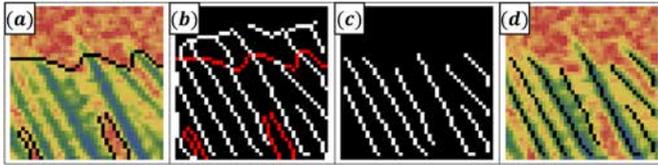


FIG. 138. An example of IW skeleton cropping using the LZ mask: (a) colorized IW image with a SZ boundary (black contours), (b) IW skeleton with the SZ boundary overlay (red contours), (c) cropped IW skeletons and (d) cropped IW skeletons overlaid on top of the colorized IW image (black lines).

Observe, however, that many of the IWs in Figure 139 exhibit multiple dendrite crossings and branching, making the orientation analysis problematic. To mitigate this issue, Algorithm 19 is applied to each of the IWs (steps visualized in [223]). Although the results have certain imperfections and a small amount of information may be lost in IW skeletons in general, it was found that the set of parameters that were selected for Algorithm

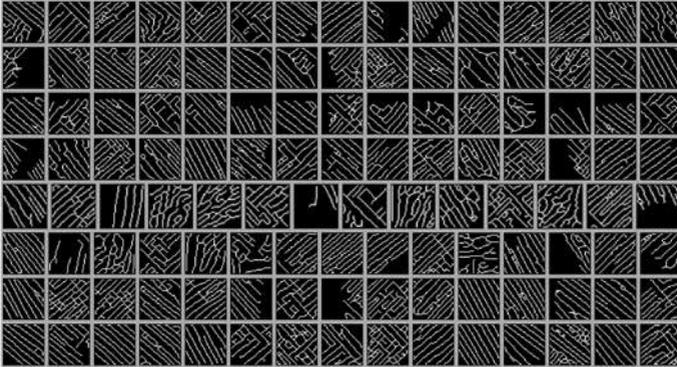


FIG. 139. An example of detected IW skeletons for an image after cropping to the SZ mask, e.g. Figure 135(e). This case corresponds to Figures 133 and 135.

19 allows recovering much more information that would otherwise be lost. This is very clearly illustrated in Figures 140 and 141 where one can see the significant difference between before and after Algorithm 19 is applied.

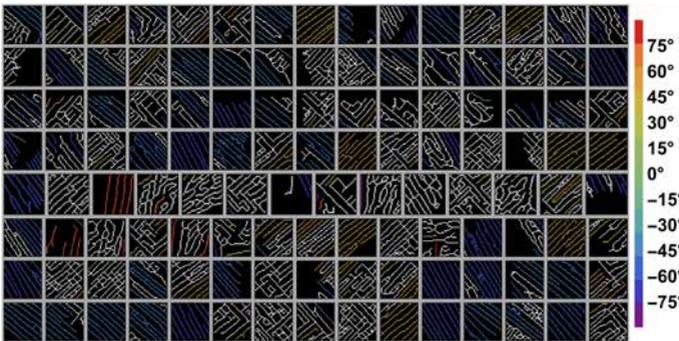


FIG. 140. IW skeletons shown in Figure 139 colored by their orientations with respect to the horizontal image axis. The white-colored segments are unoriented according to the criteria established in Algorithm 19.

As noticeable in Figure 141, some of the new segments still do not exhibit clear orientations or are otherwise ambiguous, but most of the information otherwise inaccessible from Figure 140 has been recovered with minimal losses. With this, one can now re-assemble the processed IWs into a global skeleton, which yields results showcased in Figures 142 (corresponding to Figure 141) and 143.

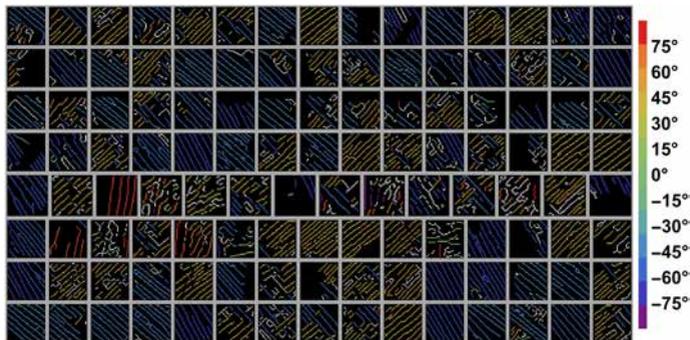


FIG. 141. The same IW skeletons as in Figure 140 after Algorithm 19 was applied to every IW. The color convention matches Figure 140.

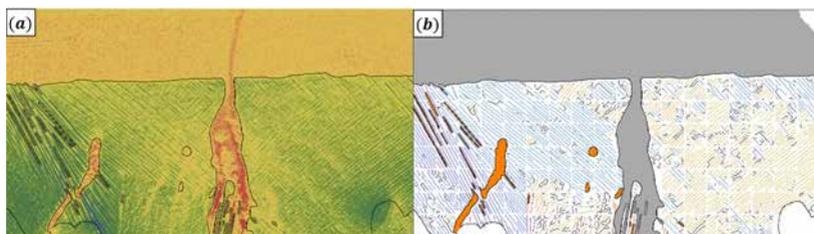


FIG. 142. SZ analysis results for an image from the case considered in 141: (a) pre-processed image with overlaid SZ/LZ boundaries (black contours, image artifact areas included) and (b) the same SZ/SZ boundaries with dendrite skeletons (colored by orientation otherwise black if unoriented) and liquid metal cavity areas (orange) within. The gray area in (b) is the bulk liquid above the SF plus the liquid metal channel connected to it. The orientation color map is as in Figure 141.

While Figures 142 and 143 are already quite informative, there is more to be extracted from the assembled skeletons. One can measure the orientation (φ) spectrum, computing the relative orientation frequency by weighing over dendrite segment lengths to account for gaps due to IW boundaries and dendrite interruptions due to other reasons. An example of this is shown in Figure 144b where the φ spectrum is computed for the image seen in Figure 143. Note that the unoriented dendrite skeletons (colored black in Figure 143b) do not count towards the spectrum. One can also compute the median φ for every IW to have a coarser but simplified overview of how dendrite φ are distributed over the FOV – this is seen in Figure 144a.

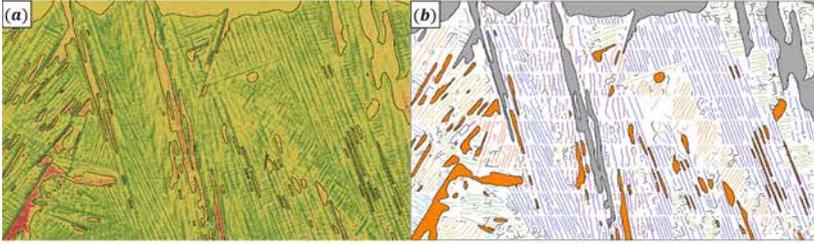


FIG. 143. Another example of the SZ analysis.

In addition, by computing φ spectra for an entire image sequence, one can observe the dynamics over time. Statistics for liquid cavity (orange areas in Figure 143b) areas, aspect ratio, orientations, etc. can also be computed per frame and their dynamics visualized. The same is true for channels extending into the SZ (e.g., Figure 142b).

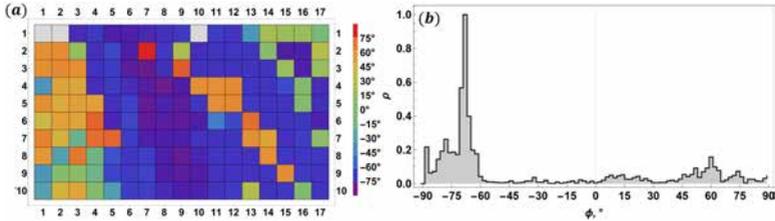


FIG. 144. (a) Median dendrite orientation φ for the IWs (the case in Figure 143) and (b) orientation spectrum for the assembled dendrite skeleton, excluding unoriented segments.

Importantly, the φ spectrum shown in 144b is later used for the DGD procedure (Algorithm 20), which is used to decompose the global dendrite skeleton seen in Figure 143 into grains. An example of this process is shown in Figure 145. Figure 145a is Figure 143b stripped of all background with only dendrites remaining – here the most intense φ spectrum peak shown in Figure 144b is considered. In (b), one can see the segments with φ within an interval of the selected φ peak (Algorithm 20) and (c) shows the result of thresholding by color-space distance. Afterward, the grain masks are constructed using the closing transform, then size-thresholded, which results in grain masks as seen in (d). These operations are performed for every peak in the global φ spectrum that survives filtering and thresholding (Algorithm 20). More examples of this are shown in [223].

When the refined scans and grain cleanup are done, one can assemble

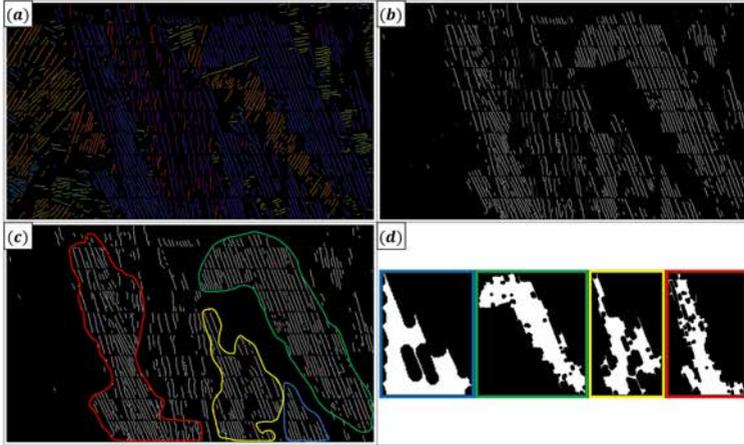


FIG. 145. The DGD process (Algorithm 20): (a) global dendrite skeleton with unoriented segments removed, colored by orientation as in Figure 143b, (b) dendrite skeleton segments detected about the main orientation spectrum peak seen in Figure 144b, with grayscale color map encoding the distance to the peak in color-space for each segment, (c) segments from (b) after color-space distance thresholding, with rough outlines for the four detected grains, (d) masks for the detected dendrite grains corresponding to the outlined in (c).

the resulting grains (both their masks and dendrite segments) within the FOV, examples of which are shown in Figures 146 and 147. It becomes visible in Figure 146, and especially (b) that the detected grains indeed constitute the major dendrite clusters with coherent (sufficiently similar) φ sets. Note also that in (a) the grains separated by the liquid cavities (e.g. upper-left corner and the lower-right part of the FOV) are correctly separated, even the ones with sufficiently similar φ . This is because liquid cavities are accounted for in the grain cleanup process. Observe that the detected dendrite skeletons match the landscape seen in the background of (b).

The same holds for Figure 147, where in (a) one can see a black-colored cluster of skeletons designated as undetermined, i.e. they do not belong to either of the two adjacent grains with certainty and are a transition zone. This is because both grains are very close to this area, and the φ of dendrites within this zone is somewhere between the mean φ for the two grains in question. It might seem that the lower part of one of the grains in (a), highlighted with light blue, should be treated as a separate smaller grain because the long diagonal cavities seem to split it in two. However,

one can clearly see in (b) that they are actually connected via one of the dendrites that bridges a narrow gap between the two liquid metal cavities. A similar situation holds for the grain highlighted with light red in Figure 146a. If one does observe some small clusters of grains with similar φ in Figures 146b and 147b that were discarded by the DGD process, then these grains are most likely below the size threshold. One can, of course, always adjust settings accordingly if capturing smaller grains is necessary.

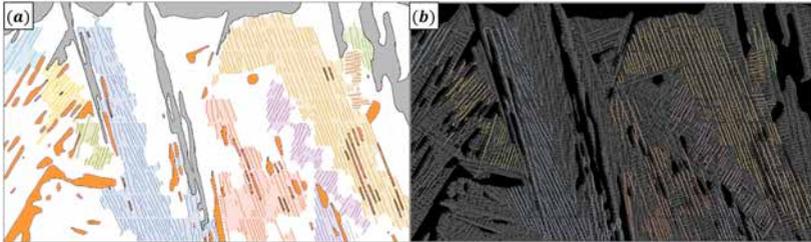


FIG. 146. Results of DGD for the case corresponding to Figure 145: (a) dendrite segments colorized by dendrite grain IDs with grain mask overlays, with orange areas representing liquid metal cavities and bulk liquid with channels highlighted as gray areas, and (b) grain dendrite skeletons overlays. The background in (b) is a post-processed raw image followed by image inversion, reference-less FFC, CTM, re-scaling and sharpening with the LZ masked.

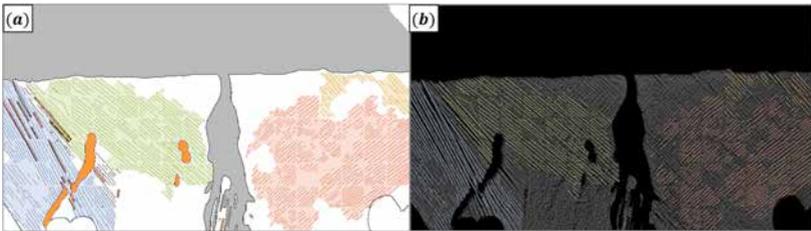


FIG. 147. Results of DGD for the case considered in Figure 142. Note the transition zone (black-colored dendrite segments) in (a) between the two grains to the left.

Once the dominant grains are detected, their relative areas (with respect to the SZ area) and φ statistics can be determined, which is demonstrated in Figure 148 for the case seen in Figure 147. Again, note that this can be done for all or selected frames in an image sequence to observe the dynamics of grain formation, fragmentation and how their statistics change.

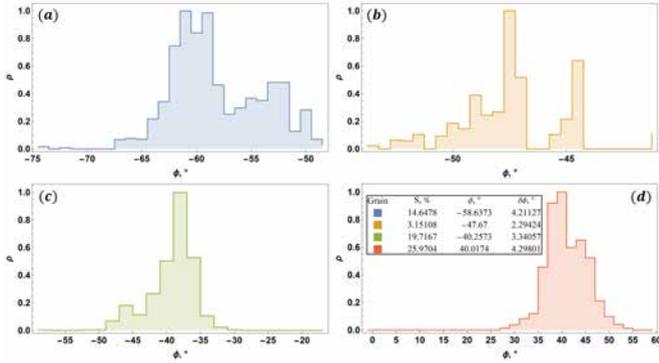


FIG. 148. Dendrite orientation φ spectra for the grains identified via the DGD process, as shown in Figure 147. Colors for plots (a-b) correspond to the grain colors in Figure 147. The legend in (d) shows the grain area fraction S (with respect to the SZ mask area), mean dendrite orientation φ and its standard deviation $\delta\varphi$.

c Solidification front dynamics

In addition to the above, having derived the SF, one can now look at its dynamics. Figure 149 shows how the SF height over the cell width changes over time for one of the cases found in [223], as the SF advances vertically upwards and the SFF increases. This also contains the shape information to some degree. If the latter is not of interest, or if a more continuous dynamics visualization is needed, one can generate an image as in Figure 150 where the color encodes the front height and variations over time and cell width can be clearly seen for an entire image sequence.

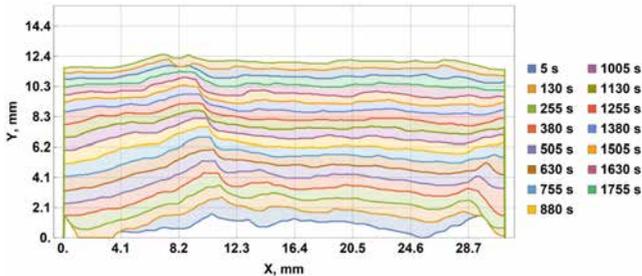


FIG. 149. SF height Y over cell width X at different time stamps (legend to the right).

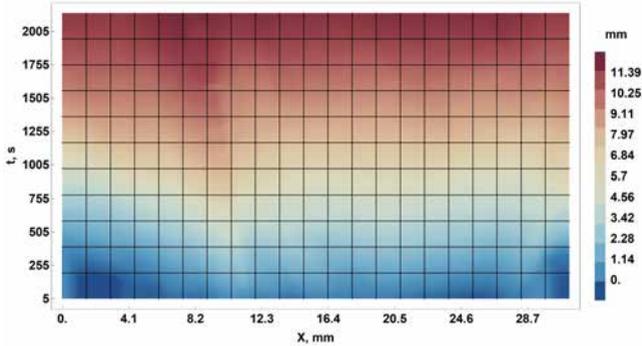


FIG. 150. SF height (color-coded, with color legend to the right) over cell width X and time t for the case shown in Figure 149.

Figure 150 is obtained by median-filtering the front height matrix. The resulting matrix can then be used to calculate the matrix of instantaneous SF propagation velocity, which is shown in Figure 151. In the case of velocity, outlier removal and bilateral filtering are performed (parameters stated in [223]). If mean dynamics are of interest, they can be readily derived from the above matrices. The corresponding results can be seen in Figure 152.

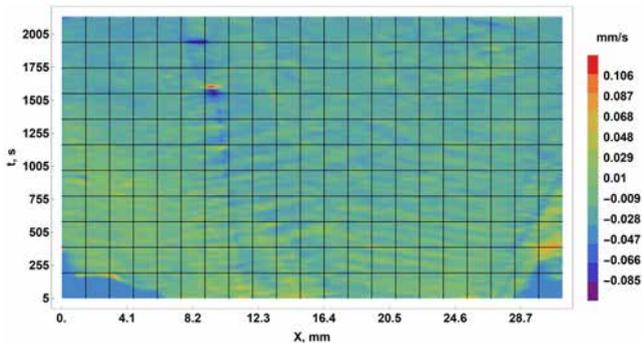


FIG. 151. SF vertical velocity (color-coded, with color legend to the right) over cell width X and time t .

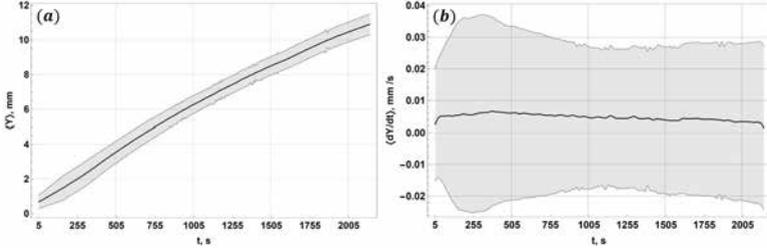


FIG. 152. Mean SF dynamics over cell width: (a) height $\langle Y \rangle$ and (b) vertical velocity $\langle dY/dt \rangle$ over time t .

Solute concentration can be measured above the SF as explained in Algorithm 22 and plotted for different locations along the cell width over time, which is showcased in Figure 153.

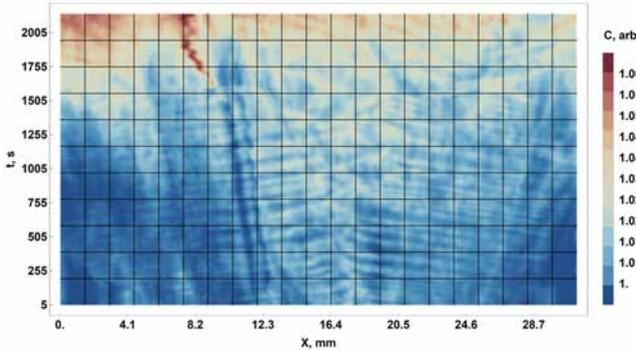


FIG. 153. Solute concentration dynamics above the SF and along the cell width X over time t . Concentration is expressed via image luminance (encoding Ga/In concentration via the Beer-Lambert law) in relative units (with respect to the initial frame after Algorithm 22 is applied) and is color-coded as shown in the color bar to the right.

Note the maximum spot within $X \in (8.2; 12.3) \text{ mm}$ and $t \in (1505; 1755) \text{ s}$, which corresponds to a rapid channel opening at the SF. Figure 22 reveals the sudden appearance of a highly concentrated "trace" in the same region as the velocity minimum/maximum in Figure 151. Once a liquid metal cavity breaches the SF and a channel forms, much greater X-ray transmission is consistently measured, as it should be, implying increased gallium concentration. The trail shift to the left is simply due to the change in the location of the channel outlet at the SF, which is caused by remelting.

Observe also the banded structure for $t \lesssim 1500$ – these are not artifacts due to data processing, but rather physical concentration oscillations above the SF, as well as the result of SF fluctuations.

d Convective plume segmentation

Finally, one can analyze what occurs above the SF in the bulk of the LZ further by examining the convection plumes using Algorithm 23. The results of its applications to example frames representing different cases are shown in Figures 154-155. As with the concentration measurements just above the SF, the buffer zone mask is intended to prevent the interference from potentially sticking out dendrite tips that in general may occur above the derived SF for low SNR/CNR images. In addition, the artifact areas must be excluded, since any information contained there is otherwise meaningless. This makes obvious that the code successfully segments the plumes in the showcased examples and largely preserves their shapes despite the clearly visible large-grained noise.

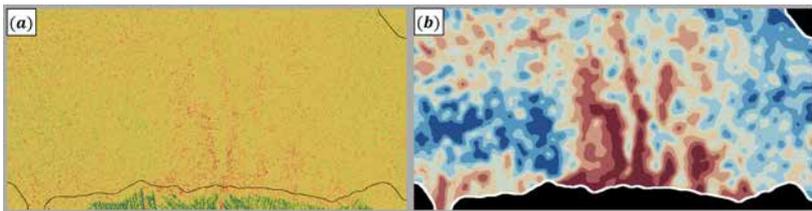


FIG. 154. Analyzing gallium-rich convective plumes in the LZ above the SF (Algorithm 23): (a) pre-processed image with the outlines of the buffer zone (black contours) extended from the SF and the artifacts (e.g., lower-left and upper-right corners), and (b) convective plumes segmented and highlighted outside the buffer zone (black areas with white boundaries). The level sets representing the plumes are colored by image luminance (representative of the solute concentration). The color scheme is as in Figure 153. This case discussed in more detail in [223]. The current frame is close to the beginning of the image sequence, where the plumes are the most intense.

This multiple level set representation, in conjunction with morphological analysis shown in Figure 135, as well as the SF dynamics and SZ analysis, should provide a wealth of details, enabling in-depth analysis and physical interpretation of solidification processes studies using experimental setups that are similar in principle to those considered here. It is worth noting that what can be seen in the above sub-figures (a) a purposely lowered SNR. Given the results seen in (b), higher SNR images should enable even better results. This, of course, also holds for the SZ analysis.

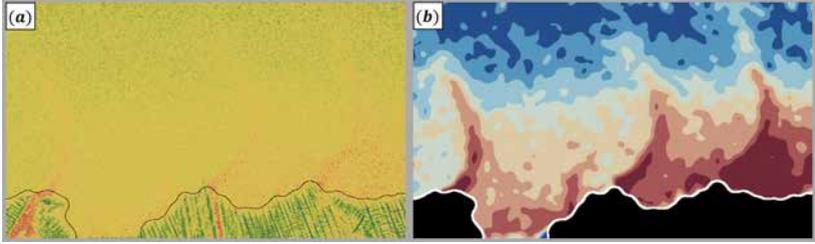


FIG. 155. Another example of convective plumes segmented in the LZ. Here, an early frame from the image sequence corresponding to Figure 146 is shown.

4. Summary

A robust and noise-resilient image processing pipeline for analyzing meso-scale metal alloy solidification processes in laboratory-scale experiments using Hele-Shaw liquid metal cells and dynamic X-ray imaging has been demonstrated. The developed methodology at present allows one to segment liquid and solidified zones within the field of view, detect the skeletons of solidified structures in the solid zone, perform orientation analysis, detect dominant dendrite grains (if any), quantify the dynamics of the solidification front and solute concentration above it, detect and separate liquid metal channels and cavities, as well as segment and characterize convective plumes in the liquid zone. Even with artificially lowered SNR, the code performed reliably and the demonstrated performance is such that in-depth physical analysis of images is feasible.

The code combines both existing and original methods, including the ones developed for bubble- and particle-laden flow imaging (Sections II A 2 b and II D), and is open-source, available on *GitHub*: Mihails-Birjukovs/Meso-scale_Solidification_Analysis.

III. Conclusions & outlook

The overarching goal of this thesis was the development of methods required for physical analysis of bubble- and particle-laden flow in liquid metal, as well as metal alloy solidification. This meant leveraging the most prospective of the currently available experimental approaches, and bridging the gap between experimental/numerical data acquisition and physical interpretation of system dynamics. The results of this thesis can be summarized as follows:

1. The previously reported experimental setups were improved upon, and it was shown that dynamic neutron radiography (NR) is indeed feasible for *in situ* imaging of bubble flow in liquid metal in thick model systems, and its *quantitative* analysis. NR allows one to observe the effects of varying the gas flow rates and applied magnetic field (MF) on bubble chain shape, and bubble trajectory and shape dynamics. The results, aside from being physically agreeable, were validated both indirectly, by comparison against numerical simulations, and directly by imaging reference objects and emulating bubble detection.
2. The above was accomplished using the developed image processing methodology, combining both existing and original methods, which enabled bubble detection in low signal-to-noise (SNR) images. It is expected that implemented code and/or separate elements thereof should be applicable well beyond this application and context.
3. At the moment, even with limited analysis of a part of the experimental data, it was shown that, depending on the MF configuration, radically different 1-bubble and bubble chain flow regimes occur. For example, for a fixed flow rate, bubbles in a chain can exhibit nearly constant, two stage or single stage non-uniform acceleration, each with significantly different trajectories, velocity time series and shape dynamics. Bubble chain shapes can also vary quite drastically. This is a clear proof-of-concept for MF control of bubble flow with industrially relevant dimensionless parameters.
4. Note that the accumulated neutron image dataset is the first in the field that systematically covers an industrially relevant region of the (Re, Eo, N) space for bubble chain flow model systems with both vertical and horizontal applied MF. It is planned to further process the data and systematically study the effects of applied horizontal and vertical MF with different strengths on bubble chain flow in a rectangular liquid gallium vessel. Bubble shape and motion

parameters will be assessed for different MF configurations and gas flow rates, and correlations with Reynolds (Re), Eötvös (EO) and Stuart (N) numbers will be obtained. This should enable to describe and predict the bubble chain instabilities that occur under different conditions.

5. Additional reference imaging experiments were also performed at extra high frame rates (up to 800 FPS), to be later used as benchmarks for numerical modelling of bubble shape dynamics, as well as for further development of tools for experimental data analysis.
6. Bubble chain flow imaging was also performed using X-ray radiography (XR) for the Hele-Shaw variants of the model systems imaged with NR, with identical MF configurations and appropriate flow rate ranges. The acquired images form a unique dataset which systematically covers geometrically confined bubble chain flow and bubble collective dynamics. An image processing code was implemented to resolve the very closely packed bubbles with significant irregular deformations. A shape analysis algorithm was implemented to quantify the shape dynamics. It is planned to process the accumulated data to systematically study bubble collective dynamics in different MF configurations, something that has not yet been done. It is also planned to combine the shape analysis algorithm with object tracking and optical flow velocimetry to measure bubble shape deformation rate, which will enable a more in-depth analysis of how bubbles interact with one another and applied MF.
7. To reconstruct bubble trajectories from image data for further physical analysis, an object tracking algorithm, MHT-X (offline multiple hypothesis tracking with Algorithm X), was developed, which was demonstrated as an efficient and versatile tool that can operate even in the presence of strong data noise, and can successfully treat cases with bubble flow where breakup and coalescence occur using weak mass and momentum conservation laws for motion prediction. It is expected that MHT-X will find many more applications in the scientific community owing to its generality and modularity. Note that this is the first reported numerically feasible implementation of MHT for offline tracking, which is perfectly suitable for scientific purposes.
8. A combination of high frame rate NR, a newly developed image processing methodology and an extended version (branch) of MHT-X has, for the first time, enabled successful detection and tracking of

particles in NR images of particle-laden liquid metal flow. The MHT-X branch developed for particle tracking and can also recover global continuous flow velocity from obtained particle tracking velocimetry data using divergence-free interpolation.

9. The above developments resulted in the first successful measurements of turbulence spatial and temporal characteristics from NR images of particle-laden liquid metal flow, and the implemented image processing and MHT-X code have been validated. Of particular importance was the reconstruction of the trajectory curvature probability density function, which is an embedding of the turbulence energy spectrum. Thus, the above methods are expected to play important roles in further research of bubble and particle flow, as well as turbulence in liquid metals, and beyond. In particular, it is planned to study bubble wake flow using particle tracking velocimetry combined with NR imaging.
10. A noise-resilient and memory-efficient dynamic mode decomposition (DMD) code was developed, and applied to numerical simulation output to compute the velocity field modes. DMD revealed several important aspects of bubble flow: despite the overall flow stabilization by MF, the system turns out to be very sensitive even to slight MF asymmetry, and flow is destabilized at lower gas flow rates than normal; bubble chain stability can be quantified by examining the associated modes, and chain coherence (stability height) can be determined; bubble wake flow can be decomposed into different momentum transfer mechanisms, i.e. vortex shedding modes can be identified; it can be seen that large scale vessel modes modulate the bubble wake flow modes, whereas the more disordered bubble wake modes can be seen to propagate instability into larger flow scales.
11. It is planned to apply DMD to higher-resolution simulations of in the future to study finer spatial and temporal scales within the flow field. In addition, DMD will be performed for bubble shape oscillations to observe how the application of MF affects the surface modes, which is of interest as it strongly determines individual and collective bubble dynamics. Note that the showcased application of DMD to (explicitly resolved) bubble flow simulations is the first of its kind.
12. The problem of a sorely lacking all-in-one open-source image processing code for analyzing X-ray images of directional solidification of metal alloys has been at least partially solved. The implemented (and open-source) code provides most of the required functionality by combining both known and new methods, and is expected to be of

great value to researchers in the field of alloy solidification, but also potentially beyond this context.

13. Finally, as promised, all the developed methods, which include image processing for dynamic NR and XR of bubble- and particle-laden liquid metal flow and metal alloy solidification, bubble and particle tracking (MHT-X), bubble shape analysis and DMD, have been made publicly available via open-source repositories on *GitHub*:

- Mihails-Birjukovs/Low_C-SNR_Bubble_Detection
- 4n515/X-ray-bubble-detection
- Mihails-Birjukovs/Bubble_shape_analysis
- Mihails-Birjukovs/Low_C-SNR_Particle_Detection
- Peteris-Zvejnieks/MHT-X
- Peteris-Zvejnieks/DivergenceFreeInterpolation
- Divergence-Free-Interpolant PyPi
- MartinKlevs/PyDMD
- MartinKlevs/MOSES-SVD
- Mihails-Birjukovs/Meso-scale_Solidification_Analysis

In-detail accounts for all the underlying methods are provided in the publications associated with this thesis [152, 153, 158, 182, 185, 213, 214, 223].

14. Using the developed methods and code, with the accumulated extensive neutron and X-ray image datasets for liquid metal systems with bubble and particle flow, as well as solidification, the aim is to apply the set of the new tools to extract physically meaningful correlations from the data and provide new insights into and systematic coverage of the properties of relevant model systems. In addition, experience and knowledge gained during the experiments and data processing over the course of this thesis will allow developing next-generation experimental methods – for instance, combined simultaneous X-ray or neutron imaging with ultrasound Doppler velocimetry and the application of contactless inductive flow tomography to bubble chain flow systems.

References

- [1] E. Baake, T. Fehling, D. Musaeva, and T. Steinberg, Neutron radiography for visualization of liquid metal processes: bubbly flow for CO₂ free production of hydrogen and solidification processes in EM field, IOP Conference Series: Materials Science and Engineering **228**, 012026 (2017), publisher: IOP Publishing.
- [2] K. Timmel, S. Eckert, G. Gerbeth, F. Stefani, and T. Wondrak, Experimental modeling of the continuous casting process of steel using low melting point metal alloys — the LIMMCAST program, ISIJ International **50**, 1134 (2010).
- [3] K. Timmel, N. Shevchenko, M. Röder, M. Anderhuber, P. Gardin, S. Eckert, and G. Gerbeth, Visualization of liquid metal two-phase flows in a physical model of the continuous casting process of steel, Metallurgical and Materials Transactions B **46**, 700 (2015).
- [4] T. Wondrak, S. Eckert, G. Gerbeth, K. Klotsche, F. Stefani, K. Timmel, A. Peyton, N. Terzija, and W. Yin, Combined electromagnetic tomography for determining two-phase flow characteristics in the submerged entry nozzle and in the mold of a continuous casting model, Metallurgical and Materials Transactions B **42**, 1201 (2011).
- [5] C. Zhang, Liquid metal flows driven by gas bubbles in a static magnetic field (2009).
- [6] E. Strumpf, Experimental study on rise velocities of single bubbles in liquid metal under the influence of strong horizontal magnetic fields in a flat vessel, International Journal of Multiphase Flow **97**, 168 (2017).
- [7] C. Zhang, S. Eckert, and G. Gerbeth, Experimental study of single bubble motion in a liquid metal column exposed to a DC magnetic field, International Journal of Multiphase Flow **31**, 824 (2005).
- [8] Z. Wang, S. Wang, X. Meng, and M. Ni, UDV measurements of single bubble rising in a liquid metal galinstan with a transverse magnetic field, International Journal of Multiphase Flow **94**, 201 (2017).
- [9] W. I. Shew, S. Poncet, and J.-F. Pinton, Force measurements on rising bubbles, Journal of Fluid Mechanics **569**, 51–60 (2006).
- [10] T. Richter, O. Keplinger, N. Shevchenko, T. Wondrak, K. Eckert, S. Eckert, and S. Odenbach, Single bubble rise in GaInSn in a horizontal magnetic field, International Journal of Multiphase Flow **104**, 32 (2018).
- [11] S. Schwarz, An immersed boundary method for particles and bubbles in magnetohydrodynamic flows (2014).
- [12] S. Schwarz and J. Fröhlich, Numerical study of single bubble motion in liquid metal exposed to a longitudinal magnetic field, International Journal of Multiphase Flow **62**, 134 (2014).
- [13] K. Jin, P. Kumar, S. Vanka, and B. Thomas, Rise of an argon bubble in

- liquid steel in the presence of a transverse magnetic field, *Physics of Fluids* **28**, 093301 (2016).
- [14] J. Zhang and M.-J. Ni, Direct simulation of single bubble motion under vertical magnetic field: Paths and wakes, *Physics of Fluids* **26**, 102102 (2014).
- [15] J. Zhang, M.-J. Ni, and R. Moreau, Rising motion of a single bubble through a liquid metal in the presence of a horizontal magnetic field, *Physics of Fluids* **28**, 032101 (2016).
- [16] D. Gaudlitz and N. Adams, Numerical investigation of rising bubble wake and shape variations, *Physics of Fluids* **21** (2009).
- [17] X. Wang, B. Klaasen, J. Degève, A. Mahulkar, G. Heynderickx, M.-F. Reyniers, B. Blanpain, and F. Verhaeghe, Volume-of-fluid simulations of bubble dynamics in a vertical Hele-Shaw cell, *Physics of Fluids* **28**, 053304 (2016).
- [18] V. Roig, M. Roudet, F. Risso, and A.-M. Billet, Dynamics of a high-Reynolds-number bubble rising within a thin gap, *Journal of Fluid Mechanics* **707**, 444 (2012).
- [19] R. May, F. Gruy, and J. Fröhlich, Impact of particle boundary conditions on the collision rates of inclusions around a single bubble rising in liquid metal, *PAMM* **18**, 1 (2018).
- [20] J.-S. Kroll-Rabotin, M. Gisselbrecht, B. Ott, R. May, J. Fröhlich, and J.-P. Bellot, Multiscale simulation of non-metallic inclusion aggregation in a fully resolved bubble swarm in liquid steel, *Metals* **10**, 10.3390/met10040517 (2020).
- [21] J.-P. Bellot, J.-S. Kroll-Rabotin, M. Gisselbrecht, M. Joishi, A. Saxena, S. Sanders, and A. Jardy, Toward better control of inclusion cleanliness in a gas stirred ladle using multiscale numerical modeling, *Materials* **11**, 10.3390/ma11071179 (2018).
- [22] M. Gisselbrecht, J.-S. Kroll-Rabotin, and J.-P. Bellot, Aggregation kernel of globular inclusions in local shear flow: application to aggregation in a gas-stirred ladle, *Metallurgical Research & Technology* **116**, 512 (2019).
- [23] G. Mougin and J. Magnaudet, Path instability of a rising bubble, *Physical Review Letters* **88**, 014502 (2002).
- [24] M. Tripathi, K. Sahu, and R. Govindarajan, Dynamics of an initially spherical bubble rising in quiescent liquid, *Nature Communications* **6**, 6268 (2015).
- [25] J. Zhang and M.-J. Ni, What happens to the vortex structures when the rising bubble transits from zigzag to spiral?, *Journal of Fluid Mechanics* **828**, 353 (2017).
- [26] D. Gaudlitz and N. Adams, Numerical investigation of rising bubble wake and shape variations, *Physics of Fluids* **21** (2009).
- [27] J. Zhang, K. Sahu, and M.-J. Ni, Transition of bubble motion from spiralling to zigzagging: A wake-controlled mechanism with a transverse magnetic field, *International Journal of Multiphase Flow* **136**, 103551 (2020).
- [28] J. Will, V. Mathai, S. Huisman, D. Lohse, C. Sun, and D. Krug, Kinematics and dynamics of freely rising spheroids at high reynolds numbers, *Journal*

- of Fluid Mechanics **912**, A16 (2021).
- [29] O. Keplinger, N. Shevchenko, and S. Eckert, Experimental investigation of bubble breakup in bubble chains rising in a liquid metal, *International Journal of Multiphase Flow* **116**, 39 (2019).
- [30] O. Keplinger, N. Shevchenko, and S. Eckert, Visualization of bubble coalescence in bubble chains rising in a liquid metal, *International Journal of Multiphase Flow* **105**, 159 (2018).
- [31] T. Ziegenhein and D. Lucas, Observations on bubble shapes in bubble columns under different flow conditions, *Experimental Thermal and Fluid Science* **85** (2017).
- [32] T. Haas, C. Schubert, M. Eickhoff, and H. Pfeifer, A review of bubble dynamics in liquid metals, *Metals* **11**, 664 (2021).
- [33] Z. Liu and B. Li, Large-eddy simulation of transient horizontal gas–liquid flow in continuous casting using dynamic subgrid-scale model, *Metallurgical and Materials Transactions B* (2017).
- [34] W. Yang, Z. Luo, N. Zhao, and Z. Zou, Numerical analysis of effect of initial bubble size on captured bubble distribution in steel continuous casting using euler-lagrange approach considering bubble coalescence and breakup, *Metals* **10**, 1160 (2020).
- [35] W. Yang, Z. Luo, Y. Gu, Z. Liu, and Z. Zou, Numerical analysis of effect of operation conditions on bubble distribution in steel continuous casting mold with advanced bubble break-up and coalescence models, *ISIJ International* **60(10)**, 2234 (2020).
- [36] M. Taborda, M. Sommerfeld, and M. Muniz, LES-Euler/Lagrange modelling of bubble columns considering mass transfer, chemical reactions and effects of bubble dynamics, *Chemical Engineering Science* **229**, 116121 (2021).
- [37] W. Lou and M. Zhu, Numerical simulation of desulfurization behavior in gas-stirred systems based on computation fluid dynamics–simultaneous reaction model (CFD–SRM) coupled model, *Metallurgical and Materials Transactions B* **45**, 1706 (2014).
- [38] Y. Liu, M. Ersson, H. Liu, P. Jönsson, and Y. Gan, A review of physical and numerical approaches for the study of gas stirring in ladle metallurgy, *Metallurgical and Materials Transactions B* **50** (2018).
- [39] Q. Cao and L. Nastac, Numerical modelling of the transport and removal of inclusions in an industrial gas-stirred ladle, *Ironmaking & Steelmaking* **45**, 1 (2018).
- [40] R. Morales, F. A. Calderon Hurtado, K. Chattopadhyay, and S. Guarneros, Physical and mathematical modeling of flow structures of liquid steel in ladle stirring operations, *Metallurgical and Materials Transactions B* **51** (2020).
- [41] E. Ramasetti, V.-V. Visuri, P. Sulasalmi, T. Palovaara, A. Gupta, and T. Fabritius, Physical and CFD modeling of the effect of top layer properties on the formation of open-eye in gas-stirred ladles with single and dual-plugs, *Steel Research International* **90** (2019).
- [42] A. Nguyen and H.-J. Schulze, *Colloidal science of flotation* (2004).
- [43] M. Fuerstenau, R.-H. Yoon, and G. Jameson, *Froth flotation: A century*

- of innovation (Society for Mining, Metallurgy, and Exploration, Littleton, Colo., Ix, 2007).
- [44] A.-E. Sommer, M. Nikpay, S. Heitkam, M. Rudolph, and S. Eckert, A novel method for measuring flotation recovery by means of 4D particle tracking velocimetry, *Minerals Engineering* **124**, 116–122 (2018).
- [45] M. Amoorezaei, S. Gurevich, and N. Provatas, Orientation selection in solidification patterning, *Acta Materialia* **60** (2012).
- [46] A. Kao, N. Shevchenko, S. He, P. Lee, S. Eckert, and K. Pericleous, Magnetic effects on microstructure and solute plume dynamics of directionally solidifying Ga-In alloy, *JOM* **72** (2020).
- [47] D. Stefanescu and R. Ruxanda, *Fundamentals of solidification* (2004) pp. 71–92.
- [48] C. Beckermann, Modelling of macrosegregation: Applications and future needs, *International Materials Reviews* **47**, 243 (2002).
- [49] R. Reed, T. Tao, and N. Warnken, Alloys-by-design: Application to nickel-based single crystal superalloys, *Acta Materialia* **57**, 5898 (2009).
- [50] J. D. Madison, Investigation of solidification defect formation by three-dimensional reconstruction of dendritic structures (2010).
- [51] P. Auburtin, T. Wang, S. Cockcroft, and A. Mitchell, Freckle formation and freckle criterion in superalloy castings, *Metallurgical and Materials Transactions B* **31**, 801 (2000).
- [52] C. Beckermann, J. Gu, and W. Boettinger, Development of a freckle predictor via rayleigh number method for single-crystal nickel-base superalloy castings, *Metallurgical and Materials Transactions A* **31**, 2545 (2000).
- [53] A. Saad, C.-A. Gandin, M. Bellet, N. Shevchenko, and S. Eckert, Simulation of channel segregation during directional solidification of In—75 wt pct Ga. qualitative comparison with in situ observations, *Metallurgical and Materials Transactions A* (2015).
- [54] N. Shevchenko, S. Boden, G. Gerbeth, and S. Eckert, Chimney formation in solidifying Ga-25wt pct In alloys under the influence of thermosolutal melt convection, *Metallurgical and Materials Transactions A* **44** (2013).
- [55] M. Asta, C. Beckermann, A. Karma, W. Kurz, R. Napolitano, M. Plapp, G. Purdy, M. Rappaz, and R. Trivedi, Solidification microstructures and solid-state parallels: Recent developments, future directions, *Acta Materialia* **57**, 941 (2009).
- [56] T. Haxhimali, A. Karma, F. Gonzales, and M. Rappaz, Orientation selection in dendritic evolution, *Nature materials* **5**, 660 (2006).
- [57] J. Strickland, B. Nenchev, and H. Dong, On directional dendritic growth and primary spacing—a review, *Crystals* **10**, 627 (2020).
- [58] J. Ramirez and C. Beckermann, Evaluation of a rayleigh-number-based freckle criterion for Pb-Sn alloys and Ni-base superalloys, *Metallurgical and Materials Transactions A* **34**, 1525 (2003).
- [59] N. Shevchenko, O. Keplinger, O. Sokolova, and S. Eckert, The effect of natural and forced melt convection on dendritic solidification in Ga-In alloys, *Journal of Crystal Growth* **417** (2015).
- [60] N. Shevchenko, H. Neumann-Heyme, C. Pickmann, E. Schaberger-

- Zimmermann, G. Zimmermann, K. Eckert, and S. Eckert, Investigations of fluid flow effects on dendritic solidification: Consequences on fragmentation, macrosegregation and the influence of electromagnetic stirring, IOP Conference Series: Materials Science and Engineering **228**, 012005 (2017).
- [61] D. Anderson and P. Guba, Convective phenomena in mushy layers, Annual Review of Fluid Mechanics **52** (2020).
- [62] P. Soar, A. Kao, G. Djambazov, N. Shevchenko, S. Eckert, and K. Pericleous, The integration of structural mechanics into microstructure solidification modelling, IOP Conference Series: Materials Science and Engineering **861**, 012054 (2020).
- [63] B. Cai, A. Kao, E. Boller, O. Magdysyuk, R. Atwood, N. T. Vo, K. Pericleous, and P. Lee, Revealing the mechanisms by which magneto-hydrodynamics disrupts solidification microstructures, Acta Materialia **196** (2020).
- [64] S. He, N. Shevchenko, and S. Eckert, In situ observation of directional solidification in Ga-In alloy under a transverse dc magnetic field, IOP Conference Series: Materials Science and Engineering **861**, 012025 (2020).
- [65] S. Eckert, P. Nikrityuk, B. Willers, D. Raebiger, N. Shevchenko, H. Neumann-Heyme, V. Travnikov, S. Odenbach, A. Voigt, and K. Eckert, Electromagnetic melt flow control during solidification of metallic alloys, The European Physical Journal Special Topics **220**, 123 (2013).
- [66] F. Stefani, T. Gundrum, and G. Gerbeth, Contactless inductive flow tomography, Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics **70**, 056306 (2004).
- [67] T. Gundrum, G. Gerbeth, F. Stefani, and M. Xu, Experimental aspects of contactless inductive flow tomography, Magnetohydrodynamics **42**, 153 (2006).
- [68] T. Vogt, A. Andruszkiewicz, S. Eckert, K. Eckert, S. Odenbach, and G. Gerbeth, Mixing enhancement in gas-stirred melts by rotating magnetic fields, Metallurgical and Materials Transactions B **43**, 1454 (2012).
- [69] O. Keplinger, N. Shevchenko, and S. Eckert, Validation of X-ray radiography for characterization of gas bubbles in liquid metals, IOP Conference Series: Materials Science and Engineering **228**, 012009 (2017).
- [70] L. Liu, O. Keplinger, T. Ziegenhein, N. Shevchenko, S. Eckert, H. Yan, and D. Lucas, Euler-euler modeling and X-ray measurement of oscillating bubble chain in liquid metals, International Journal of Multiphase Flow **110**, 218 (2018).
- [71] B. Krull, E. Strumpf, O. Keplinger, N. Shevchenko, J. Fröhlich, S. Eckert, and G. Gerbeth, Combined experimental and numerical analysis of a bubbly liquid metal flow, IOP Conference Series: Materials Science and Engineering **228**, 012006 (2017).
- [72] M. Banowski, U. Hampel, E. Krepper, M. Beyer, and D. Lucas, Experimental investigation of two-phase pipe flow with ultrafast X-ray tomography and comparison with state-of-the-art CFD simulations, Nuclear Engineering and Design **336**, 90 (2017).
- [73] F. Fischer, D. Hoppe, E. Schleicher, G. Mattausch, H. Flaske, R. Bartel,

- and U. Hampel, An ultra fast electron beam X-ray tomography scanner, *Measurement Science and Technology* **19**, 094002 (2008).
- [74] F. Barthel, M. Bieberle, D. Hoppe, M. Banowski, and U. Hampel, Velocity measurement for two-phase flows based on ultrafast X-ray tomography, *Flow Measurement and Instrumentation* **46**, 196 (2015).
- [75] V. Dzelme, M. Sarma, V. Geža, M. Ščepanskis, and A. Jakovics, Transition of different voltical structures in rotating permanent magnets agitated flow: numerical and experimental neutron radiography investigation, 10th PAMIR International Conference Fundamental and Applied MHD (Cagliari, Italy, 2016).
- [76] M. Ščepanskis, M. Sarma, P. Vontobel, P. Trtik, K. Thomsen, A. Jakovics, and T. Beinerts, Assessment of electromagnetic stirrer agitated liquid metal flows by dynamic neutron radiography, *Metallurgical and Materials Transactions B* **48**, 1045 (2017).
- [77] R. Zboray and P. Trtik, 800 fps neutron radiography of air-water two-phase flow, *MethodsX* **5**, 96 (2018).
- [78] R. Zboray, R. Adams, M. Cortesi, and H.-M. Prasser, Development of a fast neutron imaging system for investigating two-phase flows in nuclear thermal-hydraulic phenomena: A status report, *Nuclear Engineering and Design* **273**, 10–23 (2014).
- [79] R. Zboray and P. Trtik, In-depth analysis of high-speed, cold neutron imaging of air-water two-phase flows, *Flow Measurement and Instrumentation* **66**, 182 (2019).
- [80] S. Heitkam, M. Rudolph, T. Lappan, M. Sarma, S. Eckert, P. Trtik, E. Lehmann, P. Vontobel, and K. Eckert, Neutron imaging of froth structure and particle motion, *Minerals Engineering* **119**, 126 (2018).
- [81] S. Heitkam, T. Lappan, S. Eckert, P. Trtik, and K. Eckert, Tracking of particles in froth using neutron imaging, *Chemie Ingenieur Technik* **91**, 1001 (2019).
- [82] Y. Saito, K. Mishima, Y. .Tobita, T. Suzuki, and M. Matsubayashi, Measurements of liquid-metal two-phase flow by using neutron radiography and electrical conductivity probe, *Experimental Thermal and Fluid Science* **29**, 323 (2005).
- [83] Y. Saito, K. Mishima, Y. .Tobita, T. Suzuki, M. Matsubayashi, I. Lim, and J. Cha, Application of high frame-rate neutron radiography to liquid-metal two-phase flow research, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **542**, 168 (2005).
- [84] M. Akashi, O. Keplinger, N. Shevchenko, S. Anders, and M. Reuter, X-ray radioscopic visualization of bubbly flows injected through a top submerged lance into a liquid metal, *Metallurgical and Materials Transactions B* **51** (2019).
- [85] T. Lappan, M. Sarma, S. Heitkam, P. Trtik, D. Mannes, K. Eckert, and S. Eckert, Neutron radiography of particle-laden liquid metal flow driven by an electromagnetic induction pump, *Magneto hydrodynamics* **56**, 167 (2020).

- [86] M. Sarma, M. Ščepanskis, A. Jakovics, K. Thomsen, R. Nikoluškins, P. Vontobel, T. Beinerts, A. Bojarevics, and E. Platacis, Neutron radiography visualization of solid particles in stirring liquid metal, *Physics Procedia* **69**, 457 (2015).
- [87] M. Ščepanskis, M. Sarma, P. Vontobel, P. Trtik, K. Thomsen, A. Jakovics, and T. Beinerts, Assessment of electromagnetic stirrer agitated liquid metal flows by dynamic neutron radiography, *Metallurgical and Materials Transactions B* **48**, 1045 (2017).
- [88] V. Dzelme, A. Jakovics, J. Vencels, D. Köppen, and E. Baake, Numerical and experimental study of liquid metal stirring by rotating permanent magnets, *IOP Conference Series: Materials Science and Engineering* **424**, 012047 (2018).
- [89] Y. Liu, M. Ersson, H. Liu, P. Jönsson, and Y. Gan, A review of physical and numerical approaches for the study of gas stirring in ladle metallurgy, *Metallurgical and Materials Transactions B* **50** (2018).
- [90] Q. Cao and L. Nastac, Numerical modelling of the transport and removal of inclusions in an industrial gas-stirred ladle, *Ironmaking & Steelmaking* **45**, 984 (2018).
- [91] H. Kusuno and T. Sanada, Wake-induced lateral migration of approaching bubbles, *International Journal of Multiphase Flow* **139**, 103639 (2021).
- [92] J. Zhang, L. Chen, and M.-J. Ni, Vortex interactions between a pair of bubbles rising side by side in ordinary viscous liquids, *Physical Review Fluids* **4** (2019).
- [93] J. Zhang, M.-J. Ni, and J. Magnaudet, Three-dimensional dynamics of a pair of deformable bubbles rising initially in line. Part 1: Moderately inertial regimes (2020).
- [94] A. Filella, P. Ern, and R. Véronique, Interaction of two oscillating bubbles rising in a thin-gap cell: vertical entrainment and interaction with vortices, *Journal of Fluid Mechanics* **888** (2020).
- [95] R. Clift, J. Grace, and M. Weber, *Bubbles, Drops, and Particles* (1978).
- [96] T. Lappan, M. Sarma, S. Heitkam, P. Trtik, D. Mannes, K. Eckert, and S. Eckert, Neutron radiography of particle-laden liquid metal flow driven by an electromagnetic induction pump, *Magneto hydrodynamics* **56**, 167 (2020).
- [97] C. Zhang, S. Eckert, and G. Gerbeth, The flow structure of a bubble-driven liquid-metal jet in a horizontal magnetic field, *Journal of Fluid Mechanics* **575**, 57 (2007).
- [98] A.-E. Sommer, K. Ortmann, M. van Heerden, T. Richter, T. Leadbeater, K. Cole, S. Heitkam, P. Brito-Parada, and K. Eckert, Application of positron emission particle tracking (PEPT) to measure the bubble-particle interaction in a turbulent and dense flow, *Minerals Engineering* **156**, 106410 (2020).
- [99] D. Burnard, A. Caden, J. Gargiuli, T. Leadbeater, D. Parker, and W. Griffiths, A positron emission particle tracking (PEPT) study of inclusions in liquid aluminium alloy, *Advanced Materials Research* **922**, 43 (2014).
- [100] D. Burnard, J. Gargiuli, T. Leadbeater, D. Parker, and W. Griffiths, The

- application of positron emission particle tracking (PEPT) to study inclusions in the casting process (2011) pp. 25–28.
- [101] W. Griffiths, Y. Gaber Beshay, A. Caden, X. Fan, J. Gargiuli, T. Leadbeater, and D. Parker, The use of positron emission particle tracking (PEPT) to study the movement of inclusions in low-melting-point alloy castings, *Metallurgical and Materials Transactions B* **43** (2011).
- [102] A. Dybalska, A. Caden, D. Parker, J. Wedderburn, and W. Griffiths, Liquid metal flow studied by positron emission tracking, *Metallurgical and Materials Transactions B* **51** (2020).
- [103] J. Cimbala, D. Hughes, S. Levine, and D. Sathianathan, Application of neutron radiography for fluid flow visualization, *Nucl. Technol.; (United States)* **81:3** (1988).
- [104] T. Lappan, M. Sarma, S. Heitkam, D. Mannes, P. Trtik, N. Shevchenko, K. Eckert, and S. Eckert, X-ray and neutron radiographic experiments on particle-laden molten metal flows (2021) pp. 13–29.
- [105] S. Heitkam, T. Lappan, S. Eckert, P. Trtik, and K. Eckert, Tracking of particles in froth using neutron imaging, *Chemie Ingenieur Technik* **91** (2019).
- [106] S. Anders, D. Noto, M. Seilmayer, and S. Eckert, Spectral random masking: a novel dynamic masking technique for PIV in multiphase flows, *Experiments in Fluids* **60**, 68 (2019).
- [107] S. Anders, D. Noto, Y. Tasaka, and S. Eckert, Simultaneous optical measurement of temperature and velocity fields in solidifying liquids, *Experiments in Fluids* **61** (2020).
- [108] N. Chenouard, I. Smal, F. Chaumont, M. Maška, I. Sbalzarini, Y. Gong, J. Cardinale, C. Carthel, S. Coraluppi, M. Winter, A. Cohen, W. Godinez, K. Rohr, Y. Kalaidzidis, L. Liang, J. Duncan, H. Shen, Y. Xu, K. Magnusson, and E. Meijering, Objective comparison of particle tracking methods, *Nature methods* **11** (2014).
- [109] D. Schanz, S. Gesemann, and A. Schröder, Shake-the-box: Lagrangian particle tracking at high particle image densities, *Experiments in Fluids* **57** (2016).
- [110] B. Wieneke, Iterative reconstruction of volumetric particle distribution, *Measurement Science and Technology* **24**, 024008 (2012).
- [111] G. Haller, An objective definition of a vortex, *Journal of Fluid Mechanics* **525**, 1–26 (2005).
- [112] V. Holmen, Methods for vortex identification (2012).
- [113] J. Tu, C. Rowley, D. Luchtenburg, S. Brunton, and J. Kutz, On dynamic mode decomposition: Theory and applications, *Journal of Computational Dynamics* **1** (2013).
- [114] P. Schmid and J. Sesterhenn, Dynamic mode decomposition of numerical and experimental data, *Journal of Fluid Mechanics* **656**, <https://doi.org/10.1017/S0022112010001217> (2008).
- [115] C. Rowley, I. Mezic, S. Bagheri, P. Schlatter, and D. Henningson, Spectral analysis of nonlinear flows, *Journal of Fluid Mechanics* **641**, 115 (2009).
- [116] A. Seenaa and H. J. Sung, Dynamic mode decomposition of turbulent cavity

- flows for self-sustained oscillations, *International Journal of Heat and Fluid Flow* **32**, 1098 (2011).
- [117] M. R. Jovanović, P. J. Schmid, and J. W. Nichols, Sparsity-promoting dynamic mode decomposition, *Physics of Fluids* **26**, 024103 (2014), 1309.4165.
- [118] J. Rojsel, Koopman mode analysis of the side-by-side cylinder wake (2017).
- [119] J. Manning and R. Baldick, Forecasting short-term dynamics of fair-weather cumuli using dynamic mode decomposition, arXiv:1907.12980 [cs] (2019), 1907.12980.
- [120] J. Grosek and J. N. Kutz, Dynamic mode decomposition for real-time background/foreground separation in video, arXiv:1404.7592 [cs] (2014), 1404.7592.
- [121] I. Ul Haq, K. Fujii, and Y. Kawahara, Dynamic mode decomposition via dictionary learning for foreground modeling in videos, *Computer Vision and Image Understanding* **199**, 103022 (2020).
- [122] B. W. Brunton, L. A. Johnson, J. G. Ojemann, and J. N. Kutz, Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition, arXiv:1409.5496 [q-bio] 10.1016/j.jneumeth.2015.10.010 (2020), 1409.5496.
- [123] Y. Shiraishi, Y. Kawahara, O. Yamashita, R. Fukuma, S. Yamamoto, Y. Saitoh, H. Kishima, and T. Yanagisawa, Neural decoding of electrocorticographic signals using dynamic mode decomposition, *Journal of Neural Engineering* **17**, 036009 (2020), publisher: IOP Publishing.
- [124] A. B. Albidah, W. Brevis, V. Fedun, I. Ballai, D. B. Jess, M. Stangalini, J. Higham, and G. Verth, Proper orthogonal and dynamic mode decomposition of sunspot data, arXiv (2020), publisher: Sheffield.
- [125] V. C. Loukopoulos, G. C. Bourantas, and K. Miller, Study of the thermo-magneto-hydrodynamic flow of micropolar-nanofluid in square enclosure using dynamic mode decomposition and proper orthogonal decomposition, *European Journal of Mechanics, B/Fluids* **84**, 272 (2020), publisher: Elsevier.
- [126] R. Taylor, J. N. Kutz, K. Morgan, and B. A. Nelson, Dynamic mode decomposition for plasma diagnostics and validation, *The Review of Scientific Instruments* **89**, 053501 (2018).
- [127] A. A. Kaptanoglu, K. D. Morgan, C. J. Hansen, and S. L. Brunton, Characterizing magnetized plasmas with dynamic mode decomposition, *Physics of Plasmas* **27**, 032108 (2020), <https://doi.org/10.1063/1.5138932>.
- [128] A. Alessandri, P. Bagnerini, M. Gaggero, D. Lengani, and D. Simoni, Dynamic mode decomposition for the inspection of three-regime separated transitional boundary layers using a least squares method, *Physics of Fluids* **31**, 044103 (2019), publisher: American Institute of Physics.
- [129] M. Liu, L. Tan, and S. Cao, Dynamic mode decomposition of gas-liquid flow in a rotodynamic multiphase pump, *Renewable Energy* **139**, 1159 (2019).
- [130] R. Qiu, R. Huang, Y. Wang, and C. Huang, Dynamic mode decomposition and reconstruction of transient cavitating flows around a clark-y hydrofoil, *Theoretical and Applied Mechanics Letters* **10**, 327 (2020).
- [131] W. Wu, C. Meneveau, and R. Mittal, Spatio-temporal dynamics of turbulent

- separation bubbles, *Journal of Fluid Mechanics* **883** (2020).
- [132] E. Ramos, G. Darze, F. Nascimento, J. Faccini, and G. Giralaldi, Comparison of dynamic mode decomposition and deep learning techniques for two-phase flows analysis, *Flow, Turbulence and Combustion* **105** (2020).
- [133] G. Reinhart, D. Grange, L. Abou Khalil, N. Manginck-Noël, N. Niane, V. Maguin, G. Guillemot, C.-A. Gandin, and H. Nguyen-Thi, Impact of solute flow during directional solidification of a Ni-based alloy: In-situ and real-time X-radiography, (2020).
- [134] M. Becker, L. Sturz, D. Bräuer, and F. Kargl, A comparative in situ x-radiography and dnn model study of solidification characteristics of an equiaxed dendritic Al-Ge alloy sample, *Acta Materialia* **201**, 286 (2020).
- [135] P. Delaleau, C. Beckermann, R. Mathiesen, and L. Arnberg, Mesoscopic simulation of dendritic growth observed in X-ray video microscopy during directional solidification of Al-Cu alloys, *ISIJ International* **50**, 1886 (2010).
- [136] T. Werner, M. Becker, J. Baumann, C. Pickmann, L. Sturz, and F. Kargl, In situ observation of the impact of hydrogen bubbles in Al-Cu melt on directional dendritic solidification, *Journal of Materials Science* **56** (2021).
- [137] S. Karagadde, L. Yuan, N. Shevchenko, S. Eckert, and P. Lee, 3-D microstructural model of freckle formation validated using in situ experiments, *Acta Materialia* **79**, 168–180 (2014).
- [138] A. Kao, N. Shevchenko, M. Alexandrakis, I. Krastins, S. Eckert, and K. Pericleous, Thermal dependence of large-scale freckle defect formation, *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* **377** (2019).
- [139] W. Mirihanage, K. Falch, I. Snigireva, A. Snigirev, Y. Li, L. Arnberg, and R. Mathiesen, Retrieval of three-dimensional spatial information from fast in situ two-dimensional synchrotron radiography of solidification microstructure evolution, *Acta Materialia* **81**, 241–247 (2014).
- [140] H. Soltani, G. Reinhart, M. Benoudia, F. Ngomesse, M. Zahzouh, and H. Nguyen-Thi, Equiaxed grain structure formation during directional solidification of a refined Al-20wt.%Cu alloy: In situ analysis of temperature gradient effects, *Journal of Crystal Growth* **587**, 126645 (2022).
- [141] Y. Tang, Y. Wu, Y. Zhang, Y. Dai, Q. Dong, Y. Han, G. Zhu, J. Zhang, Y. Fu, and B. Sun, Intermittent nucleation and periodic growth of grains under thermo-solutal convection during directional solidification of Al-Cu alloy, *Acta Materialia* **212**, 116861 (2021).
- [142] S. Boden, S. Eckert, B. Willers, and G. Gerbeth, X-ray radioscopic visualization of the solutal convection during solidification of a Ga-30 Wt Pct In alloy, *Metallurgical and Materials Transactions A* **39**, 613 (2008).
- [143] B. Nenchev, J. Strickland, K. Tassenberg, S. Perry, S. Gill, and H. Dong, Automatic recognition of dendritic solidification structures: Denmap, *Journal of Imaging* **6**, 19 (2020).
- [144] K. Tassenberg, B. Nenchev, J. Strickland, S. Perry, and D. Weston, Denmap single crystal solidification structure feature extraction: Automation and application, *Materials Characterization* (2020).
- [145] W. Wan, D. Li, H. Wang, L. Zhao, X. Shen, D. Sun, J. Chen, and C. Xiao,

- Automatic identification and quantitative characterization of primary dendrite microstructure based on machine learning, *Crystals* **11**, 1060 (2021).
- [146] T. Hughes, A. Robinson, and S. Mcfadden, Multiple dendrite tip tracking for in-situ directional solidification: Experiments and comparisons to theory, *Materials Today Communications* **29**, 102807 (2021).
- [147] A. Viardin, K. Noth, M. Torabi Rad, and L. Sturz, Automatic detection of equiaxed dendrites using computer vision neural networks, (2022).
- [148] N. Wang, Y. Tang, Y. Wu, Y. Zhang, Y. Dai, J. Zhang, R. Zhang, Y. Xu, and B. Sun, Dynamic evolution of microstructure morphology in thin-sample solidification: Deep learning assisted synchrotron X-ray radiography, *Materials Characterization* **181**, 111451 (2021).
- [149] J. Schindelin, I. Arganda-Carreras, E. Frise, V. Kaynig, M. Longair, T. Pietzsch, S. Preibisch, C. Rueden, S. Saalfeld, B. Schmid, J.-Y. Tinevez, D. White, V. Hartenstein, K. Eliceiri, P. Tomancak, and A. Cardona, Fiji: An open-source platform for biological-image analysis, *Nature methods* **9**, 676 (2012).
- [150] C. Schneider, W. Rasband, and K. Eliceiri, NIH image to ImageJ: 25 years of image analysis, *Nature Methods* **9** (2012).
- [151] E. Lehmann, P. Vontobel, and L. Wiesel, Properties of the radiography facility NEUTRA at SINQ and its use as European reference facility, *Nondestructive Testing and Evaluation* **16**, 191 (2001).
- [152] M. Birjukovs, V. Dzelme, A. Jakovics, K. Thomsen, and P. Trtik, Argon bubble flow in liquid gallium in external magnetic field, *International Journal of Applied Electromagnetics and Mechanics* **63**, 1 (2020).
- [153] M. Birjukovs, V. Dzelme, A. Jakovics, K. Thomsen, and P. Trtik, Phase boundary dynamics of bubble flow in a thick liquid metal layer under an applied magnetic field, *Physical Review Fluids* **5**, 10.1103/PhysRevFluids.5.061601 (2020).
- [154] J. Vencels, P. Råback, and V. Geža, EOF-Library: Open-source Elmer FEM and OpenFOAM coupler for electromagnetics and fluid dynamics, *SoftwareX* **9** (2019).
- [155] G. Sapiro, *Vector (self) snakes: a geometric framework for color, texture, and multiscale image segmentation* (1996) pp. 817 – 820 vol.1.
- [156] *CurvatureFlowFilter* (2010).
- [157] R. Gonzalez and R. Woods, *Digital Image Processing* (2006).
- [158] M. Birjukovs, P. Trtik, A. Kaestner, J. Hovind, M. Klevs, D. J. Gawryluk, K. Thomsen, and A. Jakovics, Resolving gas bubbles ascending in liquid metal from low-SNR neutron radiography images, *Applied Sciences* **11**, 10.3390/app11209710 (2021).
- [159] A. Kaestner, S. Hartmann, G. Kühne, G. Frei, C. Grünzweig, L. Josic, F. Schmid, and E. Lehmann, The ICON beamline—a facility for cold neutron imaging at SINQ, *Nuclear Instruments & Methods in Physics Research Section A-accelerators Spectrometers Detectors and Associated Equipment* **659**, 387 (2011).
- [160] P. Perona and J. Malik, Scale-space and edge detection using anisotropic diffusion, *IEEE Transactions on Pattern Analysis and Machine Intelligence*

- 12**, 629 (1990).
- [161] J. Weickert, B. Romeny, and M. Viergever, Efficient and reliable schemes for nonlinear diffusion filtering, *IEEE Transactions on Image Processing* **7**, 398 (1998).
- [162] PeronaMalikFilter (2010).
- [163] T. Le, R. Chartrand, and T. Asaki, A variational approach to reconstructing images corrupted by poisson noise, *Journal of Mathematical Imaging and Vision - JMIV* **27**, 257 (2007).
- [164] Totalvariationfilter (2018).
- [165] R. Haralick, S. Sternberg, and X. Zhuang, Image analysis using mathematical morphology, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **9**, 532 (1987).
- [166] N. Otsu, A threshold selection method from gray-level histograms, *IEEE Transactions on Systems, Man and Cybernetics* **9**, 62 (1979).
- [167] Tone reproduction, in *The Reproduction of Colour* (John Wiley & Sons, Ltd, 2004) Chap. 6, pp. 47–67, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/0470024275.ch6>.
- [168] P. Getreuer, Chan-Vese segmentation, *Image Processing On Line* **2**, 214 (2012).
- [169] J. Kapur, P. Sahoo, and A. Wong, A new method for gray-level picture thresholding using the entropy of the histogram, *Computer Vision, Graphics, and Image Processing* **29**, 273 (1980).
- [170] B. Coll and J.-M. Morel, A non-local algorithm for image denoising (2005) pp. 60– 65 vol. 2.
- [171] B. Coll and J.-M. Morel, Non-local means denoising, *Image Processing On Line* **1** (2011).
- [172] Nonlocalmeansfilter (2012).
- [173] M. Sezgin and B. Sankur, Comparison of thresholding methods for non-destructive testing applications, (2004).
- [174] L. Dah-Jye, S. Antani, and L. Long, Similarity measurement using polygon curve representation and Fourier descriptors for shape-based vertebral image retrieval, *Proceedings of SPIE - The International Society for Optical Engineering* **5032** (2003).
- [175] U. Ramer, An iterative procedure for the polygonal approximation of plane curves, *Computer Graphics and Image Processing. Academic Press Inc.*, August, 1972. vol. 1: pp. 244- 256 : ill. ; diagrams. includes bibliography **1** (1972).
- [176] C. Kim, F. Li, A. Ciptadi, and J. Rehg, Multiple hypothesis tracking revisited (2015) pp. 4696–4704.
- [177] J. Rubio, J. Serrat, and A. López, Multiple target tracking and identity linking under split, merge and occlusion of targets and observations, *ICPRAM 2012 - Proceedings of the 1st International Conference on Pattern Recognition Applications and Methods* **2**, 15 (2012).
- [178] S. Blackman, Multiple hypothesis tracking for multiple target tracking, *Aerospace and Electronic Systems Magazine, IEEE* **19**, 5 (2004).
- [179] I. Cox and S. Hingorani, An efficient implementation of Reid’s multiple

- hypothesis tracking algorithm and its evaluation for the purpose of visual tracking, *Pattern Analysis and Machine Intelligence*, IEEE Transactions on **18**, 138 (1996).
- [180] D. Reid, An algorithm for tracking multiple targets, *IEEE Transactions on Automatic Control* **24**, 843 (1979).
- [181] D. E. Knuth, Dancing links, (2000), arXiv:cs/0011047.
- [182] P. Zvejnieks, M. Birjukovs, M. Klevs, M. Akashi, S. Eckert, and A. Jakovics, MHT-X: offline multiple hypothesis tracking with algorithm X, *Experiments in Fluids* **63** (2022).
- [183] G. Brasó and L. Leal-Taixé, Learning a neural solver for multiple object tracking, in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020) pp. 6246–6256.
- [184] D. Obiso, M. Akashi, S. Kriebitzsch, B. Meyer, M. Reuter, and A. Richter, CFD modeling and experimental validation of top-submerged-lance gas injection in liquid metal, *Metallurgical and Materials Transactions B* **51** (2020).
- [185] M. Birjukovs, P. Zvejnieks, T. Lappan, M. Sarma, S. Heitkam, P. Trtik, D. Mannes, S. Eckert, and A. Jakovics, Particle tracking velocimetry in liquid gallium flow around a cylindrical obstacle, *Experiments in Fluids* **63** (2022).
- [186] D. R. Lide, ed., *English CRC Handbook of Chemistry and Physics* (CRC Press, Boca Raton, 2019).
- [187] W. Martienssen and H. Warlimont, *Springer Handbook of Condensed Matter and Materials Data* (Springer, Berlin, Heidelberg, 2005).
- [188] A. P. Kaestner, S. Hartmann, G. Kühne, G. Frei, C. Grünzweig, L. Josic, F. Schmid, and E. H. Lehmann, The ICON beamline - A facility for cold neutron imaging at SINQ, *Nuclear Instruments and Methods in Physics Research Section A - Accelerators, Spectrometers, Detectors and Associated Equipment* **659**, 387 (2011).
- [189] B. Blau, K. N. Clausen, S. Gvasaliya, M. Janoschek, S. Janssen, L. Keller, B. Roessli, J. Schefer, P. Tregenna-Piggott, W. Wagner, and O. Zaharko, The Swiss spallation neutron source SINQ at Paul Scherrer Institut, *Neutron News* **20**, 5 (2009).
- [190] Inpaint (2015).
- [191] E. J. Fuselier, Sobolev-type approximation rates for divergence-free and curl-free RBF interpolants, *Mathematics of Computation* **77**, 1407 (2008).
- [192] H. Wendland, Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree, *Advances in Computational Mathematics* **4**, 389 (1995).
- [193] P. Farrell, K. Gillow, and H. Wendland, Multilevel interpolation of divergence-free vector fields, *IMA Journal of Numerical Analysis* **37**, 332 (2017).
- [194] H. Wendland, Divergence-free kernel methods for approximating the stokes problem, *SIAM J. Numerical Analysis* **47**, 3158 (2009).
- [195] B. Tapley, H. Andersson, E. Celledoni, and B. Owren, Computational geometric methods for preferential clustering of particle suspensions, *Journal*

- of Computational Physics **448**, 110725 (2021).
- [196] M. Wilczek, O. Kamps, and R. Friedrich, Lagrangian investigation of two-dimensional decaying turbulence, *Physica D: Nonlinear Phenomena* **237**, 2090 (2007).
- [197] Y. Choi, Y. Park, and C. Lee, Helicity and geometric nature of particle trajectories in homogeneous isotropic turbulence, *International Journal of Heat and Fluid Flow* **31**, 482 (2010).
- [198] R. Pandit, D. Banerjee, A. Bhatnagar, M.-E. Brachet, A. Gupta, D. Mitra, N. Pal, P. Perlekar, S. Ray, V. Shukla, and D. Vincenzi, An overview of the statistical properties of two-dimensional turbulence in fluids with particles, conducting fluids, fluids with polymer additives, binary-fluid mixtures, and superfluids, *Physics of Fluids* **29**, 111112 (2017).
- [199] W. Braun, F. Lillo, and B. Eckhardt, Geometry of particle paths in turbulent flows, *Journal of Turbulence Volume* **7** (2006).
- [200] L. Bentkamp, T. Drivas, C. Lalescu, and M. Wilczek, The statistical geometry of material loops in turbulence, *Nature Communications* **13**, 2088 (2022).
- [201] H. Xu, N. Ouellette, and E. Bodenschatz, Curvature of lagrangian trajectories in turbulence, *Physical review letters* **98**, 050201 (2007).
- [202] A. Gupta, D. Mitra, P. Perlekar, and R. Pandit, Statistical properties of the intrinsic geometry of heavy-particle trajectories in two-dimensional, homogeneous, isotropic turbulence, (2014).
- [203] B. Kadoch, D. del Castillo-Negrete, W. Bos, and K. Schneider, Lagrangian statistics and flow topology in forced two-dimensional turbulence, *Physical review. E, Statistical, nonlinear, and soft matter physics* **83**, 036314 (2011).
- [204] N. Sakaki, T. Maruyama, and Y. Tsuji, Statistics of the lagrangian trajectories' curvature in thermal counterflow, *Journal of Low Temperature Physics* (2022).
- [205] N. Ouellette and J. Gollub, Dynamic topology in spatiotemporal chaos, *Department of Physics Papers* **20** (2008).
- [206] N. Ouellette and J. Gollub, Curvature fields, topology, and the dynamics of spatiotemporal chaos, *Physical review letters* **99**, 194502 (2007).
- [207] X. He, S. Apte, K. Schneider, and B. Kadoch, Angular multiscale statistics of turbulence in a porous bed, *Physical Review Fluids* **3** (2018).
- [208] A. Bhatnagar, A. Gupta, D. Mitra, P. Perlekar, M. Wilkinson, and R. Pandit, Deviation-angle and trajectory statistics for inertial particles in turbulence, *Physical Review E* **94** (2016).
- [209] B. Kadoch, D. del Castillo-Negrete, W. J. T. Bos, and K. Schneider, Transport, flow topology and lagrangian conditional statistics in edge plasma turbulence 10.48550/ARXIV.2205.07135 (2022).
- [210] J. Pratt, A. Busse, and W.-C. Müller, Lagrangian statistics for dispersion in magnetohydrodynamic turbulence, *Journal of Geophysical Research: Space Physics* **125** (2020).
- [211] K. Alards, H. Rajaei, R. Kunnen, F. Toschi, and H. Clercx, Directional change of tracer trajectories in rotating Rayleigh-Bénard convection, *Physical Review E* **97** (2018).

- [212] L. Kaufman and P. Rousseeuw, Partitioning around medoids (program PAM) (1990) pp. 68 – 125.
- [213] M. Birjukovs, P. Zvejnieks, T. Lappan, M. Klevs, S. Heitkam, P. Trtik, D. Mannes, S. Eckert, and A. Jakovics, Particle tracking velocimetry and trajectory curvature statistics for particle-laden liquid metal flow in the wake of a cylindrical obstacle 10.48550/ARXIV.2206.11033 (2022).
- [214] M. Klevs, M. Birjukovs, P. Zvejnieks, and A. Jakovics, Dynamic mode decomposition of magnetohydrodynamic bubble chain flow in a rectangular vessel, *Physics of Fluids* **33**, 083316 (2021).
- [215] J. Roenby, H. Bredmose, and H. Jasak, IsoAdvector: Geometric VOF on general meshes (2018).
- [216] J. Roenby, H. Bredmose, and H. Jasak, A computational method for sharp interface advection, *Royal Society Open Science* **3** (2016).
- [217] I. Mezic, Spectral properties of dynamical systems, model reduction and decompositions, *Nonlinear Dynamics* **41**, 309 (2005).
- [218] S. Le Clainche and J. M. Vega, Higher order dynamic mode decomposition, *SIAM Journal on Applied Dynamical Systems* **16**, 882 (2017), publisher: Society for Industrial and Applied Mathematics.
- [219] S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley, Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition, *Experiments in Fluids* **57**, 42 (2016), version: 3, 1507.02264.
- [220] S. L. Brunton, B. W. Brunton, J. L. Proctor, E. Kaiser, and J. N. Kutz, Chaos as an intermittently forced linear system, *Nature Communications* **8**, 19 (2017).
- [221] B. Loring, H. Karimabadi, and V. Rortershteyn, A screen space GPGPU surface LIC algorithm for distributed memory data parallel sort last rendering infrastructures (2014).
- [222] J. Fröhlich, S. Schwarz, S. Heitkam, C. Santarelli, C. Zhang, T. Vogt, S. Boden, A. Andruszkiewicz, K. Eckert, S. Odenbach, and S. Eckert, Influence of magnetic fields on the behavior of bubbles in liquid metals, *The European Physical Journal Special Topics* **220**, 167 (2013).
- [223] M. Birjukovs, N. Shevchenko, and S. Eckert, An image processing pipeline for in-situ dynamic X-ray imaging of directional solidification of metal alloys in thin cells (2022).
- [224] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, Image denoising with block-matching and 3D filtering, *Proceedings of SPIE - The International Society for Optical Engineering* **6064**, 354 (2006).
- [225] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, Image denoising by sparse 3-D transform-domain collaborative filtering, *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* **16**, 2080 (2007).
- [226] M. Lebrun, An Analysis and Implementation of the BM3D Image Denoising Method, *Image Processing On Line* **2**, 175 (2012), <https://doi.org/10.5201/ipo1.2012.1-bm3d>.
- [227] Y. Makinen, L. Azzari, and A. Foi, Collaborative filtering of correlated

- noise: Exact transform-domain variance for improved shrinkage and patch matching, *IEEE Transactions on Image Processing* **PP**, 1 (2020).
- [228] L. I. Rudin, S. Osher, and E. Fatemi, Nonlinear total variation based noise removal algorithms, *Physica D: Nonlinear Phenomena* **60**, 259 (1992).
- [229] J. Canny, A computational approach to edge detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**, 679 (1986).
- [230] Brightnessequalize (2017).
- [231] C. Harris and M. Stephens, A combined corner and edge detector, *Proceedings 4th Alvey Vision Conference* **1988**, 147 (1988).
- [232] J. Sánchez, N. Monzón, and A. Salgado, An analysis and implementation of the Harris corner detector, *Image Processing On Line* **8**, 305 (2018).
- [233] Morphologicalbranchpoints (2010).
- [234] C. Tomasi and R. Manduchi, Bilateral filtering for gray and color images (1998) pp. 839–846.

Acknowledgments

The author wishes to emphasize that this thesis and all the results presented herein are not only the author's achievements. Rather, the thesis represents a milestone in a collective effort in which the author plays an important role. Without the support of the collaborators, most of what has been reported here would not have been accomplished. In no particular order, the author would like to acknowledge the colleagues whose input has been instrumental to success.

The author expresses profound gratitude to Andris Jakovics (UL), Pavel Trtik (PSI), Knud Thomsen (PSI) and Sven Eckert (HZDR), all of whom were the kind of mentors one is thankful to have encountered and wishes were there for every aspiring PhD student. Andris Jakovics has been there from the very beginning as the MSc and PhD thesis supervisor and a colleague at INM. The author is grateful to Andris for *many* things, but especially for allowing and fostering the author's creative freedom and independence throughout all the joint projects at INM, and in particular over the course of this thesis. A lesser known fact, perhaps, is that the thesis was originally supposed to have been devoted to the application of the lattice Boltzmann method (LBM) to MHD bubble flow modelling. However, when the author pointed out that it was more scientifically "profitable" and practical to instead focus on experimental and image & data processing methods while relying on the already existing numerical approaches, in order to lay the desperately missing groundwork for future physics-focused papers and next-generation experimental methods, Andris trusted the author with taking the research in a radically different direction, despite LBM being a mandatory part of an ongoing project. This thesis is proof that trust has paid off, and is perhaps a good example of what happens when senior scientists and supervisors respect the vision of their students, for which the author is very thankful. Of course, Andris has also introduced the author to many key players in the field of research to which this thesis belongs.

Due to previous collaborations between Andris Jakovics, Mihails Scepanskis (UL) and PSI, the author had the honor of meeting Pavel Trtik and Knud Thomsen. Pavel and Knud introduced the author to the world of neutron radiography and the sheer scale of research done at PSI, which was a turning point for both the thesis and the author's career. The author owes Pavel and Knud for the very first neutron imaging campaign at the PSI NEUTRA neutron beamline, which was allotted from personal time at the instrument, which was a gesture of trust and investment in itself. It was also Pavel and Knud who were ever encouraging during the first stages of the thesis, on the way to first major publications. It will not be

an exaggeration to say that much of the paper writing skill improvement can be attributed to Pavel's and Knud's mentorship and advice. Pavel has also contributed a great amount of time and effort to help prepare for and conduct the experiments at NEUTRA. The author met Sven Eckert at a later point, during one of the MHD-themed conferences – this meeting soon resulted in the first collaboration between the author and HZDR, which is central to both the success of this thesis and the author's professional growth. The author is deeply grateful to Sven for introducing him to the other colleagues at HZDR, as well as to its MHD department, which at this point the author considers almost like home away from home. Much credit is also due to Sven for helping the author further improve both paper and proposal writing skills. The author would like to thank Pavel, Knud and Sven for initiating and supporting research visits (~ 1 year in total) and remote contracts to both PSI and HZDR, which have significantly boosted the rate of research by the author, and is also thankful to Andris who helped make the transitions between these visits smooth and comfortable. The author also thanks Andris, Pavel, Knud and Sven for the provided financial support throughout the years that this thesis took to complete.

Of course, much is owed to fellow Latvian researchers at INM. Although Mihails Scepaniskis moved on before INM became, well, INM, it was him who originally introduced the author to the (former) lab. In addition, much of the work on particle-laden flow imaging is based on the first experiments performed by Mihails at PSI in collaboration with Pavel and Knud. The author would also like to recognize very significant initial support from Valters Dzelme (UL) who not only provided the first version of the numerical model for MHD bubble flow, but also greatly assisted the first bubble flow imaging experiment at PSI, in particular hardware preparations. Further down the road, the author received tremendous support from Jevgenijs Telicko (UL), who played a key role in the improvement and implementation of an improved experimental setup for the second neutron imaging campaign.

The author feels that their luck over the course of their journey to this thesis is particularly clear, considering the opportunity to work with great students of their own. The author is glad to have met Martins Klevs (UL), Peteris Zvejnieks (UL), Ansis Ziemelis (UL) and Aleksandrs Jegorovs (UL), and it has truly been a pleasure to work and grow professionally alongside them. Martins made a very significant contribution as the main developer of the DMD code and methods, while Peteris is the main driving force for the MHT-X development, and has co-developed the bubble shape analysis code with the author. In addition, both have participated in the second neutron imaging campaign at PSI and have contributed their ideas to the experiments as well. Ansis has co-developed the bubble detection

code for the X-ray imaging experiments, implementing his own ideas on top of the author's work and performing necessary optimization. Finally, Aleksandrs, while not directly involved in the X-ray and neutron imaging experiments or code development, has been very helpful during discussions and, perhaps sometimes unbeknownst to him, has often provided useful insight or alternative viewpoints. Aleksandrs is currently working on *optical* imaging of bubble flow in Hele-Shaw liquid metal cells, a spin-off with respect to the experiments considered in this thesis, while Martins, Peteris and Ansis are continuing their work with DMD, MHT-X and X-ray image analysis, respectively. The author is deeply grateful to all of their students for the sense of camaraderie and mutually reinforced learning, and for their resilience in the face of a sometimes remarkably heavy workload and university studies.

Many thanks go out to other colleagues from PSI, in particular to Anders Kaestner who provided access to the ICON beamline, lent their neutron imaging expertise and assistance during experiments, and provided many insights during discussions; Jan Hovind, who has supported the experiments at PSI at every turn; Dariusz Jakub Gawryluk, who manufactured and provided the author with the reference bodies used for experimental validation of image processing methods; Marino Corrado, Yohei Sato and Bojan Niceno, who have provided the author with an alternative numerical simulation framework for bubble flow modelling, and who are currently involved with the author in further research regarding the applications of DMD to bubble flow analysis. The author also wishes to express gratitude to their colleagues from HZDR, who were not only important contributors to this thesis, but have also generated a warm and productive atmosphere during the author's visits to HZDR. Specifically, the author is grateful to Tobias Lappan for providing the X-ray images of bubble flow and neutron images of particle flow in liquid metal, as well as for invaluable support during the X-ray imaging experiments at HZDR; Natalia Shevchenko, who has also been very supportive not only during the experiments, but has helped the author in many other ways during his stay at HZDR, in addition to being the first person to test and contribute improvement ideas for the image processing code developed for studying metal alloy solidification, as well as providing experimental data for the author to work with; Sascha Heitkam, for many lengthy and insightful discussions on particle flow physics and critically important advice and suggestions for papers; Megumi Akashi, for providing the data used to test MHT-X and the shape analysis code; Thomas Wondrak and Frank Stefani, for hosing Martins Klevs and Peteris Zvejnieks during their visits to HZDR, as well as for providing financial support and taking the time to engage in productive discussions; Max Sieger, Thomas Gundrum and Stefanie Sonntag for their assistance

with magnetometry experiments, as well as other matters; Henrik Schulz, for helping the author perform image processing using the hardware at the HZDR high-performance computing (HPC) center and assisting with troubleshooting; Sten Anders and Tom Weier for providing many insights into imaging techniques for multiphase flow.

The author would also like to thank Imants Buceniaks (UL) for assembling the MF systems used in experiments at PSI, HZDR and UL; Andrew Kao (University of Greenwich (UG), UK) and Catherine Tonry (UG) for their invaluable assistance with setting up and running *OpenFOAM*, *Elmer* and *EOF-Library* on HPC systems; Robert Zboray (Empa, Switzerland) and Tianshu Liu (Western Michigan University) for their time and discussions on multiphase flow imaging and image processing.

Many thanks also to Liene Bandeniece (UL), Inese Suija (UL) and Aida Roch (HZDR) for their help with logistics and formal matters, without which many endeavors would surely have ended in failure.

Finally, the author expresses heartfelt gratitude to their family, who have supported and propped up the author throughout this journey, at times exceedingly taxing both mentally and physically. The author is convinced they would not have made it through to the end if not for the people closest to them, and would like them to know that this work is their achievement just as much as the author's.