A comprehensive shape analysis pipeline for stereoscopic measurements of particulate populations in suspension

Author(s):
Rajagopalan, Ashwin K.; Schneeberger, Janik; Salvatori, Fabio; Bötschi, Stefan; Ochsenbein, David R.; Oswald, Martin R.; Pollefeys, Marc; Mazzotti, Marco

Publication date:
2017-11

Permanent link:
https://doi.org/10.3929/ethz-b-000192659

Originally published in:
Powder Technology 321, https://doi.org/10.1016/j.powtec.2017.08.044

Powder Technology 321 (2017) 479–493

Contents lists available at ScienceDirect

Powder Technology

journal homepage: www.elsevier.com/locate/powtec

A comprehensive shape analysis pipeline for stereoscopic measurements of particulate populations in suspension

Ashwin Kumar Rajagopalan[a,1], Janik Schneeberger[a,1], Fabio Salvatori[a], Stefan Bötschi[a], David R. Ochsenbein[a,2], Martin R. Oswald[b], Marc Pollefeys[b], Marco Mazzotti[a,*]

[a] Institute of Process Engineering, Sonneggstrasse 3, ETH Zürich, CH-8092 Zürich, Switzerland
[b] Computer Vision and Geometry Group, Universitätstrasse 6, ETH Zürich, CH-8092 Zürich, Switzerland

ARTICLE INFO

Article history:
Received 10 June 2017
Received in revised form 10 August 2017
Accepted 13 August 2017
Available online 16 August 2017

Keywords:
Crystallization
Particle size and shape distribution
Imaging-based particle sizing technique
Volumetric visual hull reconstruction
Automated crystal shape classification

ABSTRACT

A state-of-the-art, compact optomechanical setup coupled with an image analysis routine to measure multi-dimensional particle size and shape distributions (*n*D PSSDs) for crystallization processes is presented. A novel image processing pipeline to process the raw images from the cameras is presented. The pipeline consists of a stereoscopic camera calibration model, adaptive background subtraction, particle contour matching, and 3D reconstruction of the segmented crystals. The reconstructed crystals are subjected to a supervised shape classification strategy, which categorizes each detected crystal into spheres, needles, quasi-equant particles, platelets and non-convex particles. Additionally, a high-speed image capture mode, capable of monitoring processes with fast kinetics, is presented. The device discussed in this work is subjected to an experimental campaign, to validate size measurements, characterize steady state, and confirm repeatability of measurements to affirm and assess the non-invasive nature of the setup on the measurement. An experiment aimed at evaluating the enhancement in the proposed image analysis pipeline performance, more specifically the automatic shape classification, is further conducted. Finally, a dissolution process is monitored using a stereoscopic imaging setup for the first time, and the size and shape evolution of the population in a growth and dissolution phase is monitored for about 18 h.

1. Introduction

In crystallization processes, the particle size and shape of the final product is of key importance as it influences the downstream processing operations such as filtration, drying, and tableting. Crystals exhibit different shapes depending on the crystal habit, and an accurate characterization of shape is critical in the design and control of such processes. Often, commercially available crystallization process characterization tools condense shape related information of crystals into a single characteristic length [1]. Hence, a one dimensional particle size distribution is obtained leaving out the shape information of the crystals. However, due to the variety of shapes exhibited by crystals, in order to accurately quantify the population of crystals, a *multidimensional* particle size and shape distribution (PSSD), rather than a PSD, would be preferable.

As we have recently reported [2], commercially available sizing tools that rely on the assumption of a single characteristic length, such as focused beam reflectance measurement (FBRM), laser diffraction (LD), Coulter counter (CC), as well as monoscopic imaging tools, are prone to errors and misleading effects for particles that are non-spherical. For example, FBRM provides a one-dimensional chord length distribution (CLD) which is difficult to interpret; the transformation from the CLD to PSD is an ill-posed problem and can in fact simply not be done without additional tools [3]. Multi-projection imaging systems [4–7], have been proposed as remedy as they are able to tackle shape-related issues rather satisfactorily, thus reducing the ill effects encountered by commercial sizing tools for non-spherical particles.

In an earlier publication, the successful implementation of a dual projection imaging device using a stereoscopic camera setup with a sapphire glass based flow channel was demonstrated [8]. The setup has been used to monitor growth and agglomeration of βL-Glutamic acid [9–12], and the measurement device enabled modeling the phenomenon with both size and shape information, which would have been difficult with other process characterization tools. Even though the image analysis routine implemented in previous works

* Corresponding author.
E-mail address: marco.mazzotti@ipe.mavt.ethz.ch (M. Mazzotti).

[1] A.K.R. and J.S. contributed equally to this work.

[2] Current address: Janssen: Pharmaceutical Companies of Johnson & Johnson, Hochstrasse 201, CH-8200 Schaffhausen, Switzerland.

http://dx.doi.org/10.1016/j.powtec.2017.08.044
0032-5910/

was able to distinguish primary particles and non-convex particles like agglomerates, the need to implement a more accurate and robust image analysis routine was identified. For example, the classification of different primary particles was performed only based on the contour boundary pixels of the particles. But considering the boundary pixels alone, not utilizing the available information in its entirety, would definitely hinder the characterization of complex shapes in a process. Although particle shape characterization by means of supervised and unsupervised learning algorithms is not entirely novel [13], methods for shape categorization based on reconstructed volumetric 3D models have never been applied to the best of our knowledge in the past.

1.1. Contributions

A compact version of the stereoscopic camera setup, ≈5 times smaller than our previous setup, was engineered. A novel calibration procedure for the alignment of the multi-camera setup is proposed, which was missing in our previous works. In contrast to our previous work, the proposed image segmentation algorithm accounts for visual artifacts which finally leads to an improved reliability of measured particle size distributions over time. Furthermore, the novel shape classification of volumetrically 3D reconstructed particles leads to a better shape approximation and segregation of the particles observed by the stereoscopic camera setup. In particular, the more accurate particle shape description allows for a more precise and fine-grained classification of particles which allows for better real-time control of crystallization processes.

1.2. Outline

First, a detailed hardware overview of the measurement device is provided in Section 2. Second, in Section 3, a comprehensive theoretical study of the basic concepts of camera calibration, image analysis, and 3D reconstruction is presented. The measurement device validation and the application of the new particle reconstruction and shape classification is reported in Section 4. Finally, in Section 5, a discussion on the improvements of the measurement device is given along with concluding remarks.

2. Measurement device

A major challenge of imaging-based particle sizing techniques is the dependence of the observed particle size on the orientation of the particle under inspection with respect to the camera. These orientation-related issues can be mitigated by means of a multi-camera setup which provides particle projections from different angles. The previously published stereoscopic imaging setup, henceforth referred to as FTC [8], uses a dual-projection technique capable of merging particle size information provided by two cameras into *n*D PSSDs. This feature yields a more accurate measurement than what single-view setups can provide [2]. However, the major drawback of the FTC is its bulky mechanical design (126 × 126 × 90 cm), making it vulnerable to vibrations during image acquisition. Moreover, the Xenon flashes employed required additional optics to provide collimated light; also, a square flow channel assembled by gluing four sapphire glass windows held by a brass holder was used, making maintenance of the device cumbersome.

Based on the issues described above, a more compact version of the optomechanical setup (80 × 74 × 42 cm) was developed with the goal of overcoming the problems associated with the FTC. The smaller setup described, henceforth referred to as *dual imaging system for crystallization observation* (*μ*-DISCO), fits into a standard laboratory hood and is less vibration-susceptible during operation. A schematic of the new setup is shown in Fig. 1; it consists of two monochrome CMOS cameras (Point Grey Research, Canada) in an orthogonal configuration with telecentric optics (Opto Engineering, Italy) resulting in an orthographic projection with very low spatial distortions (<0.1%). The camera-lens system provides a field of view (FOV) of 2.41 × 2.02 mm at a nominal magnification of 3.5×. Two high-power, telecentric LED illuminators (Opto Engineering, Italy), which emit collimated chief rays parallel to the optical axis produce high contrast silhouettes of particles passing through the flow channel. The whole setup is mounted on an optical rail cross-construction. Manual XYZ-translation stages (Newport Corporation, USA) that allow high precision alignment of the two cameras, and a rotation stage (Newport Corporation, USA) that allows orienting the parallel illumination beam, are used. A microcontroller (Atmel, USA) running in-house software provides an external trigger signal, which enables a synchronized image acquisition from the cameras.

The introduction of the new cameras allows to operate the *μ*-DISCO either in a *standard mode*, with a constant low frame rate (1–7 Hz), or in a *burst mode*, with higher frame rates up to 75 Hz. The burst mode is particularly useful for capturing processes with fast dynamics, such as dissolution. In standard mode, the *μ*-DISCO can be operated either online, that is, the image processing is performed in real-time, or offline, where the image processing is performed after

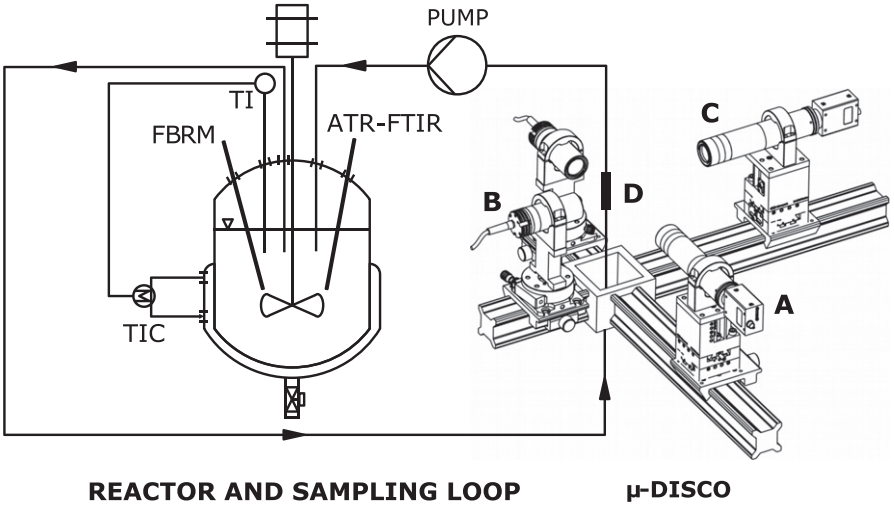

Fig. 1. Schematic of the new dual imaging system for crystallization observation (*μ*-DISCO). The suspension flowing from the reactor through the flow channel **(D)** is back-light illuminated using two telecentric illuminators **(B)**. The suspension is photographed using two digital cameras **(A)** with telecentric optics **(C)**. The camera, lens and the illuminator system are mounted orthogonally on an optical rail construction.

the image acquisition for the entire measurement period. The cameras are connected to high-speed USB 3.0 ports and controlled with custom drivers implemented using the FlyCapture SDK (Point Grey Research, Canada).

The suspension is sampled from the reactor using a peristaltic pump (Watson-Marlow, Inc., UK) and flows through a single piece quartz square channel (FireflySci, Inc., USA) embedded in an insulated sampling loop. The cross-section of the flow channel is 2×2 mm and incorporates the two transitions between the circular tubing and the square channel at both ends. A one-piece flow channel is superior to the old channel, because of a reduced risk of leakage and a completely eliminated risk of dissolution of the glue used to hold the channel.

3. Theoretical background

3.1. Camera calibration

The calibration of camera-based measurement setups is necessary to obtain accurate metric size information from images. It involves finding optical characteristics of the camera that affect the image formation (*intrinsic* parameters), as well as the position and orientation of the camera coordinate systems in space with respect to a world coordinate system (*extrinsic* parameters). Camera calibration has been a research subject for a long time [14–16], and open-source camera calibration software is available [17]. Standard off-the-shelf calibration software usually assumes the cameras to fulfill the so-called *pinhole camera model*[18], which however is not applicable when dealing with telecentric lenses since these types of lenses are designed to produce an orthographic projection of the object imaged [19]. Our calibration procedure is based on the following simplifying assumptions:

1. The optical axes of the two lenses lie on parallel planes and are orthogonal to each other;
2. Nonlinear lens distortions are negligible due to the particular telecentric lenses used.

Based on the mathematical model of a telecentric camera presented by Huiyang et al. [20], a custom camera model is derived. The world coordinate system in homogeneous form is denoted by $\mathcal{O}_\mathrm{w} = [x_\mathrm{w}, y_\mathrm{w}, z_\mathrm{w}, 1]^\mathrm{T}$, and its origin is placed in the center of the square glass channel (shown in Fig. 2). The two camera coordinate systems $\mathcal{O}_\mathrm{A} = [x_\mathrm{A}, y_\mathrm{A}, z_\mathrm{A}, 1]^\mathrm{T}$ and $\mathcal{O}_\mathrm{B} = [x_\mathrm{B}, y_\mathrm{B}, z_\mathrm{B}, 1]^\mathrm{T}$ of camera A and B, respectively, are shown in Fig. 2.

The local image plane coordinate system of each camera, $[u_\mathrm{A}, v_\mathrm{A}]$ and $[u_\mathrm{B}, v_\mathrm{B}]$, is placed at the upper-right corner of each image plane. Given that the location of both cameras is fixed by the optical rail, the remaining degrees of freedom of each camera include one rotation axis, around z_A and z_B, respectively, denoted by the roll angles, ϕ_A and ϕ_B, and three translational degrees of freedom for each camera along axes x_j, y_j, and z_j, denoted with $d_{j,x}$, $d_{j,y}$, and $d_{j,z}$, respectively, where the subscript j identifies the camera source, i.e., $j \in \{\mathrm{A}, \mathrm{B}\}$. The distance along z_j is determined by the nominal working distance of the telecentric lenses, which defines the distance between the front of the lens and the object when the lens is correctly focused ($d_{\mathrm{A},z} = d_{\mathrm{B},z} = 132.3$ mm). The image formation of the camera-lens system is modeled by a *weak perspective projection* [18], which is an orthographic projection followed by isotropic scaling, denoted by the scaling factors α_A and α_B, respectively. Following the notation proposed by Huiyang et al. [20], the relationship between a 3D point $[x_\mathrm{w}, y_\mathrm{w}, z_\mathrm{w}, 1]^\mathrm{T}$ in the world coordinate system $\mathcal{O}_\mathrm{w}$ and its image projection $[u_j, v_j, 1]^\mathrm{T}$ is given in homogeneous form as

$$\begin{bmatrix} u_j \\ v_j \\ 1 \end{bmatrix} = \mathbf{K}_j \begin{bmatrix} \mathbf{R}_j & \mathbf{T}_j \\ \mathbf{0} & 1 \end{bmatrix} \begin{bmatrix} x_\mathrm{w} \\ y_\mathrm{w} \\ z_\mathrm{w} \\ 1 \end{bmatrix} \tag{1}$$

where, the intrinsic camera matrices, $\mathbf{K}_j$, are given by

$$\mathbf{K}_j = \begin{bmatrix} \alpha_j^{-1} & 0 & 0 & \tilde{u}_0 \\ 0 & \alpha_j^{-1} & 0 & \tilde{v}_0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2}$$

and the rotation matrices, $\mathbf{R}_j \in \mathbb{R}^3$, are given by

$$\mathbf{R}_\mathrm{A} = \begin{bmatrix} -\sin(\phi_\mathrm{A}) & 0 & -\cos(\phi_\mathrm{A}) \\ -\cos(\phi_\mathrm{A}) & 0 & \sin(\phi_\mathrm{A}) \\ 0 & 1 & 0 \end{bmatrix} \quad \text{and}$$
$$\mathbf{R}_\mathrm{B} = \begin{bmatrix} 0 & \sin(\phi_\mathrm{B}) & -\cos(\phi_\mathrm{B}) \\ 0 & -\cos(\phi_\mathrm{B}) & -\sin(\phi_\mathrm{B}) \\ -1 & 0 & 0 \end{bmatrix} \tag{3}$$

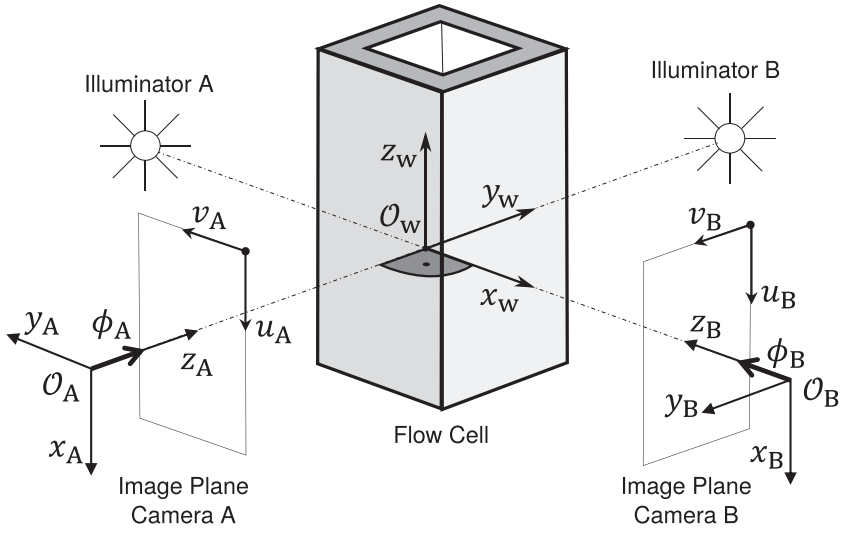

Fig. 2. Coordinate system convention used for the calibration of the stereoscopic camera setup.

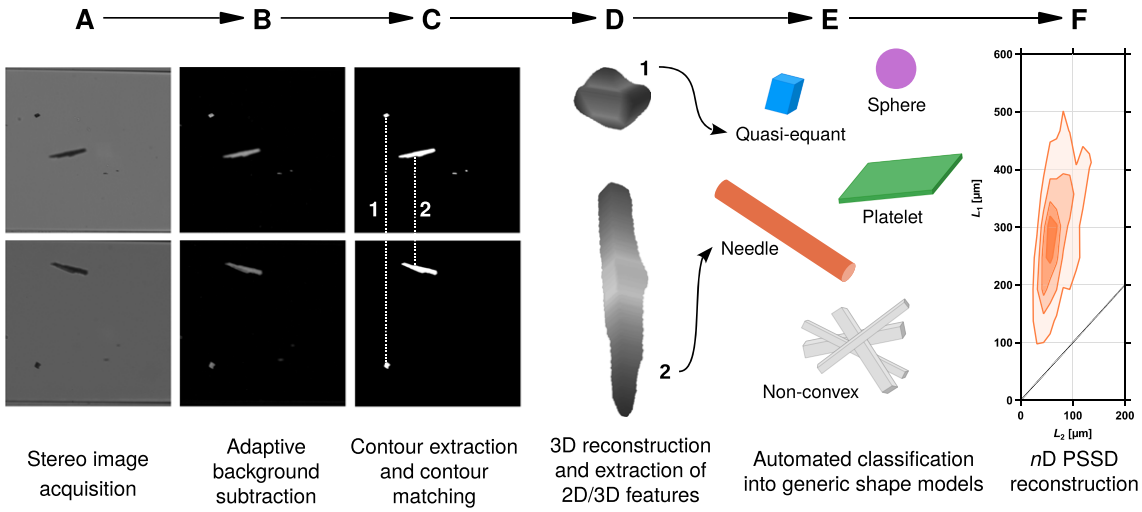

Fig. 3. Illustration of the multi-step image processing approach. This work includes the adaptive background subtraction **(B)**, a contour matching algorithm **(C)**, 3D shape reconstruction **(D)**, and the automated shape classification of reconstructed particles **(E)**. The result of this procedure is an *n*D PSSD for each generic shape model.

respectively. Here, $[\tilde{u}_0, \tilde{v}_0]^{\mathrm{T}} = [1224, 1024]^{\mathrm{T}}$ denotes the principal point in pixel coordinates, i.e., the point where the principal camera axis intersects the image plane. The image resolution of both cameras is 2448 × 2048 pixels. The translation vectors are given by $\mathbf{T}_j = [d_{j,x}, d_{j,y}, d_{j,z}]^{\mathrm{T}}$. The relative shift along the z_{w} axis between the two camera coordinate systems is defined as $\epsilon = d_{\mathrm{B},x} - d_{\mathrm{A},x}$. In the ideal case, ϵ should be equal to zero after the mechanical alignment, thus indicating that a projected 3D point in $\mathcal{O}_{\mathrm{w}}$ shares one common coordinate in the image planes of the cameras. However, in reality, $\epsilon \neq 0$ since mechanical alignment cannot be attained with sufficient precision; the exact value of the relative shift, ϵ, has to be found during the extrinsic calibration. The distances $d_{j,y}$ are mechanically adjusted, so as that the cell is centered in both images.

In summary, the five parameters to be found by means of the calibration procedure are α_{A} and α_{B} (*intrinsic* parameters), and ϕ_{A}, ϕ_{B}, and ϵ (*extrinsic* parameters).

3.1.1. Intrinsic calibration

A planar-object-based calibration method is employed to find each cameras intrinsic parameters. In order to do so, a calibration plate with a checkerboard pattern of 0.2 × 0.2 mm (Opto Engineering, Italy) is mounted perpendicular to the optical axis and an image of the checkerboard calibration pattern is captured. After the manual selection of 9 × 9 checkerboard corner points, a corner-refinement routine [17] is applied, and the exact location of the corner points on the image plane is found. The distances between neighbouring corner pixel locations are measured and plotted against the corresponding nominal checkerboard size of the pattern. The scaling factor is then taken as the slope of the line that best fits the data distribution. The scaling factors for the cameras of the setup presented in this article are $\alpha_{\mathrm{A}} = 0.993 \pm 0.002$µm/pixel, and $\alpha_{\mathrm{B}} = 0.991 \pm 0.002$µm/pixel.

3.1.2. Extrinsic calibration

The extrinsic calibration procedure is carried out in two steps, namely a rough mechanical alignment of the optical components followed by a high-precision software calibration. A thin nylon filament is used as a reference object for the rotational alignment of the two cameras; one end of the filament is attached to the frame of the setup and the other end is attached to a weight so as the filament is aligned with the direction of gravity and centered in the FOV of the cameras. After an initial, coarse mechanical alignment of the cameras, multiple images of the shadow created by the nylon filament are captured. A simple image gradient-based algorithm is used to determine the exact location of a set of center points of the nylon filament with sub-pixel accuracy. The rotation angles are calculated from the slope of a linear function found using linear regression of the set of center points. The values for the roll angles determined in this way are $\phi_{\mathrm{A}} = -0.86 \pm 0.02^{\circ}$ and $\phi_{\mathrm{B}} = 0.15 \pm 0.01^{\circ}$.

The remaining extrinsic parameter, ϵ, is determined by means of a statistical calibration measurement involving spherical latex beads (CC Size Standard L90, Beckman Coulter, USA) with a nominal diameter of 90 µm. For this, the bounding box centroid difference along the z_{w}-direction, $\Delta x = \hat{x}_{\mathrm{B}} - \hat{x}_{\mathrm{A}}$, is reported for every matched contour pair, where $\hat{x}_j$ are the x-components of the particle bounding box centroid. Assuming a perfect rotational correction of the cameras, the mean of Δx over all contour-matches can be taken as an approximation for ϵ.

3.2. Image processing and analysis

A multi-step, in-house developed image processing approach (illustrated in Fig. 3) for extracting contours of particle silhouettes (Sections 3.2.1 and 3.2.2), reconstruction of particle shapes in 3D (Section 3.2.4), and shape classification of the particles (Section 3.3) is presented.

3.2.1. Adaptive background subtraction

After the acquisition of a stereoscopic image pair, an adaptive background subtraction step [21,22] is applied to both images in order to remove stationary content from the images such as dust, dirt or scratches on the flow cell, which tend to increase as the cell ages. The idea behind this technique is to generate a foreground mask of pixels belonging to moving objects (see III in Fig. 4) by estimating the stationary background of the scene (see II in Fig. 4), which can be subtracted from successive images. A simple yet effective way to model the background, $B(x, y, t_k)$, at the pixel location (x, y) and time step, t_k is used to compute the temporal mean over a series of images. Using this approach, the background is updated with a new frame, $I(x, y, t_k)$, according to the recursive update equation

$$B(x, y, t_k) = (1 - \gamma)\, B(x, y, t_{k-1}) + \gamma I(x, y, t_k) \tag{4}$$

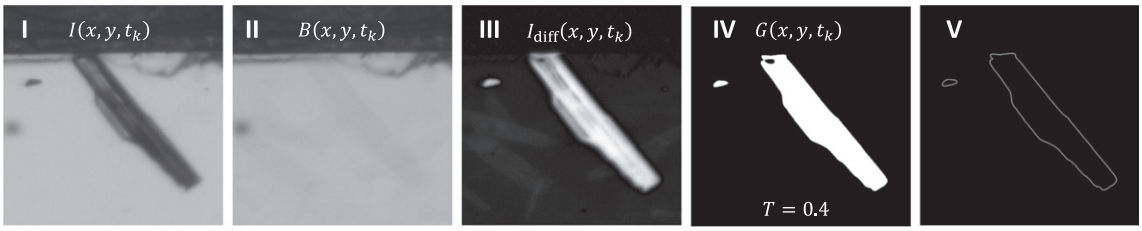

Fig. 4. **(I)**: new frame at time t_k showing a crystal; **(II)**: static background; **(III)**: difference frame (intensity normalized); **(IV)**: segmented difference frame; **(V)**: extracted contour of the crystal.

where, $\gamma \in [0,1]$ is the *learning rate* of the update equation, which determines the contribution from a new frame to the background. For our application, we used a small value of $\gamma = 0.05$. The subtraction of a new frame with the estimated background model results in a difference frame, $I_{\text{diff}}(x,y,t) = |I(x,y,t_k) - B(x,y,t_k)|$, containing only moving objects, i.e., objects that appear new in the image.

3.2.2. Image segmentation

This processing step identifies whether pixels in the image either depict a crystal (shown as bright pixels in image IV of Fig. 4) or if they belong to the background (shown as dark pixels). Given the difference image $I_{\text{diff}} : \Omega \subset \mathbb{R}^2 \times T \rightarrow [0,1]$ from the previous background subtraction step, we compute a binary image $G : \Omega \subset \mathbb{R}^2 \times T \rightarrow \{0,1\}$ that labels every pixel as either foreground or background. $T \subset \mathbb{R}$ denotes the time domain. Being the simplest and fastest segmentation algorithm, we use thresholding and classify all pixels by means of a global threshold $\tau \in [0,1]$ as

$$G(x,y,t_k) = \begin{cases} 1 & \text{if } \; I_{\text{diff}}(x,y,t_k) > \tau \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

In our application, the global threshold value, is found using a photometric calibration involving spherical latex beads (Beckman Coulter, USA) with a nominal diameter of 90μm. It is a well-known fact that thresholding is prone to noise because the classification of each pixel is independent of its direct neighbors. Therefore, the segments might not be contiguous or contain holes. In order to cope with this problem, a median filter with window size of 5 × 5 pixels is applied to $G(x,y,t_k)$ after segmentation, followed by a morphological opening and closing operation with a rectangular structuring element with the size of 13 pixels [23]. These operations reduce noise and close thin concavities in the binary mask. In a successive step, the contour of the particle silhouettes is extracted using a boundary tracing algorithm [24] (illustrated in image V of Fig. 4). Contours containing less than 50 connected pixels are considered to be below the optical resolution of the camera-lens system and are thus excluded from further analysis. Contours intersecting the image boundaries are discarded as they cannot provide an accurate size measurement of the corresponding particle. All image processing operations are implemented using the open source computer vision library OpenCV [25].

In the course of this evaluation, also a variational segmentation approach known as Total Variation Segmentation (TVSeg) [26–28] was extensively tested and carefully evaluated for the application of this work. This method overcomes the drawbacks of simple segmentation techniques and produces excellent segmentation outcomes also for images with intensity inhomogeneity and Poisson noise [27]. However, when compared with binary thresholding, the computation time of TVSeg is significantly higher which impedes its application for real-time measurements. Hence, it was decided to implement binary thresholding for the context of this work.

3.2.3. Stereoscopic contour matching

The aim of this processing step is to find spatial correspondences between silhouettes on the two stereoscopic images that originate from the same crystal. Assuming perfect alignment of both cameras and a segmentation without errors, the matching would be performed by comparing the upper and lower coordinates of the contours from the two stereoscopic images along the common z_w-coordinate. In reality, the particles from the two stereoscopic images do not share identical upper and lower coordinates due to calibration errors, mechanical micro-vibrations, or image segmentation errors.

A particle matching algorithm based on the Euclidean norm of the upper and lower coordinates of the bounding box, and centroid of the different particles with a fixed matching threshold from the two stereoscopic projections was reported earlier [6,8]. Instead of assigning a fixed threshold for the contour matching procedure, a statistical assignment approach to find the correspondence between contoured particles from the two segmented images is employed. The matching is performed based on conditions imposed on (i) the absolute difference in the z_w-coordinate of the contour centroids, and (ii) absolute or relative differences in the upper and lower z_w-coordinates of the bounding box of the contour.

In the first step, pairs of contours that satisfy the centroid threshold from the two matched stereoscopic images are checked for difference in the z_w-coordinate of the bounding box. On the one

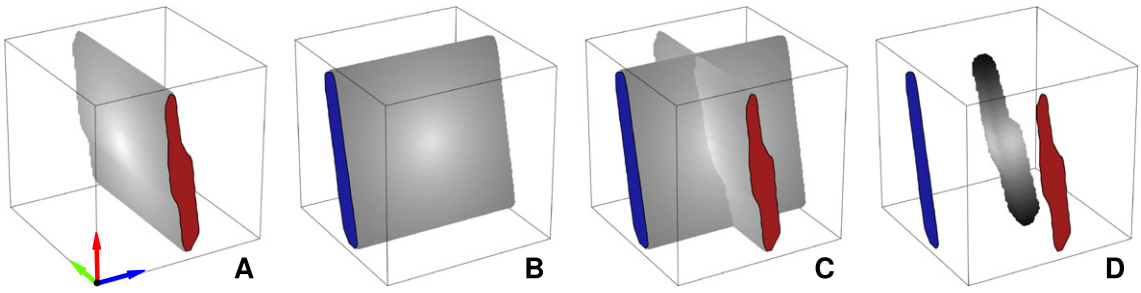

Fig. 5. Volume intersection method used for the 3D reconstruction of particles being imaged by the two cameras. **(A)** and **(B)** show the volumes spanned by the two silhouettes, $\mathcal{V}_A$ (red silhouette) and $\mathcal{V}_B$ (blue silhouette); **(C)** illustrates the intersection of $\mathcal{V}_A$ and $\mathcal{V}_B$; **(D)** shows the final reconstructed visual hull $\mathcal{V} = \mathcal{V}_A \cap \mathcal{V}_B$.

hand, if a pair of contour satisfies the threshold for absolute difference of the z_w-coordinate, it is considered to be matched. While on the other hand, if the absolute difference threshold is not met, but the relative difference threshold is satisfied, the contour pair is considered to be matched and is subjected to further image processing steps.

3.2.4. Visual hull-based 3D reconstruction

For the classification of crystals based on shape, a 3D volumetric model of each crystal is required from which shape attributes can be extracted and used in the context of shape classification. Shape-from-Silhouette (SFS) [29,30] is a popular technique for the shape estimation of an object from a set of calibrated silhouette images. The reconstruction obtained with this method, henceforth referred to as *visual hull* (VH), is the intersection of the projective sets (silhouette cones) obtained by unprojecting all camera-aligned silhouette images into 3D space. Since object concavities can only be recovered if they are visible in at least one of the silhouettes, the VH can be regarded as an upper bound of the actual object shape.

Two methods for the construction of the VH are found in literature, namely volume- and surface-based methods. The first category includes methods that approximate the VH by collections of cubic elementary cells, so-called voxels [31], and the latter method estimates the surface of the VH as a polyhedron by intersecting silhouette cones [32]. Due to its efficient implementation, the volume-based approach is the selected method for the application in this work. The visual hull is approximated by a set of N points inside the VH, denoted by the data matrix, $\mathcal{V} = [x_1, \ldots, x_N] \in \mathbb{R}^{3\times N}$, where each column x_i represents a data point in $\mathbb{R}^3$ defining the centroid location of a cubic voxel with edge length Δs. With the two orthogonally aligned cameras, $\mathcal{V}$ can be computed by the volume intersection $\mathcal{V} = \mathcal{V}_A \cap \mathcal{V}_B$ of the two silhouette cones, $\mathcal{V}_A$ and $\mathcal{V}_B$, each representing a discretized volume spanned by the two projected silhouettes as illustrated in Fig. 5. The volume intersection is performed separately for each pair of silhouette contours that have been found in the previous contour matching step. This avoids many reconstruction artifacts such as so-called "ghosts" that appear when particles occlude each other in the camera views [33]. The accuracy of the approximation of the VH is inherently determined by the edge length of a voxel, Δs. An adaptive voxel size approach is pursued in order to keep the computation time for the reconstruction constant, independent of particle volume. As a consequence, larger particles are approximated by larger voxels and vice-versa. An upper-bound estimate of the volume of a reconstructed particle is given by $V_{VH} = N\Delta s^3$. A limitation of the exploited method in the given application is the dependence of the reconstructed VH on the orientation of the particle with respect to the cameras. With only two cameras, the shape of the VH might be ambiguous, e.g., a needle-like particle cannot be distinguished from a platelet-like particle when it is oriented in a 45°-angle with respect to both cameras (see Fig. 6). This ambiguity also affects the reconstruction of flat, platelet-like particles, where certain orientations can lead to non-platelet shaped reconstructions. These inherent ambiguities could be resolved by adding more cameras to the setup, providing additional silhouette projections of particles flowing through the cell.

3.3. Automated shape classification

The measurement of the particle size by means of imaging methods requires a size metric capable of adequately representing the characteristic size of the particle to be measured. Finding such a suitable metric requires prior 3D shape information of the particle before particle sizing is performed. In the context of this work, the task of classifying the shape of a crystal consists of assigning a generic shape model to each particle recorded by the stereoscopic camera setup, and hence categorizing crystals into shape groups that share certain geometric properties.

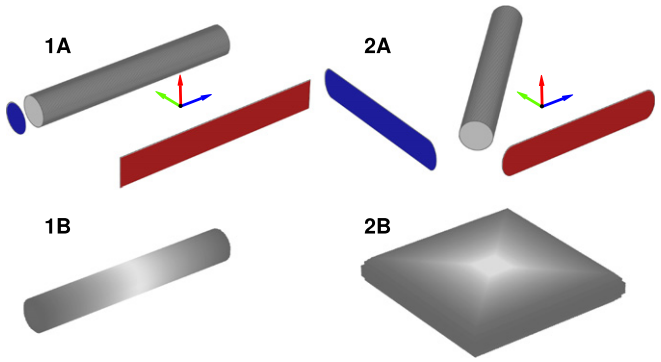

Fig. 6. Favourable orientation of a cylinder **(1A)** leads to a correct VH-reconstruction **(1B)**. In an ambiguous cylinder orientation **(2A)** one cannot reconstruct the true shape of the cylinder with only two cameras since both shapes **(2A)** and **(2B)** generate the same silhouettes. The virtual cameras are in both cases aligned along the green and blue coordinate direction.

3.3.1. Definition of generic shape models

In general, one can distinguish between physically based morphological particle models and generic particle models. The first class of models uses crystallographic data and aims at reflecting the morphology of an ideal crystal as accurately as possible [8]. In contrast, generic particle models aim at simplifying the shape of the crystals as much as is sensible for a given application. The faces of crystals are not necessarily respected using such an approach and only the macroscopic form is considered. Following the definition of Schorsch et al. [7], the generic model particles *sphere* and *needle* are distinguished, along with the new classes *quasi-equant*, *platelet*, and *non-convex*. Note that the non-convex class strongly resembles the class denoted by the term *agglomerate* in the work by Ochsenbein et al. to describe agglomerates of needle-like crystals [12]. In this article, however, the more general term *non-convex* is used since this label is restricted not only to the agglomeration of needle-like crystals. The *non-convex* label in the context of this work might also include overlapping particles which exhibit non-convex shape, thus making it difficult to distinguish them from real agglomerates.

The set of generic shapes distinguished in this work is illustrated in process step E of Fig. 3. Members of the *sphere* class are spherical objects (spheroids), which can be described with a diameter, L_1, only. Elongated, needle-like particles are described by the *needle* class whose characteristic dimensions are approximated by a cylindrical shape with length, L_1, and width, L_2, where $L_1 \gg L_2$. The class *quasi-equant* consists of regular, non-spherical particles, for which $L_1 \geq L_2 \geq L_3$ can be assumed, as for instance in the case of cuboid shaped crystals. The *platelet* class describes flat-shaped, tabular objects like platelet crystals, for which $L_1 \geq L_2 \gg L_3$ is assumed. The *non-convex* class is intended for the detection and classification of crystal agglomeration events yielding irregularly shaped particles.

3.3.2. 2D and 3D shape descriptors

The design and the selection of shape descriptors is of central importance to supervised learning approaches. A common way for characterizing particle shapes is to use 2D shape factors that are computed from the contour of a particle silhouette. Shape factors are dimensionless quantities that numerically describe the shape of a particle, independent of its size. Ochsenbein et al. used a set of shape factors for a supervised machine learning strategy using a non-linear Support Vector Machine (SVM) algorithm for the separation of primary (needle-like, spheres, cuboids, etc.) particles and

agglomerates [12]. Following up on their work, new descriptors are investigated in this article. We measure circularity (roundness) of a particle contour using the isoperimetric quotient (reciprocal of the P2A compactness measure [34]). In the context of our application, the isoperimetric quotient, Q_j, of a closed contour associated with camera j is defined as the ratio of the projected contour area $A_{\mathrm{proj},j}$ to the area of a circle with same perimeter as the contour, denoted by $P_{\mathrm{proj},j}$, i.e.,

$$Q_j = \frac{4\pi A_{\mathrm{proj},j}}{P^2_{\mathrm{proj},j}} \tag{6}$$

It is easy to show that Q_j belongs to the interval $(0,1]$ and becomes 1 for a perfect circle. In order to combine the circularity from both projected contours, the feature transform $\omega = Q_A Q_B$ is applied, where ω is referred to as the *circularity product* of a particle contour pair. Besides circularity, also perimeter-based convexity [35], κ_j, is calculated for each particle contour pair:

$$\kappa_j = \frac{P_{\mathrm{conv},j}}{P_{\mathrm{proj},j}} \tag{7}$$

where $P_{\mathrm{proj},j}$ and $P_{\mathrm{conv},j}$ are the perimeters of the silhouette contours and the convex hull of the silhouette, respectively, associated with camera j. The same concept for combining shape descriptors from multiple projections into a single-valued number is applied to the convexity measurement, i.e., $\beta = \kappa_A \kappa_B$ is referred to as the *convexity product*. 3D shape descriptors are extracted using an approach for the analysis of point clouds based on the covariance matrix of the VH. For this, a sub-set, $\hat{\boldsymbol{\mathcal{V}}} \subset \boldsymbol{\mathcal{V}}$ with $\hat{\boldsymbol{\mathcal{V}}} = [y_1, \ldots, y_n] \in \mathbb{R}^{3\times n}$, is sampled uniformly at random without replacement for each reconstructed particle. The sample covariance matrix $\boldsymbol{\Sigma} = \mathrm{Cov}(\hat{\boldsymbol{\mathcal{V}}})$ for the voxels represented in the rows of matrix $\hat{\boldsymbol{\mathcal{V}}}$ is calculated as

$$\boldsymbol{\Sigma} = \frac{1}{n}\sum_{i=1}^{n} (y_i - \bar{y})\,(y_i - \bar{y})^{\mathrm{T}} \in \mathbb{R}^{3\times 3} \qquad \text{where} \qquad \bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i \in \mathbb{R}^3 \tag{8}$$

contains the means of all voxels in the corresponding dimension. In order to construct geometric covariance-descriptors as proposed by Blomley et al. [36], the eigenvalues of the covariance matrix are sorted in such a way that $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$. A scale normalization is achieved by dividing all eigenvalues by the largest eigenvalue, i.e., $\lambda^*_i = \lambda_i/\lambda_1$ where $i \in \{1,2,3\}$. The covariance-sphericity, $S_\lambda = \lambda^*_3/\lambda^*_1$, and the covariance linearity, $L_\lambda = (\lambda^*_1 - \lambda^*_2)/\lambda^*_1$, were identified as promising descriptors for the given shape classification problem. The resulting 4-dimensional feature vector, including shape descriptors from both the 2D and 3D domain, takes the form $\vec{x} = [\omega, \beta, S_\lambda, L_\lambda]^{\mathrm{T}}$, and is used within the framework of a supervised learning scheme.

3.3.3. Classifier selection

In this context, the term classifier refers to a function inferred from labeled training data, which maps the feature vector representing the shape of a particle to a shape category, called class. Due to its computational efficiency and its capability of performing multi-class classification on a data set, the *decision tree* algorithm [37,38] is chosen as a classifier. Additionally, discriminant functions for the binary classification of the *sphere* and the *non-convex* class are found using linear support vector machines. A cascaded classification scheme with two instances is thus proposed. The first classification instance performs two binary classifications using the convexity and circularity product, ω and β, in order to detect non-convex particles and spheroids. Particles belonging to none of the two classes are classified by the second classification instance which distinguishes *needle*, *quasi-equant* or *platelet* based on a decision tree inferred from simulated training data. It is worth noting that the discussed classification model does not include a rejection option, i.e., a particle is assigned to one of the five shape classes in any case.

3.3.4. Training data sets and supervised learning

3.3.4.1. Training set of spheres. A training set was obtained from real experiments containing 962 particles out of which 12.6 % were labeled as *spheres* (positive) and the rest as *others* which build the set of the remaining four classes (negative). The positive examples were spherical latex beads (Beckman Coulter, USA) of different nominal diameters. The optimal linear class-separating decision boundary (shown in Fig. 7a) was calculated from the circularity product of the training data set by means of a linear Support Vector Machine (SVM) using the `fitcsvm` function from the Statistics and Machine Learning Toolbox of MATLAB [39].

3.3.4.2. Training set of non-convex particles. The same approach was followed in the case of non-convex shaped particles like agglomerates. A training set consisting of 973 particles, out of which 11.8 % were labeled as non-convex shaped by an independent expert, were used to identify the decision boundary using a linear SVM. After the calculation of the convexity product, the optimal linear decision boundary, shown in Fig. 7b, was calculated by the linear SVM reported above.

3.3.4.3. Simulated training set. On the basis of the reconstruction ambiguity (mentioned in Section 3.2.4), the second classifier instance was trained with in silico data so as feature vectors are not compromised by the limitations of the reconstruction method using only two cameras. For this, the *extended virtual test bench* (eVTB) [2] for simulating crystal morphologies was used. The eVTB is a simulation framework which provides an idealized in silico environment for the simulation of a variety of different crystal morphologies. A training set including 2400 cuboids and 1200 cylinders served as basis for learning a decision tree model to distinguish the shape classes *quasi-equant*, *platelets*, and *needles*. Feature vectors of all simulated objects were collected and the `fitctree` function from the Statistics and Machine Learning Toolbox of MATLAB was used to build a decision tree model from the generated training data. The covariance linearity, L_λ, and the covariance sphericity, S_λ, along with the decision boundaries are shown in Fig. 7c. The complexity (or depth) of the decision tree is determined using 10-fold cross validation. Using this method, the training samples are split into 10 smaller sets and decision trees ensembles are generated while the complexity parameter is increased. For each ensemble, the cross-validated training error is reported and the parameter yielding a small training error is chosen. It is important to mention that the characteristic size of simulated particles for the generation of training data was sampled from a user-defined distribution. Therefore, the decision boundaries found using the decision tree algorithm were implicitly determined by the parameters of this distribution. The chosen parameters for generating the distribution were related to outcomes from real experiments.

3.3.4.4. Shape classification validation. In order to validate the classification performance of the implemented model, three independent experts were asked to classify 524 samples from a data set not used for training of the classifier. Classifier labels are compared to the majority label given by the experts using the *confusion matrix*[40], which is commonly used for calculating precision and sensitivity of multi-class classifiers. The validation revealed an overall agreement of 69.0 % of the classification model labels and the majority label by the experts.

The reader is referred to the supplementary material of this article for a detailed assessment of the multiclass validation results.

A few examples of shaded 3D visual hulls, along with the predicted shape class label of the classification model, are shown in Fig. 8. The figure shows corresponding stereoscopic image pairs captured by camera A and B, respectively, and the resulting reconstructed VH.

4. Experimental results

An experimental campaign was executed for assessing the performance of μ-DISCO and of the image analysis routine presented in this article. In Section 4.1, the stability of the measurements, as well as their repeatability is shown by measuring a population of γD-Mannitol crystals in 2-Propanol at saturated conditions. In Section 4.2, the accuracy of μ-DISCO in estimating the particle size is assessed by measuring populations of standardized spherical latex beads of three different nominal diameters. In Section 4.3, the capability of the image analysis routine to perform shape classification discussed in Section 3.3 and the characteristic response time of the device for a measurement of a suspension consisting of a mixture of particles exhibiting three different shapes is discussed. Also, an experiment with a growth and dissolution phase for βL-Glutamic acid is presented in Section 4.4, with the aim of demonstrating the potential of the device to provide time resolved changes in size and shape of the crystals during an experiment, even for processes with fast dynamics. It is to be noted that every measurement performed using μ-DISCO produces five different distributions for the five different classes. In this article the number- and volume-based distributions are shown for the class with the highest contribution, unless stated otherwise. The reader is directed towards the supplementary material for information regarding the calculation of the characteristic sizes, nD PSSDs, and moments that will be used in the following subsections.

4.1. Steady state and measurement repeatability

In order to establish that the system is at steady state, i.e., there is no change in the underlying population in the sampling loop or the channel, the PSSD of needle-like γD-Mannitol suspended in a saturated solution ($S = 1.0$) in 2-Propanol (Sigma-Aldrich, Buchs, Switzerland) at 25 °C was measured over a period of 5 h by circulating the suspension at a flow rate of 100 mL min^{-1} through the channel of μ-DISCO, capturing images at a frequency of 5 Hz. Crystals of the γ polymorph of D-Mannitol were produced by crash cooling from 25 °C to 5 °C, a saturated solution of D-Mannitol suspended in deionized and filtered (filter size of 0.22 µm) water obtained from a Milli-Q Advantage A10 system (Millipore, Zug, Switzerland). The crystals thus obtained were washed with 2-Propanol, filtered and dried for 24 h. In the case of appearance or disappearance of particles, a change in the average characteristic size of the population or in the percentage of particles belonging to the different classes would be observed. For example, during a measurement with needle-like particles, a breakage event would lead to particles exhibiting a prismatic shape, which in turn would lead to an increase in the percentage of quasi-equant particles and decrease in the percentage of needle-like particles; this would most likely be associated to a decrease in the average characteristic length along the length (L_1) direction.

Fig. 9a shows the time-resolved evolution of number-weighted characteristic size, along the mean length ($L_{1,\mathrm{mean}} = \mu_{10}/\mu_{00}$) and mean width ($L_{2,\mathrm{mean}} = \mu_{01}/\mu_{00}$) of the needle-like crystals. The measured time average $L_{1,\mathrm{mean}}$ and $L_{2,\mathrm{mean}}$ for repetition 1 was 157 µm and 19 µm, respectively, while for repetition 2, these measures were 156 µm and 19 µm, respectively. The generic particle model assumes a needle-like particle to be approximated by a cylinder. It is to be

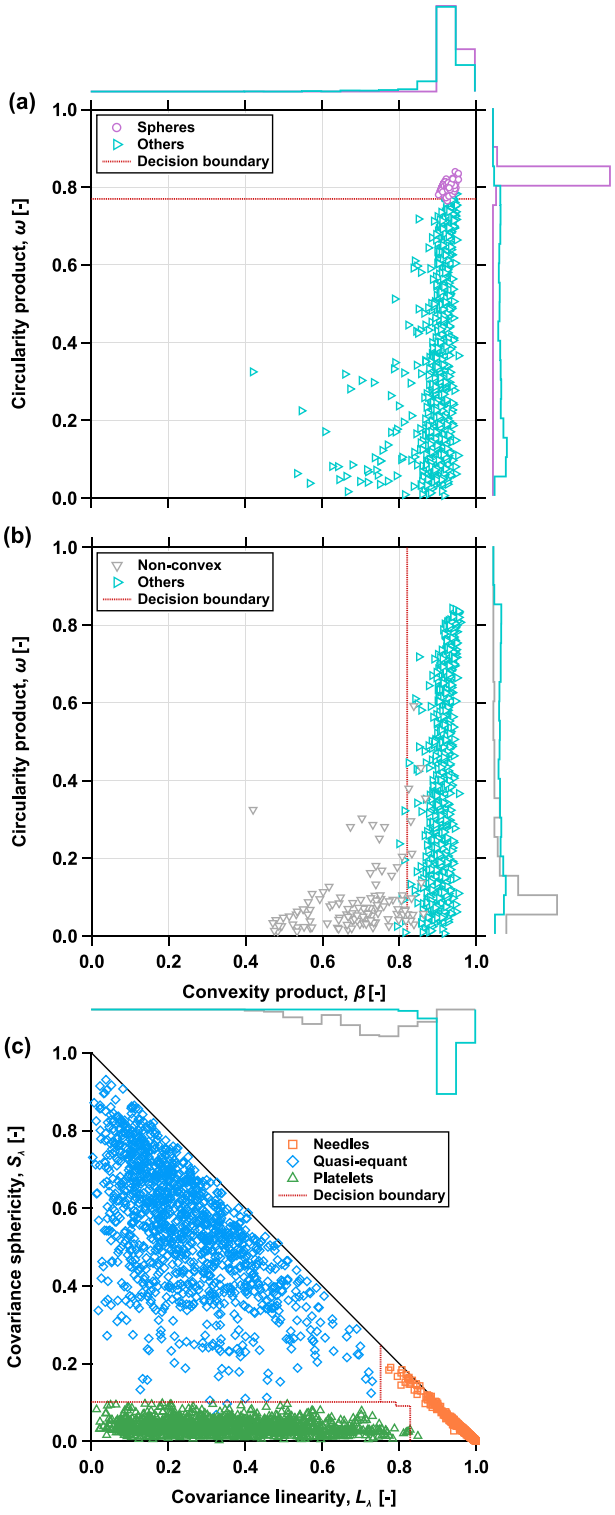

Fig. 7. Feature space with projected histograms of convexity product, β, and circularity product, ω, to distinguish **(a)** spheres (purple circles) from others (needles, quasi-equant, platelets and non-convex – cyan right-pointing triangle); **(b)** non-convex (grey downward-pointing triangle) from others (spheres, needles, quasi-equant, and platelets – cyan right-pointing triangle). **(c)** feature space of covariance linearity, L_λ, and sphericity, S_λ, to distinguish between needles, quasi-equant particles and platelets. The brown dotted lines indicate the decision boundaries for the different classes.

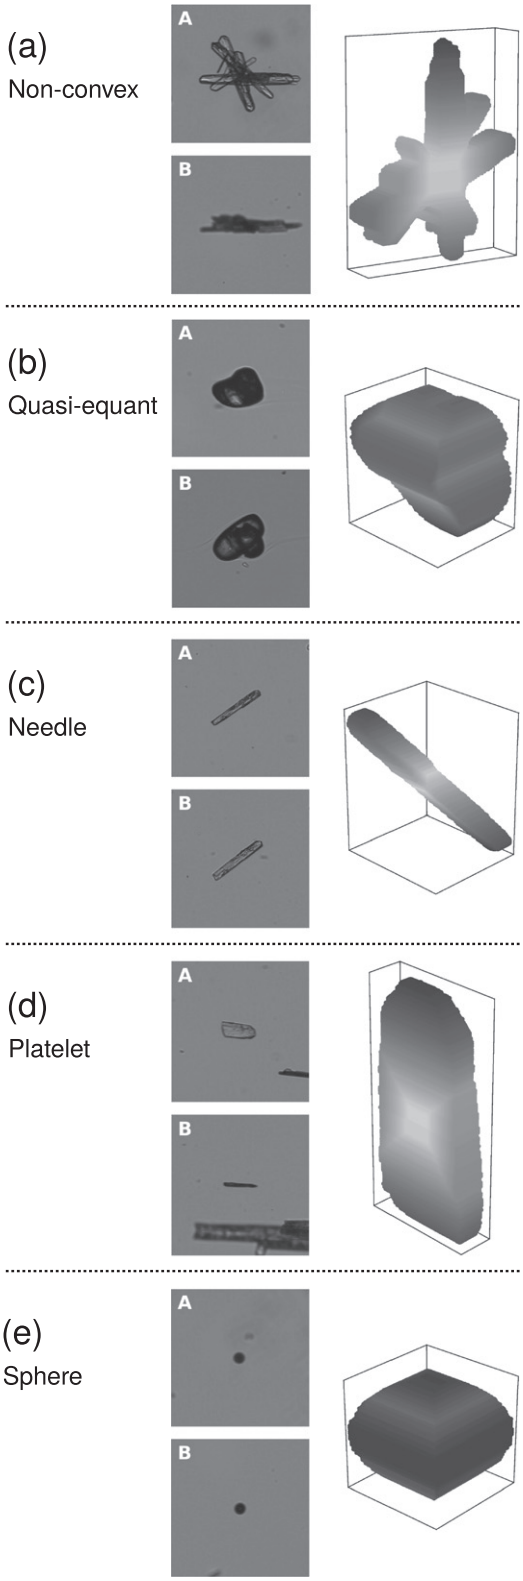

Fig. 8. Examples of reconstructed visual hulls of crystals belonging to the generic shape classes distinguished by the classifier. **(a)**: non-convex shaped agglomerate of needle-like βL-Glutamic acid; **(b)**: quasi-equant shaped Aspirin crystal; **(c)**: needle-like βL-Glutamic acid crystal; **(d)**: broken crystal segment of βL-Glutamic acid which takes the shape of a platelet; **(e)**: spherical latex bead with a nominal diameter of 90 μm (Beckman Coulter, USA).

noted that this approximation introduces a geometrical error independent of the optical error reported for the measurement of a population of spheres in Section 4.2. The quantification of this error is beyond the scope of this work. Apart from the characteristic lengths, the time-resolved broadness of the number-weighted distribution, σ_i, along the two directions are shown in Fig. 9b. The broadness measure of the 2D distribution used in this article is calculated as

$$\sigma_1 = \left[\frac{\mu_{20}}{\mu_{00}} - \left(\frac{\mu_{10}}{\mu_{00}}\right)^2\right]^{\frac{1}{2}} \quad \text{and} \quad \sigma_2 = \left[\frac{\mu_{02}}{\mu_{00}} - \left(\frac{\mu_{01}}{\mu_{00}}\right)^2\right]^{\frac{1}{2}} \tag{9}$$

The metric also exhibited a deviation with the same order of magnitude as that of the two characteristic lengths. The near-constant broadness of the distribution further confirms the fact that there were no significant changes in the suspension, over the entire duration of the measurement. A negligible change in the number and fraction of particles in the different shape classes was observed over the duration of the experiment, where, $87.8 \pm 1.6\%$ and $90.4 \pm 0.4\%$ of the detected particles accounted for the needle-like shape in the two repetitions, respectively.

The constant and repeatable trends observed from Fig. 9 are good indicators that the hardware setup has negligible to minimal impact on quantifying the population in a crystallizer. In addition, the observation of the minor differences in number and percentage of the predicted shapes by the new shape classification algorithm has further strengthened our confidence in using μ-DISCO for quantitative and reliable measurements.

4.2. 1D particle size validation

Spherical latex beads (Beckman Coulter, Nyon, Switzerland) with nominal diameters of 20 μm, 65 μm, and 90 μm were suspended in 400 mL of deionized and filtered (filter size of 0.22 μm) water obtained from a Milli-Q Advantage A10 system (Millipore, Zug, Switzerland). The size of the particles was measured using μ-DISCO at a frequency of 5 Hz in order to evaluate the sizing accuracy of the device. The suspension is circulated for a duration of 15 min at a constant flow rate of 100 mL min^{-1} through the channel. After each experiment, the sampling loop and the channel were carefully cleaned and thoroughly flushed with purified water, ethanol, and acetone.

On average, $86.5 \pm 1.0\%$ of all detected particles were classified as spheres by the shape classifier over the three independent experiments. The second largest shape contribution was the quasi-equant particles, which contributed to $13.2 \pm 1.0\%$ to the overall shape composition. Particles of other shapes were present in negligible amounts (<0.3%). The discretized normalized 1D number- and volume-weighted distributions for the three independent repetitions of the measurement are shown in Fig. 10, where the particle size distribution (PSD) is shown only for particles classified as spheres. The distributions nicely overlap, again indicating an excellent repeatability of the three independent measurements. In the smaller particle size range (below 60 μm), a lower measurement accuracy was observed. This can be attributed to the limited depth of field and the optical resolution of the telecentric lenses, which lead to an impairment of the measurements as the particles are likely to be out of focus. This is in line with observations from similar imaging devices [4]. For larger particles (above 60 μm), the measurement error between the mode of the number- or volume-weighted distributions, and the nominal diameters of the latex beads was less than 5 %.

An interesting observation from the microscope images of the latex beads, shown in Fig. 10c, was the presence of broken beads and of beads with nominal diameter larger than the one obtained from the manufacturer. The broken beads would explain the contribution from quasi-equant-particles, while the presence of the larger beads would explain the peak at around 110 μm, clearly visible in the volume-weighted distribution shown in Fig. 10b.

4.3. Characterization of spheres, needle-like, and quasi-equant particles

The automated shape classification model was tested with particles of three distinct shapes, namely spheres, needles, and quasi-equant particles, to evaluate its efficacy in tracking sudden changes in the suspension during the online measurements. Two different polymorphs of L-Glutamic acid (L-Glu), namely the prismatic metastable α polymorph, and the needle-like, stable β polymorph, exist. The two polymorphs combined with the spherical latex beads (Beckman Coulter, Nyon, Switzerland) make an excellent system to test the performance of the shape classification algorithm during a dynamic measurement with sudden changes in the population of suspension of particles. Crystals of the α polymorph were produced by mixing equimolar amounts of monosodium L-Glu monohydrate (NAGLU, Sigma-Aldrich, Buchs, Switzerland, purity >99 %) and hydrochloric acid (HCL, Sigma-Aldrich, Buchs, Switzerland, 37–38 %) in deionized and filtered water with continuous stirring at 5.0 °C for 1 h. The α crystals obtained were filtered and dried. To obtain the β polymorph, a saturated solution with respect to the α form of L-Glu at 45.0 °C was created by mixing equimolar amounts of NaGlu and HCl, and the α crystals obtained from the previous step were allowed to undergo a polymorphic transformation over 36 h. The transformed crystals were subsequently filtered and dried [41].

The measurement was initiated by the addition of spherical latex beads (phase I), with a nominal diameter of 90 µm, to 300.0 g of pure ethanol (Merck KGaA, Darmstadt, Germany, purity >99.5 %), followed by the addition of 0.5 of αL-Glu after 10 min (phase II). Finally, 10 min after the addition of αL-Glu, 0.1 g of βL-Glu was added to the suspension (phase III). Since, L-Glu is insoluble in ethanol, neither growth nor dissolution of the crystals during the experiment occurred. The suspension was stirred at 300 rpm and was pumped through the μ-DISCO at a flow rate of 100 mL min^{-1}. The addition of particles of different shapes at short frequent intervals renders the standard mode of operation, i.e., obtaining images at a constant, low frequency over the entire duration of the measurement, suboptimal. Hence, the burst mode described in Section 2 with a frequency of image capture, which is an order of magnitude higher than in the standard mode, would be efficient in capturing the sudden changes in the process.

The suspension was analyzed by μ-DISCO operating in the burst mode and capturing 600 images every minute at a frame rate of ≈70 Hz for a total measurement duration of 30 min. A total of 9 × 10^4 particles were observed and analyzed. The normalized number of particles (defined as the ratio of number of particles of a desired class at a specific time to the total number of particles of the desired class over the entire measurement duration), and the time resolved fraction of the five particle classes are shown in Fig. 11. During phase I of the measurement, as expected, the sphere class has the highest contribution to the fraction of particles observed, while the needles, platelets and non-convex particles have a negligible contribution. The higher fraction of the quasi-equant class during phase I can be attributed to the presence of broken spherical latex beads, which were classified as quasi-equant particles. During phase II, an increase in the number and fraction of quasi-equant particles due to the presence of the prismatic (classified as quasi-equant particles) α polymorph was observed. Interestingly in phase II, also an increase of non-convex particles can be observed, which can be explained by the fact that crystals of the sample can stick together. In phase II, the delay in the stabilization of the normalized number of particles and of the fraction of classes is understandable due to the time required for proper mixing of the newly added particles and to the dead volume of the sampling loop.

During phase III, more than 60 % of the particles were classified as needles. With the presence of needle-like particles, an increase in the number and contribution from the platelet class was also evident due to the ambiguities arising from the camera setup as described in Section 3.2.4. The higher fraction of non-convex particles compared to phase-I can be explained by the presence of crystals exhibiting non-convex shapes, and overlapping needle-like crystals.

One interesting point to note is that even though the three phases of the measurement last only for 10 min each, the number and fraction of different classes in each of these individual phases exhibit only a small deviation around their mean, thus confirming that the system was indeed at steady state during the individual phases.

4.4. Growth and dissolution of βL-Glutamic acid

The potential of the stereoscopic imaging setup to monitor the growth of a population of βL-Glu crystals has been reported in our previous publications [9,10]. However, to the best of our knowledge, dissolution has never been monitored using a stereoscopic imaging setup like the one described in this article, although dissolution kinetics have been reported in the past using other monoscopic imaging techniques [42,43]. Hence, an experiment was set up consisting of growth followed by dissolution, to demonstrate the ability of μ-DISCO to track the evolution of a population of βL-Glu seeds during the two distinct phases.

Needle-like βL-Glu crystals produced by the protocol described in Section 4.3 were milled at 10,000 rpm in a wet mill to produce quasi-equant shaped particles. An amount of 0.165 g of the seeds were suspended in 2000 g of saturated solution of βL-Glu in water at 30.1 °C, and the population was characterized using μ-DISCO for 0.6 h (steady state phase). The suspension was then crash cooled (from

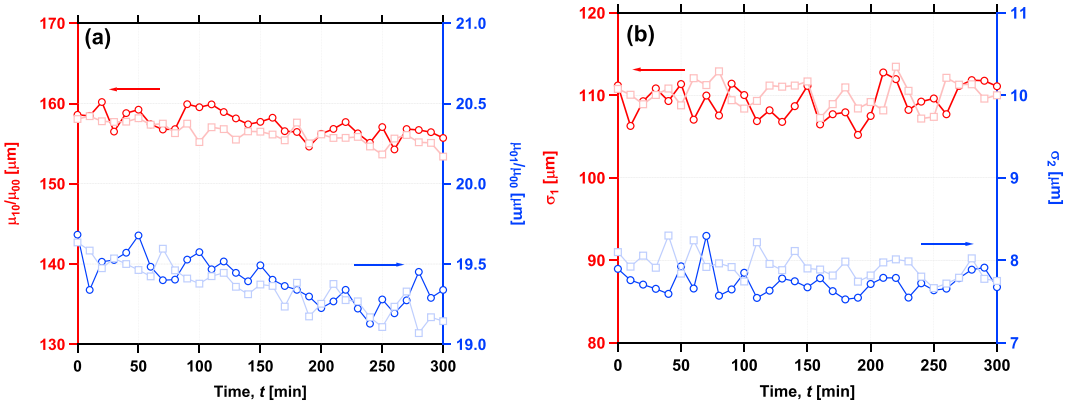

Fig. 9. Time-resolved **(a)** number-weighted characteristic lengths, μ_{10}/μ_{00} (left ordinate) and μ_{01}/μ_{00} (right ordinate) and the **(b)** number-weighted broadness of the distribution, σ_1 (left ordinate) and σ_2 (right ordinate) of γD-Mannitol in 2-Propanol at saturated conditions ($S = 1.0$) for the two repetitions (repetition 1 - circles and repetition 2 - squares).

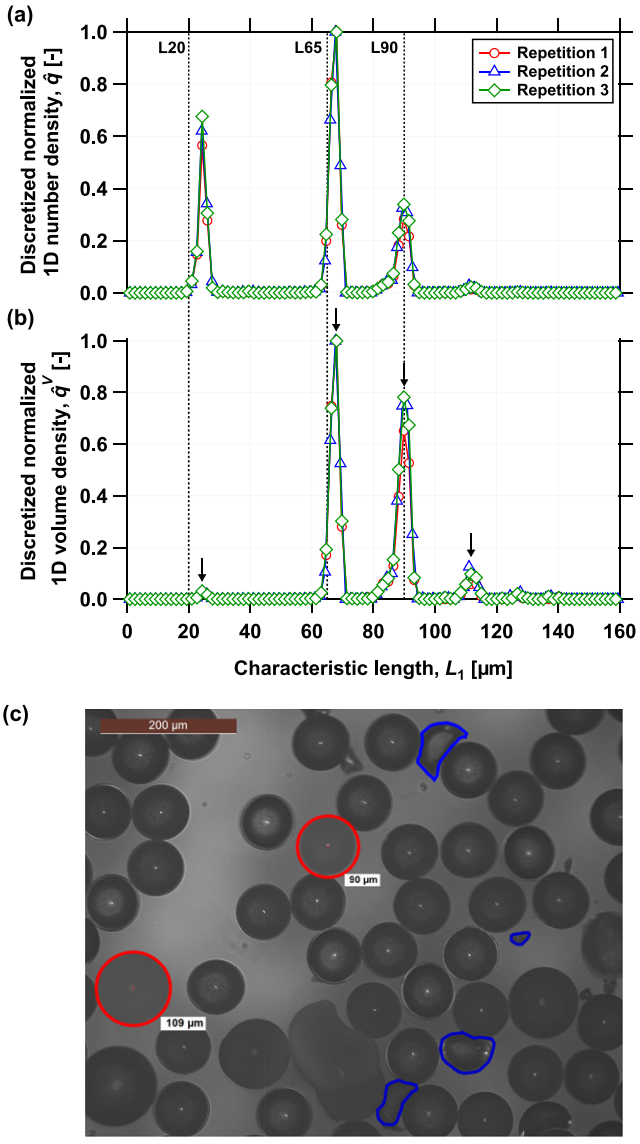

Fig. 10. Discretized normalized 1D **(a)** number density distribution, $\hat{q}$, and **(b)** volume density distribution, $\hat{q}^V$ (solid lines) from three repetitions for a mixture of spherical latex beads. Spacing between the numbers and units is not consistent in the page proof version of the document. Needs to be fixed. **(c)** Microscope image of the standard spherical latex beads with a nominal diameter of 90μm. Particles highlighted with red contours would be classified as spheres and broken particles highlighted with blue contours would be classified as quasi-equant particles. The spherical particles highlighted (red) shows two different diameters, 90μm and 109μ m.

30.1 °C to 25.0 °C) to induce supersaturation ($S = 1.20$), while being continuously stirred at 300 rpm. The population was allowed to grow for 11.4 h (growth phase), while being monitored by μ-DISCO operating in intermittent burst mode, capturing 800 images at a frame rate of ≈70 Hz every 6 min. After the growth phase, a linear heating ramp (25.0 °C to 33.1 °C at 2.7 °C h^{-1}) aiming at a complete dissolution of the population was applied. Since the dissolution kinetics of βL-Glu are significantly faster than growth kinetics, it was decided to operate the device in the burst mode capturing 800 images at a frame rate of ≈70 Hz every 2 min to obtain a better time-resolved measurement in this phase. It is worth noting that in the experiment considered in this section, i.e., in the temperature range and at the pressure explored here, L-Glu is present only in the β form, which is the thermodynamically stable one. The supersaturation generated during the growth phase in the experiment discussed above does not cause any observable nucleation and agglomeration, as previously confirmed by other independent studies [9,12,41].

Over the entire duration of the experiment, a total of 1.3×10^6 particles were measured by μ-DISCO. Fig. 12a and b show the normalized number of particles and fraction of different classes, respectively. Spheres, platelets, and non-convex particles are lumped together as *others*, as the classes of interest for this experiment were the needle-like and quasi-equant particles. The normalized total reconstructed VH volume, which, ideally should be proportional to the true volume of crystals being observed by the measurement device, is shown in Fig. 12c.

The quasi-equant class is characterized by a 3D PSSD in (L_1, L_2, L_3) space. In order to facilitate a comparison with the needle class, a reduced 2D PSSD for the quasi-equant class combined with the needle class in the ($\tilde{L}_1$, $\tilde{L}_2$) plane was used, where $\tilde{L}_1$ and $\tilde{L}_2$ are the length and width of the needle-like particles, respectively. The reduced 2D PSSD for the quasi-equant class was generated by assuming $\tilde{L}_1 = L_1$, and $\tilde{L}_2 = \sqrt{4L_2L_3/\pi}$, which preserves the volume of the particle and transforms it into a needle-like particle with length, $\tilde{L}_1$, and width, $\tilde{L}_2$. The number-weighted characteristic length along the $\tilde{L}_1$ direction is shown in Fig. 13b, and the aspect ratio, $\tilde{\mu}_{21}/\tilde{\mu}_{12}$, for the reduced 2D PSSD is shown in Fig. 13c.

During the steady state phase, the normalized number of particles, the fraction of classes, as well as the total reconstructed VH volume remained unchanged as expected. Similar values for the fraction of needle-like and of quasi-equant shaped particles for the milled seeds of βL-Glu were observed. The volume-weighted seed distribution shown in Fig. 13a has an aspect ratio of 3.3. The seed

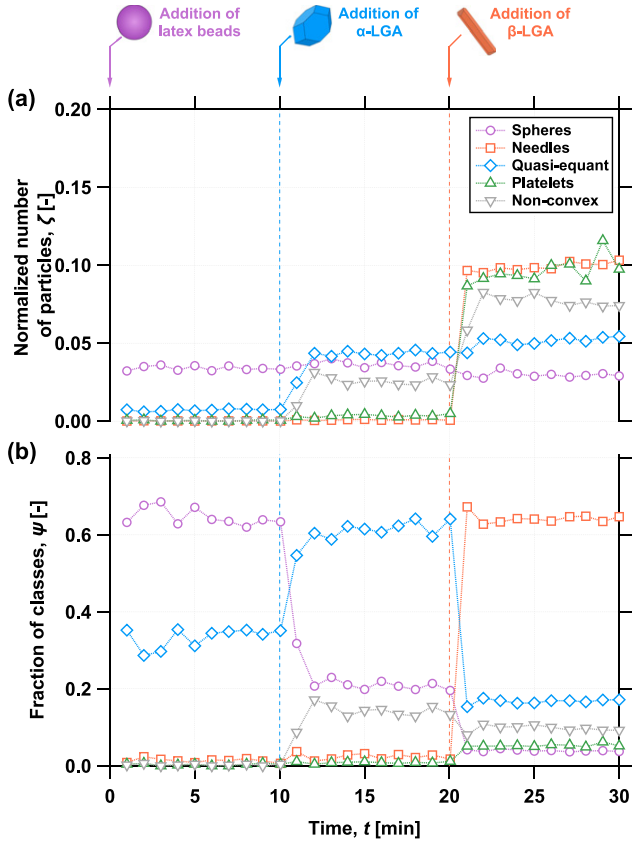

Fig. 11. Time-resolved **(a)** normalized number of particles (ζ), and **(b)** fraction of classes (Ψ), for a mixture of spherical latex beads, αL-Glu, and βL-Glu in pure ethanol. At $t = 0$ min, spherical latex beads were added in ethanol, followed by the addition of prismatic shaped αL-Glu at $t = 10$ min, and needle-like βL-Glu at $t = 20$ min.

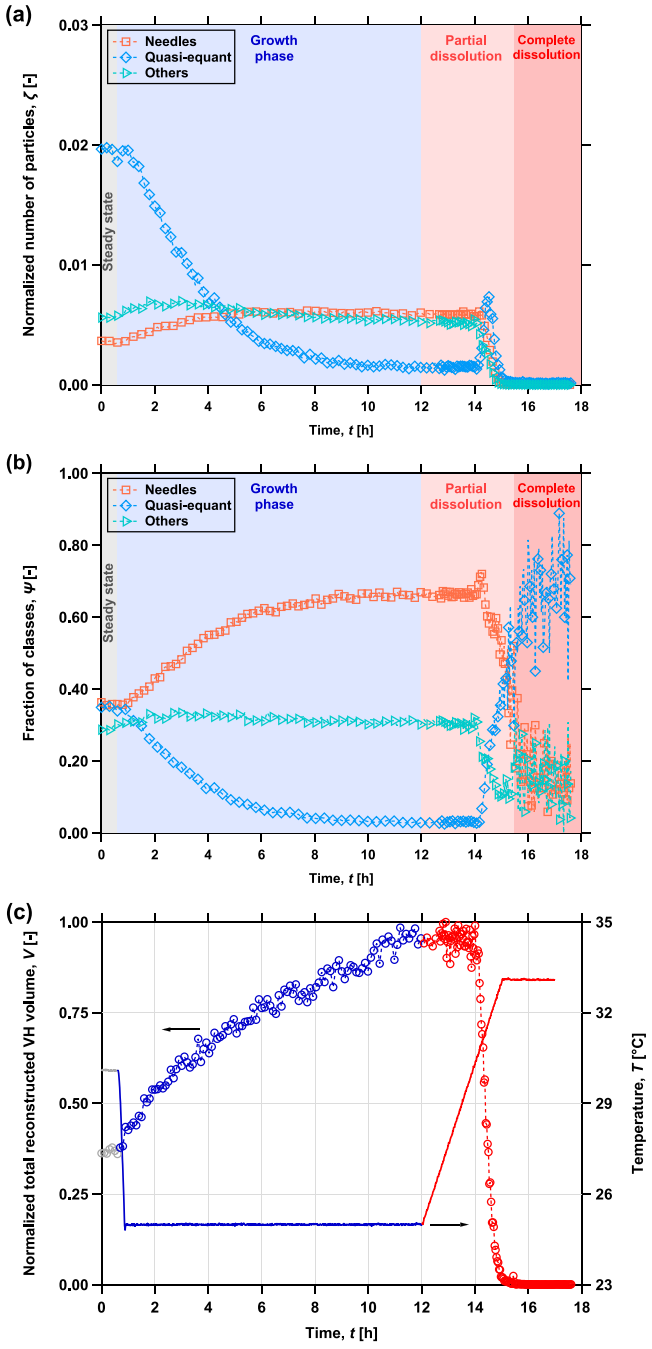

Fig. 12. Time resolved **(a)** normalized number of particles (ζ), **(b)** fraction of classes (Ψ), and **(c)** normalized reconstructed VH volume ($\bar{V}$) of all particles analyzed by the measurement device. Seeds of βL-Glu (grey shaded region) were grown at a supersaturation of $S = 1.20$ for 11.4 h (growth phase - blue shaded region), followed by a complete dissolution of the grown crystals (red shaded region). The 2D PSSD for the seeds and the product at the end of the growth phase, time-resolved number-weighted characteristic length and aspect ratio of the reduced 2D PSSD is shown in Fig. 13.

population has a number-weighted characteristic length of 84 µm. During the growth phase, the kinetics of βL-Glu along $\tilde{L}_1$ direction are faster than along the $\tilde{L}_2$ direction. The preferential growth along one characteristic direction suggests that βL-Glu crystals exhibiting quasi-equant shape develop into a needle-like shape, hence a decrease in the number and fraction of quasi-equant particles, and an increase in the needles would be expected; this hypothesis was indeed confirmed by the experiment. The population at the end of the growth phase had an aspect ratio of 5.9 with a number-weighted characteristic length of 179 µm. It is interesting to note that the aspect ratio during the growth phase is not constant; the increase in the aspect ratio can be explained by the change in shape of the quasi-equant βL-Glu particles. The increase in the reconstructed VH volume during this phase serves as an additional aid to visualize crystal growth of the population of crystals.

The final phase of the experiment is divided into two sub-phases, namely a partial dissolution phase and a complete dissolution phase for better clarity. No visible change in the number or fraction of classes, in the total reconstructed VH volume, or in the number-weighted characteristic length was observed in the first 2 h of the dissolution phase, which could be explained by the incomplete consumption of supersaturation during the growth phase. The number of needles drop significantly from this moment on, while the number of quasi-equant particles attains a maximum followed by a gradual drop to zero. Upon examination of the images obtained, it could be confirmed that the long needle-like particles undergo a shape change to quasi-equant particles, followed by a complete dissolution. During the final dissolution phase, from Fig. 12a, it is clear that the particle count for all the three classes drops to zero, while the fraction of classes shown in Fig. 12b has a high variance, as expected from the low particle count. The number-weighted characteristic length and aspect ratio also exhibit a high variance in the complete dissolution phase, hence for better clarity it was decided to plot these quantities only during the partial dissolution phase. The particles in the suspension dissolve, leading to a decrease in the total reconstructed VH volume and number-weighted characteristic length; and also the distribution evolves towards a lower aspect ratio, which might indicate faster dissolution kinetics along the length direction when compared to the width direction.

The versatility of μ-DISCO to provide information at two different levels is clearly demonstrated in this experiment. On the one hand, Fig. 12 provides a qualitative high-level information concerning the evolution of the population during the experiment, which can be used for process monitoring. On the other hand, Fig. 13 provides a quantitative detailed (low-level) information about the experiment, which can be used for process modeling and development. Note that, the results shown in this section concerning the dissolution phase represent merely a proof-of-concept experiment, intended to demonstrate the ability of μ-DISCO to monitor dissolution; a quantitative interpretation of such results should be made only with great caution.

5. Concluding remarks

5.1. Key outcomes

In this work, we have presented an improved dual-camera measurement device, along with enhanced image processing and analysis algorithms, which are used to extract 3D metric information from stereoscopic images acquired using a back-light illumination technique. A novel calibration procedure for multi-camera systems with telecentric lenses was proposed which aids in the correction of mechanical misalignment, and thereby increases the accuracy of spatial matching of particle projections. 3D visual hull reconstruction of particle shapes observed by the two stereoscopic cameras and in turn, the ability of geometric shape descriptors to classify particles into generic shape classes was shown.

The compact optomechanical camera setup, alongside improved optics, calibration model, background subtraction for image segmentation, and statistical contour matching from the stereoscopic images, paves the way towards an even more accurate characterization of the particle size, and hence the PSSDs when compared to our previous works. The characterization of a suspension with

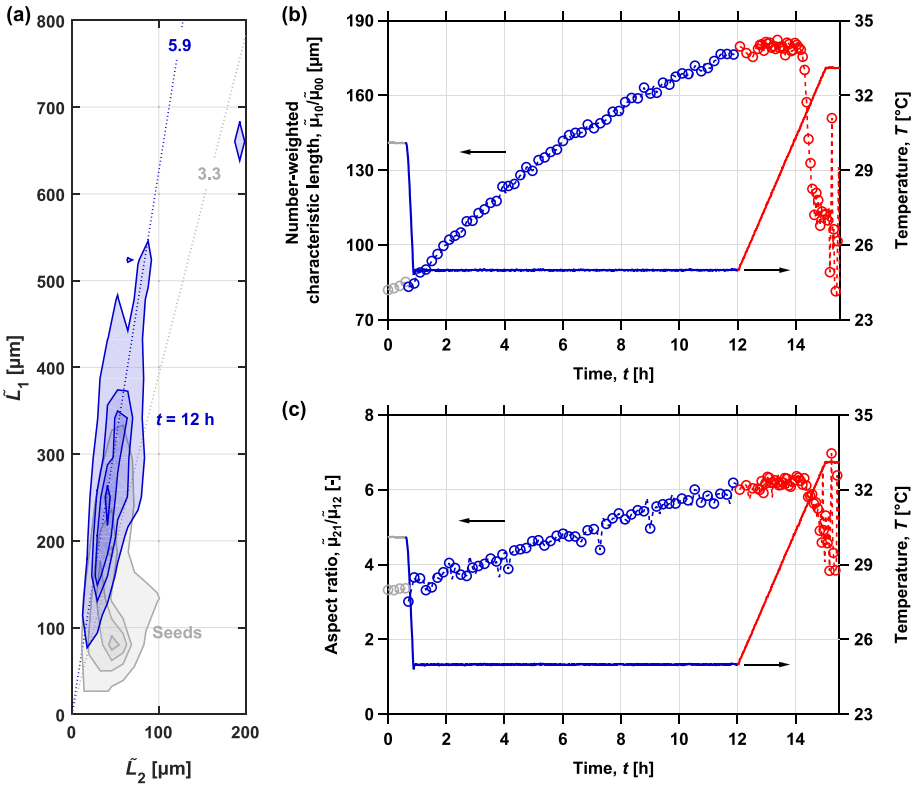

Fig. 13. **(a)** Volume weighted 2D PSSD, $\hat{q}^V$, in the reduced $(\tilde{L}_1, \tilde{L}_2)$ plane for the seeds (grey) and products (blue) at the end ($t = 12$ h) of the seeded growth phase of βL-Glu. The contour line values correspond to 0.2, 0.5, 0.7, and 0.9 of the maximum of the PSSD. A linearized grid with 20 bins at a spacing of 46 µm for seeds and 54 µm for products along the $\tilde{L}_1$ direction, and a spacing of 13 µm for seeds and 12 µm for products along the $\tilde{L}_2$ direction was used. The constant aspect ratios $(\tilde{\mu}_{21}/\tilde{\mu}_{12})$ for the two distributions are given by the dotted lines. Time-resolved **(b)** number-weighted characteristic length $(\tilde{\mu}_{10}/\tilde{\mu}_{00})$, and **(c)** aspect ratio $(\tilde{\mu}_{21}/\tilde{\mu}_{12})$ of the reduced 2D PSSD.

three characteristic particle shapes presented in this article indicated the ability of μ-DISCO to track shape changes with a time resolution an order of magnitude higher when compared to the FTC. Also, we are optimistic that the positive outcome from the experiment concerning the dissolution of βL-Glu is the first step towards monitoring, characterizing, and modeling a dissolution process using a stereoscopic imaging setup.

5.2. Key limitations

We acknowledge the fact that the long sequence of image processing steps going from the two raw images to the final size and shape distributions introduces several uncertainties. The camera-lens system introduces an optical error, while the assumption of a generic particle model introduces a geometrical error. Alas, the quantification of such errors is a nontrivial task, which is beyond the scope of this article. We are aware that the VH reconstruction method presented here merely yields a rough approximation of the three-dimensional particle shape as long as only two projections are available. It is clear that configurations with more cameras would alleviate the problems associated with this aspect. However, the accompanying, necessary changes in cost, (flow cell) geometry and the additional effort w.r.t. the calibration, imaging cost, etc. imply that the choice of the number of cameras represents a trade-off in this application. It is our strong belief that the boost in information obtained from having a second camera is such that adding it is hugely advantageous, whereas every subsequent camera only offers incremental improvement. Furthermore, there is no doubt that the output from the shape classification model is a qualitative result, which is not necessarily linked to physical quantities. However, it may be closely correlated with such. The work carried out is conceived to make meaningful process indicators available, which provide an insight into the predominant shape of crystal populations during a crystallization process. Considering the mentioned limitations, the output from the measurements should be analyzed with care, to avoid over-interpretation.

5.3. Way forward

The techniques outlined in this work could be a useful aid in other fields of research dealing with particulate matter where shape related information is critical. The compact design enables the device to be set up in a standard laboratory hood or cart, which makes us feel confident that we can use it for industrial applications in particular during late-stage development and for the trouble-shooting of launched products, in synergy with available commercial characterization tools to provide better insights into various processes. We have made a significant step forward towards making size and shape feedback control of a process possible in the near future.

Nomenclature

Greek symbols

α_j	j-th isotropic scaling factor [µm/pixel]
β	convexity product [–]
ϵ	relative translational misalignment of cameras along the z_w axis [µm]
γ	learning rate of the adaptive background subtraction [–]
κ_j	convexity pertaining to camera j [–]

λ_n	n-th eigenvalue of $\hat{\boldsymbol{\mathcal{V}}}$ [µm]
λ_n^*	n-th normalized eigenvalue of $\hat{\boldsymbol{\mathcal{V}}}$ [–]
μ_{ij}	ij-cross moment of shape distribution [µm^{i+j-3}]
Ω	image domain
ω	circularity product [–]
ϕ_j	roll angle of j-th camera [rad]
Ψ	normalized fraction of classes [–]
σ_n	broadness measure of a 2D distribution in direction n[µm]
τ	global intensity threshold for binarization
$\tilde{\mu}_{ij}$	ij-cross moment of reduced shape distribution [µm^{i+j-3}]
ζ	normalized number of particles [–]

Roman symbols

$[\tilde{u}_0, \tilde{u}_0]$	principal point offset coordinates [pixel]
$[u_j, u_j]$	image plane coordinates camera j [pixel]
$[x, y]$	discrete pixel coordinates on Ω [pixel]
$\hat{\boldsymbol{\mathcal{V}}}$	data matrix holding voxels of the subsampled visual hull [µm]
$\boldsymbol{\mathcal{V}}$	data matrix holding voxels of the visual hull [µm]
Δs	edge length of a cubic voxel in the visual hull [µm]
Δx	bounding box centroid difference along the z_w-direction [µm]
$\hat{q}$	discretized normalized nD number density distribution [–]
$\hat{q}^V$	discretized normalized nD volume density distribution [–]
$\hat{x}_j$	bounding box centroid x-coordinate in coordinate system $\mathcal{O}_j$[µm]
$\mathcal{O}_w$	world coordinate system
$\mathcal{O}_j$	j-th camera coordinate system
$\mathbf{K}_j$	j-th camera matrix
$\mathbf{R}_j$	j-th camera rotation matrix
$\mathbf{T}_j$	j-th translation vector
$\bar{V}$	normalized reconstructed visual hull volume [–]
$\tilde{L}_n$	reduced characteristic particle size [µm]
$\vec{x}$	4-dimensional feature vector [–]
$B(x, y, t_k)$	background model at time step t_k
$G(x, y, t_k)$	binary segmentation mask at time step t_k
$I(x, y, t_k)$	new frame at time step t_k
$I_{\text{diff}}(x, y, t_k)$	new difference frame at time step t_k
L_n	characteristic particle size, $L_1 \geq L_2 \geq L_3$ [µm]
N	number of voxels in the visual hull [–]
Q_j	isoperimetric quotient (circularity) pertaining to camera j [–]
S	supersaturation [–]
T	temperature [°C]
t_k	time at discrete time step k [s]

Abbreviations, subscripts and superscripts

j	camera identifier $j \in \{A, B\}$
k	discrete time step
x, y, z	subscripts denoting direction in Euclidean space
A, B	subscript pertaining to primary and secondary camera

Acronyms

1D, 2D, 3D	n-dimensional, $n = 1, 2, \ldots$
CC	Coulter counter
CLD	chord length distribution
CMOS	complementary metal-oxide-semiconductor
DISCO	dual imaging system for crystallization observation
eVTB	extended virtual test bench
FBRM	focused beam reflectance measurement
FOV	field of view
fps	frames per second
FTC	flow through cell
LD	laser diffraction
LED	light-emitting diode
PAT	process analytical technology
PSD	particle size distribution
PSSD	particle size and shape distribution
SDK	software development kit
SFS	shape-from-silhouette
SVM	Support Vector Machine
TVSeg	total variation segmentation
VH	visual hull

Acknowledgments

The authors are thankful to the Swiss National Science Foundation for their support (CrystoCAM 2.0, project number - 155971). The authors thank Daniel Trottmann for the discussions and his assistance in building the μ-DISCO.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.powtec.2017.08.044.

References

[1] M.R. Singh, J. Chakraborty, N. Nere, H.-H. Tung, S. Bordawekar, D. Ramkrishna, Image-analysis-based method for 3D crystal morphology measurement and polymorph identification using confocal microscopy, Cryst. Growth Des. 12 (7) (2012) 3735–3748.
[2] I. d. Albuquerque, M. Mazzotti, D.R. Ochsenbein, M. Morari, Effect of needle-like crystal shape on measured particle size distributions, AIChE J. 62 (9) (2016) 2974–2985.
[3] A. Ruf, J. Worlitschek, M. Mazzotti, Modeling and experimental analysis of PSD measurements through FBRM, Part. Part. Syst. Charact. 17 (4) (2000) 167–179.
[4] B. Bujak, M. Bottlinger, Three-dimensional measurement of particle shape, Part. Part. Syst. Charact. 25 (4) (2008) 293–297.
[5] X. Wang, J.C. De Anda, K. Roberts, Real-time measurement of the growth rates of individual crystal facets using imaging and image analysis: a feasibility study on needle-shaped crystals of L-glutamic acid, Chem. Eng. Res. Des. 85 (7) (2007) 921–927.
[6] M. Kempkes, T. Vetter, M. Mazzotti, Measurement of 3D particle size distributions by stereoscopic imaging, Chem. Eng. Sci. 65 (4) (2010) 1362–1373.
[7] S. Schorsch, T. Vetter, M. Mazzotti, Measuring multidimensional particle size distributions during crystallization, Chem. Eng. Sci. 77 (2012) 130–142.
[8] S. Schorsch, D.R. Ochsenbein, T. Vetter, M. Morari, M. Mazzotti, High accuracy online measurement of multidimensional particle size distributions during crystallization, Chem. Eng. Sci. 105 (2014) 155–168.
[9] D.R. Ochsenbein, S. Schorsch, T. Vetter, M. Mazzotti, M. Morari, Growth rate estimation of βL-glutamic acid from online measurements of multidimensional particle size distributions and concentration, Ind. Eng. Chem. Res. 53 (22) (2014) 9136–9148.
[10] D.R. Ochsenbein, S. Schorsch, F. Salvatori, T. Vetter, M. Morari, M. Mazzotti, Modeling the facet growth rate dispersion of βL-glutamic acid-combining single crystal experiments with nD particle size distribution data, Chem. Eng. Sci. 133 (2015) 30–43.
[11] D.R. Ochsenbein, T. Vetter, M. Morari, M. Mazzotti, Agglomeration of needle–like crystals in suspension. II. Modeling, Cryst. Growth Des. 15 (9) (2015) 4296–4310.
[12] D. Ochsenbein, T. Vetter, S. Schorsch, M. Morari, M. Mazzotti, Agglomeration of needle-like crystals in suspension: I. Measurements, Cryst. Growth Des. 15 (4) (2015) 1923–1933.
[13] H. Hundal, S. Rohani, H. Wood, M. Pons, Particle shape characterization using image analysis and neural networks, Powder Technol. 91 (3) (1997) 217–227.
[14] Z. Zhang, Flexible camera calibration by viewing a plane from unknown orientations, ICCV, 1999, pp. 666–673.
[15] R. Tsai, A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses, IEEE J. Robot. Autom. 3 (4) (1987) 323–344.
[16] M. Li, J.-M. Lavest, Some aspects of zoom lens camera calibration, IEEE Trans. Pattern Anal. Mach. Intell. 18 (11) (1996) 1105–1110.
[17] J.Y. Bouguet, Camera Calibration Toolbox for Matlab, 2008, available online: http://www.vision.caltech.edu/bouguetj/calib_doc/ (Accessed: 2016-03-14).
[18] R. Hartley, A. Zisserman, Multiple View Geometry in Computer Vision, Cambridge University Press. 2004.
[19] M. Watanabe, S.K. Nayar, Telecentric optics for computational vision, Proceedings of the 4th European Conference on Computer Vision-Volume II, Springer-Verlag, London, UK, 1996, pp. 439–451.

[20] L. Huiyang, C. Zhong, Z. Xianmin, Calibration of camera with small FOV and DOF telecentric lens, 2013 IEEE International Conference on Robotics and Biomimetics (ROBIO), 2013, pp. 498–503.
[21] N. Friedman, S. Russell, Image Segmentation in Video Sequences: A Probabilistic Approach, Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1997, pp. 175–181.
[22] Y. Benezeth, P.-M. Jodoin, B. Emile, H. Laurent, C. Rosenberger, Comparative study of background subtraction algorithms, J. Electron. Imaging 19 (2010).
[23] R.M. Haralick, S.R. Sternberg, X. Zhuang, Image analysis using mathematical morphology, IEEE Trans. Pattern Anal. Mach. Intell. PAMI 9 (4) (1987) 532–550.
[24] S. Suzuki, K. Abe, Topological structural analysis of digitized binary images by border following., Comput. Vis. Graph. Image Process. 30 (1) (1985) 32–46.
[25] G. Bradski, The OpenCV library, Dr. Dobb's J. Softw. Tools (2000).
[26] M. Unger, T. Pock, H. Bischof, Continuous globally optimal image segmentation with local constraints, Computer Vision Winter Workshop, Moravske Toplice, Slovenija, 2008.
[27] Q. Chen, C. He, Variational segmentation model for images with intensity inhomogeneity and Poisson noise, EURASIP J. Image and Video Process. 2013 (1) (2013) 1–11.
[28] M. Donoser, M. Urschler, M. Hirzer, H. Bischof, Saliency driven total variation segmentation, 2009 IEEE 12th International Conference on Computer Vision, 2009, pp. 817–824.
[29] B. Baumgart, Geometric Modeling for Computer Vision, Memo (Stanford University. Stanford Artificial Intelligence Laboratory), University Microfilms. 1974.
[30] A. Laurentini, The visual hull concept for silhouette-based image understanding, IEEE Trans. Pattern Anal. Mach. Intell. 16 (2) (1994) 150–162.
[31] D. Schneider, Visual Hull 2009, available Online: http://homepages.inf.ed.ac.uk/rbf/CVonline/LOCAL_COPIES/AV0809/schneider.pdf (Accessed: 2016-06-14).
[32] J.-S. Franco, E. Boyer, Exact polyhedral visual hulls, British Machine Vision Conference (BMVC'03), Vol. 1, Norwich, United Kingdom, 2003, pp. 329–338.
[33] L. Guan, J.S. Franco, M. Pollefeys, Multi-object shape estimation and tracking from silhouette cues, IEEE Conference on Computer Vision and Pattern Recognition, 2008, pp. 1–8.
[34] R.S. Montero, E. Bribiesca, State of the art of compactness and circularity measures, Int. Math. Forum 4 (27) (2009) 1305–1335.
[35] J. Zunic, P.L. Rosin, A new convexity measure for polygons, IEEE Trans. Pattern Anal. Mach. Intell. 26 (7) (2004) 923–934.
[36] R. Blomley, M. Weinmann, J. Leitloff, B. Jutzi, Shape distribution features for point cloud analysis - a geometric histogram approach on multiple scales, ISPRS Ann. Photogramm. Remote. Sens. Spat. Inf. Sci. (2014) 9–16.
[37] J.R. Quinlan, Induction of decision trees, Mach. Learn. 1 (1) (1986) 81–106.
[38] T.M. Mitchell, Machine Learning, first ed., McGraw-Hill Education. 1997.
[39] MATLAB, version 8.6.0.267246 (R2015b), The MathWorks Inc., Natick, Massachusetts, 2015.
[40] T. Fawcett, An introduction to ROC analysis, Pattern Recogn. Lett. 27 (8) (2006) 861–874.
[41] J. Cornel, C. Lindenberg, M. Mazzotti, Experimental characterization and population balance modeling of the polymorph transformation of L-glutamic acid, Cryst. Growth Des. 9 (1) (2009) 243–252.
[42] M.R. Singh, N. Nere, H.-H. Tung, S. Mukherjee, S. Bordawekar, D. Ramkrishna, Measurement of polar plots of crystal dissolution rates using hot-stage microscopy. Some further insights into dissolution morphologies, Cryst. Growth Des. 14 (11) (2014) 5647–5661.
[43] H. Eisenschmidt, A. Voigt, K. Sundmacher, Face-specific growth and dissolution kinetics of potassium dihydrogen phosphate crystals from batch crystallization experiments, Cryst. Growth Des. 15 (1) (2015) 219–227.