SN Computer Science (2026) 7:733
https://doi.org/10.1007/s42979-026-05300-1

ORIGINAL RESEARCH

Automated Monitoring and Active Control of Jet-Breakup Ceramic-Pebble Fabrication

Markus Vogelbacher[1] · **Miao Zhang**[1] · **Oliver Leys**[2] · **Regina Knitter**[2] · **Jörg Matthes**[1]

Received: 22 April 2026 / Accepted: 12 August 2026

Abstract

To meet the stringent dimensional tolerances required for lithium-enriched ceramic granules, the Karlsruhe Institute of Technology (KIT) has devised the KALOS (KArlsruhe Lithium OrthoSilicate) technique, in which a steady molten jet is deliberately destabilised by an adjustable excitation frequency. Lithium-enriched ceramic granules are a cornerstone for the tritium-breeding blankets that will be employed in future fusion power plants, and their quality hinges on precise control of the jet-breakup process. In this work we present a novel high-speed camera-based measurement and control platform that automatically monitors and regulates pebble production. Image-processing algorithms extract droplet size, position, and inter-droplet distance distributions in real time. For this purpose, both classical image-processing methods and deep neural networks are investigated. Experimental results demonstrate that the system delivers accurate measurements and reliably adjusts the driving frequency to maintain target pebble dimensions. Thus, the proposed computer-vision system enables closed-loop control of the KALOS process, paving the way for robust, large-scale ceramic pebble fabrication.

Keywords Ceramic pebbles manufacturing · Control system · High-speed camera-based measurement system · Melt-based process · Real-time measurement and control

Introduction

Over the last several decades, nuclear fusion has emerged as a promising route to sustainable energy for future generations, owing to its inherent safety features and the very limited production of long-lived radioactive waste. The reactor's fuel cycle depends on two components: deuterium, which can be extracted from seawater, and tritium, which must be generated on-site to achieve self-sufficiency and continuous operation. To breed tritium, lithium-rich ceramic pebbles are incorporated into the reactor walls, forming pebble beds within solid breeder blankets [1, 2]. When high-energy neutrons from the deuterium-tritium fusion reaction strike these pebbles, lithium atoms transmute into helium and tritium. The tritium is then processed and recirculated to react with deuterium.

A variety of manufacturing routes have been explored worldwide to meet the demand for tritium-breeding ceramics. Lulewicz and Roux employed an extrusion-spherodisation technique [3], whereas Park *et al.* produced Li_2TiO_3 pebbles by means of a slurry-droplet wetting process [4]. Cai *et al.* introduced a piezoelectric micro-droplet jetting method that yields green Li_2TiO_3 pebbles subsequently

M. Vogelbacher, M. Zhang, O. Leys, R. Knitter and J. Matthes authors contributed equally to this work.

✉ Jörg Matthes
joerg.matthes@kit.edu

Markus Vogelbacher
markus.vogelbacher@kit.edu

Miao Zhang
miao.zhang@kit.edu

Oliver Leys
oliver.leys@kit.edu

Regina Knitter
regina.knitter@kit.edu

1 Institute for Automation and Applied Informatics, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, Eggenstein-Leopoldshafen 76344, Germany

2 Institute for Applied Materials, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, Eggenstein-Leopoldshafen 76344, Germany

Published online: 02 September 2026

sintered [5]. Other concepts include Hoshino's emulsion-based approach [6] and the use of additive manufacturing for constructing breeder structures.

At the Karlsruhe Institute of Technology (KIT), the KALOS (KArlsruhe Lithium OrthoSilicate) process was established to fabricate Advanced Ceramic Breeder (ACB) pebbles by fragmenting a molten laminar jet and instantly solidifying the droplets with liquid nitrogen [7, 8]. These pebbles consist of lithium orthosilicate reinforced by lithium metatitanate and are designated as the EU reference material for tritium breeding. The melt-based approach offers distinct benefits: it can be scaled up to satisfy the large-scale requirements of future reactors, and it enables pebble recycling without the need for wet-chemical treatments [9]. Compared with the earlier melt-spraying technique [10], the controlled breakup of a laminar jet yields a more uniform droplet size, which directly translates into a narrower pebble size distribution.

Achieving a high packing factor for the pebble bed is crucial to maximize the bulk lithium density and thus the overall efficiency of the breeder blanket. At KIT, pebble diameters are produced in the range of 250 to 1250 μm, a window that is expected to provide the necessary packing performance in the intricate geometries of blanket modules. The breakup of the molten jet follows the Plateau–Rayleigh instability: surface-tension forces eventually dominate over viscous forces, causing a droplet to separate from the jet. In the absence of external perturbations, these instabilities arise from random ambient fluctuations. By deliberately imposing a defined driving frequency, the instability wavelength can be steered, allowing precise control over the droplet and consequently pebble size distribution. This targeted control has led to higher material yields and improved monodispersity.

To investigate and quantify the jet-breakup behavior, a high-speed camera system coupled with a dedicated image-processing algorithm was developed. The coefficient of variation (CV) of droplet spacing was introduced as a quantitative measure of jet regularity. A lower CV reflects a more periodic breakup and therefore tighter size tolerances for the resulting pebbles. Adjusting the excitation frequency enables the system to maintain a low CV, ensuring stable jet fragmentation and the production of pebbles within the prescribed size limits. The present study identifies the optimal frequency window for jet-breakup and integrates a feedback loop into the process control to preserve the low-CV regime.

Recent publications in *SN Computer Science* demonstrate the relevance of computer vision and deep learning for automated visual analysis and quality inspection. Trampert *et al.* presented an AI-driven visual quality inspection system for production environments and investigated different deep-learning architectures with respect to detection performance and computational requirements [11]. Sharma *et al.* reviewed deep-learning-based object instance segmentation [12]. Suwannaphong *et al.* demonstrated the use of transfer learning for automated object detection in real-world images and compared different pretrained CNN architectures [13]. Furthermore, Montaha *et al.* investigated U-Net for automated image segmentation [14]. These studies demonstrate the relevance of deep learning, image segmentation, and automated visual inspection within the scope of *SN Computer Science*. However, the application of these methods to the real-time monitoring and closed-loop control of a high-temperature ceramic pebble production process has not been addressed.

Consequently, this paper presents a fully automatic high-speed camera based monitoring and control system for the real-time production of ceramic pebbles, first introduced in [15]. The main contributions of this work are (i) the development of a high-speed-camera-based measurement system for non-contact monitoring of molten-jet breakup under real KALOS process conditions, (ii) the development and evaluation of an image-processing pipeline for extracting droplet diameter, inter-droplet spacing, and the coefficient of variation (CV), (iii) the systematic evaluation of different DNN architectures for droplet segmentation and detection using real process data, and (iv) the integration of the vision-based measurements into a closed-loop frequency-control strategy based on the CV and Simulated Annealing. The complete system is experimentally validated by comparing the measured droplet sizes with the size distribution of the subsequently produced ceramic pebbles. The contribution therefore goes beyond the evaluation of an individual deep-learning architecture by combining high-speed imaging, image-based process monitoring, deep-learning-based droplet detection, and automatic process control within a single production system. Section "Fabrication of ceramic pebbles" outlines the theoretical background of jet-breakup and the KALOS fabrication procedure. Section "Computer-vision based measurement system for monitoring the pebbles fabrication" describes the measurement hardware and the image-processing pipeline. Section "Automatic Control System" details the closed-loop control of the droplet-generation frequency. Section "Results and discussion" reports and discusses the experimental results of the monitoring-control system, and Section Conclusions provides concluding remarks.

Fabrication of Ceramic Pebbles

As outlined in the introduction, the ceramic pebbles are generated by the KALOS process, which exploits the controlled breakup of a molten laminar jet. This section describes the advanced pebble-fabrication workflow together with the theoretical background of jet-breakup.

The KALOS method relies on the growth of Plateau–Rayleigh instabilities in a continuous molten jet. These instabilities appear as sinusoidal surface waves. Once the wave amplitude exceeds a critical value, surface-tension forces dominate viscous forces and a droplet detaches from the jet. In practice, the jet is deliberately perturbed by vibrating the process-pressure at selected driving frequencies, thereby governing the breakup dynamics and droplet formation. Because the jet velocity (v) and radius (r) remain constant during operation, the imposed frequency translates directly into a wavenumber, and a regular droplet spacing corresponds to a narrow droplet size distribution. Only frequencies within a specific band affect the breakup positively, with an optimal frequency that maximizes the growth rate of the most unstable mode. Within this band the imposed disturbances suppress ambient noise, yielding a more regular and uniform jet fragmentation. At the optimal frequency the disturbance waves amplify most rapidly, producing the lowest coefficient of variation (CV) of droplet spacing. The CV is defined as the normalized standard deviation of the spacing s:

$$CV = \frac{\sigma(s)}{\mu(s)} \tag{1}$$

where $\sigma(s)$ and $\mu(s)$ denote the standard deviation and the mean of the droplet spacing, respectively. A small CV signals a stable fabrication process and high product quality. Droplet size constitutes a second key quality indicator.

In the KALOS process the feedstock powders are first pre-reacted to obtain a composition of 70 mol% Li_4SiO_4 and 30 mol% Li_2TiO_3. The mixture is loaded into a platinum-alloy crucible, which is heated to roughly 1400°C (Figure 1) to generate a melt. To ensure stable melt flow conditions, the nozzle region is temperature-controlled at approximately 1360°C. A pressure of 320 mbar drives the melt through a 300 µm nozzle at the crucible base, forming a laminar jet. The jet disintegrates into droplets that exit the furnace and enter a cooling tower, where liquid nitrogen instantly solidifies them. Based on the high-speed recordings and the known nozzle diameter as spatial reference, the average droplet velocity can be estimated to approximately 3.4 ms^{-1}.

The solidified pebbles are collected at the bottom of the tower and transferred to the laboratory for characterization. An example of the final pebbles is shown on the right side of Figure 1.

During operation, the nozzle temperature is maintained at 1360°C to ensure stable melting conditions.

A high-speed camera captures the jet and droplets in the center of Figure 1, demonstrating the system's capability to resolve individual droplets an essential prerequisite for the measurement routine. After cooling, the pebbles are analyzed with a particle analyzer (HAVER CPA 2-1, Haver & Boecker, Germany) to obtain size distributions. These measurements provide a quantitative benchmark for validating the performance of the presented high-speed camera based monitoring system.

Computer-Vision Based Measurement System for Monitoring the Pebbles Fabrication

As introduced in the previous section, the quality of the jet-breakup based ceramic-pebble fabrication is evaluated by two key metrics: the droplet size and the coefficient of variation (CV) of the droplet spacing. To guarantee production

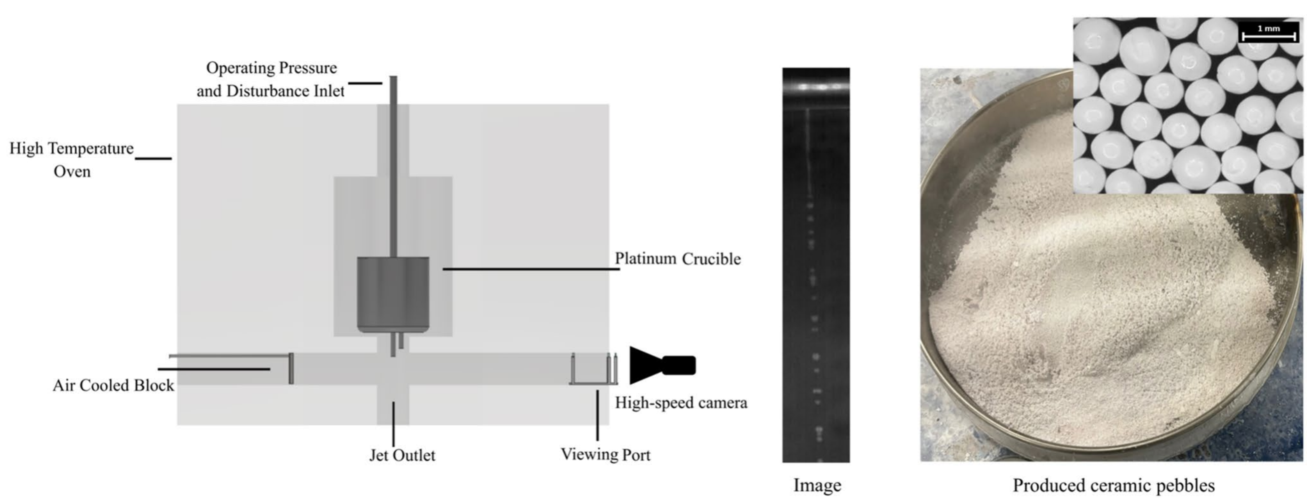

Fig. 1 High-temperature KALOS production experimental set-up and the gathered pebbles [16]

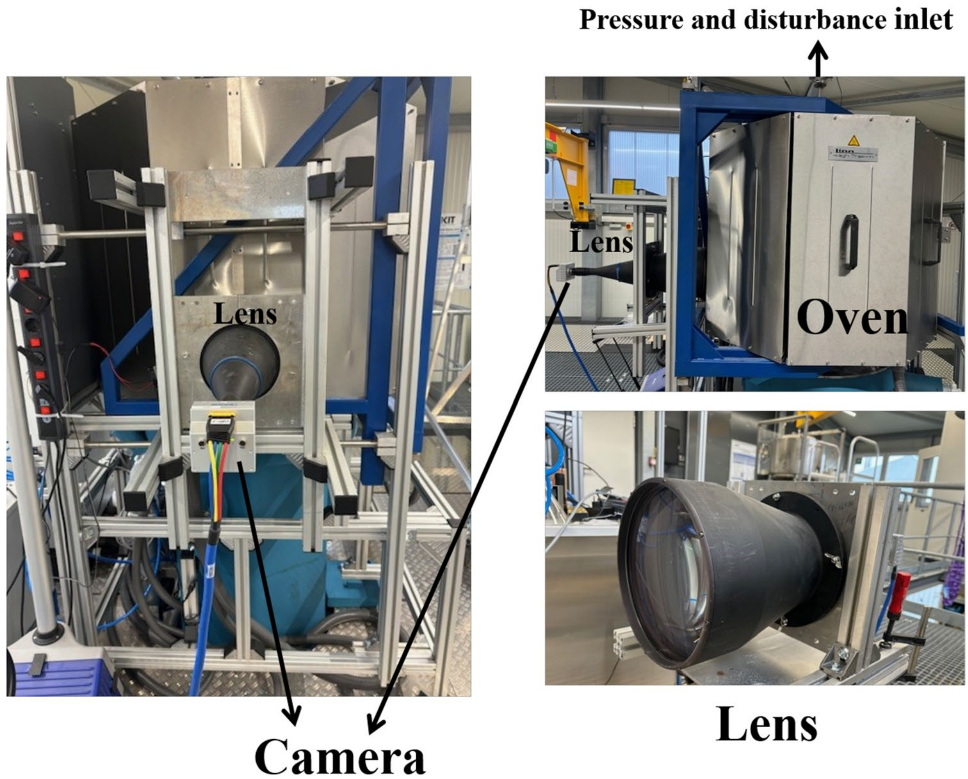

Fig. 2 High-temperature test facility and the applied high-speed camera system. [15]

quality, a dedicated measurement system capable of real-time droplet monitoring has been developed. This system comprises a hardware setup and a suite of image-processing algorithms, which are described below.

High-Speed Camera System

The core of the measurement hardware is a high-speed camera (Optronis CP70-1-M-1000) that can acquire up to 1000 frames per second (fps) at its full resolution of 1280× 1024 pixels. By defining a region of interest and discarding irrelevant image areas, the frame size can be reduced, thereby increasing the attainable frame rate. The camera together with a telecentric lens is mounted on the side of the high-temperature oven and observes the droplet formation through a window protected by specialized heat-shielding glass (Figure 2). The camera is positioned at the lower part of the oven so that the nozzle of the crucible is in direct view.

A cooled plate placed behind the jet enhances the contrast between the jet, the droplets and the background, facilitating robust segmentation. For the KALOS process the sensor active area measures 8.448 mm × 6.758 mm, with a pixel size of 6.6 µm × 6.6 µm. Calibration shows that one pixel corresponds to approximately 71.94 µm in both lateral directions. The camera supports exposure times as short as 2 µs, which suppresses motion blur during the rapid jet-breakup. For the experiments the exposure was set to 160 µs to obtain sufficient brightness while preserving the spherical appearance of the droplets. The telecentric lens (model TC23172) maintains a constant magnification over a range of object distances and exhibits minimal distortion, making it ideal for high-precision measurement of micron-sized droplets. Detailed technical specifications of the camera and lens are listed in Table 1.

Table 1 Technical details of the high-speed camera and lens. [15]

Camera	Optronis CP70-1-M-1000
Max. resolution	1280x1024 px
Frame rate for max. resolution	1040 fps
Frame rate for 1/4 max. resolution	3800 fps
Pixel size	6.6 µm
Exposure time	2 µs-1/frame rate
Lens	TC23172
Working distance	526.9 mm
Depth of field	159.2 mm

During operation, the camera records a continuous image sequence that constitutes a single measurement cycle. By analyzing the entire sequence rather than a single frame, the system mitigates random fluctuations of individual images and yields robust estimates of droplet size and spacing.

Image-Processing

After acquiring an image sequence, each frame is processed with the image-processing pipeline to extract the production-quality parameters. The measurement system performs three main tasks, determining the nozzle position, detecting the jet, and identifying the generated droplets, as illustrated in Figure 3. Since the nozzle position remains essentially constant during a sequence, it is detected only once at the beginning of the sequence to improve computational efficiency. In contrast, jet and droplet detection are carried out on every frame to guarantee accurate measurement of the droplet size, spacing, and CV.

Consequently, the essential task is to determine each droplet's size, count, and spatial position. To localize droplets accurately and suppress background noise, the nozzle and jet are detected first, thereby defining a region of interest (ROI) for droplet detection. By restricting the search to the vicinity of the jet extension, non-droplet detections are reduced and processing speed is increased, enabling real-time operation. In our implementation the jet is modelled as a straight line deviating up to $\pm 10°$ from the vertical. The ROI is a quadrilateral whose top and bottom edges are horizontal and the left and right edges run parallel to the jet line. The top edge coincides with the end of the detected jet, the bottom edge aligns with the bottom of the image, and the side edges are offset by a fixed pixel distance from the jet on each side (Figure 3). The jet line is extracted with the Hough transform [17]. The nozzle position is taken as the intersection between the jet line and the lowest horizontal line also identified by the Hough transform, which serves as the jet's starting point. Once the ROI is established, droplets inside it are detected using the appropriate methods for the given experimental conditions.

The detection of droplets inside the identified ROI is examined in greater detail below. Several approaches can be pursued. Firstly, droplets can be detected by means of conventional image-processing techniques that exploit their circular appearance. This is the method currently implemented in the measurement system. Alternatively, machine-learning methods such as deep neural networks (DNNs) may be employed, provided that a sufficient amount of ground-truth data is available for network training.

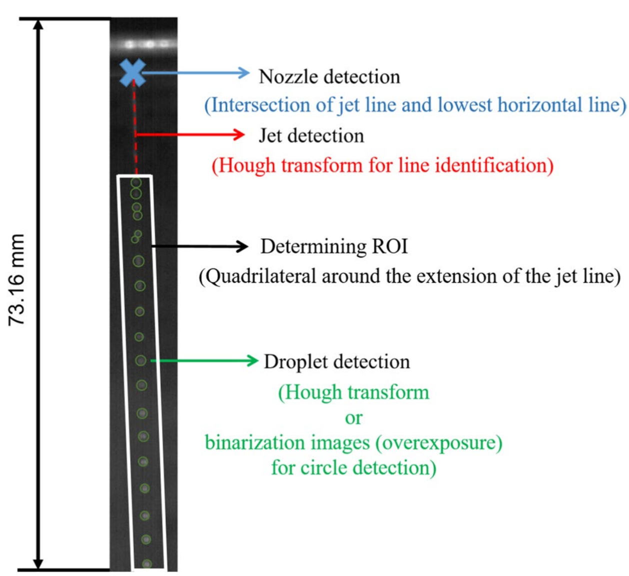

Fig. 3 Image-processing tasks of the high-speed camera-based measurement system. [15]

In the following, we first focus on detection using classical image-processing methods as used in the measurement system. Within the detected ROI, droplets can be identified either by a Hough transform or by image binarization. The Hough transform is a feature-extraction technique that locates objects through a voting process in a parameter space [17]. Although originally developed for line-segment detection, the classical Hough transform has been extended to recognize other shapes such as circles [18], which makes it suitable for droplet detection. Two detection pathways are integrated into the measurement system. For most operating conditions the Hough transform is preferred because it is robust against variations in individual pixel intensities and therefore yields high accuracy. However, its localization precision depends on the regularity of the target shape. Over-exposed or deformed droplets can bias the circle centre and radius estimates, which in turn corrupts the CV calculation. To mitigate this limitation, an alternative binary-image method is provided for cases with irregular droplet shapes. Each frame is first binarized using Otsu's automatic threshold selection [19], which adapts to the image's gray-value distribution. Circles are then extracted from the binary foreground by analyzing connected components [20]. This approach compensates for the Hough transform's sensitivity to shape deformation caused by long exposure times. Consequently, the system automatically switches to the binary-image pipeline when the exposure exceeds 500 µs.

Based on the results of the classical image-processing pipeline, we investigate deep neural networks (DNNs) as an alternative droplet-detection approach. While the currently deployed Hough-Circle method provides robust and computationally efficient detection under standard operating conditions, it strongly relies on the assumption of nearly circular droplet shapes. In practice, however, droplets may become elongated, fragmented, or partially overexposed due to motion blur and illumination variations. Such effects can reduce the reliability of purely geometry-based methods. DNN-based semantic segmentation offers greater flexibility, as it can learn more complex droplet morphologies directly from image data without relying on strict geometric assumptions. Therefore, the DNN-based approach is investigated in this work as an exploratory extension of the measurement system, with the aim of assessing its potential for future applications involving non-ideal droplet shapes and more challenging imaging conditions. To this end, a dedicated dataset was created from two independent experimental campaigns. From each campaign, 1,000 grayscale images were randomly selected, resulting in a total dataset of 2,000 images. Each image was cropped to a fixed resolution of 688 $\times$ 224 pixels and contains between 3 and 16 droplets under varying jet-breakup conditions. Representative examples are shown in Figures 4(a) and 4(b).

The dataset was randomly divided into 90 % training data (1,800 images), 5 % validation data (100 images), and 5 % test data (100 images). To improve generalization and reduce overfitting, online data augmentation was applied during training, including random rotations within $\pm 30°$ and horizontal translations within ± 50 pixels. This increases the variability of droplet orientations and positions and prevents the networks from learning spurious geometric features specific to the training data.

Ground-truth labels for training and testing are derived from the Hough-Circles detections described earlier: the circle centers and radii are converted into pixel-wise masks for semantic segmentation. In the masks, black pixels represent background and white pixels belong to a droplet (Figure 4(c)).

Four neural-network architectures were evaluated:

1. *Simple custom network* (simple DNN) consisting of six 2-D convolutional layers.
2. *U-Net*, a standard encoder-decoder segmentation network with an encoder depth of four levels, including symmetric skip connections between encoder and decoder stages [21].
3. *DeepLabV3+* with a **ResNet18** backbone (established for classifying image data) [22].
4. *DeepLabV3+* with a **ResNet50** backbone (established for classifying image data) [22].

For the DeepLabV3+ architectures, transfer learning was employed using ImageNet-pretrained backbone networks. The pretrained encoder weights were fine-tuned on the droplet dataset, while the segmentation heads were adapted to the two-class problem (background and pebble). In contrast, the custom DNN and U-Net were trained from the ground up without any pre-training. To ensure a fair comparison between all architectures, identical training settings were applied in all experiments. Training was performed in MATLAB using stochastic gradient descent with momentum (SGDM), a mini-batch size of 16, an initial learning rate of 0.01, and a piecewise learning-rate schedule. Validation was carried out every five iterations. All models were trained and evaluated on a workstation equipped with an Intel Core i7-8750H CPU (6 cores, 12 threads, 2.20 GHz) and 32 GB RAM, without GPU acceleration.

For semantic segmentation performance on the test set we report *Accuracy* (ratio of correctly classified pixels to total

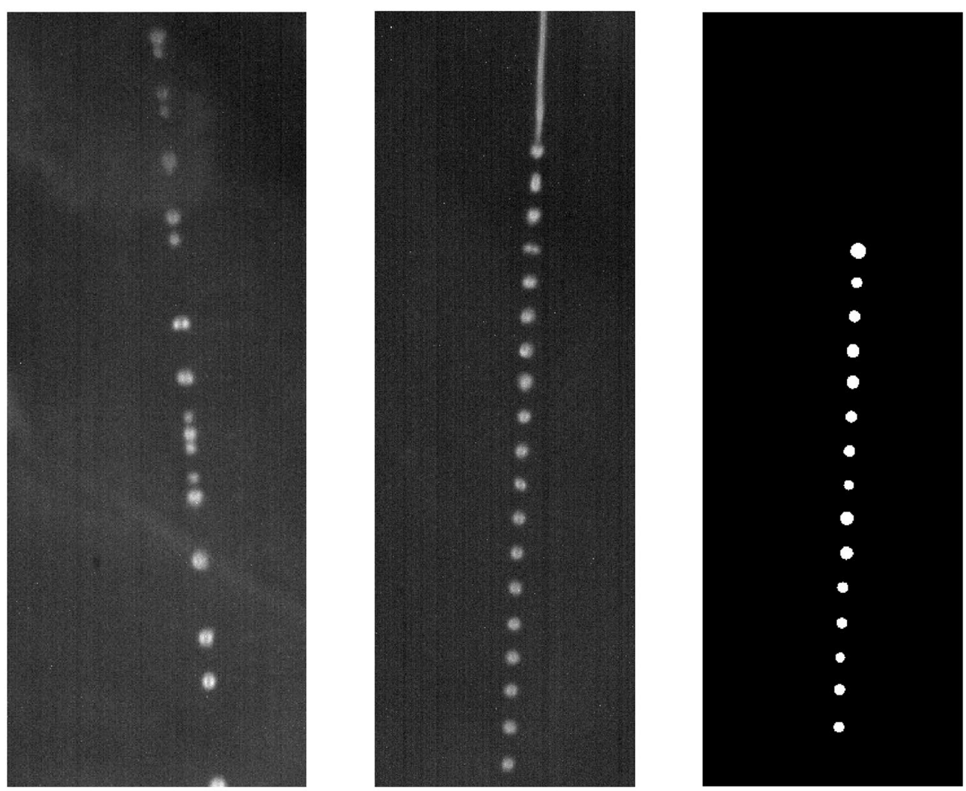

(a) (b) (c)

Fig. 4 **a** and **b** - example images from the two experimental series that are used as training and test data for the DNNs. (c) - example of a ground truth mask generated with the Hough Circles detection

Table 2 Semantic segmentation results of the investigated DNNs

	Simple DNN	U-Net	Deeplabv3+ ResNet18	Deep-labv3+ ResNet50
Mean Accuracy	0.796	0.684	0.880	0.883
Mean IoU	0,717	0.662	0.836	0.831

pixels) and *Intersection over Union* (IoU) (overlap between predicted droplet region and ground truth region divided by their union). U-Net performs worse than the DeepLab-based architectures but outperforms the simple custom DNN in detection robustness. The semantic segmentation results of all investigated architectures are summarized in Table 2. DeepLabV3+ with ResNet50 achieves the highest mean accuracy (0.883), closely followed by DeepLabV3+ ResNet18 (0.880), while ResNet18 provides the highest IoU (0.836). The U-Net architecture yields lower segmentation performance than both DeepLab variants and also under-performs the simple custom DNN in terms of pixel-wise accuracy and IoU.

For the measurement system, however, a pixel-wise classification is unnecessary. Only droplet centroids are required for downstream CV calculations. Consequently, we further analyze the individual droplet detections produced by the DNNs.

When inspecting the semantic-segmentation outputs of the DNNs, it becomes evident that droplet regions are sometimes highly fragmented (cyan rectangles in Figure 5 center) or contain internal holes (green rectangles in Figure 5 center). These imperfections arise, for example, from droplet motion during exposure or from illumination inhomogeneities across the field of view. To recover the compact, nearly circular shape that is expected for a droplet and that is

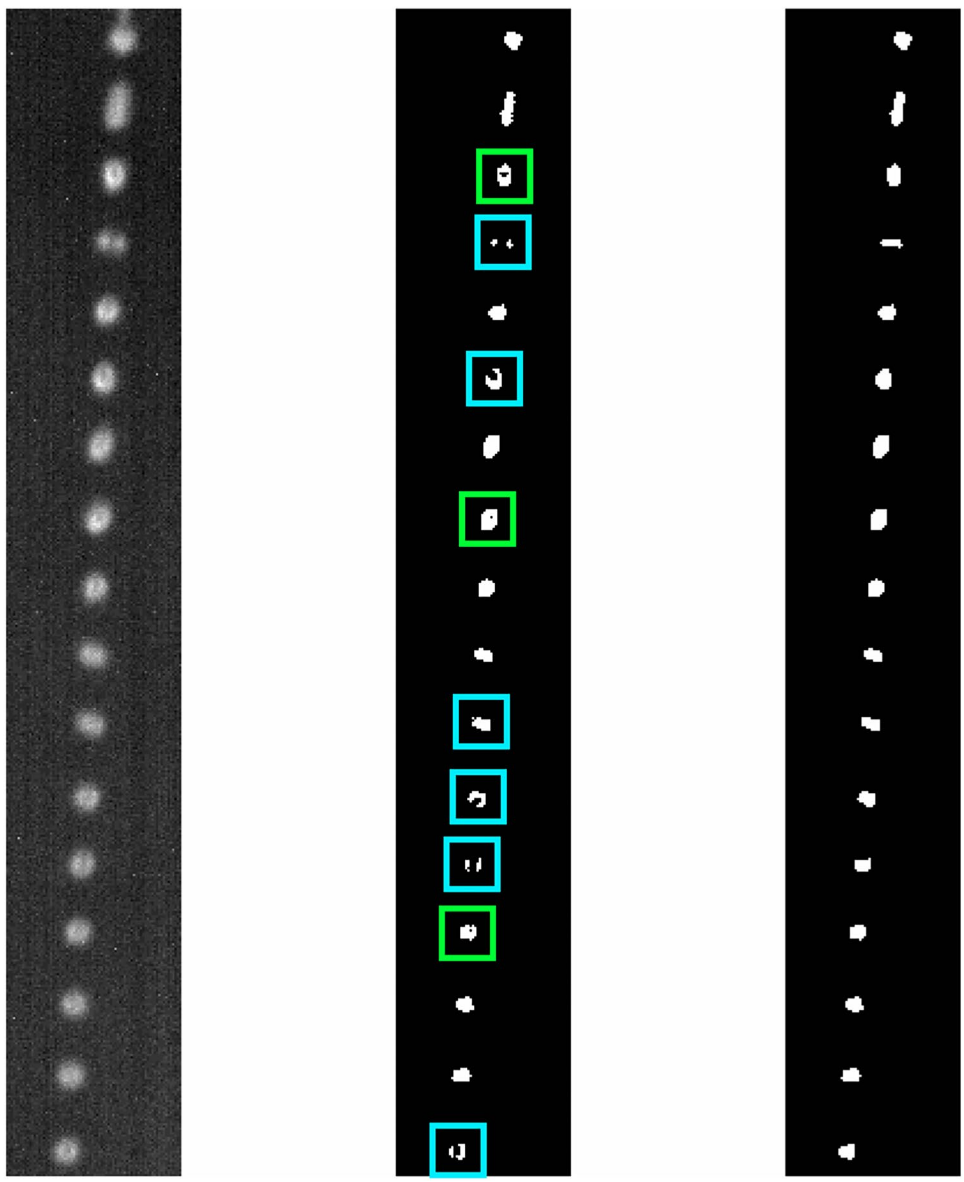

Fig. 5 Results of the morphological post-processing applied to the semantic-segmentation outputs of the DNNs (left: original image, center: raw output of the simple DNN, right: output after morphological post-processing)

required for subsequent quantitative analysis, a morphological post-processing step is applied. The binary masks are first subjected to a closing operation to bridge small gaps, followed by a hole-filling procedure that eliminates interior voids. After this processing the droplets exhibit the desired morphology (Figure 5 right). Connected components are then extracted, their centroids computed, and the resulting positions are treated as droplet detections.

A detection from a DNN is counted as a true positive (tp) if its centroid lies within the radius of a corresponding Hough Circle detection. Otherwise, it is a false positive (fp). Conversely, any Hough Circle detection without a matching DNN centroid is a false negative (fn). With these definitions, the subsequent evaluation metrics for the DNNs can be calculated:

- Precision:

$$P = \frac{tp}{tp + fp} \tag{2}$$

- Recall:

$$R = \frac{tp}{tp + fn} \tag{3}$$

- F-Score:

$$\text{F-Score} = 2 \cdot \frac{P \cdot R}{P + R} \tag{4}$$

The results of the DNNs on the test dataset are summarized in Table 3.

It is evident that all investigated architectures achieve a high mean recall, ranging from 0.952 to 0.975, indicating that most droplets are successfully detected. In contrast, the precision values show larger differences between the architectures. The simple DNN attains the lowest precision of 0.596, which leads to a comparatively low F-score of 0.727 and indicates a substantial number of false-positive detections. Among all investigated architectures, DeepLabV3+ ResNet18 achieves the highest F-score of 0.907, representing the best overall balance between precision and recall. DeepLabV3+ ResNet50 shows a comparable performance with an F-score of 0.890, but the slightly lower precision indicates that the increased network depth does not necessarily improve droplet localization. U-Net reaches an intermediate F-score of 0.816, outperforming the simple DNN despite its weaker semantic-segmentation metrics. This suggests that U-Net generates more coherent droplet regions, which benefits centroid extraction for the downstream CV calculation. Overall, the gap between precision and recall across all architectures points to a remaining proportion of false-positive detections, especially for the simpler network structures.

Table 3 Droplet-detection results of the investigated DNNs on the test set

	Simple DNN	U-Net	Deeplabv3+ ResNet18	Deeplabv3+ ResNet50
Mean Precision	0.596	0.715	0.883	0.847
Mean Recall	0.959	0.975	0.952	0.962
F-Score	0.727	0.816	0.907	0.890

The observed behavior can be assessed by inspecting an example detection result from the test set, as shown in Figure 6.

The example shows that the DNNs detect more droplets than are present in the ground truth masks. These additional detections correspond to droplets whose shapes deviate from perfect circles and therefore were missed by the Hough Circle method, which was used to generate the ground truth. Because the training labels inherit this bias, the networks learn to reproduce it. The shallow architecture of the simple DNN limits its capacity to fully fit the training data, resulting in many non-circular detections that are erroneously counted as false positives. With a manually refined training set, the DeepLabV3+ ResNet18 could achieve comparable results.

Overall, DNN-based droplet detection shows the potential to replace the Hough Circle approach, especially for non-circular droplets. However, the computational cost per frame limits real-time applicability. Figure 7 shows the computation times required for all 2,000 images of the combined training and test datasets on a CPU-based workstation equipped with an Intel Core i7-8750H processor (2.20 GHz, 6 physical cores, 12 logical processors) and 32 GB RAM with no GPU acceleration.

The Hough Circle method remains the fastest (baseline). The simple DNN requires roughly twice the processing time, while DeepLabV3+ ResNet18 needs almost seven times longer. The runtime comparison is limited to these architectures, as they represent the two most relevant trade-offs between computational cost and detection performance among the investigated models. Consequently, DeepLabV3+ ResNet18 is not feasible for the current online requirement. Given its superior detection performance and still acceptable runtime, the simple DNN constitutes a viable alternative to Hough Circles for droplet detection in the measurement system.

Production Quality Measures

In the measurement system, the droplet diameter is obtained directly by a circle-detection algorithm, whereas the CV

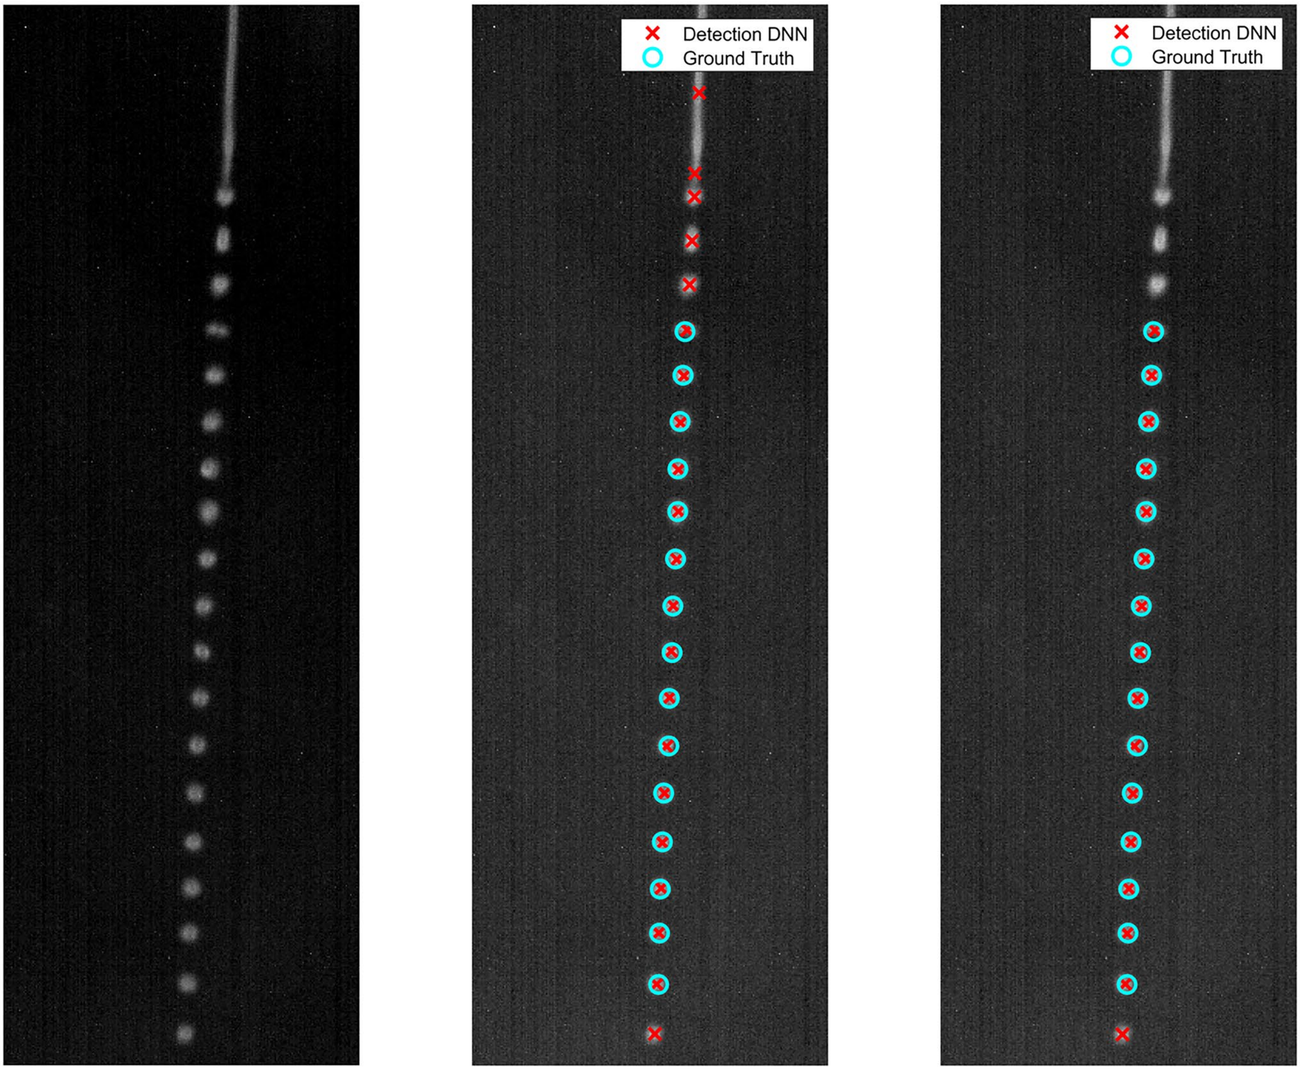

Fig. 6 Comparison of detection results on a test-set example (left: original image, center: simple DNN, right: DeepLabV3+ ResNet18)

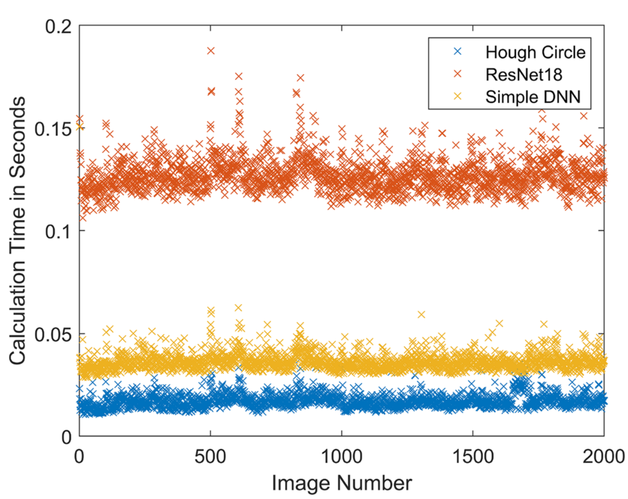

Fig. 7 Computation times for the different droplet detection methods

value is calculated according to the definition given in equation 1. The standard deviation of the droplet spacing, $\sigma(s)$, is expressed as

$$\sigma(s) = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(s_i - \mu(s))^2}, \tag{5}$$

and the mean spacing, $\mu(s)$, is computed by

$$\mu(s) = \frac{1}{N}\sum_{i=1}^{N} s_i \tag{6}$$

Prior to evaluating the CV, the detected droplets are ordered from top to bottom according to their vertical positions. The distance between two consecutive droplets is then

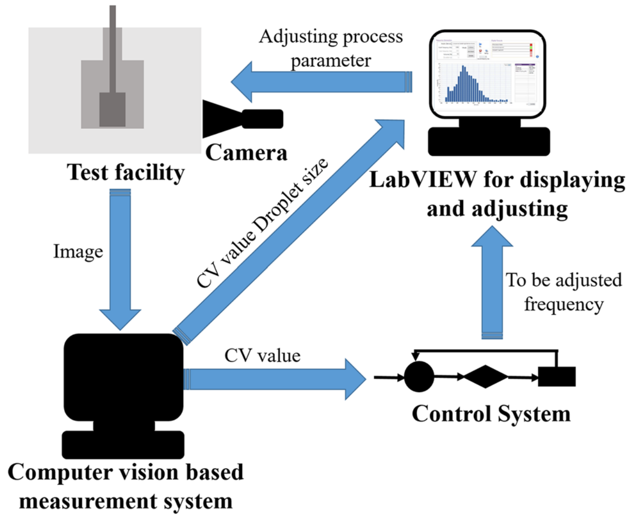

Fig. 8 Schematic of the introduced measurement systems and control systems for monitoring and controlling the KALOS process. [15]

determined and denoted as s_i. Here, i indexes the ith spacing out of the total of N spacings that correspond to $N+1$ detected droplets. To reduce the influence of fluctuations that arise from individual frames, the system reports the average CV over an image sequence comprising M frames. The averaged standard deviation is calculated as

$$\sigma = \frac{\sum_{j=1}^{M} \sigma(j)}{M}, \tag{7}$$

where $\sigma(j)$ denotes the CV computed for the jth image.

Automatic Control System

As introduced in the previous sections, the coefficient of variation (CV) quantifies the regularity of the molten-jet breakup and therefore represents the most critical parameter for ceramic-pebble production. Under constant jet diameter and flow conditions, a low CV directly corresponds to a narrow droplet size distribution. The present control system uses the CV as its control basis. Although theory predicts that frequencies within a certain range should improve jet-breakup regularity (and thus lower the CV), experimental studies have shown that the CV's response to the driving frequency can be irregular [23]. Even inside the effective frequency band, the CV may increase abruptly at several frequencies, which we attribute to resonances within the system that disrupt the jet's regularity. To guarantee a stable droplet generation and high-quality pebble fabrication, a real-time control loop that adjusts the driving frequency is required. Since the thermophysical process dynamics evolve on a significantly slower timescale than the image acquisition and processing chain, the computational delay of the measurement system can be neglected for control purposes. Whenever the CV exceeds a predefined threshold, the controller adapts the frequency to restore droplet-generation stability. By integrating this controller, the KALOS process can be monitored and regulated automatically in real time. The complete schematic of the system is shown in Figure 8.

The measured CV value is transmitted from the measurement subsystem to the controller, which then determines the appropriate driving frequency. This frequency command is sent to LabVIEW, which implements the adjustment of the process parameter and simultaneously displays both the parameter and the measurement results in real time. Consequently, LabVIEW serves as the visualization and operation platform for the entire production line. The measurement system, the controller, and LabVIEW together realize a closed-loop monitoring and control of the droplet-generation process. The following paragraphs focus on the controller.

Figure 9 illustrates the control logic. A low CV indicates satisfactory droplet quality. When the measured CV exceeds the predefined threshold, the current driving frequency must be modified. Theory predicts that the optimal frequency for a given setup does not deviate strongly from the principal optimum. Typically, each setup possesses several nearby optimal frequencies, and adjustments are modest. Therefore, the controller defines a permissible variation range around the current set point.

The controller operates in two modes. In standard mode (Figure 9), frequency adjustment stops as soon as the CV falls below the threshold. In optimization mode, iterative frequency adjustments continue until the global optimum

Fig. 9 Schematic of one iteration of control logic in standard mode. Note: [15]

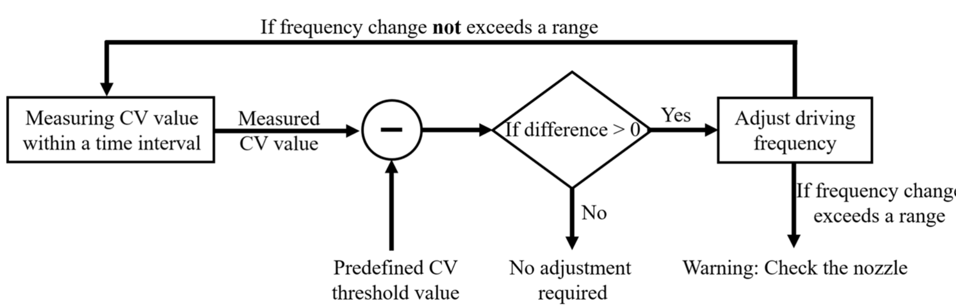

is found or a maximum number of iterations is reached, even if the CV has already dropped below the threshold. If frequency adjustments fail to reduce the CV to the desired level, the system alerts the operator to inspect the nozzle. In such cases the problem is usually caused by melt wetting or a partially blocked outlet rather than by the driving frequency.

At startup, the user selects an initial frequency based on theory and prior experiments. For the KALOS process this value is 1000 Hz. The measurement system records the CV at this frequency, then evaluates a second frequency between 750 Hz and 1000 Hz and a third between 1000 Hz and 1250 Hz. Using these three (frequency, CV) pairs, a quadratic polynomial is fitted to approximate the CV-frequency relationship. The polynomial's minimum, i. e. the smallest estimated CV, is taken as a new driving frequency. After measuring the CV at this frequency, the polynomial is refitted with the four data points, and the process repeats. The iteration terminates when the change in frequency becomes negligible, indicating that further polynomial refinement would yield only marginal benefit. To refine the optimum, the controller subsequently applies a global search algorithm based on Simulated Annealing (SA) [24]. Compared with other global optimization techniques such as the Burg algorithm [25, 26] (suited to spectral estimation) and Particle Swarm Optimization [27, 28] (prone to local-optimum traps), SA offers strong global search capability, high adaptability, and straightforward implementation.

SA exploits the analogy between function minimization and the annealing process in statistical mechanics [29]. After the polynomial fit yields an initial frequency, SA generates a candidate frequency in the neighborhood of the current solution. The CV associated with the candidate is measured and compared with the CV of the incumbent solution. If the candidate CV is lower, the new frequency is accepted. Otherwise, acceptance occurs with probability

$$P = \exp\left(-\frac{\Delta CV}{f}\right), \tag{8}$$

where ΔCV is the CV difference and f the current frequency, following the Metropolis criterion [30]. The new frequency is then updated by

$$f_{new} = \alpha \cdot f \tag{9}$$

with α drawn from a normal distribution centred at 1 and a standard deviation of 0.01. In standard mode, the SA iteration stops as soon as the CV falls below the threshold. In optimization mode, the iteration continues until either the global minimum is identified or a predefined maximum number of iterations is reached. The global minimum is considered found when the best CV does not improve for 12 consecutive iterations (approximately 100 seconds of experiment time). The controller finally outputs the frequency that yielded the lowest CV. Because SA alone can be computationally intensive for real-time control, we combine it with the preceding polynomial fit. The fit limits the search space, enabling SA to converge rapidly while still escaping local minima [31].

If the controller is initially idle and later triggered by a CV exceeding the threshold, the polynomial fit is initialized at the current frequency, and a new candidate frequency is randomly chosen within ±250 Hz. Finally, to guarantee that the KALOS process can respond to frequency changes, the measurement subsystem records several image sequences before each CV evaluation. This acquisition is synchronized by the measurement trigger signal.

Results and Discussion

The image-processing pipeline and the control algorithm were described in the preceding sections. This section presents and discusses the results obtained with these methods. For the computer-vision based measurement system, the two key metrics, coefficient of variation (CV) and droplet diameter, are extracted from the high-speed camera images. To validate the measurement accuracy, the solidified pebbles are measured after cooling and their diameters are then compared with the values computed by the vision system. Concerning the control system, the evolution of the driving frequency and the corresponding CV values are visualized, allowing a clear view of the frequency adjustment process.

Computer-Vision Based Measurement System

At first, droplets were generated by applying driving frequencies ranging from 0 to 5000 Hz to obtain an overview of the frequency's influence. Extensive experimental data and theoretical considerations show that frequencies above 3000 Hz have only a negligible effect on the jet-breakup. Therefore the full 0 to 5000 Hz sweep guarantees that the entire active range is covered. Figure 10 presents high-speed camera images of the jet-breakup captured at 500 fps.

With increasing driving frequency the average droplet diameter first rises and then falls, while the CV trend is not immediately apparent from the images. The smallest CV is observed near 1000 Hz. To visualize the frequency dependence of both CV and diameter more clearly, the processed results are displayed in Figure 11. The horizontal axis denotes the driving frequency and the vertical axes show CV (Figure 11(a)) and median diameter (Figure 11(b)). Each green symbol represents a single-frame measurement,

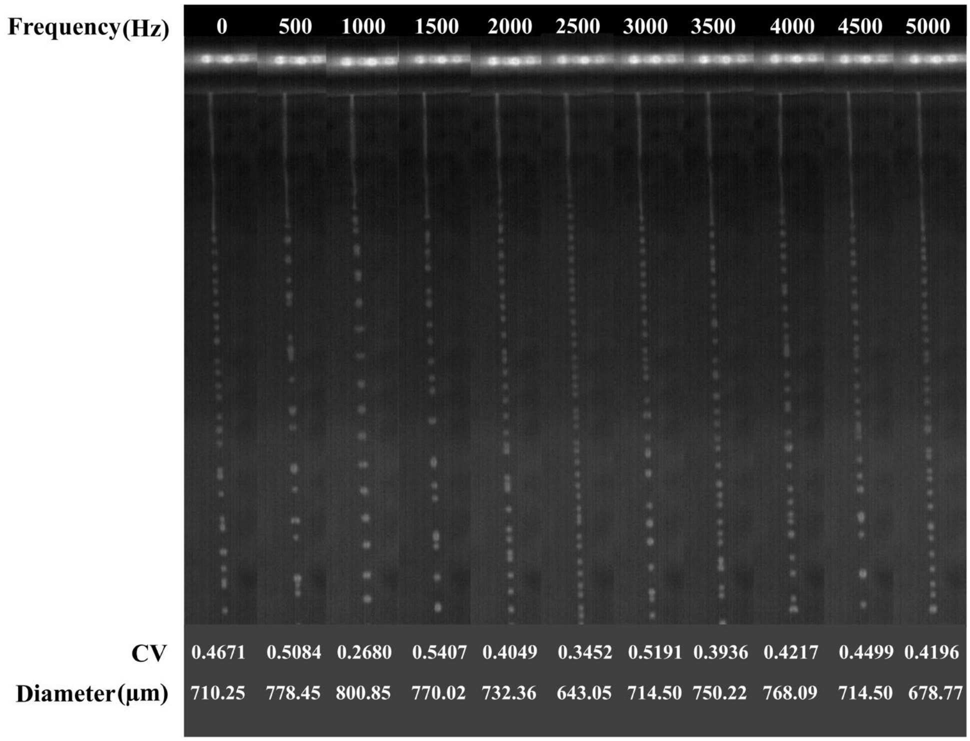

Fig. 10 Effect of the driving frequency on the appearance of the molten ceramic jet-breakup. The corresponding CV values and droplet diameters (in μm), calculated by the image-processing algorithms, are shown at the bottom [16]

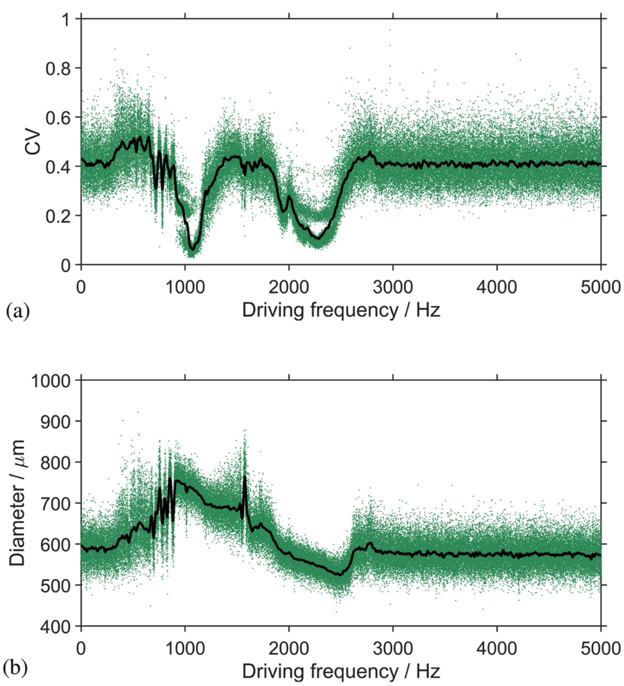

Fig. 11 Output of the measurement system. **a** The CV value and **b** The median diameter of the droplets. [15]

and a robust linear regression [32] over a sliding window of 200 frames produces the black smoothing curves.

Figure 11(a) reveals that frequencies between approximately 800 Hz and 2600 Hz affect the breakup, and two distinct minima appear around 1000 Hz and 2300 Hz. Between these minima irregular jet behavior is visible (e.g., at 1500 Hz and 2000 Hz in Figure 10), indicating internal resonances. Regarding droplet size, the median diameter is about 750 μm at 1000 Hz and reduces to 550 μm at 2500 Hz, following an approximately linear decline. Again, the resonance region around 1500 Hz shows larger deviations from the smoothed trend. Because the CV optimum occurs near 1000 Hz, this frequency was chosen as the initial set point for the control experiments. Thereafter, we also conducted experiments to investigate the droplet generation at the driving frequency of 1000 Hz with respect to the droplet diameter.

Droplet diameters measured by the high-speed camera and the diameters of the solidified pebbles obtained from the particle analyzer are compared in Figure 12. In Figure 12(a) a green histogram (bin width 5 μm) displays the measured droplet sizes together with an approximated normal probability density curve. The black dashed line marks the droplet size with the highest probability at 760 μm to 765 μm, while the orange dashed line indicates the mean of the fitted normal distribution (732.05 μm). The majority of droplets lie between 500 μm and 1000 μm.

Figure 12(b) shows the pebble-size distribution measured after solidification. The histogram and the red cumulative-distribution line indicate that most pebbles fall between 600 μm and 800 μm, with a modal size around 750 μm, which matches the droplet sizes observed in Figure 12(a). Considering solidification shrinkage and possible deformation during free fall, a slight size shift between droplets and final pebbles is expected. The close agreement between the

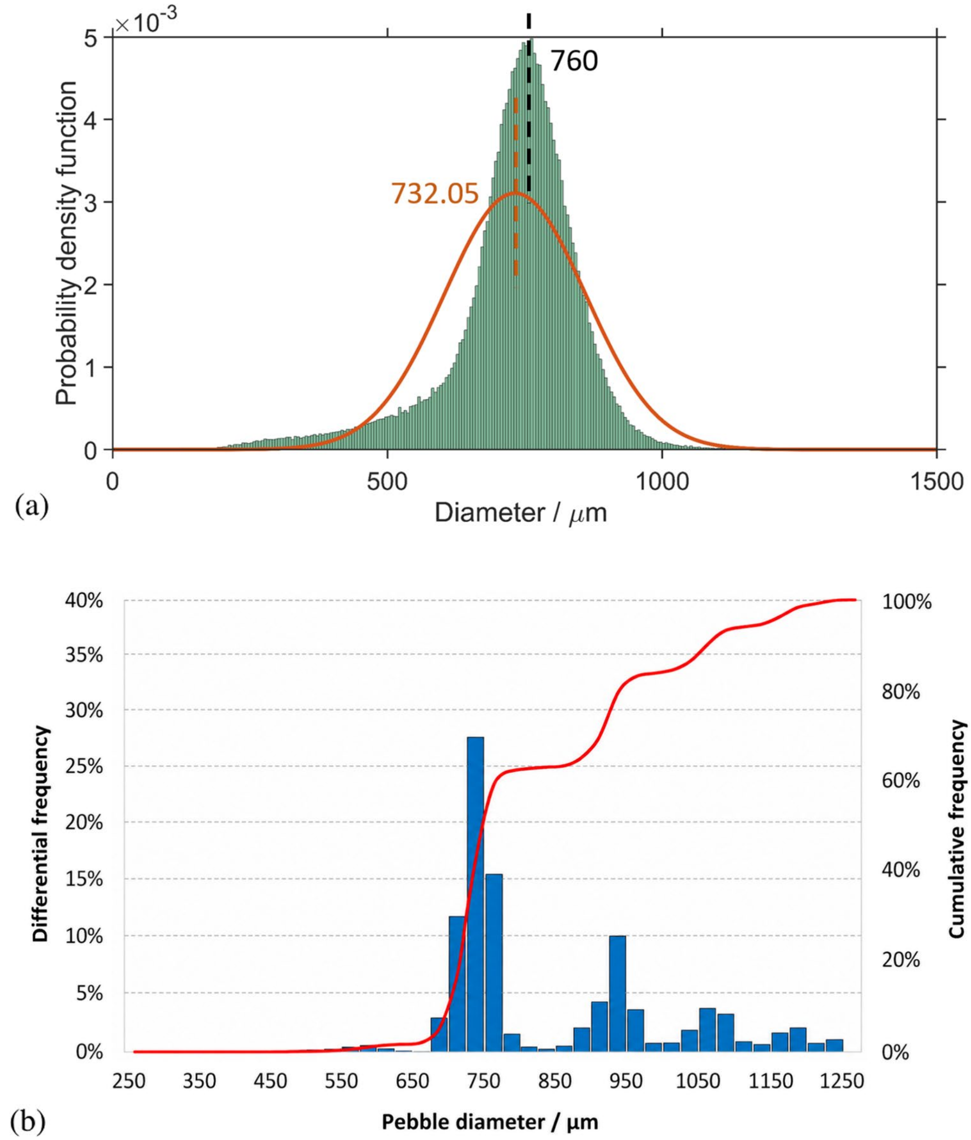

Fig. 12 Diameter of droplets and gathered pebbles produced at the driving frequency of 1000 Hz. **a** Measured droplet diameters and its distribution, **b** Measured diameter distribution of the pebbles provided by particle analyzer HAVER CPA 2-1. [15]

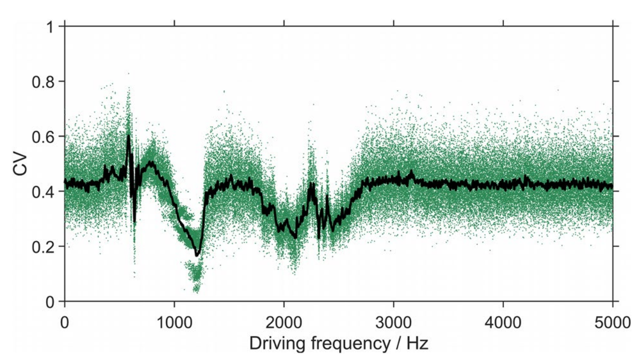

Fig. 13 Plot of CV value versus driving frequency for one experiment. Frequency ramp from 0-5000 Hz. [15]

image-processing results and the independent particle analysis confirms the accuracy and reliability of the high-speed camera measurement system.

Control System

To illustrate the regulation behavior of the controller, a dedicated validation experiment was performed. The driving frequency was ramped from 0 to 5000 Hz while images were recorded at 500 fps. The resulting frequency versus CV curve for this experiment can be seen in Figure 13. As in Figure 11(a), green symbols denote the CV of individual frames and the black line shows a 200-frame robust regression smoothing. Compared with the experiment shown in Figure 11, this experiment exhibits a broader CV distribution and larger fluctuations, making modulation more demanding. Two local minima around 1000 Hz and 2000 Hz are still visible, but the CV near 1000 Hz appears fragmented because internal resonances cause an irregular jet-breakup.

To visualize this effect, Figure 14 presents pairs of images captured at adjacent frequencies. At 1200 Hz the white box

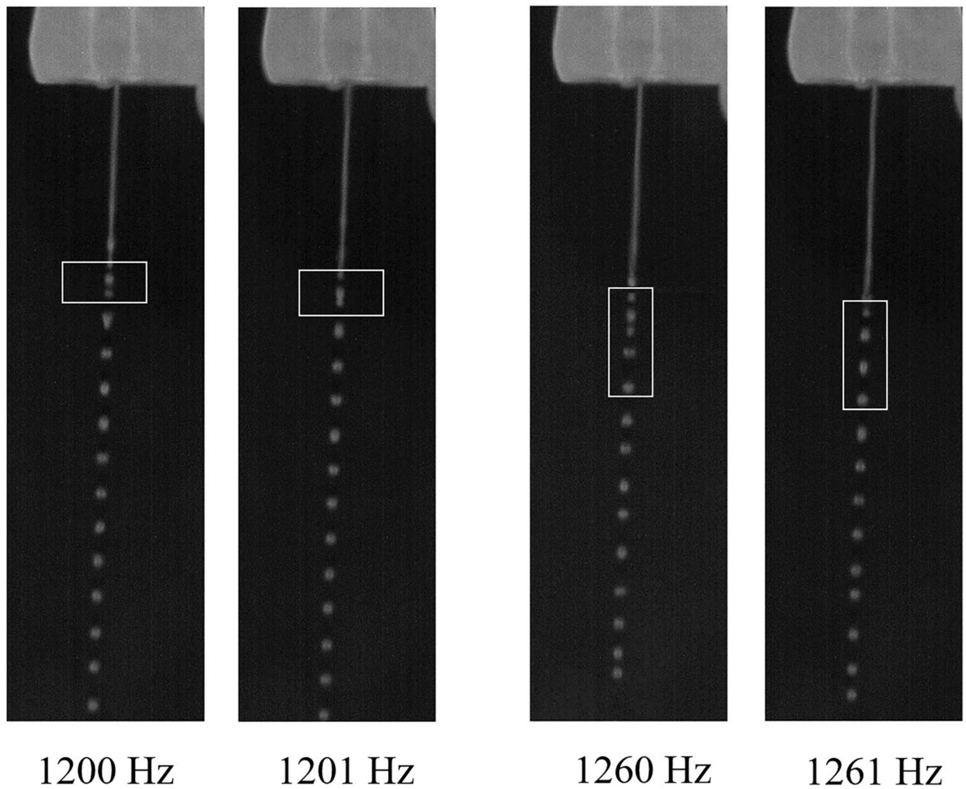

Fig. 14 Captured images of adjacent frequencies. [15]

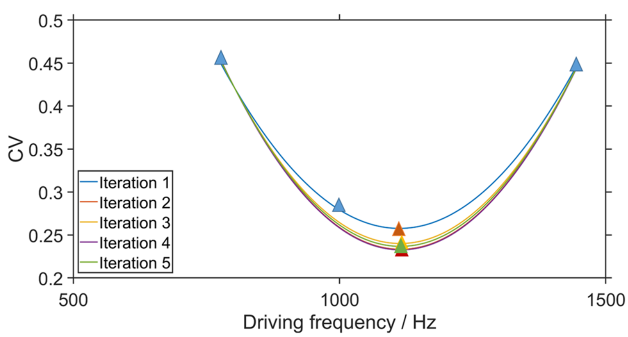

Fig. 15 Process of polynomial fitting. Different curves represent different iterations. The triangles in the image are the data points for each fitting, whose colors also identify the number of the iteration. Note: This content is cited from [15]

contains two detectable droplets, whereas at 1201 Hz the same region shows a single, larger droplet. A more pronounced instability is observed between 1260 Hz and 1261 Hz, where the jet-breakup becomes erratic, and the CV rises sharply. These observations underline the necessity of a control system capable of compensating unknown resonances. The delayed response of the controller can partly attenuate random irregularities.

For the validation the controller was operated in optimization mode, i. e. the iterative search continued until the optimal frequency was reached even if the CV threshold had already been under-cut. As described in Section Automatic Control System, three data points around 1000 Hz were first selected and a quadratic polynomial fitted to them. The fitted minimum was then used as a new input, and the procedure was repeated. Figure 15 displays five successive polynomial fits. The iteration stopped when successive curves differed only marginally. The final curve's minimum served as the starting point for the subsequent global optimization.

Using the polynomial output as an initial guess, the controller applied Simulated Annealing (SA) to refine the optimum. The temporal evolution of the SA search is shown in Figure 16. Green dots indicate measured points, while red asterisks mark frequencies actively probed by the regulator. Table 4 lists the iteration number, the experimentally applied frequency, the resulting CV, and the internal updates performed by the algorithm.

Initially the controller explored frequencies around 1114 Hz, but the CV did not improve significantly, prompting a reduction to 1050 Hz. The algorithm then shifted the search toward 1200 Hz, where a markedly lower CV was observed, signaling proximity to the optimum. The global minimum was identified in iteration 23 (frequency = 1174 Hz, CV = 0.041). Since no improvement occurred over the following 12 iterations, the optimization terminated after iteration 35. The identified optimum can be retained for subsequent production until the CV exceeds the predefined threshold, at which point a new optimization cycle is triggered.

Despite the fact that a rough estimate of the local minimum is obtained by the polynomial fitting, the complete

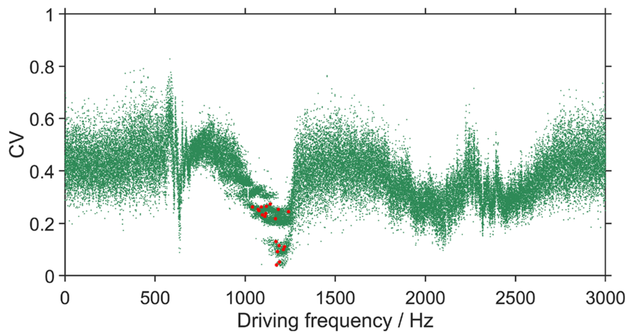

Fig. 16 Temporal results of the control output by the SA algorithm. Green points denote measurement points and red asterisks indicate the data points that are scanned by regulation. [15]

Table 4 Driving frequency and CV value of each control step. [15]

Iteration	Experimented frequency (Hz)	Experimented CV	Updated frequency (Hz)	Updated CV
1	1114	0.238	1114	0.238
2	1109	0.227	1109	0.227
3	1072	0.253	1072	0.253
4	1051	0.309	1072	0.253
5	1124	0.310	1072	0.253
6	1188	0.115	1188	0.115
7	1245	0.2138	1188	0.115
8	1180	0.251	1188	0.115
9	1195	0.116	1195	0.116
10	1189	0.194	1195	0.116
11	1201	0.095	1201	0.095
12	1235	0.100	1235	0.100
13	1268	0.301	1235	0.100
14	1300	0.439	1235	0.100
15	1222	0.095	1222	0.095
16	1213	0.103	1222	0.095
17	1199	0.110	1199	0.110
18	1172	0.2049	1199	0.110
19	1180	0.2506	1199	0.110
20	1173	0.092	1173	0.092
21	1137	0.208	1173	0.092
22	1195	0.115	1173	0.092
22	1158	0.242	1173	0.092
22	1183	0.084	1183	0.084
23	**1174**	**0.041**	**1174**	**0.041**
24	1156	0.250	1174	0.041
25	1166	0.214	1174	0.041
26	1171	0.131	1174	0.041
27	1212	0.100	1174	0.041
28	1190	0.052	1190	0.052
29	1185	0.247	1190	0.052
30	1175	0.216	1190	0.052
31	1211	0.096	1190	0.052
32	1199	0.111	1190	0.052
33	1204	0.118	1190	0.052
34	1226	0.140	1190	0.052
35	1179	0.190	1190	0.052

Table 5 Number of iterations of the SA algorithm with and without polynomial fitting. [15]

Start frequency (Hz)	800	1000	1200
Iteration count with fitting	34	35	32
Iteration count without fitting	137	50	76

control loop still requires several dozen iterations. Each iteration involves a waiting period of 2 to 3 s for processing an image sequence, resulting in a total optimization time of roughly 3 min, which is acceptable for the KALOS process. Without the polynomial pre-fit, the regulation would take considerably longer and would adversely affect production throughput. Table 5 quantifies the contribution of the polynomial fitting by comparing the number of SA iterations needed with and without this pre-processing step.

As shown in Table 5, the iteration count with polynomial fitting varies only slightly with different starting frequencies because the fitted curve places the SA start point close to the true optimum. This not only enhances control efficiency but also reduces the random bias associated with the initial frequency. In contrast, SA without a preceding fit requires many more iterations, and the required search time strongly depends on the chosen start frequency. For example, an initial frequency of 800 Hz (far from the optimum) leads to a staggering 137 iterations, whereas a start near the optimum needs far fewer steps. Consequently, the polynomial fitting stage is essential: it shortens the regulation time and markedly diminishes the controller's sensitivity to the initial frequency.

Conclusions

The paper introduces a novel high-speed camera based measurement and control system for the automatic production of lithium-rich ceramic pebbles, which are essential for tritium-breeding blankets in fusion reactors. By employing dedicated image-processing algorithms, the system monitors droplet generation in real time and extracts two key metrics: the droplet diameter and the coefficient of variation (CV) of the droplet spacing. Because the CV directly reflects the regularity of the jet-breakup, it serves as the primary indicator of fabrication performance and as the feedback signal for adjusting the driving frequency.

For droplet detection, we investigated several deep-learning-based semantic segmentation architectures in addition to the conventional Hough-Circle pipeline: a lightweight six-layer CNN (simple DNN), U-Net, and DeepLabV3+ with both ResNet18 and ResNet50 backbones. All networks were trained on a curated dataset of 2,000 grayscale images, with ground-truth masks derived from the Hough Circle detections and enhanced by online data augmentation.

Among the investigated architectures, DeepLabV3+ ResNet18 achieved the best overall droplet-detection performance with an F-score of 0.907, while DeepLabV3+ ResNet50 reached the highest segmentation accuracy (0.883). U-Net showed intermediate detection performance (F-score 0.816), whereas the simple DNN achieved the lowest F-score (0.727) but remained computationally the most efficient among the neural-network approaches. The comparison shows that increasing architectural complexity improves segmentation and detection robustness, but does not necessarily yield better localization performance for the downstream control task. Importantly, the imperfect ground-truth labels, biased toward nearly circular droplets due to their derivation from Hough Circle detections, limit the ability of the networks to fully learn irregular droplet morphologies. This explains why the DNN-based methods did not surpass the classical Hough-based baseline in overall robustness. From a real-time perspective, the lightweight CNN required approximately twice the inference time of the Hough Circle method, whereas DeepLabV3+ ResNet18 was about seven times slower, rendering it unsuitable for the current online control requirements. Nevertheless, the results demonstrate that deep-learning-based detectors can provide a viable alternative for future upgrades, particularly for handling non-circular droplet geometries.

According to the jet-breakup theory presented in Section "Fabrication of ceramic pebbles", a low CV corresponds to a stable, periodic breakup—conditions that are required for high-quality pebble production. Using the measurement system, we investigated how the driving frequency influences both CV and droplet size. Experiments were carried out with a frequency sweep (to obtain a broad overview) and with a fixed frequency to study the optimal operating point. The results show that the driving frequency strongly affects production within a limited band, with the optimum for the employed nozzle located at approximately 1000 Hz.

A detailed study at the optimal frequency (1000 Hz) was performed, after which the solidified pebbles were collected and analyzed with a particle-size analyzer. The measured droplet diameters and the analyzed pebble diameters agree closely, confirming the accuracy and reliability of the computer-vision measurement chain.

Since the driving frequency is the critical control variable, the measurement system is indispensable for selecting an appropriate initial frequency and for providing continuous feedback. When the CV exceeds a predefined threshold, the control module adapts the frequency. The adaptation proceeds in two stages: first, a quadratic polynomial fit identifies a rough minimum, which serves as the initial guess for a global optimization based on Simulated Annealing (SA). The SA algorithm rapidly converges to the global optimum, enabling real-time frequency regulation. Due to stochastic fluctuations in the droplet formation process, individual CV measurements may vary even at constant operating conditions. Therefore, the optimization is based on temporally averaged CV values to obtain a robust estimate of the optimal driving frequency.

Overall, the experimental campaign demonstrates that the integrated measurement-and-control framework reliably monitors and stabilizes ceramic-pebble fabrication. Future work will assess the robustness of the system under varying process conditions and will extend the image-processing pipeline to quantify droplet generation rates and velocities.

Acknowledgements This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Program (Grant Agreement No 101052200 - EUROfusion). OpenAI's language model (gpt-oss:120b) was used solely for copy-editing purposes, including refining, correcting, and clarifying portions of the authors' original text.

Author Contributions MV, MZ, OL, RK and JM carried out the conception of the work. MV, MZ, OL and JM contributed to design, analysis, and interpretation of the data. MV and JM drafted the manuscript. MZ, OL and RK critically revised the manuscript for important intellectual content and gave final approval of the version to be published. All authors have read and approved the final manuscript and agree to be accountable for all aspects of the work.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability Not applicable.

Materials Availability Not applicable.

Code Availability Not applicable.

Declarations

Conflict of interest The authors have no competing interests to declare that are relevant to the content of this article.

Ethical approval and consent to participate Not applicable.

Consent for publication Not applicable.

References

1. Knitter R, Chaudhuri P, Feng YJ, Hoshino T, Yu I-K. Recent developments of solid breeder fabrication. J Nucl Mater. 2013;442(1, Supplement 1):420–4. https://doi.org/10.1016/j.jnucmat.2013.02.060. 15th International Conference on Fusion Reactor Materials
2. Hernández FA, Arbeiter F, Boccaccini LV, Bubelis E, Chakin VP, Cristescu I, et al. Overview of the hcpb research activities in eurofusion. IEEE Trans Plasma Sci. 2018;46(6):2247–61. https://doi.org/10.1109/TPS.2018.2830813.
3. Lulewicz J, Roux N. Fabrication of li 2tio 3 pebbles by the extrusion-spheronisation-sintering process. Journal of Nuclear Materials - J NUCL MATER. 2002;307:803–6. https://doi.org/10.1016/S0022-3115(02)00981-9.
4. Park Y-H, Cho S, Ahn M-Y. Fabrication of li2tio3 pebbles using pva–boric acid reaction for solid breeding materials. J Nucl Mater. 2014;455(1):106–10. https://doi.org/10.1016/j.jnucmat.2014.05.027. Proceedings of the 16th International Conference on Fusion Reactor Materials (ICFRM-16).
5. Cai J, Xu G, Lu H, Xu C, Hu Y, Cai C, et al. Preparation of tritium breeding li2tio3 ceramic pebbles via newly developed piezoelectric micro-droplet jetting. J Am Ceram Soc. 2022. https://doi.org/10.1111/jace.18320.
6. Hoshino T. Trial examination of direct pebble fabrication for advanced tritium breeders by the emulsion method. Fusion Eng Des. 2014;89: https://doi.org/10.1016/j.fusengdes.2013.12.052.
7. Heuser JM, Kolb MHH, Bergfeldt T, Knitter R. Long-term thermal stability of two-phased lithium orthosilicate/metatitanate ceramics. J Nucl Mater. 2018;507:396–402. https://doi.org/10.1016/j.jnucmat.2018.05.010.
8. Leys O, Leys JM, Knitter R. Current status and future perspectives of eu ceramic breeder development. Fusion Eng Des. 2021;164:112171. https://doi.org/10.1016/j.fusengdes.2020.112171.
9. Leys O, Bergfeldt T, Kolb MHH, Knitter R, Goraieb AA. The reprocessing of advanced mixed lithium orthosilicate/metatitanate tritium breeder pebbles. Fusion Eng Des. 2016;107:70–4. https://doi.org/10.1016/j.fusengdes.2016.04.025.
10. Knitter R, Alm B, Roth G. Crystallisation and microstructure of lithium orthosilicate pebbles. J Nucl Mater. 2007;367–370:1387–92. https://doi.org/10.1016/j.jnucmat.2007.04.002. Proceedings of the Twelfth International Conference on Fusion Reactor Materials (ICFRM-12)
11. Trampert P, Mantowsky S, Schmidt F, Masiak T, Schneider G. Ai-driven toolbox for efficient and transferable visual quality inspection in production. SN Computer Science. 2025;6(5):481.
12. Sharma R, Saqib M, Lin C-T, Blumenstein M. A survey on object instance segmentation. SN Computer Science. 2022;3(6):499.
13. Suwannaphong T, Chavana S, Tongsom S, Palasuwan D, Chalidabhongse TH, Anantrasirichai N. Parasitic egg detection and classification in low-cost microscopic images using transfer learning. SN Computer Science. 2023;5(1):82.
14. Montaha S, Azam S, Rakibul Haque Rafid A, Hasan MZ, Karim A. Brain tumor segmentation from 3d mri scans using u-net. SN Computer Science 2023;4(4):386
15. Zhang M, Leys O, Vogelbacher M, Knitter R, Matthes J. A novel automatic monitoring and control system for induced jet breakup fabrication of ceramic pebbles. In: Proceedings of the 22nd International Conference on Informatics in Control, Automation and Robotics, 2025;25–36. https://doi.org/10.5220/0013665600003982. 37.12.01; LK 01
16. Zhang M, Leys O, Vogelbacher M, Knitter R, Matthes J. A high-speed camera-based measurement system for monitoring and controlling the induced jet break-up fabrication of advanced ceramic breeder pebbles. The International Journal of Advanced Manufacturing Technology. 2024;136:1253–65. https://doi.org/10.1007/s00170-024-14887-9.
17. Hart PE. How the hough transform was invented [dsp history]. IEEE Signal Process Mag. 2009;26(6):18–22. https://doi.org/10.1109/MSP.2009.934181.
18. Ballard DH. Generalizing the hough transform to detect arbitrary shapes. Pattern Recogn. 1981;13(2):111–22. https://doi.org/10.1016/0031-3203(81)90009-1.
19. Otsu N. A threshold selection method from gray-level histograms. IEEE Trans Syst Man Cybern. 1979;9(1):62–6. https://doi.org/10.1109/TSMC.1979.4310076.
20. Haralick RM, Shapiro LG. Computer and Robot Vision, Volume I, pp 28–48. Addison-Wesley, Boston 1992
21. Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. CoRR **abs/1505.04597**, 2015 arxiv: abs/1505.04597
22. Chen L-C, Zhu Y, Papandreou G, Schroff F, Adam H. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation: 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part VII, 2018;833–851. https://doi.org/10.1007/978-3-030-01234-2_49
23. Leys O, Waibel P, Matthes J, Knitter R. Ceramic pebble production from the break-up of a molten laminar jet. In: Proc of ILASS Europe 2019: 29th Conference on Liquid Atomization and Spray Systems. 2019.
24. Laarhoven PJM, Aarts EHL. Simulated Annealing: Theory and Applications. Dordrecht: Kluwer Academic Publishers; 1987.
25. Orfanidis SJ. Optimum Signal Processing: an Introduction. New York: Macmillan; 1985.
26. Vos K. A fast implementation of burg's method. 2013. p. 2–4.
27. Pedersen MEH. Good parameters for particle swarm optimization. 2010. https://api.semanticscholar.org/CorpusID:7496444
28. Bonyadi MR, Michalewicz Z. Particle swarm optimization for single objective continuous space problems: a review. Evol Comput. 2017;25(1):1–54. https://doi.org/10.1162/EVCO_r_00180.
29. Banchs RE. Simulated annealing. Technical report, The University of Texas at Austin; 1997.
30. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equation of state calculations by fast computing machines. J Chem Phys. 1953;21(6):1087–92. https://doi.org/10.1063/1.1699114.
31. Delahaye D, Chaimatanan S, Mongeau M. Simulated annealing: From basics to applications, 2018;1–35 https://doi.org/10.1007/978-3-319-91086-4_1
32. Andersen R. Modern methods for robust regression. Sage University Paper Series on Quantitative Applications in the Social Sciences. 2008;1654–62. https://doi.org/10.4135/9781412985109.