https://doi.org/10.2352/ISSN.2169-4451.2017.32.94

Reprinted from
Journal of Imaging Science and Technology® 60(4): 040401-1–040401-9, 2016.

Measurement of Inkjet Drop Volume—The Role of Image Processing

Graham D. Martin▲
Department of Engineering, University of Cambridge, United Kingdom
E-mail: gdm31@cam.ac.uk

William C. Price
Emmanuel College, University of Cambridge, United Kingdom

Ian M. Hutchings
Department of Engineering, University of Cambridge, United Kingdom

Abstract. *The measurement of the volumes of small* (10–100 μm) *liquid drops is important in a number of fields including inkjet printing, liquid dispensing and spraying. This article explores the use of synthetic, constructed images, representing shapes with precisely known volumes, and real photographic images of inkjet drops to compare a number of image processing methods designed to estimate drop volume. The synthetic images were generated with a range of sizes, background gray levels and degrees of blur and noise. The image processing methods were chosen to represent a range of approaches, some very simple and some more complex. A comparison of the results from these methods shows that they responded differently to various image features. The process described in this article could be used to compare other existing or new processing methods, and the results should be valuable in the development of standard methods for drop measurement. © 2016 Society for Imaging Science and Technology.*

INTRODUCTION

Inkjet drop volume is an important property which is measured by inkjet print head manufacturers and some users. Manufacturers need to make these measurements during the development phase of the head, during the development of inks and during manufacturing. Some users will also make these measurements for initial set-up, to detect problems and also for precision tuning where the deposited volume is critical to the process. The human eye is very good at detecting lines or stripes in a printed image caused by variation in drop volume between different nozzles, and some processes (for example OLED screen manufacture) require precise volume control for each deposited drop.[1] Drop volume is also an important measurement during research on inkjets to study, for example, the fluid dynamics of drop formation or variations in drop volume over time ("first drop" effects). This measurement is also important where accurate dispensing is required, for example in pharmaceutical manufacture[2] or drug testing calibration.[3] Hence, it can be necessary to measure the performance of each nozzle in an array at least once and often repeatedly.

Our purpose in this work is to develop a process to compare different approaches to measuring the volume of individual inkjet drops using both simulated images of drops and real photographic images. We then use these images to compare several different image processing techniques.

In this article we address measurements using conventional optical imaging, although there are other means to measure drop volume. For example, an average drop volume can be measured by collecting and weighing a known number of drops, making an allowance for any evaporation loss.[3] This technique can be very accurate but requires great care and is perhaps not suited to measuring many thousands of nozzles, nor is it useful to assess drop reproducibility which is important in some applications. Using digital holography, drop position and diameter can be accurately measured simultaneously for many drops over a field of view at least equal to the size of the imaging sensor.[4,5] However, while providing very accurate measurements (sub-micrometer in drop radius and position) it is a technique that is experimentally difficult and expensive, and the computation involved in processing the results can be time-consuming.

The data most desired from images of inkjet drops are the direction of the drop motion, its speed and its volume.[6] The volume can be obtained from a single, calibrated image as discussed below. Drop velocities are sometimes calculated from a single image by knowing the time between the drop ejection trigger and the image exposure, and measuring the distance of the drop from the nozzle. This provides an average over the distance measured and is subject to errors arising from the time taken for the drop to emerge from the nozzle after the trigger signal, but it may be adequate, particularly if only a comparison between nozzles is required. Better estimates of velocity can be made from two or more successive images of the same drop (or different drops at different phases if consistency is assumed). Similarly, drop

▲ IS&T Member.

Received Mar. 9, 2016; accepted for publication Apr. 29, 2016; published online June 10, 2016. Associate Editor: Kye-Si Kwon.

direction can be estimated from a single image if the nozzle position is known, or from at least two images showing the same (or consistent) drops. Practically, the velocity is often measured by tracking the tip of the drop, and this may well be adequate particularly for nozzle to nozzle comparison, but a more consistent and precise measure would be to track the center of mass of the total drop.

IMAGING

Nearly all direct imaging is carried out by illuminating drops from behind and viewing the drops in silhouette with a magnifying lens and digital camera.[7] The illumination most commonly used is a short duration flash synchronized with drop ejection. This is combined with a camera (single frame or video) which will capture several flashes per frame, and hence the captured image shows an overlay of several individual drops. When combined with video this has the advantages that the effect of changes can be observed immediately, and that different stages in the drop formation can be observed by changing the timing of the flash relative to the drop trigger. The disadvantages are that any variation from drop to drop will degrade the image, and that transient effects and any satellites, which tend to vary event to event, will be difficult to observe. Alternatively, if the flash intensity and sensor sensitivity are high enough, a single flash can be used per frame, providing a record of a single jetting event.[8] If multiple images of the same event are desired then a high speed camera is required with continuous illumination (at least for the event duration).[9]

A high magnification is desirable to improve the estimation of drop edge positions and reduce the effects of pixelation (although sub-pixel measurements are possible, as discussed below). This has to be balanced against the need to observe drops at different stages of their evolution, to see drops from more than one nozzle at the same time and, in some cases, observe drops at least twice within the frame size to estimate velocity. The flash duration (or high speed camera shutter speed) needs to be short enough to minimize motion blur. For example, a drop with a velocity of 5 ms^{-1} will travel 5 μm during a flash of duration 1 μs. Hence, for accurate size and volume estimations, flash durations significantly less than 1 μs are needed.

One important issue for volume estimation is that, during their evolution, ink drops are seldom spherical. After it detaches from the nozzle, the drop will possess a tail, or ligament, which, over time, collapses into the drop or, commonly, breaks up into separate satellite drops, which may, or may not, recombine with the main drop.[10] The dynamics of breaking up, collapsing and re-joining can also cause the drops to oscillate about a spherical shape. Fortunately, observation suggests that the drop shapes are locally (over the height of at least one pixel) circularly symmetric about an axis perpendicular to the nozzle plate. This assumption, together with the assumption that the image vertical axis is aligned with this axis, is made in all the methods discussed below.

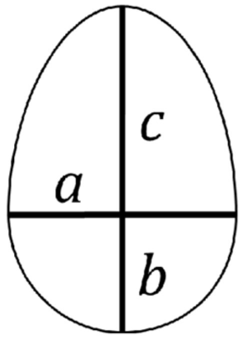

Figure 1. The oval shape created by joining two half-ellipsoids (ellipses in projection).

Table I. The set of test ovals of different sizes: *a*, *b* and *c* are the dimensions of the objects in pixels, *V* is the volume in voxels and R_{eq} is the radius of an equivalent sphere with the same volume.

a	*b*	*c*	*V*	R_{eq}
2	3	4	58.64	2.41
4	5	8	435.63	4.70
9	10	15	4241.15	10.04
18	20	31	34,607.78	20.22
54	60	93	9,34,410.19	60.65
90	100	150	4,241,150.08	100.41

TEST IMAGES

To test the selected processing techniques, some synthetic, simulated drop images were generated and real photographic drop images were also captured.

A Matlab program was written to create synthetic images. A non-spherical "drop" with a single axis of symmetry was constructed by joining two half-ellipsoids (ellipses in projection) symmetrically about a vertical axis, as shown in Figure 1. Pixels within the projected oval shape were given a foreground level and pixels outside the shape were given a background level. Pixels crossing the edge of the shape were given a level between foreground and background in proportion to the areas inside and outside the shape. The point at which the axes cross was not centered on a pixel.

Fig. 1 illustrates the construction of the ellipsoids. This constructed image could be considered to be the most accurate digital image possible for its dimensions with the given foreground and background levels. The volume (V) of the resulting solid, where the surfaces are perfectly smooth curves, is given by

$$V = \frac{2\pi a^2(b+c)}{3}, \qquad (1)$$

where a is the radius of the common axis, and b and c are the radii of the other axes.

Several synthetic images were generated, as listed in Table I with their dimensions in pixels and volumes in voxels

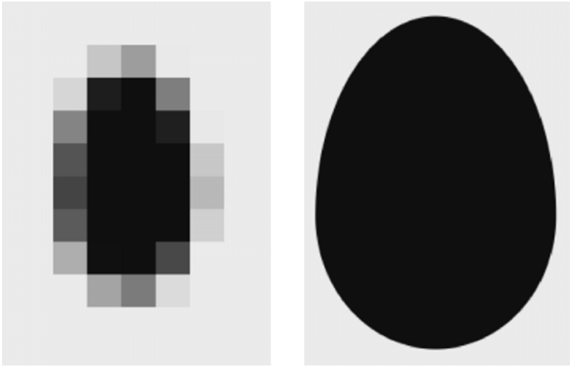

(a) (b)

Figure 2. The smallest ($a = 2$, $b = 3$, $c = 4$, $R_{eq} = 2.4$) and largest ($a = 90$, $b = 100$, $c = 150$, $R_{eq} = 100$) simulated objects.

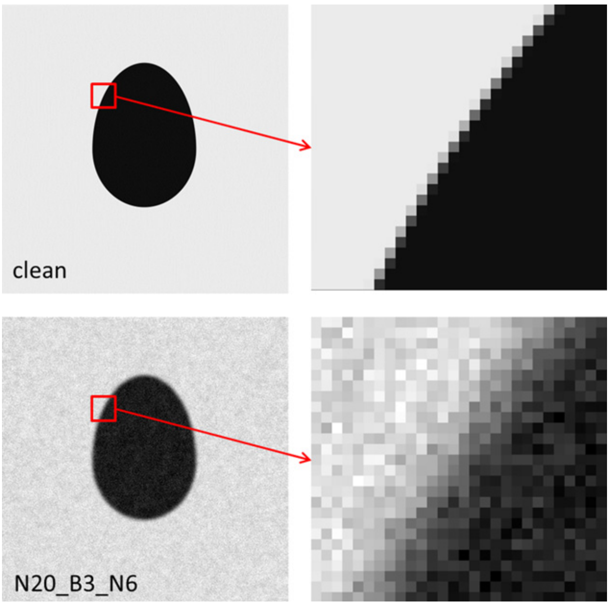

Figure 3. An object ($a = 90$, $b = 100$, $c = 150$, $R_{eq} = 100$) before (top) and after (bottom) the filter sequence N20_B3_N6.

(cubic pixels); the foreground level was set at 15 (near black) and the background level at 235 (near white) within the range 0–255. Figure 2 compares the smallest and largest images.

To investigate the effects of blur and image noise, two of the images ($a = 18$, $b = 20$, $c = 31$ and $a = 90$, $b = 100$, $c = 150$) were processed using standard Photoshop filters (Photoshop v12.1). Blur was added with the Gaussian Blur filter, which has a strength measured in pixels, and noise was added with the Gaussian Noise filter, with a percentage strength value. A variety of images were created in this way, with filters applied in sequence in some cases. For example, an image created by applying a 20% noise filter followed by a 3-pixel blur filter followed by a 6% noise filter (abbreviated as N20_B3_N6) is illustrated in Figure 3.

Photographic images of liquid drops were captured to test the processing techniques. Figure 4 shows some examples

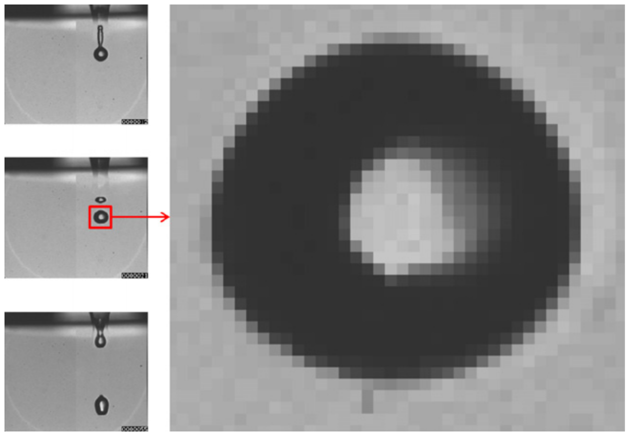

Figure 4. Examples of real images obtained during the capture of a single event using a high speed camera.

from a set of images of a drop ejection event captured using a high speed camera (Shimadzu HPV-1), with a shutter open time of approximately 500 ns and illuminated with a high power, long duration (1 ms) flash lamp (Adept Electronics Ltd, model 1002). The images were captured at a rate of 10^6 s^{-1}. The liquid was pure water ejected from a MicroFab single nozzle inkjet print head with a nominal nozzle diameter of 40 μm.

These photographic images track the development of a single ejection event, starting with a detached drop with a ligament, which separates to become a satellite drop, which then recombines with the main drop again. At some points during this event the drops appear to be approximately spherical, but for much of the time they clearly are not. Because water is transparent, there are bright areas in the centers of the drops or ligaments caused by light refraction. When a satellite is separate then volume estimates are made for each separate drop and satellite and summed to give the total volume. One advantage of using a single event captured by a high speed camera is that the total volume of the ejected liquid should remain the same during the sequence of images (if evaporation is negligible), providing a useful test of the method of estimating the total liquid volume.

IMAGE PROCESSING TECHNIQUES

To estimate the volume of a drop it is first necessary to differentiate between the drop and the image background as well as possible, and then derive a volume from the two-dimensional edge information. There is a vast literature on image segmentation and edge detection.[11–14] However, by making some reasonable assumptions and using the particular features of drop images we can focus on methods most likely to provide the necessary information. The methods investigated here do not provide an exhaustive study of all possible methods. We assume that:

- images are 8-bit gray level images (256 levels),
- liquid regions are bounded by a continuous boundary,
- they do not contain physical holes or bubbles,

- the foreground (inside) near the edges of the drop images is significantly darker than the background (outside) near the edges and
- any other significant features in the image can be masked out of the assessment region.

Several of the methods discussed below rely on an initial thresholding step to differentiate between the drop and the background. In this study we have used a technique, available as a standard Matlab function, based on Otsu's method for finding the appropriate threshold.[15] The algorithm assumes that the image has a bi-modal histogram (foreground pixels and background pixels); it then calculates the optimum threshold by separating the two classes so that their combined spread (intra-class variance) is minimized.

We used five different methods to evaluate drop volume, which are described in turn.

Thresholding and Sum Slice Volumes (Designated as Method T)

In this method the drop area in the image is isolated using Otsu's thresholding method, and the volume of each line of horizontal pixels in the drop is calculated by assuming that it represents a circular disk; disk volumes are then summed to find the total volume of the drop.[8] This technique allows the volume of non-spherical drops to be calculated provided that they are circularly symmetrical about the vertical axis. This is the method used in this work. A variant of this method calculates the volume of truncated conical slices formed by joining the center points of adjacent row edges. A further variation would assume that drop edges were smooth and that interpolation between pixel rows would improve the overall result. With these points, sub-pixel slices could be chosen to estimate volumes.[16] If necessary, a calibration could be carried out (for example, by measuring the weight of a number of drops) and the value of the threshold level adjusted appropriately.

Threshold with Sub-pixel Linear Interpolation (TSP)

This method is the same as method T with the addition of a sub-pixel estimate of each left and right edge by using the values of the pixels on each side of the initial edge and interpolating a better edge position. Figure 5 illustrates the calculation. If E is the position of the right-hand edge of the object determined by thresholding to the nearest pixel, T is the threshold value, P_i is the value at the center of the pixel inside the drop, P_o is the value at the center of the pixel outside the drop and x is the offset of the interpolated edge in pixels from E, then

$$x = \frac{2T - P_o - P_i}{2(P_o - P_i)}. \tag{2}$$

As defined,

$$P_i \le T \le P_o. \tag{3}$$

Hence,

$$-0.5 \le x \le 0.5. \tag{4}$$

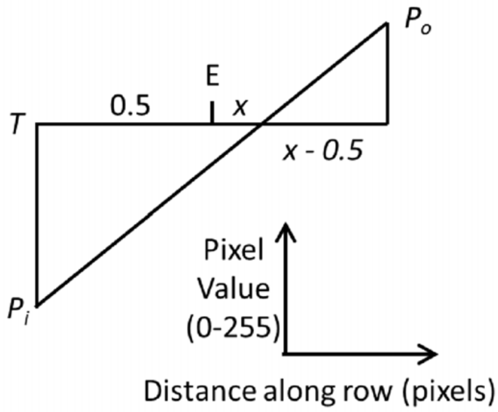

Figure 5. Sub-pixel linear interpolation.

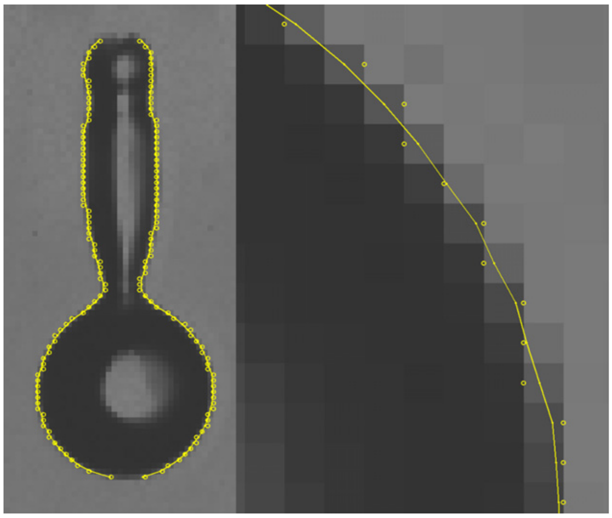

Figure 6. Pixel (light circles, yellow online) and sub-pixel (light line, yellow online) fits to one of the images of the real drop event using the T and TSP techniques.

Having found the sub-pixel edges, they are used to estimate the diameter and calculate the volume of the pixel disk. The volume of the whole drop is calculated, as in method T, by summing slice volumes.

Figure 6 presents an example of a fit to one of the images captured with the high speed camera, showing the edges estimated by using the initial thresholding and the sub-pixel process.

Pappus (P)

The second theorem of Pappus states that the volume, V, of a solid of revolution generated by the revolution of a lamina about an external axis is equal to the product of the area of the lamina, A, and the distance traveled by the lamina's geometric centroid:

$$V = 2\pi r A, \tag{5}$$

where r is the perpendicular distance from the axis of revolution to the lamina's centroid.

To implement this, the following sequence was developed as an algorithm in Matlab:

- threshold using Otsu's method,
- find centroid of drop image—x value of centroid defines location of vertical axis,
- remove pixels containing vertical axis,
- use the Pappus theorem to find volumes generated by left and right halves of object and then add in part pixel half-cylinders around the vertical axis and
- sum these volumes to obtain the volume of the whole drop.

Sub-pixel Edge Detection Using Curve Fitting (SPED)

Edge detection techniques using a variety of edge models have been proposed.[17,18] We have chosen to use an edge model based on hyperbolic tangent functions.[19] In this special case there will always be a left and right edge, hence a general function is proposed of the form

$$G = \frac{b_L + b_R}{2} - \frac{b_L - f}{2} \tanh\left(\frac{x - E_L}{s_L}\right) + \frac{b_R - f}{2} \tanh\left(\frac{x - E_R}{s_R}\right), \qquad (6)$$

where G is the gray level, b_L is the left background level, b_R is the right background level, f is the foreground level, x is the pixel location along a row of pixels passing through both edges of the drop, E_L is the left edge position, E_R is the right edge position, s_L is the slope associated with the left edge and s_R is the slope of the right edge. The edge is assumed to be at the (sub-pixel) location of the gray level halfway between the background and foreground levels (E_N), although it could be chosen elsewhere, as a result of calibration for example.

A Matlab program was written, making use of Matlab's extensive curve fitting capability, which had the following features:

- an initial edge was found using thresholding,
- at each point on the edge a fit on both left and right edges using Eq. (6) was carried out if possible,
- otherwise a fit was carried out on one edge only,
- the fitted edges were not constrained by the location of the initial thresholding,
- if the slope of the edge was greater than 45° from the vertical the fit was carried out along a vertical line of pixels and the edges at the ends of the horizontal pixel line were interpolated from those values,
- bright areas within transparent drops were detected and removed from the fit,
- nearby drops (e.g., separating or merging satellites) were detected and removed from the fit,
- if the fit failed then the edge used reverted to the original thresholded edge and
- slice volumes were summed as in method T.

Figure 7 shows an example of a fit along one line of pixels in one of the photographic images.

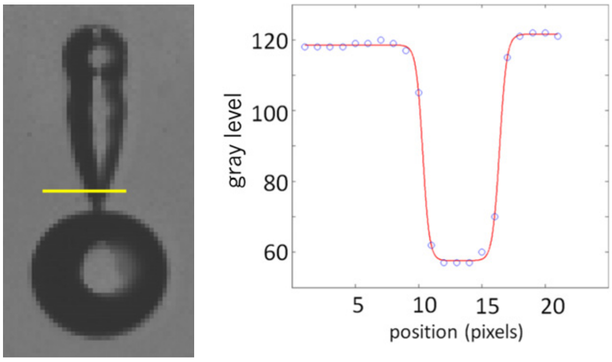

Figure 7. An example of fitting Eq. (6) to one row of pixels (indicated by a line on the image) from an image of a real drop. The open circles are the pixel values and the line is the fit.

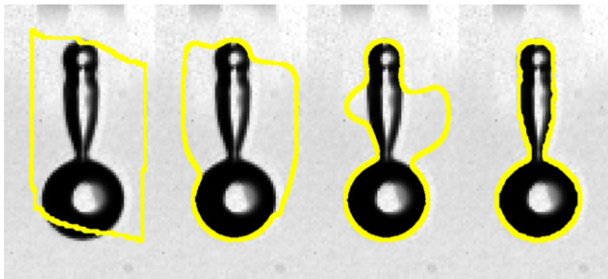

Figure 8. An example of the movement of the contour during the LCV iteration.

Localized Chan–Vese (LCV)

The Chan–Vese method[20] is a form of active contour boundary determination where image segmentation is based on "energy" minimization. The energy is defined for a given closed curve in the image (the contour), and by deforming the contour, a local energy minimum is found. At this local energy minimum it is deemed that the contour has converged on the edge of the object and the image is segmented. In the localized formulation the image intensity on the contour is compared with the local region averages, which allows intensity gradients within the image to be accommodated.[21,22] A version of this formulation was implemented in Matlab and used to find edges. An initial guess at the contour is input to the program which then iterates until the energy is minimized. Figure 8 illustrates the movement of the contour as the calculation converges on a result.

To calculate the drop volume the coordinates of the estimated (sub-pixel) edges are used to calculate pixel slice volumes as in the TSP method.

RESULTS

First, we compare the five methods using a variety of synthetic ovoid images, and then use the same methods to estimate the volumes from a sequence of high speed camera images of a single jetting event.

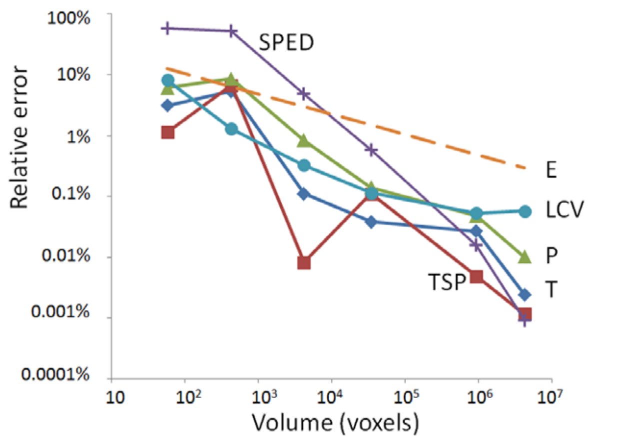

Figure 9. Log–log plot of the relative error as a function of object volume.

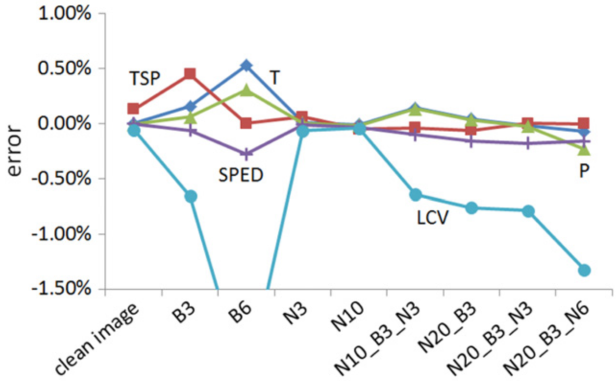

Figure 10. Error arising from analyzing images of the largest synthetic ovoid ($a = 90$, $b = 100$, $c = 150$) after the application of various noise and blur filters.

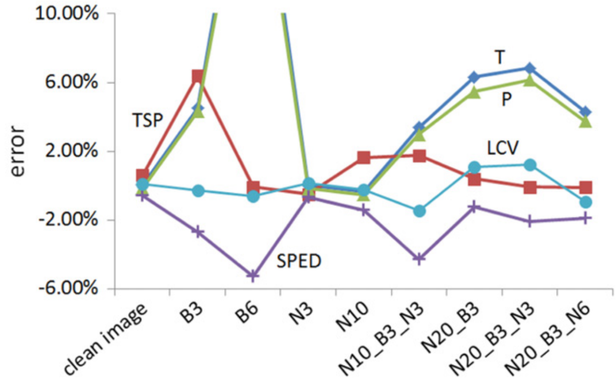

Figure 11. Error arising from analyzing a smaller ovoid ($a = 18$, $b = 20$, $c = 31$) after the application of various noise and blur filters.

Different Volumes

In Figure 9 each line represents the results from one of the methods discussed above. For the synthetic images listed in Table I, the absolute percentage error in the volume measurement is plotted as a function of the true object volume in voxels calculated from Eq. (1).

Because of the range of results, both the volume and the error are plotted logarithmically. An additional line (dashed) is plotted showing the error in the estimation of a spherical volume V that would result from a systematic error in the radius (Δr) of 0.1 pixels:

$$\frac{\Delta V}{V} = \frac{3\Delta r}{\sqrt[3]{3V/4\pi}}. \qquad (7)$$

This plot reveals a number of interesting points. Not surprisingly, the percentage error reduces as the size of the object (i.e., the number of voxels) increases. For the smallest object (volume 58.6 voxels) the equivalent spherical radius (R_{eq}) is 2.4 pixels and for the largest object (4.24×10^6 voxels) the equivalent spherical radius is 100 pixels. For the smallest volume the errors range between 1% and 10% except for the SPED method where the error is much higher. This is perhaps not surprising, as SPED relies on having a significant number of background and foreground pixels around the edge to achieve a good fit. However, for the largest volume SPED gives one of the best results. The other techniques all give similar results, with perhaps TSP showing marginally the best results overall.

Noise and Blur

Images corresponding to two sizes of object (radius equivalents of 20 and 100 pixels) were processed as described in the Test Images section to produce a number of synthetic images with various levels of noise, blur or their combinations applied in sequence. These sequences were then processed using each of the methods. Figure 10 shows the results for the larger object. This shows that the errors were nearly all less than ±0.2% except for the image with a large amount of blur (6 pixels) and for the LCV method which gave poor results for blurred images.

Figure 11 shows results for the smaller synthetic image modified in the same way. In this case the large amount of blur (6 pixels) also presented a challenge and the overall errors were much larger for most of the methods and images (with errors of ±2% for the best methods). In this case the SPED technique consistently underestimated the volume, which perhaps could be corrected with calibration. Both the T and the P methods struggled with the most degraded images, while the results from LCV and TSP remained more consistent over most of the images.

Background Level

To investigate the effects of changing the image background level, synthetic images were generated with background levels ranging from 245 (almost white) to 35 (just above the foreground level of 15). In these images the object size was $a = 18$, $b = 20$, $c = 31$. A second set with the same series of background levels was also processed with the filter sequence N10_B3_N3. Examples of these images are shown in Figure 12.

Figure 13 shows the results obtained with the clean (unfiltered) images. For most of the background level range the results from all of the methods were consistent within a few tenths of a percent although offset from each other. The LCV method seemed completely unaffected by the

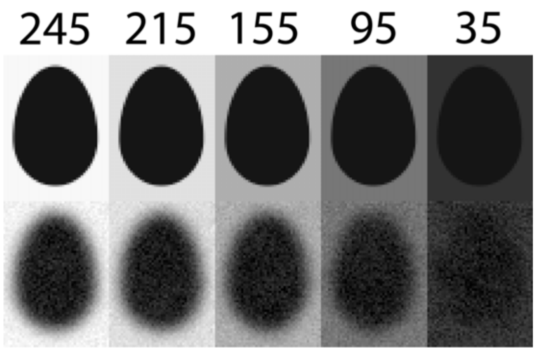

Figure 12. Examples from the series of images with varying background levels ranging from 245 (near white) to 35 (near foreground). Object size: $a = 18$, $b = 20$, $c = 31$. The top row shows unfiltered images and the bottom row show images filtered using the sequence N10_B3_N3.

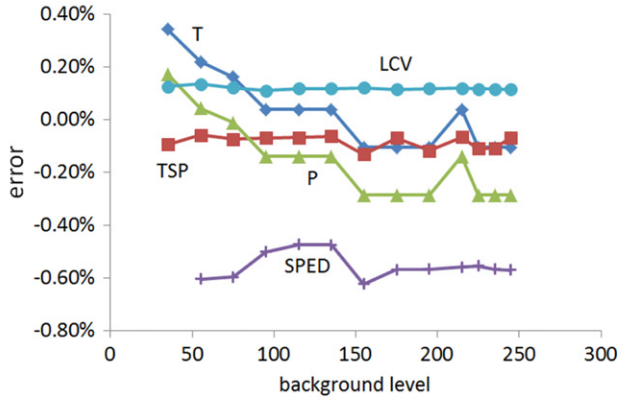

Figure 13. Error plotted against background level for the set of clean images ($a = 18$, $b = 20$, $c = 31$).

background level over the entire range. The T and P methods were least consistent, and the SPED method showed the largest offset from the actual value.

A change in the consistency and relative performance was seen with the filtered images, as shown in Figure 14. For nearly all of the techniques no results were obtained for the two lowest (darkest) background levels. Even the most consistent technique (TSP) varied over a 3% range. Others ranged more widely, the worst being LCV (9%).

Photographic Images

A set of 57 real photographic images was analyzed using the various methods. The images were frames, taken at 1 μs intervals, which follow a single jetting event from close to break-up from the nozzle, where a substantial ligament is present, through detachment of the ligament to form a satellite drop, which then catches up and merges with the main drop near the bottom of the frame. Some examples are illustrated in Fig. 4.

Subjectively, these images appear to have significantly less noise and blur than, for example, the synthetic images processed using the N10_B3_N3 filter sequence. The results for T, TSP and P were very similar, and so to make the graph less cluttered the results for the TSP, SPED and LCV methods only are displayed in Figure 15.

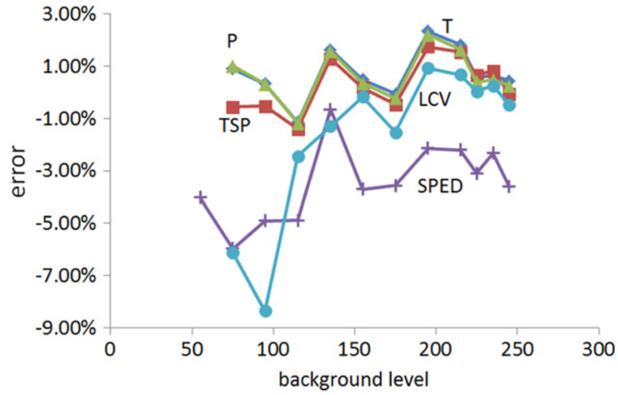

Figure 14. Error plotted against background level for the set of filtered (N10_B3_N3) images ($a = 18$, $b = 20$, $c = 31$).

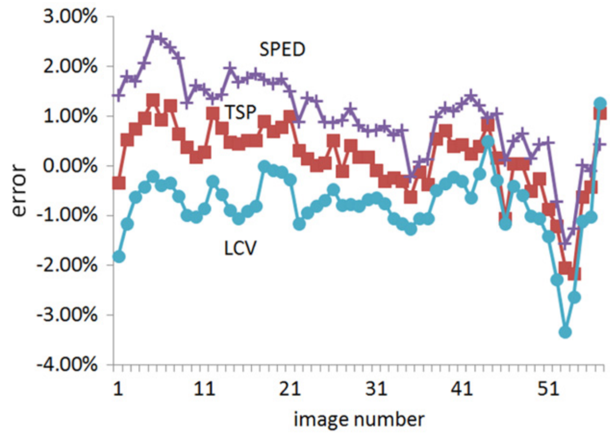

Figure 15. Error plotted against image in the sequence of images taken at 1 μs intervals of a single drop ejection event.

The mean of all of the results for each method was used to calculate the percentage error, from this mean, for each individual measurement. As discussed above, the total volume of the liquid present in the drop and ligament in each of the images was assumed to be the same. As can be seen in Fig. 15, the volumes from each of the techniques fluctuate over a range of 2% for most of the sequence of images, with a clear offset between the three methods. There appears to be some correlation between the results, which suggests that there are features in the images that affect the volume analysis in the same way. These may include variation of background level, lack of circular symmetry about the vertical axis (e.g., the presence of some vibrational modes), hidden features (e.g., an inward dip at the top of the drop following break-up), a vertical axis not completely planar with the image, and image distortion. Factors associated with the image processing include sensitivity to background level, an offset in the detection of the "real" edge, errors arising from noise and blur, and issues with extreme drop geometries.

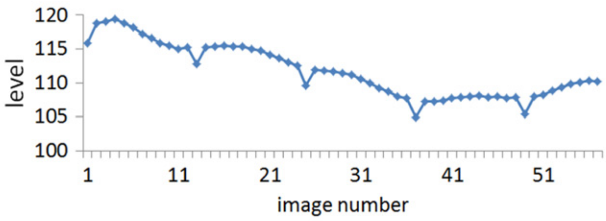

Figure 16. Image background level during the single drop ejection event.

Table II. Typical computation times taken by the various methods.

Technique	Time (sec)
T	0.0922
TSP	0.0906
P	0.0731
SPED	29.3000
LCV	1.5010

The background level of the images varied during the sequence. Figure 16 shows the background level for all of the images in this set; the foreground level remained consistent at around 15 in all of the images. The background was estimated from the mean value of pixels in a 30-pixel square box in the same place on each image.

The graph shows a smooth variation over the time of the event, probably arising from the variation in the illumination intensity from the beginning to the end of the event. There is also a regular dip every 12 frames, which is probably an artifact introduced by the camera. The apparent slight reduction in volume from the beginning of the event to the end could be caused by sensitivity of the methods to background level (although the measurements with the synthetic images suggest that this range should not produce a significant change), or by slight geometrical misalignment or optical distortion.

Processing Speed

Each of the analysis algorithms was implemented using Matlab, and no attempt was made optimize the processing speed. However, there were large differences in the typical computation times taken by each method, as shown in Table II, which lists the times taken to process one of the synthetic images ($a = 90$, $b = 100$, $c = 150$, N10_B3_N3).

DISCUSSION AND CONCLUSIONS

Synthetic images have been constructed representing images of ovoid objects with an accurately known volume. These images were used to compare the performance of five different processing methods in estimating the volume of the synthetic object. A series of differently sized objects showed, not surprisingly, that the more pixels there are, the better the estimates made by all of the methods. The SPED method did not cope well with very small objects with an equivalent diameter $R_{eq} = 2.4$ pixels, but performed very well at high pixel counts ($R_{eq} = 100$ pixels).

At $R_{eq} = 100$ pixels, with a clean image (no noise, no blur), all of the techniques gave results within 0.1% of the true volume for images with R_{eq} of 20 pixels or more. Even with added noise and blur, the volume of a large object ($R_{eq} = 100$ pixels) was, in the majority of cases, estimated within ±0.2% of the correct value except for the LCV technique which performed badly with blurred images in particular.

With a smaller image ($R_{eq} = 20$ pixels), the results were significantly poorer except for the LCV method which produced similar errors to those produced for the larger image. In this case LCV and TSP produced the best results. Varying the background level for the $R_{eq} = 20$ pixels object seemed to produce little variation where there was significant difference between foreground and background levels (>85 on an 8-bit scale); below this difference, significant changes were seen for T and P, with TSP and LCV remaining stable throughout the range. The introduction of noise and blur (N10_B3_N3) caused considerably more variation in results. T, TSP and P gave much the same results in the range −1%–+2%. The LCV method again deteriorated significantly in the presence of noise and blur, and produced a wide range of results.

Overall, it is clear that one of the simplest approaches, the TSP method, provided the most consistent results across the range of object sizes, noise and blur levels and background levels. Other techniques provided very good results in some circumstances but failed in others (e.g., for LCV blur seems to be a problem). For TSP the processing time was one of the shortest, so that for real time analysis this technique would be a good choice.

The measurements from real photographic images reveal additional sources of variation not present in the synthetic images. These may be caused by non-optimum optical and geometric set-up and by lack of circular symmetry and hidden geometry within the drop(s) as they change shape during the event.

The literature contains many examples of imaged inkjet drops.[6,16,23] Drop-on-demand inkjet drops are typically 10–50 μm in diameter and continuous inkjets would typically involve drops of up to 120 μm; satellite drops are smaller. A maximum resolution of 0.5 μm/pixel, which is typical of the limit of optical resolution with visible light, would imply drop image diameters from a few pixels up to 240 pixels. Our synthetic images cover much of this range. Our results suggest that there are several techniques, including a simple method with short calculation time (TSP), which, when processing high quality images, should achieve volume measurement precision of 0.1%–1%, significantly better than that expected from a precision of ±0.1 pixels in radius (Fig. 9).

A set of non-spherical simulated images avoids the uncertainty related to a physical experiment when comparing processing methods, and would be a useful tool for the comparison of other drop image processing algorithms, either existing or to be developed. The use of synthetic images and the means to investigate processing methods discussed in this article could assist in the development of new international (ISO/IEC) standards, for example those being drafted by the working group IEC TC119 WG3 for (printed electronics) inkjet equipment drop measurement methods (ISO/IEC 62899-302-01, drop speed, and ISO/IEC 62899-302-02, drop volume).

ACKNOWLEDGMENTS

The authors thank Steve Hoath, Department of Engineering, University of Cambridge for helpful discussions. This work was supported in part by the United Kingdom Engineering and Physical Sciences Research Council and industrial partners in the "Innovation in Industrial Inkjet Technology" project (Grant No. EP/H018913/1).

REFERENCES

[1] M. Bale, J. C. Carter, C. J. Creighton, H. J. Gregory, P. H. Lyon, P. Ng, L. Webb, and A. Wehrum, "Ink-jet printing: the route to production of full-color P-OLED displays," J. Soc. Inf. Disp. **14**, 453 (2006).

[2] R. Daly, T. S. Harrington, G. D. Martin, and I. M. Hutchings, "Inkjet printing for pharmaceutics—a review of research and manufacturing," Int. J. Pharm. **494**, 554–567 (2015).

[3] R. M. Verkouteren and J. R. Verkouteren, "Inkjet metrology: high-accuracy mass measurements of microdroplets produced by a drop-on-demand dispenser," Anal. Chem. **81**, 8577–8584 (2009).

[4] G. D. Martin, J. R. Castrejón-Pita, and I. M. Hutchings, "Holographic measurement of drop-on-demand drops in flight," *Proc. IS&T's NIP27: Int'l. Conf. on Digital Printing Technologies and Digital Fabrication 2011* (IS&T, Springfield, VA, 2011), pp. 620–623.

[5] F. Soulez, L. Denis, C. Fournier, E. Thiébaut, and C. Goepfert, "Inverse-problem approach for particle digital holography: accurate location based on local optimization," J. Opt. Soc. Am. A **24**, 1164–1171 (2007).

[6] K.-S. Kwon, M.-H. Jang, H. Y. Park, and H.-S. Ko, "An inkjet vision measurement technique for high-frequency jetting," Rev. Sci. Instrum. **85**, 1–11 (2014).

[7] R. Castrejón-García, J. R. Castrejó-Pita, G. D. Martin, and I. M. Hutchings, "The shadowgraph imaging technique and its modern application to fluid jets and drops," Rev. Mex. Fis. **57**, 266–275 (2011).

[8] I. M. Hutchings, G. D. Martin, and S. D. Hoath, "High speed imaging and analysis of jet and drop formation," J. Imaging Sci. Technol. **51**, 438 (2007).

[9] S. D. Hoath, W. Hsiao, G. D. Martin, S. Jung, S. A. Butler, N. F. Morrison, O. G. Harlen, L. S. Yang, C. D. Bain, and I. M. Hutchings, "Oscillations of aqueous PEDOT:PSS fluid droplets and the properties of complex fluids in drop-on-demand inkjet printing," J. Non-Newton. Fluid Mech. **223**, 28–36 (2015).

[10] H. Dong, W. W. Carr, and J. F. Morris, "Visualization of drop-on-demand inkjet: drop formation and deposition," Rev. Sci. Instrum. **77**, 085101 (2006).

[11] N. R. Pal and S. K. Pal, "A review on image segmentation techniques," Pattern Recognit. **26**, 1277–1294 (1993).

[12] R. C. Gonzalez, R. E. Woods, and S. L. Eddins, *Digital Image Processing using Matlab*, 2nd ed. (Gatesmark Publishing, 2009).

[13] F. Da and H. Zhang, "Sub-pixel edge detection based on an improved moment," Image Vis. Comput. **28**, 1645–1658 (2010).

[14] A. Fabijańska, "A survey of subpixel edge detection methods for images of heat-emitting metal specimens," Int. J. Appl. Math. **22**, 695–710 (2012).

[15] N. Otsu, "A threshold selection method from gray-level histograms," IEEE Trans. Syst. Man Cybern. **9**, 62–66 (1979).

[16] A. van der Bos, M.-J. van der Meulen, T. Driessen, M. van den Berg, H. Reinten, H. Wijshoff, M. Versluis, and D. Lohse, "Velocity profile inside piezoacoustic inkjet droplets in flight: comparison between experiment and numerical simulation," Phys. Rev. Appl. **1**, 014004 (2014).

[17] M. Hagara and P. Kulla, "Edge detection with sub-pixel accuracy based on approximation of edge with Erf function," Radioengineering **20**, 516–524 (2011).

[18] F. Bouchara and S. Ramdani, "Subpixel edge refinement using deformable models," J. Opt. Soc. Am. A **26**, 820–832 (2009).

[19] V. S. Nalwa and T. O. Binford, "On detecting edges," IEEE Trans. Pattern Anal. Mach. Intell. **8**, 699–714 (1986).

[20] T. F. Chan and L. A. Vese, "Active contours without edges," IEEE Trans. Image Process. **10**, 266–277 (2001).

[21] X.-F. Wang, D.-S. Huang, and H. Xu, "An efficient local Chan–Vese model for image segmentation," Pattern Recognit. **43**, 603–618 (2010).

[22] S. Lankton and A. Tannenbaum, "Localizing region-based active contours," IEEE Trans. Image Process. **17**, 2029–2039 (2008).

[23] S. D. Hoath, J. R. Castrejó-Pita, W. Hsiao, S. Jung, G. D. Martin, I. M. Hutchings, T. R. Tuladhar, D. C. Vadillo, S. A. Butler, M. R. Mackley, C. Mcilroy, N. F. Morrison, O. G. Harlen, and H. N. Yow, "Jetting of complex fluids," J. Imaging Sci. Technol. **57**, 040403 (2013).