Detection of cavitation pits on steel surfaces using SEM imagery*

Jeffery R. Price†, Kathy W. Hylton, Kenneth W. Tobin, Jr.,
Philip R. Bingham, John D. Hunn, and John R. Haines

Oak Ridge National Laboratory, Oak Ridge, TN, USA

ABSTRACT

We describe an automated image processing approach for detecting and characterizing cavitation pits on stainless steel surfaces. The image sets to be examined have been captured by a scanning electron microscope (SEM). Each surface region is represented by a pair of SEM images, one captured before and one after the cavitation-causing process. Unfortunately, some required surface preparation steps between pre-cavitation and post-cavitation imaging can introduce artifacts and change image characteristics in such a way as to preclude simple image-to-image differencing. Furthermore, all of the images were manually captured and are subject to rotation and translation alignment errors as well as variations in focus and exposure. In the presented work, we first align the pre- and post-cavitation images using a Fourier-domain technique. Since pre-cavitation images can often contain artifacts that are very similar to pitting, we perform multi-scale pit detection on each pre- and post-cavitation image independently. Coincident regions labeled as pits in both pre- and post-cavitation images are discarded. Pit statistics are exported to a text file for further analysis. In this paper we provide background information, algorithmic details, and show some experimental results.

Keywords: SEM image analysis, Spallation Neutron Source (SNS), phase correlation, Fourier-based image alignment, image-based defect detection

1. INTRODUCTION

The Spallation Neutron Source (SNS) is an accelerator-based neutron source being constructed by the U.S. Department of Energy in Oak Ridge, Tennessee (USA). In the SNS, a liquid mercury target, contained in a stainless steel vessel, is bombarded with a pulsed ion beam to produce corresponding pulses of neutrons. When the high-energy (on the order of 1GeV) proton beam strikes the liquid mercury, rapid heating of the mercury results in a pressure wave. This pressure wave is followed by a rarification wave that can produce bubbles (cavitation) in the mercury. The internally reflected positive pressure wave can cause these bubbles to collapse. When this collapse occurs near the stainless steel vessel wall, the resulting impact upon the wall can lead to pitting. The goal of the effort underlying this work[1,2] is to identify the most suitable material, target geometry, and damage mitigation scheme for the mercury target. Towards this goal, multiple material samples are imaged with a SEM tool prior to testing (pre-cavitation images). These samples are then subjected to cavitation-inducing experiments and then again imaged with a SEM tool (post-cavitation images). In Fig. 1 we show an example pre-cavitation image in (a) and an example post-cavitation image in (b). For brevity, we henceforth refer to the pre-cavitation images and post-cavitation images as simply pre- and post- images, respectively.

In this paper, we develop a system to batch process sets of pre- and post- images in order to quantify pitting damage. The desired output of this system is simply a statistical characterization of the pits. This characterization simply comprises a listing of the pits found and their size, for each image in a batch as well as the for the entire batch. In this endeavor, we were faced with several challenges. The inspection of these images was originally planned to be performed manually; all available images were manually acquired, oftentimes by different operators, over widely interspersed time intervals. This manual acquisition process resulted in significant variations in focus and exposure (see Fig. 1) as well as translation and rotation alignment differences.

*Prepared by Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

†Correspondence: J. Price, pricejr@ornl.gov, (865) 574-5743.

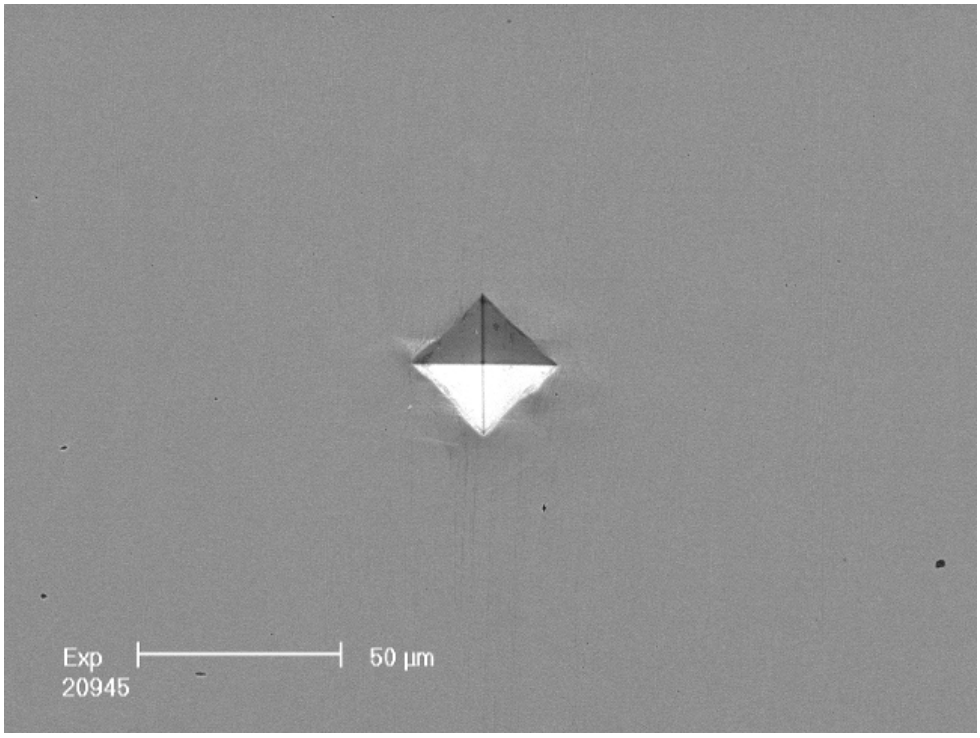

(a) Pre-cavitation image.

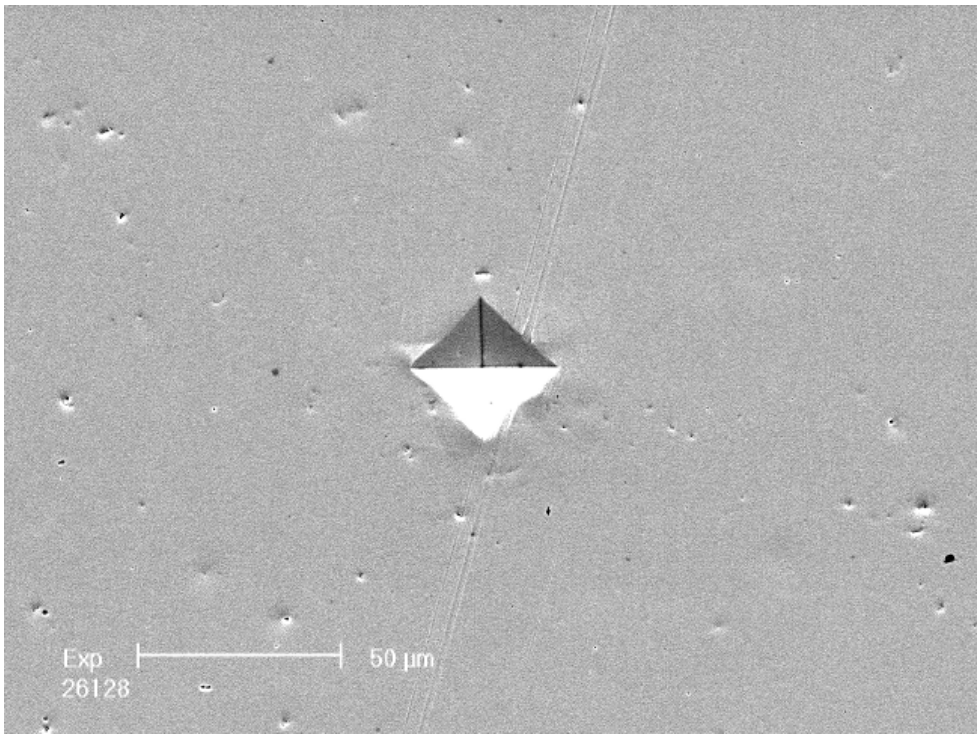

(b) Post-cavitation image.

Figure 1. Example of corresponding pre-cavitation (a) and post-cavitation (b) SEM images. Note the pitting apparent in (b) and also the diamond-shaped indent in the center of both images. These images are 1292×968 pixels and represent an area of 240×180 microns. The indent is approximately 36 microns across.

Furthermore, material decontamination procedures had to be performed after the cavitation-inducing process (between pre- and post- imaging); these procedures often altered the post- image appearance in a way unrelated to pitting. Note, for example, the diagonal linear scratches in the post- image of Fig. 1(b) that does not appear in the pre- image of (a). Finally, due to time and budget constraints, it was imperative that the system be completed quickly.

In summary, the system operates (through a MATLAB graphical user interface) as follows. A batch of corresponding pre- and post- images is selected for processing by the user and certain known parameters (i.e., image resolution, SEM detector orientation) are entered. Note that, after image selection and parameter entry, the remaining processing is performed automatically with no user input. Next, the post- image is aligned to the pre- image using a Fourier-domain, phase correlation technique. Multi-scale pit detection is then performed on both the pre- and post- images independently. Any corresponding regions that are labeled as pits in both the pre- and post- images are discarded from the post- image statistics. Finally, the post- image pit statistics for every image in the selected batch are exported to a single comma-separated-value (CSV) file for further use, beyond the scope of this paper, such as cavitation modeling, simulations, and material selection.

The remainder of this paper is organized as follows. In Section 2, we describe the method used to align the pre- and post- images to one another. We then describe the multi-scale approach adopted to detect pits in Section 3. In Section 4, we provide some representative results and we finally conclude with a few summary remarks in Section 5.

2. IMAGE ALIGNMENT

As noted above, variations in the manual placement of the samples on the SEM tool produces pre- and post-images that are not aligned. The images generally differ by small translations (within ± 30 pixels) and rotations (within $\pm 4^\circ$). We employ well-known Fourier-domain relationships[3,4] to estimate the translation and rotation. Defining an image i to be a rotated, translated copy of another (reference) image r

$$i(x,y) = r(x\cos\theta + y\sin\theta - \Delta x, -x\sin\theta + y\cos\theta - \Delta y), \tag{1}$$

their Fourier transforms are then related by

$$I(u,v) = e^{-j2\pi(u\Delta x + v\Delta y)} R(u\cos\theta + v\sin\theta, -u\sin\theta + v\cos\theta). \tag{2}$$

It is evident from Eq. (2) that the magnitudes, $|I(u,v)|$ and $|R(u,v)|$, are just rotated copies of one another. Also, if their is no rotation or it is corrected (i.e., $\theta = 0$), then only translational misalignment remains and we can see that

$$\frac{I(u,v)R^*(u,v)}{|I(u,v)R(u,v)|} = e^{-j2\pi(u\Delta x + v\Delta y)}, \tag{3}$$

whose inverse Fourier transform is just an impulse, $\delta(x-\Delta x, y-\Delta y)$, centered on the translational misalignment.

In our case, each image – which corresponds to a specific region of a specific material sample – is marked by a diamond-shaped indent (about 36 microns across) near the image center. This indent can be seen in the center of both images of Fig. 1. The alignment process is performed using sub-images from the center of both the pre- and post- images. We extract sub-images of three times the size of the indent; examples of such sub-images for one pre- and post- image pair are shown in Fig. 2. Such a large sub-region is required due to the possible variations in the indent location that result from the manual sample placement and image acquisition. From each of these extracted sub-images, we compute a gradient image given by

$$g_f = |f * d_x| + |f * d_y| \tag{4}$$

where $*$ represents two-dimensional convolution, f is the original image, and d_x and d_y represent gradient kernels in the horizontal and vertical directions, respectively. We employ 9×9 first derivative of Gaussian filters for d_x and d_y. For the original images of Fig. 2, the gradient images resulting from this process are shown in Fig. 3. We now let r represent the gradient of reference (pre-) image and i represent the gradient of the image (post-) to be aligned, where these gradients are computed according to Eq. (4). Furthermore, let R and I represent the

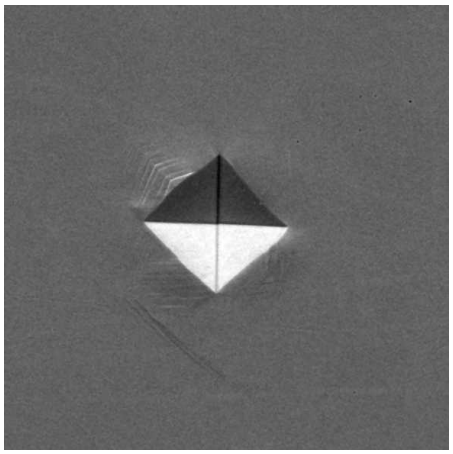

(a) Pre- image.

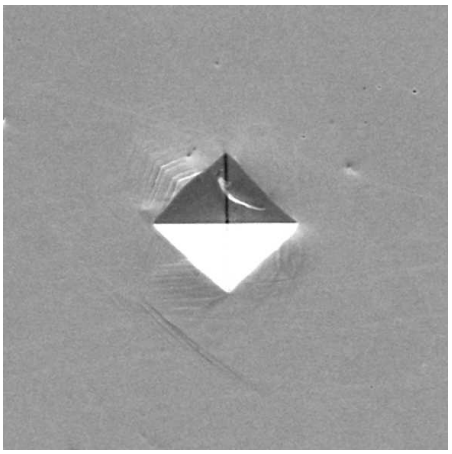

(b) Post- image.

Figure 2. Center sub-images extracted for image alignment. These images are 583×583 pixels and approximately 108 microns square.

Fourier transforms of r and i, respectively. Let $\tilde{R}$ and $\tilde{I}$ represented the log magnitudes of R and I, respectively, multiplied by a window function to remove very high and low frequencies. For the gradient images of Fig. 3, $\tilde{R}$ and $\tilde{I}$ are shown in Fig. 4. Notice how the windowing has set to zero the very low and high frequencies, which may contain disturbances from noise or pitting. To find the rotation angle θ necessary to align i to r, we compute a correlation measure, C_θ, between $\tilde{R}$ and $\tilde{I}_\theta$, where $\tilde{I}_\theta$ is $\tilde{I}$ rotated by θ degrees. The correlation measure is given by

$$C_\theta = \frac{1}{N^2} \sum_{m,n} R(m,n) I_\theta(m,n), \tag{5}$$

where I_θ is computed using computationally inexpensive nearest-neighbor interpolation. We check nine initial points of θ over the range $\{-4, -3, \ldots, 3, 4\}$. Letting θ_o represent the angle of maximum C_θ over this range, we then compute C_θ over refinement points $\theta_o + \{-0.75, -0.5, -0.25, 0.25, 0.5, 0.75\}$. We select as the rotation angle θ that maximizes C_θ over all of these angles. In Fig. 5, we show a plot of C_θ versus θ for the example $\tilde{R}$ and

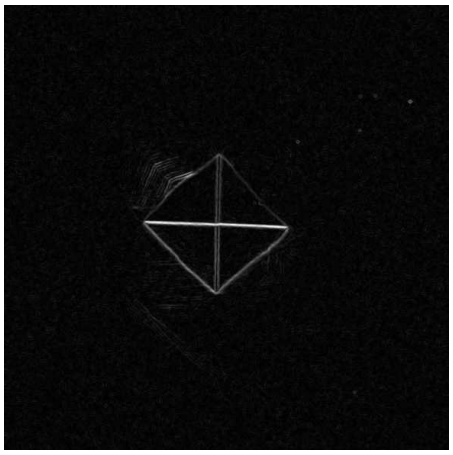

(a) Pre- image gradient.

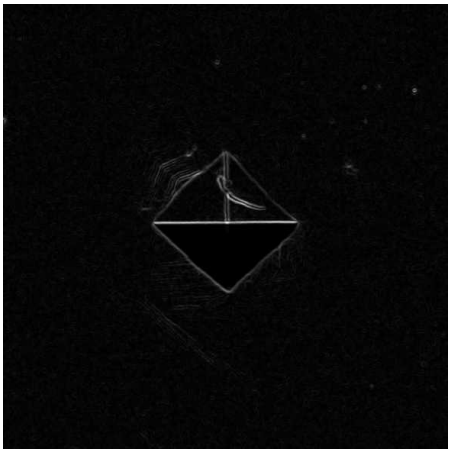

(b) Post- image gradient.

Figure 3. Gradient images, computed according to Eq. (4), for the images of Fig. 2. Images are shown in inverse grayscale so that black represents the highest value and white the lowest.

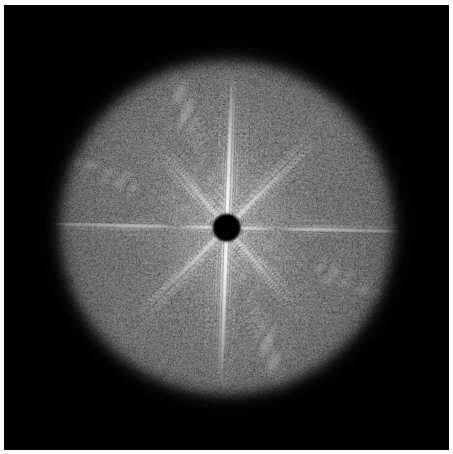

(a) Windowed, log-magnitude of the Fourier transform of the pre- gradient image from Fig. 3(a) This is the reference image (R).

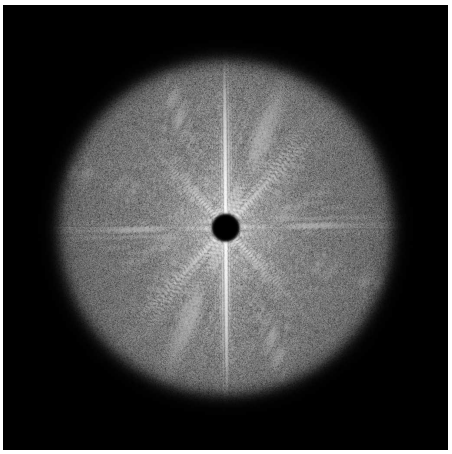

(b) Windowed, log-magnitude of the Fourier transform of the post- gradient image from Fig. 3(b) This is the image to be aligned (I).

Figure 4. Windowed, log-magnitude Fourier transforms of the images from Fig. 3.

$\tilde{I}$ from Fig. 4; initial and refinement points are indicated along with the selected point of $\theta = -2.25°$ (negative implies clockwise rotation of I). The quantization of rotation angle to $0.25°$ increments has been found to be adequate for our purposes.

Once the rotational misalignment has been computed, we rotate the post- gradient image, using bicubic interpolation, to produce i_θ whose Fourier transform is denoted I_θ. Recalling Eq. (3), we then compute the cross-power spectrum

$$S(u,v) = \frac{I_\theta(u,v)R^*(u,v)}{|I_\theta(u,v)R(u,v)|} \approx e^{-j2\pi(u\Delta x + v\Delta y)}. \tag{6}$$

We then compute the inverse Fourier transform of $S(u,v)$, which is approximately an impulse at the location corresponding to the translation misalignment:

$$s(x,y) \approx \delta(x - \Delta x, y - \Delta y). \tag{7}$$

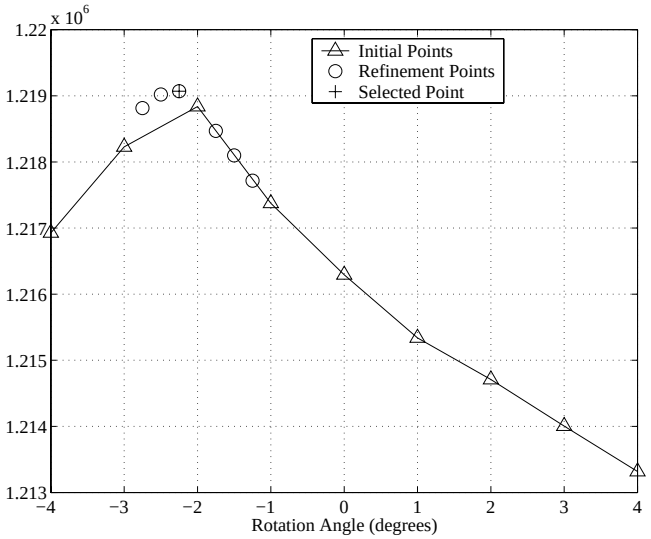

Figure 5. Plot of correlation measure C_θ vs. angle θ. Note the maximum is found at $-2.25°$.

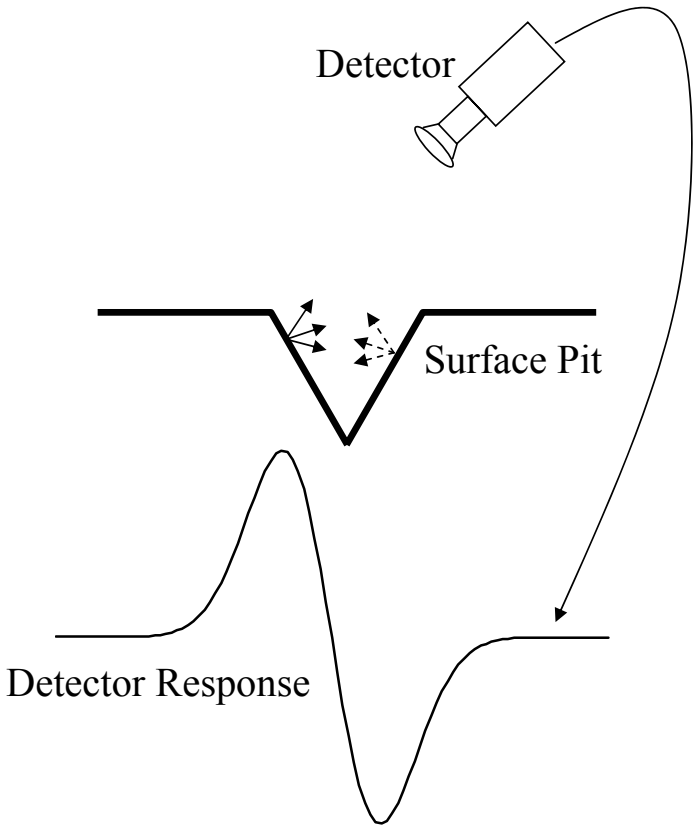

Figure 6. One-dimensional illustration of image response to a pit relative to SEM detector location. The small, solid lines with arrows indicate electrons the detector may capture, while the small, dashed lines with arrows indicate electrons the detector does not see.

We select the maximum (whole pixel) value of $s(x, y)$ as the translation $(\Delta x, \Delta y)$ between i and r. Now that the rotation angle θ and translation $(\Delta x, \Delta y)$ have been found, we align the full-size, post- image to the pre-image using these parameters and bicubic interpolation and then proceed to the pit detection stage.

3. PIT DETECTION

In the SEM imagery, the pits are generally characterized by dark-to-bright transitions, where the orientation of the transition is determined by the geometry of the SEM detector relative to the electron beam and the scanned surface. A one-dimensional analog of this idea is illustrated in Fig. 6. The side of a pit that is visible to the detector will appear brighter in the image, while the "shadowed" side will appear darker. In the two-dimensional case, the orientation of this transition can vary between $0 - 360°$. Note that fine angular resolution for the SEM detector location is not required. The user is only asked to enter the approximate SEM detector orientation as a number between 1 o'clock and 12 o'clock. An example pit and its average profile (i.e., gray values projected onto the axis oriented in the direction of the SEM detector) are shown in Fig. 7.

Pits of interest exist at multiple scales, ranging from about 0.5 microns to 80 microns or more in diameter. This corresponds to a range of about 1-400 pixels in diameter, depending upon the SEM magnification used. Some examples of the variations in pit size can be seen in Fig. 8(b). Because of this broad range of possible pit sizes, we adopt a multi-scale, pyramid-like approach for pit detection. Letting $i_k(x, y)$ represent an image at scale k (larger k implying lower resolution, i.e., smaller images), we perform pit detection, described shortly, and store the set of detected pixels, P_k. We then blur $i_k(x, y)$ with a Gaussian filter and downsample by a factor of two to produce the next lower resolution image $i_{k+1}(x, y)$. Pit detection is performed at this scale and the

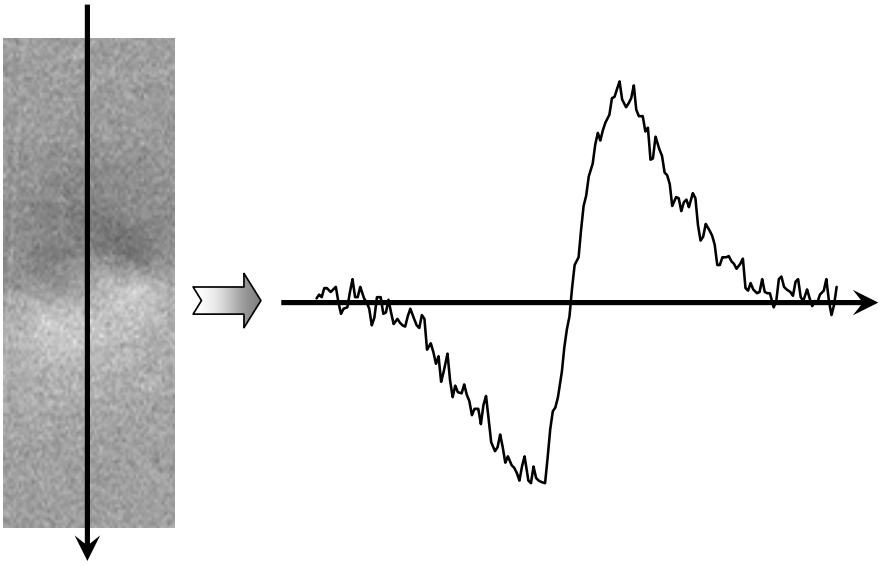

Figure 7. Image and profile (summed along line as indicated) of an example pit. In this image, the SEM detector is looking onto the image from approximately 12 o'clock in-plane (the top), in the direction of the arrow.

resulting detections, P_{k+1} are stored. This process is repeated until the downsampled image reaches a minimum size that is still greater than four times the size of the Gaussian blurring filter. All labeled pit pixels are mapped back to the original, highest resolution image.

At each resolution stage, the same process for pit detection is performed. Three simple measures are computed for each pixel. Recalling Fig. 7, we note that the profile across the pit dark-to-bright transition is similar to the first derivative of the Gaussian. For this reason, the first measure used for pit detection is correlation with a directional derivative-of-Gaussian filter, $d_\theta(x, y)$, oriented in the appropriate direction θ as determined by the SEM tool used. The result of this image correlation, which we call $c(x, y)$, is scaled so that all values lie in the range $[0, 1]$. High values of $c(x, y)$ are indicative of pits, but can also be caused by any strong bright or dark spot due to debris or scratches, hence we also compute two additional measures, $p(x, y)$ and $n(x, y)$, as follows. Assume an image pixel at location (x_o, y_o) and let the local image mean be denoted by $\mu(x_o, y_o)$. We denote the positive region of support of $d_\theta(x - x_o, y - y_o)$ – where $d_\theta(x - x_o, y - y_o) > 0$ – by $\mathcal{D}_+$ and the number of

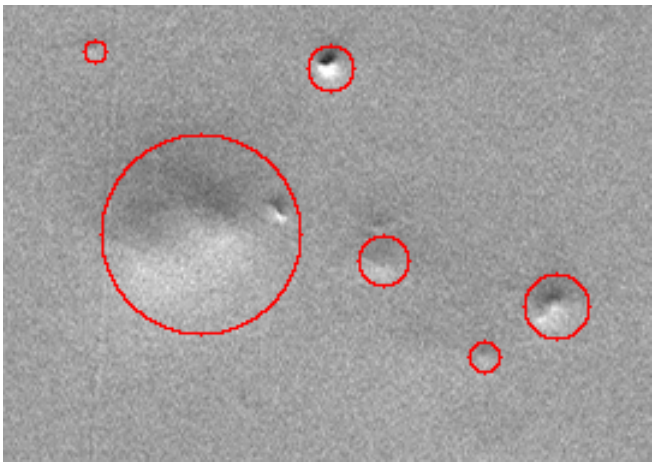

Figure 8. Example of variation in pit size. Note that the variation can even be significantly greater than displayed here.

pixels in $\mathcal{D}_+$ by $\mathcal{Q}_+$. We furthermore denote the number of image pixels in $\mathcal{D}_+$ that are greater than $\mu(x_o, y_o)$ by $\mathcal{P}$. The quantity $p(x_o, y_o)$ is simply $\mathcal{P}/\mathcal{Q}_+$. The quantity $n(x_o, y_o)$ is computed in the exact same fashion for the negative region of support. Noting that all three measures – $c(x, y)$, $p(x, y)$, and $n(x, y)$ – lie in the range $[0, 1]$, pixels are labeled pits only when all three measures exceed a user-defined threshold. We have found that a threshold of 0.4 performs suitably across our available data set.

We perform the described pit detection process on both the pre- and post- images. We then find all pits in the post- image that intersect a pit in the pre- image. If the intersecting pits are of approximately the same size, we discard the pit from the post- image. Quantities of interest (total pit area, number of pits, pit diameters) are stored in a text file and the detected pits are indicated in an output image file for manual examination, if desired.

4. RESULTS

We have applied the proposed system to several thousand images. Each material sample was imaged at two magnifications – generally 0.186 microns per pixel and 0.744 microns per pixel – in order to detect both small pits (in the high resolution image) and very large pits (in the low resolution image). Example results are shown in Figs. 9 and 10. The images of Fig. 9 each represent an area of about 59.5 microns $\times$ 74.4 microns (0.186 microns per pixel). The images of Fig. 10 each represent an area of about 298 microns $\times$ 238 microns (0.744 microns per pixel).

We note here a few shortcomings in the current implementation. The first is the requirement for a user-selectable threshold. Ideally, a more robust classifier approach could be employed if training data (i.e., manually labeled pits) and more time were available. Despite the selectable threshold, however, a threshold of 0.4 has worked adequately for all the data encountered to date. Also, as evident in Figs. 8- 10, we currently do not account for pits that lie within the boundaries of larger pits. For example, the largest pits in Fig. 8, Fig. 9(c), and Fig. 10(c) also contain smaller pits.

5. CONCLUSIONS

In this paper, we have presented a method for detecting cavitation pits on steel surfaces from SEM images. This effort is motivated by the need to identify an appropriate surface material, geometry, and damage mitigation scheme to contain the liquid mercury target in the Spallation Neutron Source. Pre- and post-cavitation images are first geometrically aligned using a Fourier-domain approach. Pit-like areas are detected in both the pre- and post- images by measuring three features at each image pixel; pixels where all three measures are high are labeled pits. Pits or pit-like areas that are coincident in the aligned images are discarded; we are interested in only those pits that resulted from the cavitation-inducing process. The process has been performed successfully on several thousand images.

REFERENCES

1. B. Riemer, J. Haines, J. Hunn, D. Lousteau, T. McManamy, and C. Tsai, "SNS target tests at the LANSCE-WNR in 2001 - part i," *Journal of Nuclear Materials* , 2003. to appear.
2. J. Hunn, B. Riemer, J. Haines, D. Lousteau, T. McManamy, and C. Tsai, "SNS target tests at the LANSCE-WNR in 2001 - part ii," *Journal of Nuclear Materials* , 2003. to appear.
3. L. G. Brown, "A survey of image registration techniques," *ACM Computing Surveys* **24**, pp. 325–376, December 1992.
4. B. S. Reddy and B. Chatterji, "An FFT-based technique for translation, rotation, and scale-invariant image registration," *IEEE Transactions on Image Processing* **5**, pp. 1266–1271, August 1996.

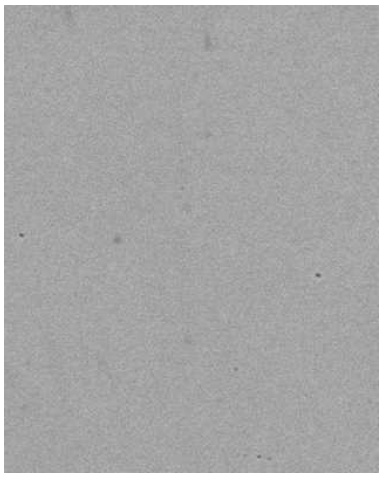

(a) Small region of a pre- image. Note some small pit-like areas.

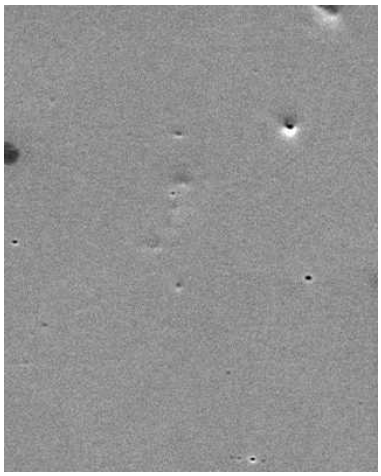

(b) Region of a post- image corresponding to (a).

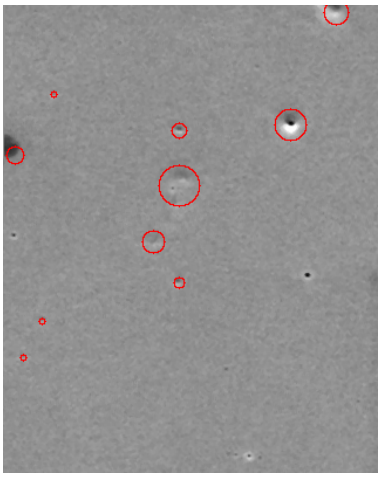

(c) Pits detected (circled) using proposed approach.

Figure 9. Example of pit detection process. Note how pit-like spots that exist in the pre- image (a) are not labeled in post- image (c). Note also the one false detection at the dark spot on the very left side of (c), near the top. These images represent an area of approximately 59.5 microns × 74.4 microns at 0.186 microns per pixel.

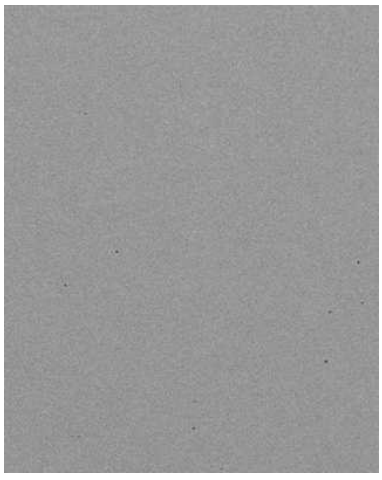

(a) Small region of a pre- image. Note some small pit-like areas.

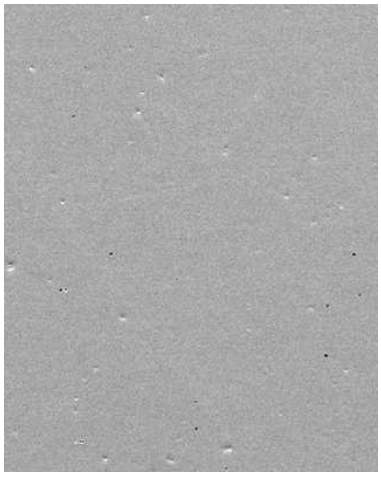

(b) Region of a post- image corresponding to (a).

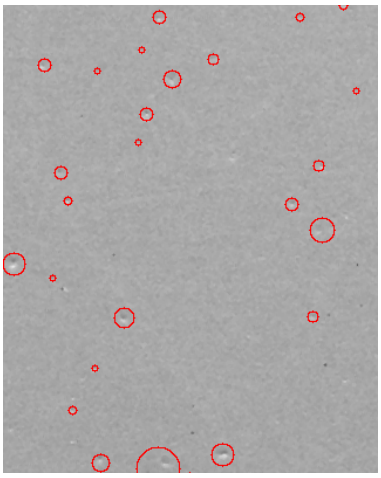

(c) Pits detected (circled) using proposed approach.

Figure 10. Another example of pit detection process. These images represent an area of approximately 298 microns × 238 microns at 0.744 microns per pixel.