OPTIMIZING BLOCKER USAGE ON NIF USING IMAGE ANALYSIS AND MACHINE LEARNING*

L Mascio Kegelmeyer, JG Senecal, AD Conder, LA Lane, MC Nostrand, PK Whitman LLNL, Livermore, CA 94550, USA.

Abstract

To optimize laser performance and minimize operating costs for high-energy laser shots it is necessary to locally shadow, or block, flaws from laser light exposure in the beamline optics. Blockers are important for temporarily shadowing a flaw on an optic until the optic can be removed and repaired. To meet this need, a combination of image analysis and machine learning techniques have been developed to accurately define the list of locations where blockers should be applied. The image analysis methods extract and measure evidence of candidate sites and their correlated downstream hot spots and this information is passed to machine learning algorithms which calculate the probability that candidates are flaws that require blocking. Results show that the machine learning helps to significantly reduce false alarms, while correctly classifying true sites, compared to the image analysis methods alone. Ten-fold cross validation of the refined training set shows about 99% of the detected candidate sites are rejected, leaving only 1% to be brought forward for review (300 sites brought forward vs. 30,000 detected candidates); about a fifth of those brought forward now are false alarms compared to many times the true positives using image analysis alone. In practice, this amounts to between 0 and 3 false alarms per image with over 98% true positive detection.

INTRODUCTION

Maintaining the quality of the optics on a high-energy laser, such as the National Ignition Facility (NIF), requires continuous monitoring of their condition. Automated analyses [1,2,3] have been developed at NIF to track and report damage on optics throughout the facility. These
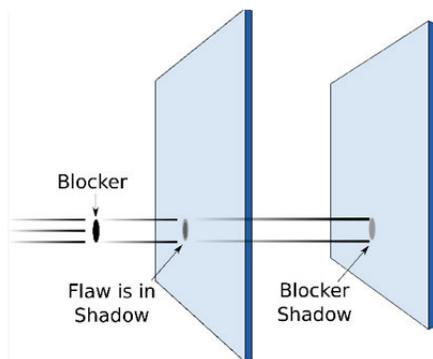
Figure 1: Blockers are used to temporarily shadow identified sites from high energy laser exposure.

sites can be blocked from further interaction with the high energy laser light until they can be removed and repaired for re-use. One inspection camera system resides at target chamber center. It can inspect the nearest (final) optics, which have individual edge-injected illumination, by focusing on each optic directly, in sequence. Images of optics further upstream and illuminated with bright field backlighting are more complex. For these images, damage sites can be confirmed, or found indirectly, by looking for a diffraction pattern "signature" on image planes closer to the camera.

The signature of interest is created when the illumination-laser plane wave interacts with a defect, modeled as a circular opaque scattering site.
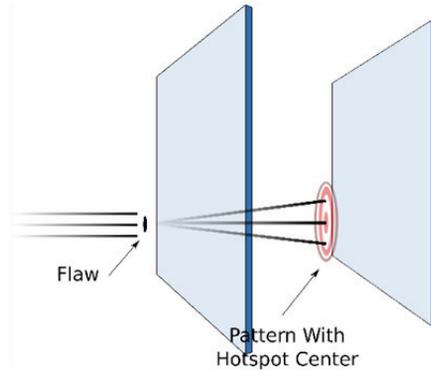
Figure 2: A plane wave from the illumination laser interacts with a scattering site to create diffraction rings downstream.

The interaction is described mathematically as:

$$I(r,R,z) = \frac{r^2 + R \cdot J\left[1, \frac{2\pi r R}{\lambda \cdot z}\right]\left(R \cdot J\left[1, \frac{2\pi r R}{\lambda \cdot z}\right] - 2r \sin\left[\frac{\pi r^2}{\lambda \cdot z}\right]\right)}{r^2}$$

where $J[1,x]$ is a Bessel function of the first kind, and λ is the wavelength of the illuminating light. In general the larger the source defect radius R, the higher the intensity amplitude of the resulting rings. The greater the distance z, the larger the rings.

Working backwards, then, the characteristics of such a ring pattern can be used to predict the size and location of the site that generated it (prior examples from NIF and elsewhere are [4,5,6]) and this in turn can be used to report and track sites on optics that may need to be blocked, or shadowed, from high energy laser light until they can be repaired.

METHODS

Detecting the diffraction ring patterns in bright field images required a custom image analysis algorithm based on the physics principles specific to the creation of diffraction rings. Because of the complexity of the images, this algorithm also finds a large number of false positives. Machine learning techniques are used to distinguish the true from false positives and significantly enrich the number of true actionable sites (candidates for blocking) brought forward for review.

Detection Algorithm

The primary mechanism for the production of diffraction ring patterns is depicted in Figure 2 and shown in the associated equation. The equation generates the "signature" shown in Figure 3 when using relevant parameters (e.g. 300 micron scattering site imaged 9.6 meters away). This can then be used as a template to find diffraction rings of interest.
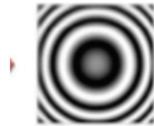
Figure 3: The resulting diffraction ring "signature" using parameters for a 300 micron scattering site imaged 9.6 meters away.

Two initial detection methods are applied and their results combined. One method locates bright centers (the peaks that remain after dividing by local average) and calculates a peak-to-mean ratio in those neighborhoods. Sites with a peak-to-mean ratio of at least 1.2 (operationally determined for sites of interest) are kept. The second method uses the template from Figure 3 and computes a cross-correlation with the image.
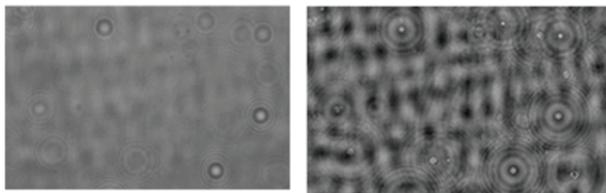
Figure 4: Cross correlating the original image (left) with the expected diffraction pattern from Eq. 1 emphasizes the locations of candidate diffraction patterns.

Locations with intensity at least 3 sigma above the mean were recorded and combined with the sites found with the first method.

The combined locations were further evaluated by measuring a number of characteristics in their local neighborhood. In addition, the co-occurrence matrix of the candidate regions provided measures of contrast, correlation and energy.

Initially, we tried tuning the image analysis parameters to select the most likely rings of interest. As is often the case with handpicked cutoffs, this resulted in either too many false alarms or not enough true positives depending on the values selected. To avoid missing true positives, we set these parameters to include many thousands of false alarms in addition to the true positives, and applied machine learning techniques to correctly classify the candidate sites.

Machine Learning

In order to reduce the overwhelming number of false alarms without reducing the true positives, we turned to machine learning techniques, which have been applied successfully at NIF in the past [7, 8]. Specifically, we used the Avatar Machine Learning Suite of Tools [9] by first training the classifier and then applying it to compute the probability of relevance for new candidates.

Training Data

Training the classifier adequately requires discovering and collecting salient measurements. During the image analysis, we measure a few dozen features in the area of each candidate site. Some measurements are generically useful while others are more specific to the expected properties of these rings. Examples include area, aspect ratio, size of the bounding box, solidity, Euler number, peak-to-mean, disk radius, radii ratio, fringe contrast and similarity to the ideal template. Once measurements are calculated and stored, we assigned a "truth" value to each. Since labeling every single one of the more than 30,000 candidate sites is a daunting task, we let the machine learning tool do some of the work for us. Several hundred sites were known and had previously been tracked manually. We found those among the candidates and labeled them as hotspots (of-interest). The rest were automatically labeled not-of-interest and we used this to initially train the classifier.

We used Avatar to generate an ensemble of decision trees (using options for bagging, Hellinger split method, non-weighted voting and its own self-stopping criterion which resulted in about 90 trees.) The first iteration of training brought to light a number of errors in the original training set, as intended. By manually reviewing and re-labeling sites that the classifier brought forward as high probability of being of-interest, we were able to add some previously untracked sites of interest to the training set. After a few iterations, Avatar accurately reported enormous numbers of false alarms and greatly enriched the percentage of true positives brought forward.

RESULTS

Evaluating the classifier results with ten-fold cross validation of the refined training set showed about 99% of the detected candidate sites were rejected, leaving only 1% to be brought forward for review (300 sites brought forward vs. 30,000 detected candidates); about a fifth of

those brought forward are false alarms. In practice, this amounts to between 0 and 3 false alarms per image (down from hundreds or even thousands in some scenarios). Evaluation suggests the remaining false alarms are primarily out-of-range (i.e. not of interest) diffraction rings and additional training may help prune these out as well.

The image analysis and machine learning combination yields an enriched collection of candidates to review for blocking, as shown in Figures 4 and 5. In addition to graphical presentations, we provide an automated list of candidates with associated probabilities as computed by the ensemble of decision trees. Candidates for blocking can then be reviewed in rank order.
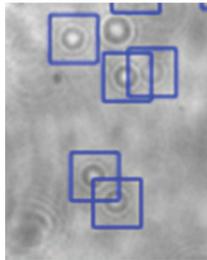
Figure 4: This region illustrates patterns that both do and do not meet the signature of the sites of interest. The blue squares indicate rings that were correctly found. Other prominent diffractions ring patterns are visible but are born from out-of-range sources, not of interest, and thus are not marked with squares (true negatives).
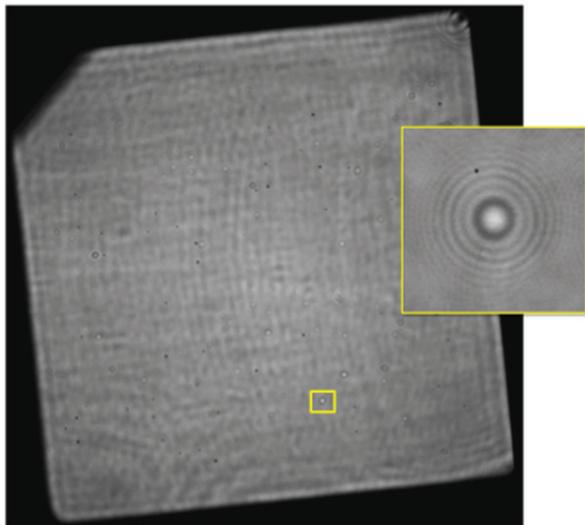
Figure 5: Full image of a backlit optic with one highlighted diffraction ring pattern showing evidence of damage on an upstream optic.

REFERENCES

[1] Laura M. Kegelmeyer, Raelyn Clark, Richard R. Leach, David McGuigan, Victoria Miller Kamm, Daniel Potter, J. Thad Salmon, Joshua Senecal, Alan Conder, Mike Nostrand, Pamela K. Whitman. Automated optics inspection analysis for NIF. "Fusion Engineering and Design", Volume 87, Issue 12, December 2012, Pages 2120-2124 http://dx.doi.org/10.1016/j.fusengdes.2012.09.017

[2] Conder, J. Chang, L. Kegelmeyer, M. Spaeth, and P. Whitman, "Final optics damage inspection (FODI) for the National Ignition Facility," in Optics and Photonics for Information Processing IV, SPIE Proceedings Vol. 7797 (2010).

[3] Kegelmeyer, L. M., Fong, P., Glenn, S. M., and Liebman, J., 2007, "Local Area Signal-to-Noise Ratio (LASNR) algorithm for Image Segmentation," SPIE: Applications of Digital Image Processing XXX (OP312, 6696-85), San Diego, August.

[4] Chen, B. Y., Kegelmeyer, L. M., Liebman, J. A., Salmon, J. T., Tzeng, and J., Paglieroni, D. W., 2006, "Detection of Laser Optic Defects Using Gradient Direction Matching," SPIE—8th International Conference on Laser Beam Control and Applications, **6101**, San Jose, CA, Manuscript #6101, L1011.

[5] C. M. G. Heffels, D. Heitzmann, E. D. Hirleman, B. Scarlett, "The Use of Azimuthal Intensity Variations in Diffraction Patterns for Particle Shape Characterization," Part. Part. Syst. Charact. Vol. 11, 1994, pp. 194-199.

[6] H. Wang, R. Valdivia-Hernandez, "Laser Scanner and Diffraction Pattern Detection: A Novel Concept for Dynamic Guaging of Fine Wires," Meas. Sci. Technol. Vol. 6, 1995, pp. 452-457.

[7] L.M. Kegelmeyer. "Applying Avatar Machine Learning to NIF Optics Inspection Analysis," Conference on Intelligent Data Understanding, Oct 6, 2010, p. 42.

[8] Abdulla G, Kegelmeyer L, Liao Z, Carr W. "Effective and efficient optics inspection approach using machine learning algorithms." SPIE Laser Damage Conference, Boulder, CO. (Sept 2010)

[9] N.V. Chawla, L.O. Hall, K.W. Bowyer, and W.P. Kegelmeyer. Learning ensembles from bites: A scalable and accurate approach. Journal of Machine Learning Research 5 (2004), 421-451.