LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-797551

Machine Learning Algorithms for Automated NIF Capsule Mandrel Selection

K. J. Boehm, R. C. Blake, A. Garcia, K. Sequoia, S. Sundram, W. Sweet

November 22, 2019

Fusion Science and Technology

Disclaimer

Machine Learning Algorithms for Automated NIF Capsule Mandrel Selection

K.-J. Boehm[1], Y. Ayzman[2], R. Blake[2], A. Garcia[1], K. Sequoia[1], S. Sundram[2], W. Sweet[1]

[1]General Atomics, P.O. Box 85608, San Diego, California 92186-5608
[2]Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94550

Abstract

Small shells, ca. 2mm in diameter, made from Poly(α–methylstyrene) (PAMS) are used as mandrels in the production of glow-discharge polymerization (GPD) capsules located at the center of inertial confinement fusion experiments. The visual inspection process of microscope images of these shell mandrels, including detection of micron-sized defects on the shell surface as well as the handling and sorting, is a very labor intensive, repetitive, and highly subjective process, which stands to benefit greatly from automation.

As part of an effort to decrease the number of labor-hours spent in capsule handling, inspection and metrology, the development of robotic systems was presented in earlier publications [1]. The current work expands the automated image acquisition systems developed previously, and adds the use of convolutional neural networks (CNN) to select capsules best suited for use in the downstream production process. Through the use of these machine learning algorithms, the selection process becomes robust, repeatable, and operator independent. As an added benefit the system developed as part of this work is able to provide defect statistics on entire shell batches and feed this information upstream to the production team.

Primary Author: Kurt Boehm, boehm@fusion.gat.com, ph 858-455-3959

1 Introduction

A decade ago, engineers and designers started developing concepts for the design of cryogenic targets to be shot the National Ignition Facility. While considering manufacturability, these targets were very specialized one-of-a-kind designs [2], hand crafted and shot at a rate of a one or two per week. Since then, the target designs were organized into target platforms [3,4] with common design features and exchangeable components to allow for batch-production of components and subassemblies to keep up with the call for more frequent shots on the NIF [5].

General Atomics, as a supplier for target components and subassemblies is constantly challenged to upgrade the production processes and metrology recipes to help keep up with target demand, which has increased ~5 fold within the last 10 years. The required increase in productivity is achieved in part through the development of automated systems [1].

Modern production facilities across the board, from consumer goods, electronics to automotive and medical device industries are adapting computer controlled robotic arms, machine vision, image recognition and machine learning algorithms in their processes to increase facility throughput, reduce man-hours spent on a task and to provide consistent and operator – independent results. This gives the developer in a rather small scale operation such as component

manufacturing for NIF targets a vast and rich pool of already developed robotic tools and numeric algorithms that can be applied to small-batch production volumes. Compared to most industries mentioned above, the batch size required to keep up with the demand of NIF targets (~400 shots per year in the most recent years) is rather small, challenging the economics of the somewhat costly development automated systems.

The current production processes in place to produce target components were examined to identify candidates, in which a modest investment in automation would yield significant cost benefit. Typically, we are looking for processes, in which rather large quantities of similar parts need to be handled, inspected, assembled or otherwise treated.

A small capsule (~2mm diameter) is located at the center of most NIF inertial confinement fusion (ICF) experiments. During the experiment energy is deposited on the capsule surface, either through direct engagement of laser beams onto the surface or through indirect exposure of the shell to an X-ray bath inside a hohlraum. The goal is to compress the sphere to fusion relevant temperatures, densities, and pressures.

Due to the basic physics of laser driven inertial confinement fusion, this capsule requires manufacturing and inspection to meet nanometer-scale out-of-roundness and surface roughness specifications as well as being virtually defect-free. There are three most commonly used capsule wall materials, also called the ablators as the material is getting ablated outwards during laser irradiation on the target, ultimately providing the propulsion force compressing the capsule. The production process of two out of these three ablator materials, namely Glow-Discharge-Polymer (GDP) and Beryllium starts by selecting PAMS mandrels, onto which the GPD and Beryllium is deposited. These mandrels have similar specifications for surface defects and shape tolerance as the final product since they determine the quality of the inner surface of the shell and solely dictate the roundness of the final product.

The effects of the PAMS mandrel surface imperfections on the quality of the final shells, whether it be GDP or Beryllium, were identified early on as potential show-stoppers of this production process [5, 6] and for laser-driven implosion research overall. With the typical yield of production-worthy PAMS mandels being in the single digit %-age points for a production run start to finish, we found the sorting, selecting and handling of these mandrels to be process that would hugely benefit from the application of an automated system. A process was developed in which a robotic arm handles the shells, machine vision algorithms image the shell surface and machine learning algorithms trained on past production data evaluate the production-worthiness of the shells. The focus of this paper is the application of the image acquisition and machine learning algorithms.

2 Motivation

About 6 years ago, General Atomics started deploying the first automated systems to help streamline the capsule inspection process [1]. Increasing the systems' throughput while decreasing the operator involvement promised to deliver better shells since more units can be analyzed at the same labor cost, allowing to pick the best shells from a bigger pool. However, there was still an operator involved in rating the quality of the shell, leading to some level of inconsistency in the final determination.

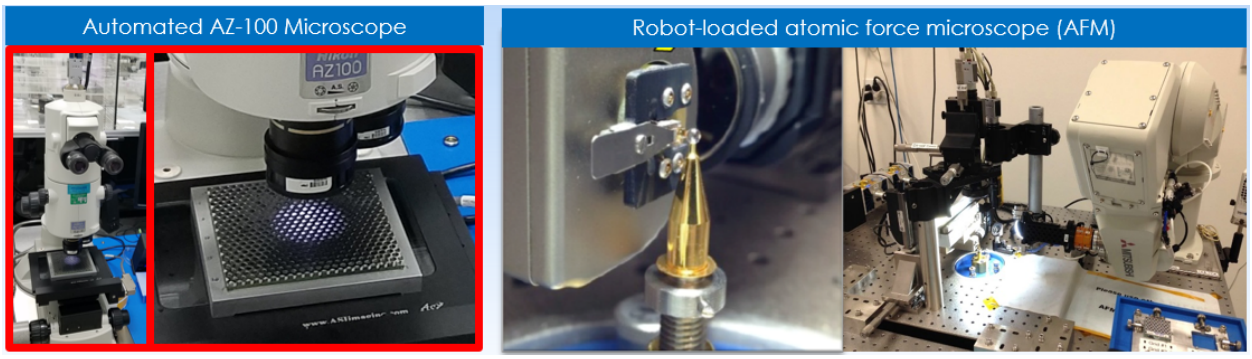

Figure 1: Automated stages, robot arms and image processing systems were first introduced into the capsule production process in 2013. (Left) The AZ-100 Microscope can image a shell and produce two images showing the capsule surface of the upper and lower hemispheres in less than one minute. Up to 300 shells can be imaged in one batch. (Right) The automated sphere mapper consists of a robotic arm to load and unload shells to a rotating spindle chuck, which presents the capsule to an Atomic Force Microscope (AFM) to measure shell Out-of-Roundness. Batch sizes of up to 30 can processed in fully automated fashion in ~4 hours.

The work presented in this paper builds on the lessons learned from the development and operation of the systems mentioned above. Methods were evaluated to remove the operator from the loop by having a computer decide whether a shell passes the optical surface inspection. In addition, it was recognized that data on the failure modes of a shell is a helpful mechanism to provide feedback to the PAMS production process upstream. Specifically, knowing the frequency and type of certain features on the shell surface or shell failure mechanisms as they appear in a whole batch would provide valuable insights into the effect certain production control parameters might have on the overall yield of a batch. Collecting this data manually would be prohibitively expensive.

Currently, up to ~650 shell mandrels are inspected each week both optically for surface defects as well as through an AFM for Out-of-Roundness. With the automated AZ-100 microscope and the auto-AFM in place, this results in 8-12 man hours of labor. Most of this time is spent loading shells into inspection grids and picking out high quality shells for further processing, as well as the analysis of the images from the AZ-100 microscope. In this process step, the goal is to pick 24 PAMS shells per week that meet the quality specifications for defect and shape for a mandrel. These 24 shells go into a GDP coater and are further processed into GDP shells. Roughly two out of that batch meet the criteria to be built into Capsule Filltube Assemblies (CFTA). From these numbers, it is clear that the automation of shell inspection and handling only yields a high rate of return when applied to the capsule mandrels. Further downstream in the process, the production numbers are too low to justify the development cost of automated systems.

3 Methodology

The different steps of the capsule inspection process are shown in figure 2. The top image shows the laboratory hardware setup, while the bottom shows are cartoon of the inspection process. There are three inspection stations that the shells are passed through, each of which inspects the shells based on different criteria. The AZ-100 inspection microscope is shown in figure 1. The first station (top left) consists of a 6-axis robotic arm mounted to an optical table. The shells, usually delivered and stored in glass vials, each containing 1000s of shells, are

poured into a tray. The robot arm is programmed to pick each shell individually from the tray and to present it to an inspection camera on the table. Since the depth of field of the camera is less than the radius of the shell, multiple images are acquired and stitched together into a "flattened" 2-D image of one hemisphere of the shell. The shell is rotated by 120 degrees twice and thus imaged from three different angles.

These three images are fed into a first automated classification algorithm trained to recognize individual defects on the shell surface through a machine learning algorithm. The details of this step are outlined in section 4. If defects of a certain type appear the shell is removed from the production line, otherwise it is placed into a tray for further processing.

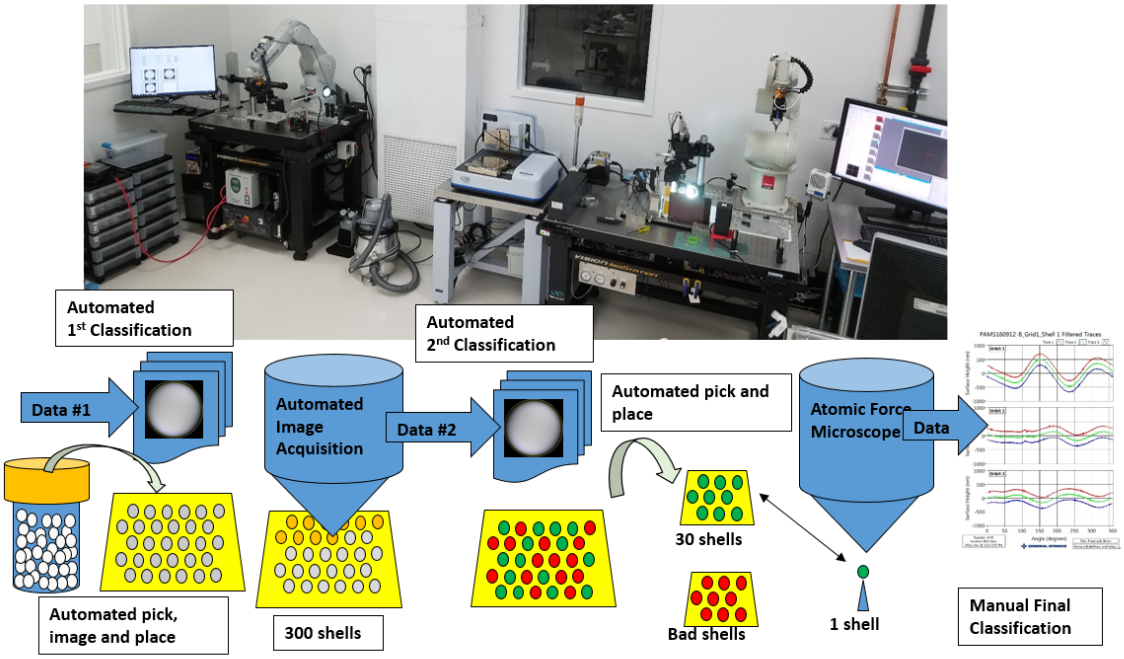

Figure 2: A photograph of the shell inspection line is shown on the top. Below is a schematic of the different inspection stations, each of which eliminates shells from the production line based on different criteria.

The robot cell described above delivers a tray of up to 300 shells, which are free from defects such as cracks and star cracks, but might still contain a large number of vacuoles or other defects, which do not automatically disqualify the shell from being used in production.

As part of the second step, each shell in the tray is imaged using a Nikon AZ-100 microscope. The resulting images are processed through a second machine learning algorithm, which is trained to look at the shell as a whole and determines whether the total number of otherwise acceptable surface defects renders the shell as unusable. The details of this algorithms are described in section 5.

In the following step, the robot cell mentioned at the beginning of the section is used to move the high quality shells into smaller grids, each of which contains up to 30 shells.

These trays of 30 shells are transferred to the auto-AFM, which loads the shells onto an AFM for Out-of-Roundness inspection as described in [1]. As a final step, an operator reviews the data from the Auto-AFM, and has the opportunity to review the surface data for only those shells that pass all three inspection steps. As a result, the operator is liberated from the tedious task of handling and inspecting shells individually, and is now only tasked with handling shells in trays

and grids, which are much easier to handle. In the following sections, the machine learning algorithms for the selection of shells is presented.

4 Inspection #1: Defect Recognition

Over the past two decades, members of the General Atomics shell production team have been manually inspecting shells, analyzing their quality and trying to set standards for minimum quality requirements. When attempting to set pass/ fail criteria, it was recognized that there is a wide variety of defect types, some of which are believed to have a greater impact on the final shell quality than others. Being able to distinguish between different failure types and recording statistics on an entire batch would provide a huge benefit in the study of how different production parameters might affect the quality of the batch as a whole, see fig 3 for examples.

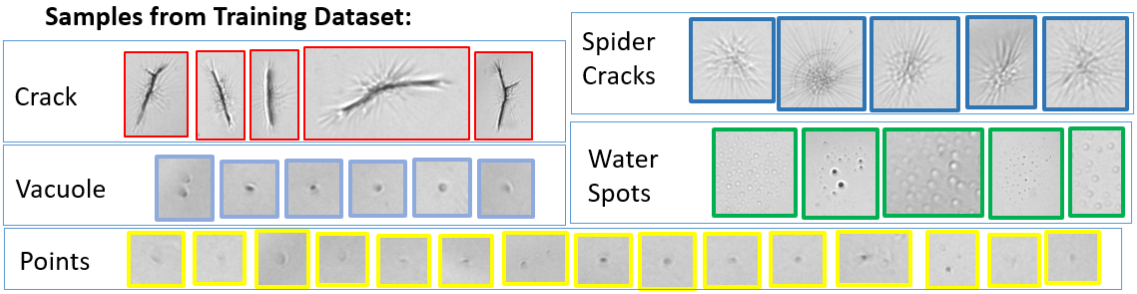

Figure 3: Different defect types are observed during the capsule surface inspection. The challenge was to develop an image recognition algorithm to detect and identify the defect types and report batch statistics.

We explored the capabilities of state-of-the-art machine learning algorithms. The problem of this type can be generalized as a “Region-based Convolutional Neural Network” R-CNN type problem, in which several objects within the same image frame are to be identified. This is a very common problem found in, among others, self-driving vehicle technology, and out-of-the-box solutions are available in all common mathematics libraries [7].

Several attempts to use algorithms from the MatLab [8] package to train a network to first recognize the regions of interest within an image and then classify the object inside the region of interest did not deliver satisfactory results. High background noise levels were observed in the images. Inconsistent contrasts in the feature edges added to the problem, along with inhomogeneous backlighting of the shell. Lastly, some features, such as water spots, appear in clusters, whereas vacuoles and point defects should be considered isolated defects, even if multiple instances of the same defect are observed in close proximity to each other.

Typically, an R-CNN type algorithm consists of two separate deep learning networks, one to find regions of interest (ROI), and one to identify the object.

We separated the two problems, and developed an NI LabView / Vision [9] image processing algorithm to identify regions of interest, crop it out and save it as an image of its own. This algorithm detects areas in the image that sand out from the background and searches the surrounding area of the shell surface to determine the extent of the feature or cluster of features. This process worked much more reliably than the deep-learning based ROI-finders in the R-CNN networks. Furthermore, the parameters of the ROI-search routine could be manually adjusted, such as to ignore or include faint features by setting the background noise threshold level, whereas the details of a deep-learning-based algorithm are much harder to adjust.

We imaged ~1000 shells and ran them through the ROI-finder, which lead to a library of ~15000 defects. These were manually sorted into five separate bins (see figure 3) to start building a labelled data set to train a deep-learning algorithm. Once a labelled dataset was assembled, a number of different deep learning networks were tested.

Due to the limited size of the training data set available and the large effort involved in developing a new type of deep learning architecture, the capabilities of existing deep learning algorithms were explored first. A number of network architectures can be found online, trained to recognize a large number of every-day objects. Since these defect types are not part of any library, the concept of transfer learning is applied.

This approach in itself is not novel, but the application of using a pre-trained network and use transfer learning to fine-tune the algorithm's capability to detect shell defects provides some interesting results and applications.

We used the libraries of pre-trained networks available for the Matlab platform. These pre-trained networks, such as Alexnet, VGG16, VGG19, GoogleNet, etc. are trained using substantial computing power on millions of images as part of the ImageNet database and the associated challenge [10]. The resulting network accuracy is the state-of-the-art in image recognition for 1000 categories of everyday objects (e.g. cats, dogs, etc). Through the transfer-learning process, these networks can be re-trained to fulfill very specific and customized image recognition tasks such as the identification of defects in a capsule image. Only small modifications to the network structures are required and they can be re-trained on high end desktop machines within a few hours to reach satisfactory accuracy.

We found that for our case, the VGG19 network worked the best in accurately determining the defect types. For this purpose, the images are re-sized to fit into the required input size for the VGG19 networks (226 by 226 pixels). In a typical training run, an accuracy of >95 % was achieved, with the biggest confusion coming from the "point defect" vs. "vacuole", which at times are hard to distinguish visually as well.

In its current state, the system acquires 3 images of the capsule. As a second step, a list of defects and their location within the image is produced, and finally the type of each defect is determined using a VGG19 network trained on the manually sorted data. This produces a list of defects for each shell, including the location and the type of defect observed.

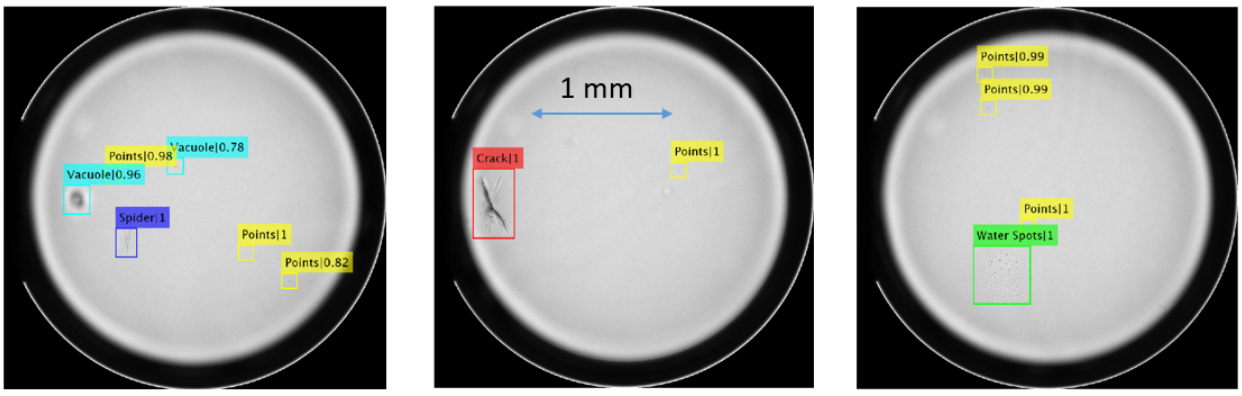

Figure 4: Three example images of three different shells are shown with their defects classified. The number next to the classification is the %-age certainty to which the machine learning algorithm is certain to have assigned the correct classification.

In a preliminary pilot run, we gathered the defect statistics for 10 batches of shells, each containing 100 shells. We counted the number of defects of each type as they occurred in each batch. Clearly, there is a significant variance in the occurrence of the different defect types. It will take a lot more data to draw meaningful conclusions as to which process parameter in the production of the PAMS shell might affect the frequency of which defect type, but the tools are developed to collect the data efficiently. It is also worth noting that the number of "Point defects" in the graph below is scaled as "instances/shell", while "Water Spots" and "Vacuoles" are scaled as "10 instances/ shell", and "Crack", "Star Cracks" and "Good shells" are instances per batch or 1 instance / shell. Overall there is a scaling factor of up to 100 applied to fit the data on the same plot.

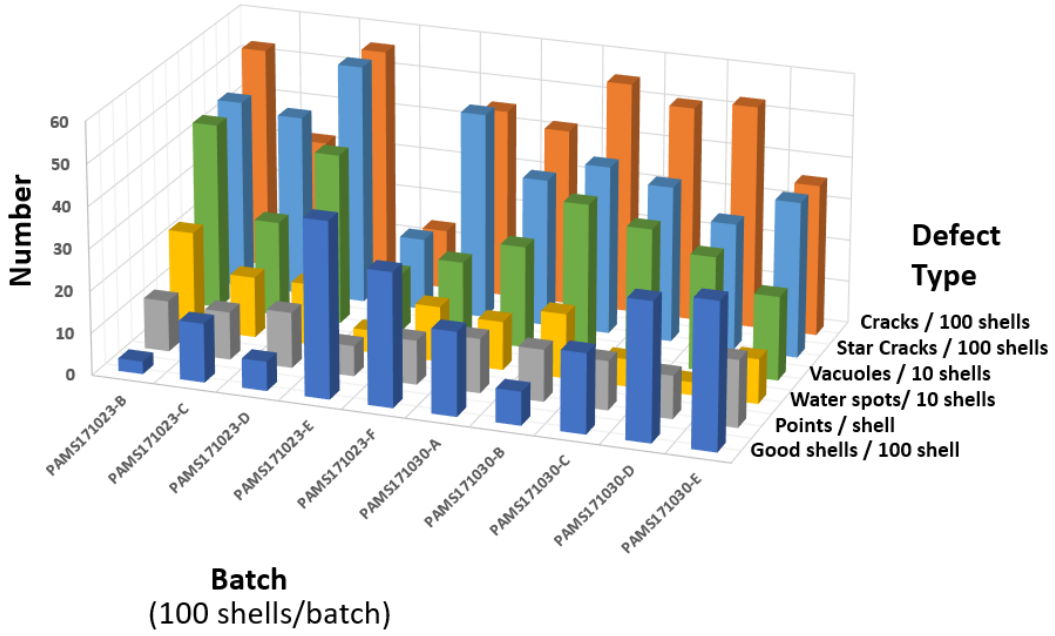

Figure 5: Defect statistics data was collected on ten batches of 100 shells each. For this particular trial, the ultimate determination of usable / not usable shell was done manually according to the previously established process. Clearly, the number of defects found by this algorithm tracks inversely with the number of good shells found in the batch.

We noticed that the occurrence of a single or even a dozen point defects or vacuoles does not automatically classify a shell as bad. However, shells with the defect type "crack" and "star crack" should be removed from the production line as they are definitely unusable.

We decided that we needed a second classification algorithm to analyze the shells based on the overall number and severity of features other than "cracks" and "star cracks" apparent on the shell surface.

5 Inspection #2: Shell Classification

The high contrast images taken by the robotic inspection station described previously show a large number of point defects and vacuoles in the shell surface. However, historical knowledge tells us that shells with these defects, if they are not too great in number or size, can be used in production. Attempts to set a criterion for size and number based on the algorithm described above did not yield satisfactory results when compared to the shells manually selected for production. The high contrast images created by that system are very useful to identify shells with star cracks and crack-type defects, which are sometimes undiscernible in the images taken with the AZ-100 microscope. In return, however, the AZ-100 microscope inspection has been part of the manual inspection process since it was first introduced in 2013, and was for years the only optical inspection step.

As a result, we have a wealth of images taken on the AZ-100 system since 2013. In this section, we describe how the historical information gathered on the AZ-100 inspection system was used to make predictions of the capsule quality of future capsules, when imaged on the same system.

The shell images of past production runs were available along with information as to whether a specific shell was used for production. This information can be gathered to serve as a label to classify the shell image data into "good" and "bad" shells. We were able to assemble a data set containing ~5000 shells, consistently imaged and manually classified by an operator within a one year time frame. Shells selected outside of this time window were classified using slightly different selection criteria. Within the last year, however, the shell selection criteria are considered stable.

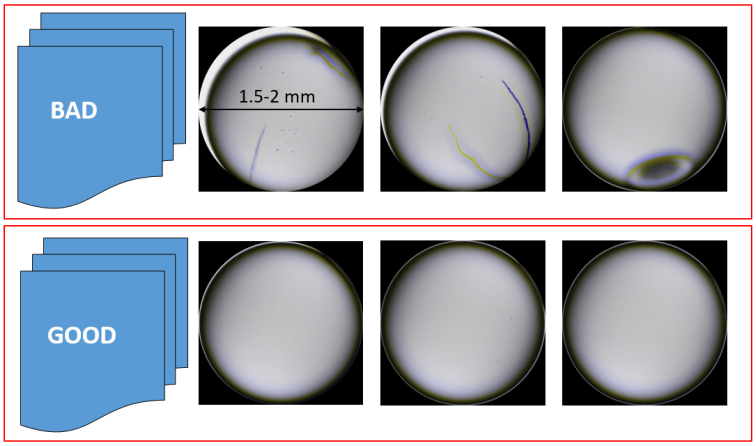

Figure 6: Schematic depiction of the labelled data set with some example images for "good" vs. "bad" shells.

In difference to the defect classification algorithms described above, we are now looking for a deep-learning algorithm, which will take the information of the entire capsule image and make a final determination as to whether this shell is usable or not. To some extent, this is much more similar to the way an operator determines the quality of the shell, by looking both at the shell surface as a whole as well as individual defects. The expectation is that deep learning algorithms can be trained to consider information on individual defects (small scale) and frequency thereof as well as presence of other defects (larger scale) using different layers of the algorithms.

In principle, image recognition of objects (e.g. cats vs. dogs) works in a similar way in that certain small features need to be considered and weighted against the frequency of occurrence of these features and the location in the image overall.

Similar to the approach applied to determine the defect classification described above, we looked for a pre-trained network from the Matlab database that would be best suited to determine the production-worthiness of a shell based on training over the past production data by applying transfer learning.

	TOTAL	% age accuracy	BAD	% age	GOOD	% - age
AlexNet	176	0.85	84	0.82	92	0.89
VGG16	179	0.87	86	0.83	93	0.90
GoogleNet	177	0.86	87	0.84	90	0.87
VGG19	177	0.86	76	0.74	99	0.96
Operator Reclassify	155	0.75	70	0.68	85	0.83

Table 1: Different pre-trained networks were fine-tuned through transfer learning using the same training set and tested against the same validation set. In total, 206 shells (103 "good" and "103" bad ones) were used in a validation set. The second and third columns show the total number of correctly classified shells and the percentage. The accuracy of the models to predict each class is shown in columns 4 and 5 ("bad") and 6 and 7 ("good"). It is worth noting that some networks are more accurate in one class over the other.

At the first look, the overall accuracy of these networks seems rather low at ~85 %, considering that we are looking at a two-class type of problem. One reason for this, however, can be found in the way the training data set is assembled. The classification of each shell is done through an operator's judgement as to whether the shells are believed to be of sufficient quality to move on to production or not. On the other hand, the quality of the shells would be better judged on a wider range, say a scale from 1-10. So, while the quality of the shell should really be graded on a sliding scale, the operator is forced to make a binary decision, which inevitably poses challenges in the "gray area".

In an attempt to test for this, a randomly assembled validation set was presented to an operator a second time. We found that the operator was only able to re-classify the shells with an accuracy of 75%. All of the networks tested as part of this study performed better than that, indicating that this process suffers from significant ambiguity and operator dependability.

Visual inspection of the images for which the algorithm misclassified a shell showed that most shells in that category could be judged either "good" or "bad", further strengthening the case that these algorithms are very likely not going to hugely impact the downstream production yield, while they carry the potential to relieve the operator from sifting through 1000's of images per month looking at shell surfaces.

The VGG19 network was selected for our production system due to its high accuracy in finding "good" shells. There might be a higher yield loss as shells that might have been good enough for production are thrown out; however, in this particular application it is considered more important to not let "bad" shells be moved into the downstream production as opposed to throwing a "potentially useable" shell out. This algorithm has been applied to the PAMS

selection process starting in early 2019, and so far no significant drop in downstream yield was reported.

6 Summary

Several techniques were combined in this body of work to automate the handling and selection of PAMS mandrels for the production of high quality capsules used at the center of NIF ICF experiments.

Two separate sets of images are taken of each capsule using different lighting conditions and camera systems. Through the application of image processing routines, images are acquired, stitched together and assembled into a data set.

One algorithm detects regions of interest within these images and determines the occurrence and frequency of different defect types. The other one analyzes images of the capsules as a whole to automatically determine the capsule quality based on the overall “appearance” of the shells.

By combining both complimentary algorithms, the system shows a 91% accuracy compared to the manual selection process with an estimated 10-20% drop in yield. This drop in yield is considered an acceptable price to pay to completely remove the operator from the selection process. With these systems in place, we expect to get more consistent, operator independent results. Combined with the increased throughput, the goal is to deliver shells at an increased overall quality level as larger and large batches of shells can be sifted through in search of the best shells. As an added benefit, the system is able to deliver quantitative statistical data on the defects observed in each batch.

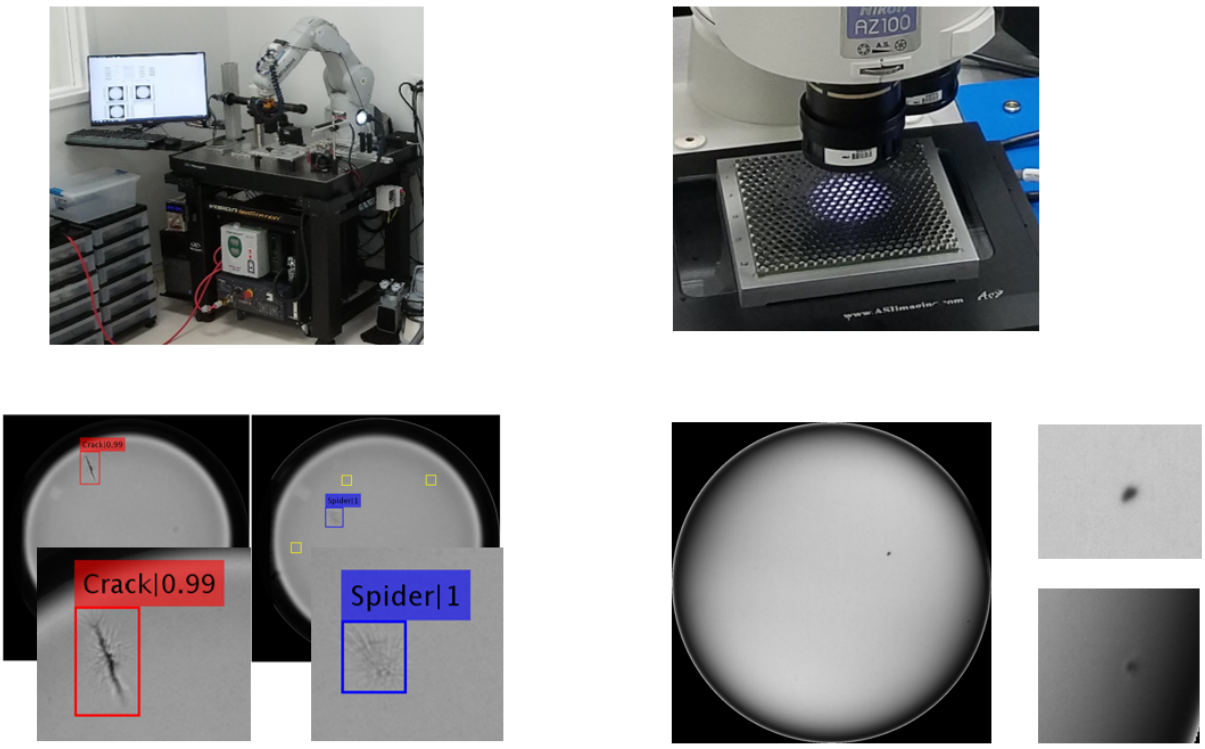

Figure 8: The culling station (left) is developed to load the tray with shells and to look for individual defects on the shell surface. In the secondary inspection an AZ-100 Microscope is used to look at the entire surface area of the shell to determine whether an accumulation of smaller defects disqualifies the shell from being used in production.

In addition to the automated data generation and shell selection, a robot arm was introduced to handle these fragile shells to load them into production batches. These previously manual pick and place operations of shells were labor intensive and repetitive in nature.

Systems like these could be developed for upstream and downstream inspection of shells (e.g. PAMS shells before they are dried) or GDP or Beryllium shells after they are coated and the mandrels are burnt out. The application of machine learning algorithms in the detections of features in the inspection of components might find other applications in the fusion technology field.

References

[1]L. C. CARLSON, H. HUANG, N. ALEXANDER, J. BOUSQUET, M. FARRELL, A. NIKROO, "Automation in Target Fabrication" *Fusion Science and Technology,* **70:2**, 274-287, (2016)

[2]J. KLINGMANN et al., "Design and Fabrication of a Novel Cryogenic Laser-Driven Ignition Target" Presentation to the *European Society of Precision Engineering,* (2007).

[3]E. DZENITIS et al., "Target Production for the First NIF Cryogenic Experiments," presented at 19th Target Fabrication Specialists' Mtg., Orlando, Florida, February 21–26, 2010.

[4]J. KROLL et al., "Design of a Cryogenic Target for Indirect Drive Ignition Experiments on NIF," presented at 19th Target Fabrication Specialists' Mtg., Orlando, Florida, February 21– 26, 2010

[5]M.TAGAKI, R.COOK, et al., "Decreasing Out-of-Round in Poly-Alpha-Methylstyrene Mandrels by increasing Interfacial Tension" *Fusion Technology,* **38,** 46, (2000)

[6]A. NIKROO, et al., "Coating and Mandrel Effects on Fabrication of Glow Discharge NIF Scal Indirect Drive Capsules" *Fusion Technology*, **41:3**, (2002)

[7]J.ESPINOSA, S.VELASTIN, "Vehicle Detection Using Alex Net and Faster R-CNN Deep Learning Models: A Comparative Study", *Advances in Visual Informatics,* **vol. 10645,** (2017)

[8]LabView and Vision are products distributed by National Instruments (2018)

[9]Matlab and various tool boxes are products developed and distributed by MathWorks (2018)

[10]Image-net.org

Declaration: Work supported by General Atomics IR&D and by LLNL under Contract DE-AC52-07NA27344. LLNL-ABS-767246