Evolution of machine learning for NIF Optics Inspection (OI)

LLNL Data Science Workshop

Laura M Kegelmeyer
Team Lead, NIF Optics Inspection Analysis

Aug 7-8, 2018

LLNL-PRES-754951
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

NIF re-creates the conditions inside of stars, giant planets; it routinely operates at energies that can damage optics

- The National Ignition Facility routinely operates at 1.8 MJ (8 J/cm2), twice the fluence that damages ordinary fused silica optics

- An optics recycle loop includes identification and repair of damage sites on optics, so specialty optics can be re-used, enabling continued high fluence operation of the laser

- Automated optics inspection (OI) informs and enables an efficient recycle loop


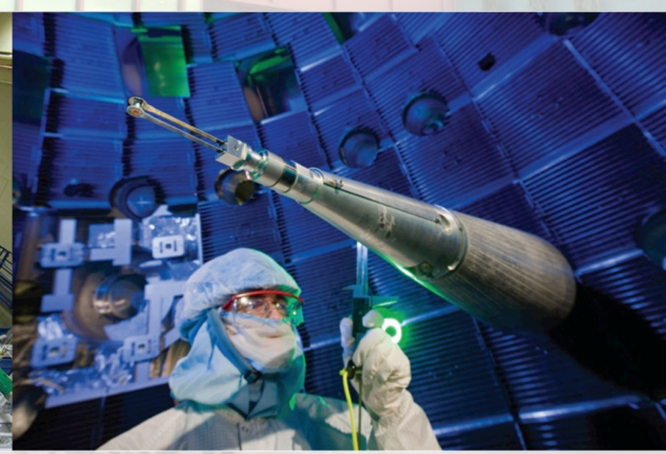
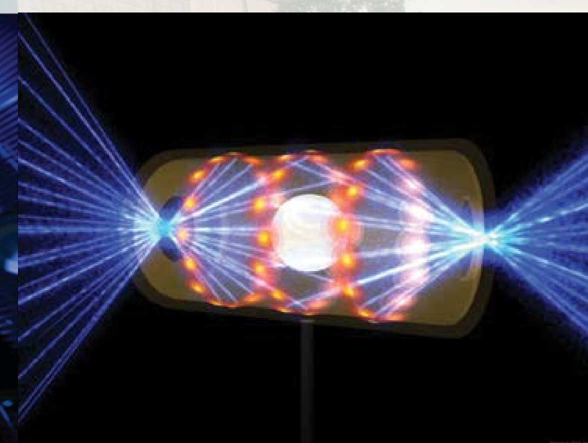

Since 2007, we’ve used machine learning to improve analysis accuracy, automation and quality control to inform and enable the NIF Optics Recycle Loop.

Several custom camera systems inspect optics *in situ* (on the NIF Beamline) to constantly monitor each and every damage site on thousands of optics

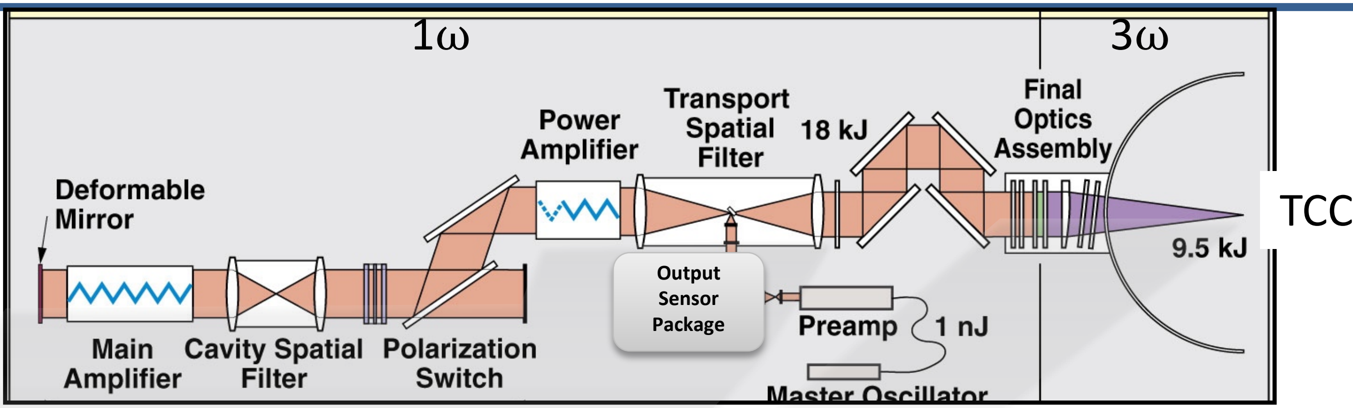

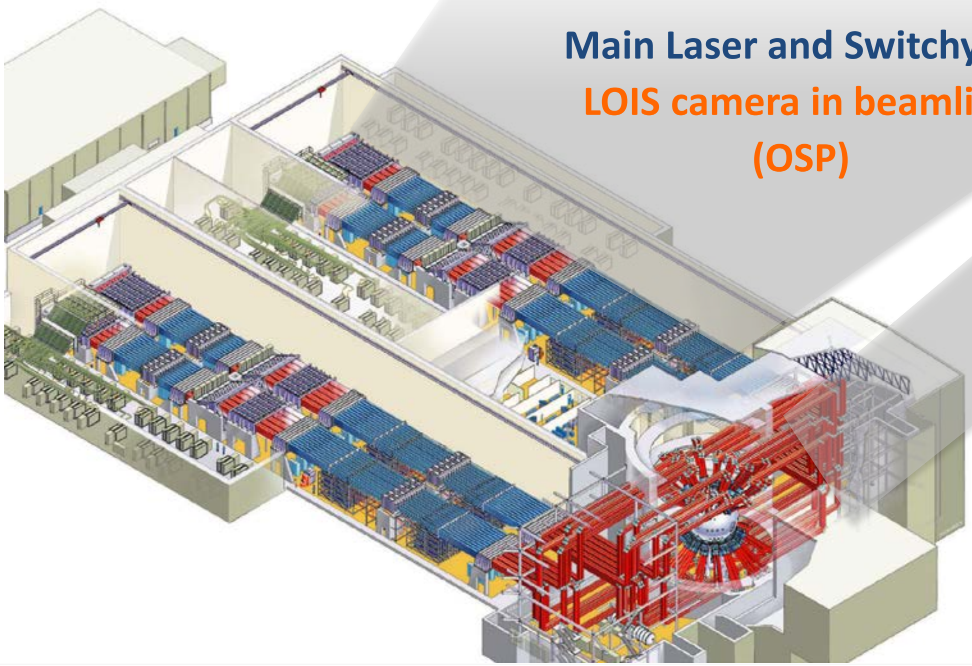

Main Laser and Switchyard
LOIS camera in beamline (OSP)

Final Optics Damage Inspection
FODI camera at target chamber center (TCC)

Side-Illuminated Damage Evaluation
SIDE camera in beamline
Inspects TCVW


Laser Bay
00ABC/abc • NIF-0508-14729s2

Backlit laser illumination travels through all the optics, picking up and carrying information along the way -- damage sites scatter light and leave a shadow

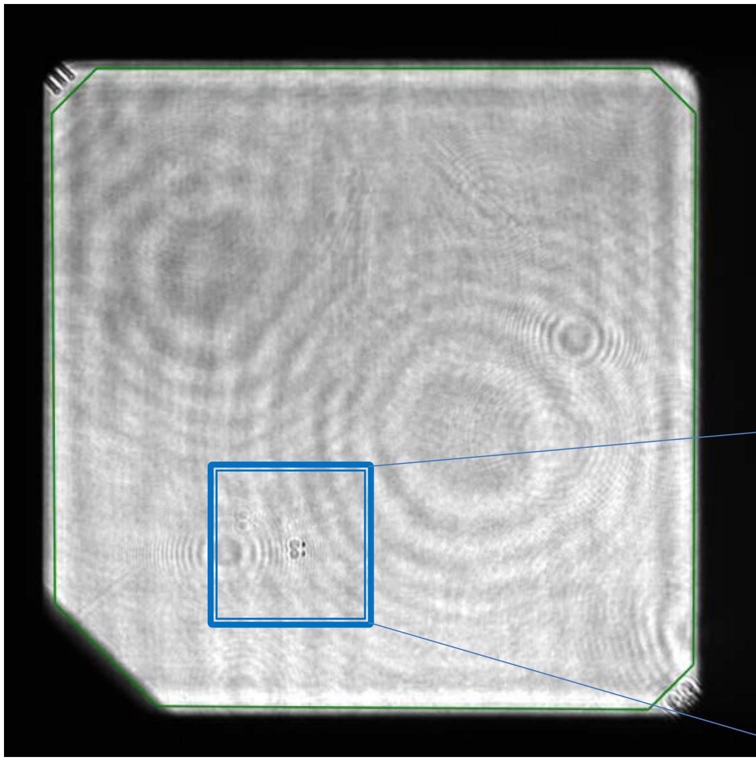

B266 SF3_LowZ N180206-002-000_180206_104329

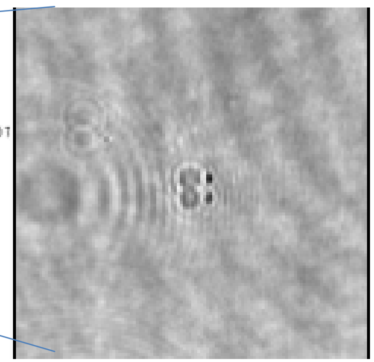

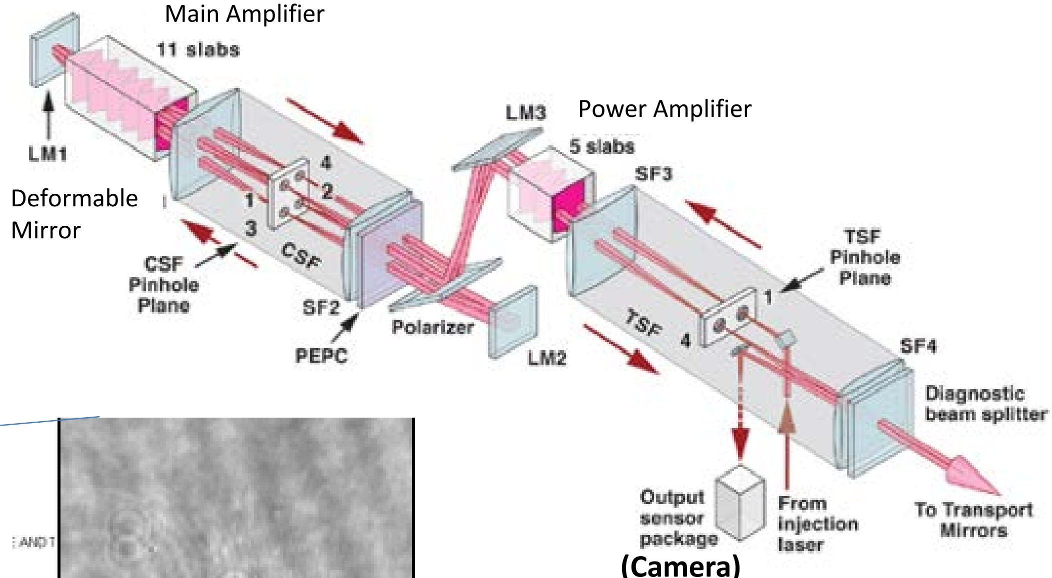

For 1 year (~2007) a dozen experts classified each site found.

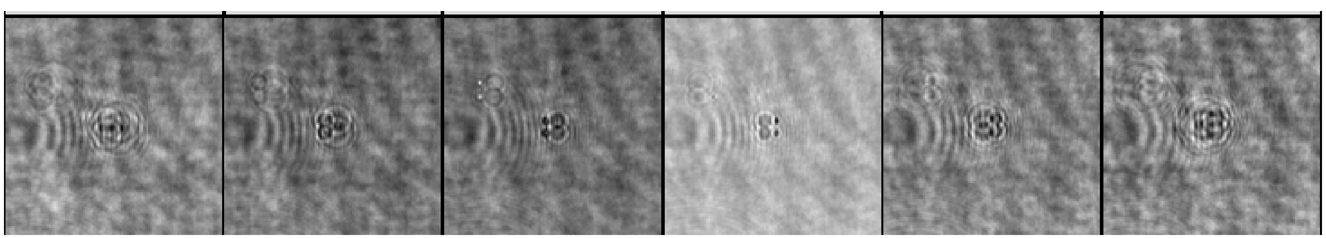

Supervised machine learning method, ensemble of decision trees (random forests), requires we analyze the images and extract information about the damage sites

Input a data file (text) that associates truth labels with results of analysis & feature extraction

Training sample	Measurements, attributes	Expert Truth
Candidate1	**Size1, OpticType1, Brightness1…**	**"Defect"**
Candidate2	**Size2, OpticType2, Brightness2…**	**"Camera Flaw"**
Candidate3	**Size3, OpticType3, Brightness3…**	**"Reflection"**

2-D Feature Space

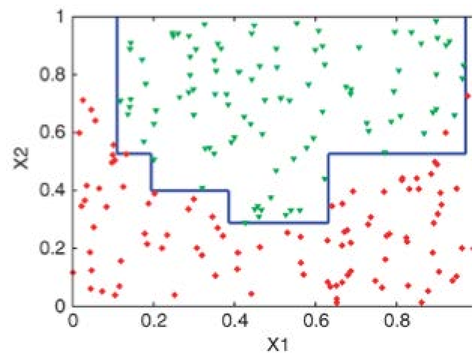

Decision Tree

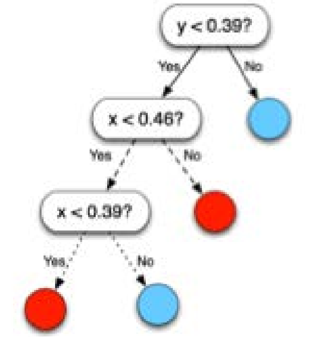

Grow decision trees

- Divide up features space with straight lines
- Use a different subset of data for each tree

Multiple trees provide a diversity of rules (partitions in feature space) and the final vote has higher accuracy than a single tree.

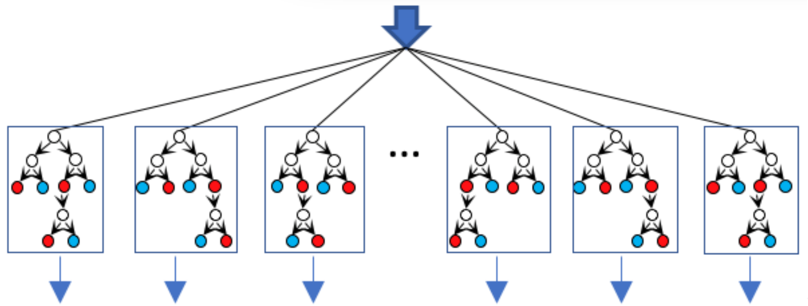

Steps for applying supervised machine learning (ensemble of decision trees) to NIF main laser optics inspections

Training phase →

Experts "Log Defects" to train data → OI Analysis team collects logged defects and "cleans" the data → Use cleaned data to grow an ensemble of decision trees

Test or prediction phase →

Operations: Send new images for analysis to find defect candidates → Candidates are "fed" to decision trees →

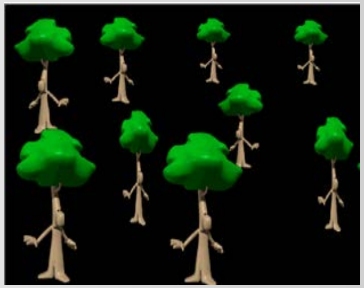

→ Trees vote for predicted class

OI software tracks defects through history, so if a candidate defect was classified during any one inspection....

Shot Id	Image Id	Defect Name Id	Beamline	System	Defect Id	Comment Id	Classification	Thumbnail
N050719-001-999	44129	663592	314	SHOTCYCLE_SF3	697238			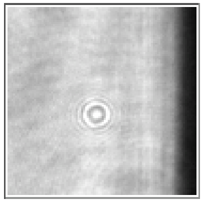
N050729-001-999	44364	663592	314	SHOTCYCLE_SF3	698434			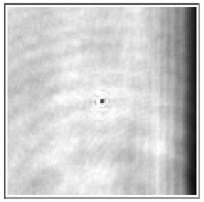
N050801-001-999	45114	663592						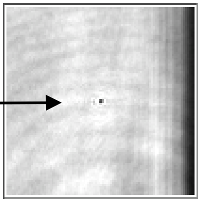
N050803-002-999	45235	663592	314	SHOTCYCLE_SF3	703542			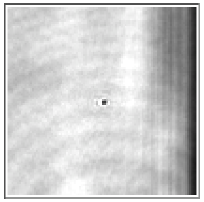
N050804-001-999	45320	663592	314	SHOTCYCLE_SF3	703906			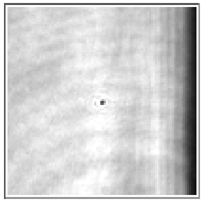

1. Label as “defect” once

... we could apply the same "expert truth" label to each instance in history to get nearly 6000 data points!

Shot Id	Image Id	Defect Name Id	Beamline	System	Defect Id	Comment Id	Classification	Thumbnail
N050719-001-999	44129	663592	314	SHOTCYCLE_SF3	697238		defect	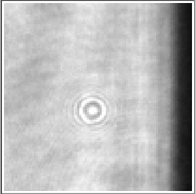
N050729-001-999	44364	663592	314	SHOTCYCLE_SF3	698434		defect	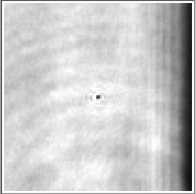
N050801-001-999	45114	663592						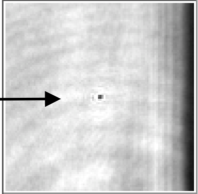
N050803-002-999	45235	663592	314	SHOTCYCLE_SF3	703542		defect	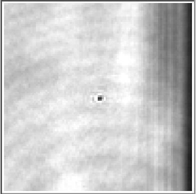
N050804-001-999	45320	663592	314	SHOTCYCLE_SF3	703906		defect	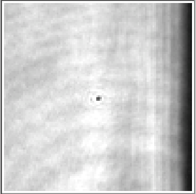

1. Label as "defect" once

2. Apply label to all

With traditional image analysis alone, the Inspection Summary Chart had too many false alarms to follow-up within the time constraints

Before

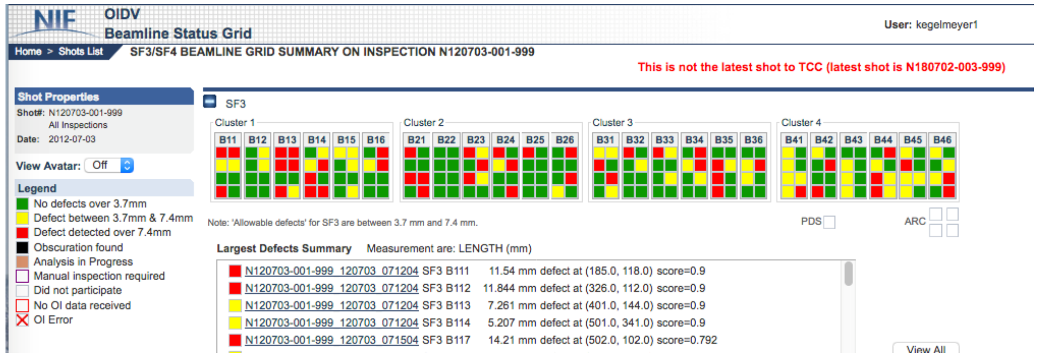

After

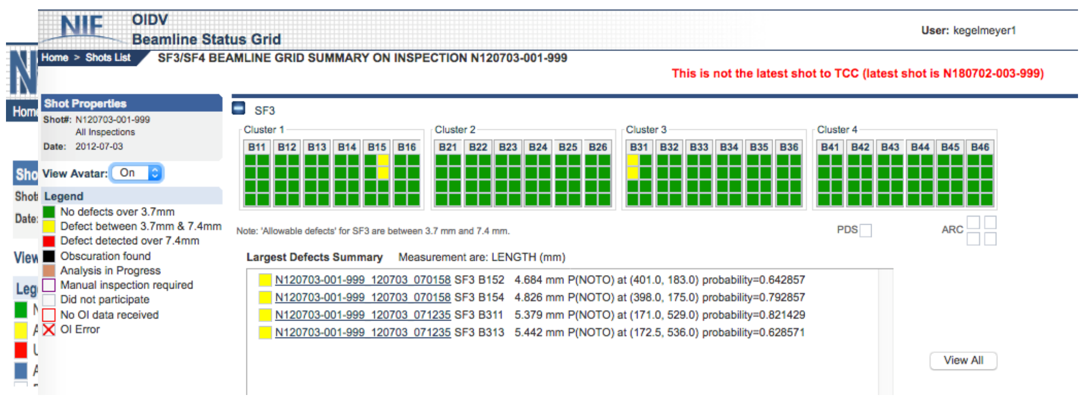

Reduction of false alarms allows operators to focus on the most relevant subset of the optics from 192 beamlines.

The Final Optics Damage Inspection (FODI) system includes a high resolution camera on a hexapod positioner inserted at Target Chamber Center

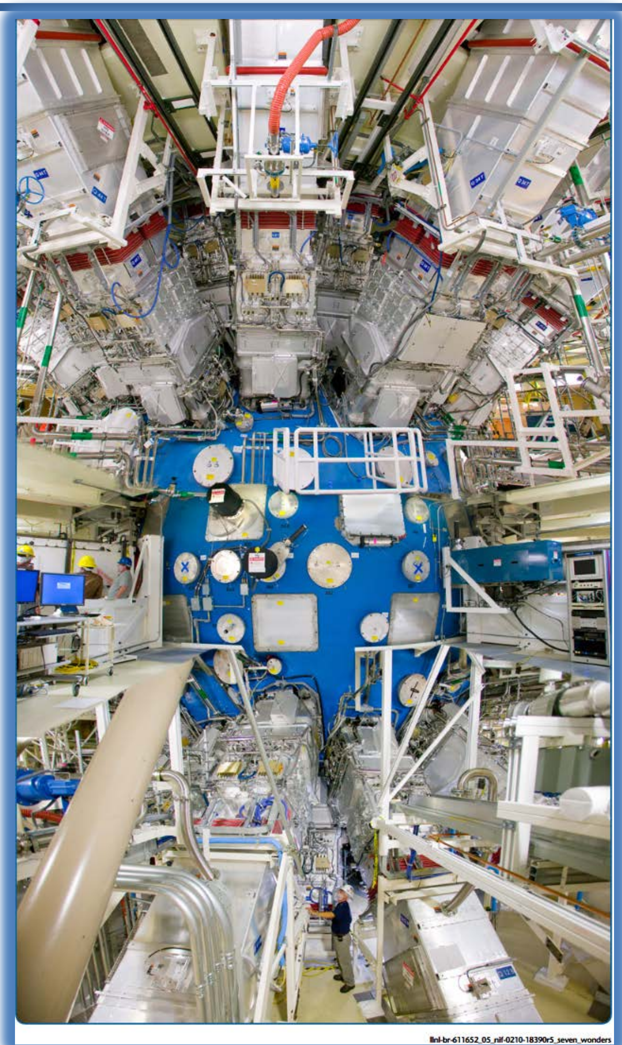

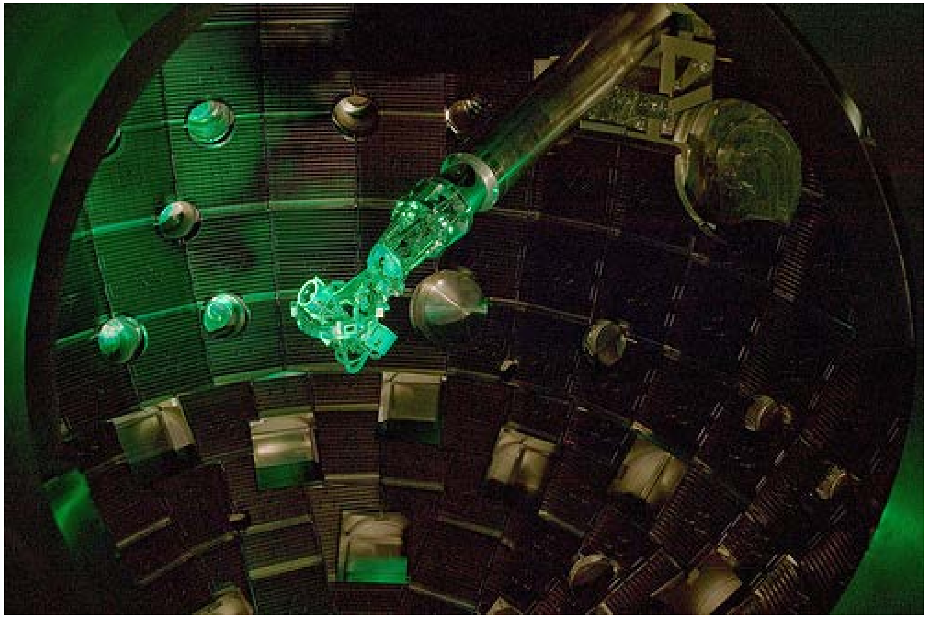

This system can image tiny damage sites (~20 microns) on optics from ~6.9 to 60 meters away (Debris Shield through LM4).

We again used ML-EDT to bring forth relevant damage sites for tracking

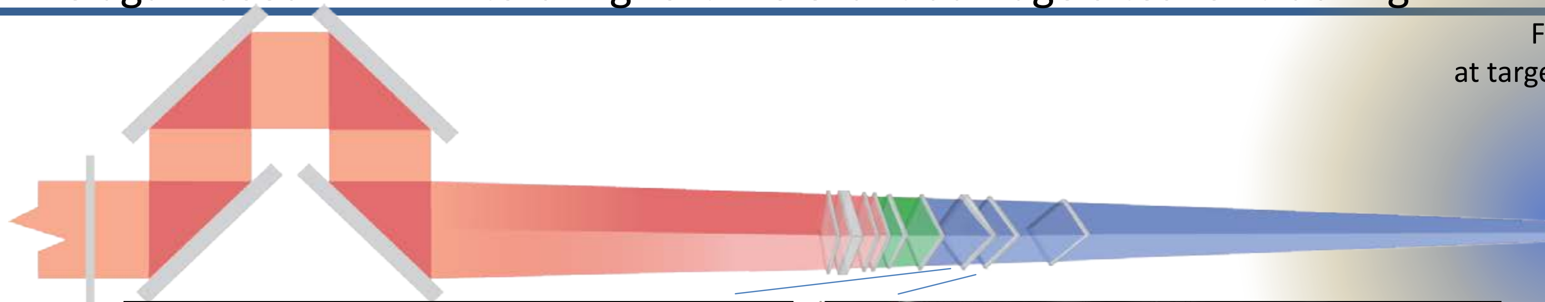

FODI Camera
at target chamber center

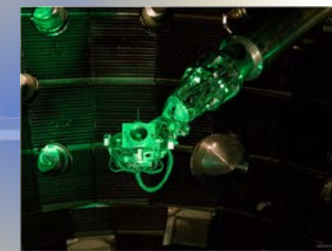

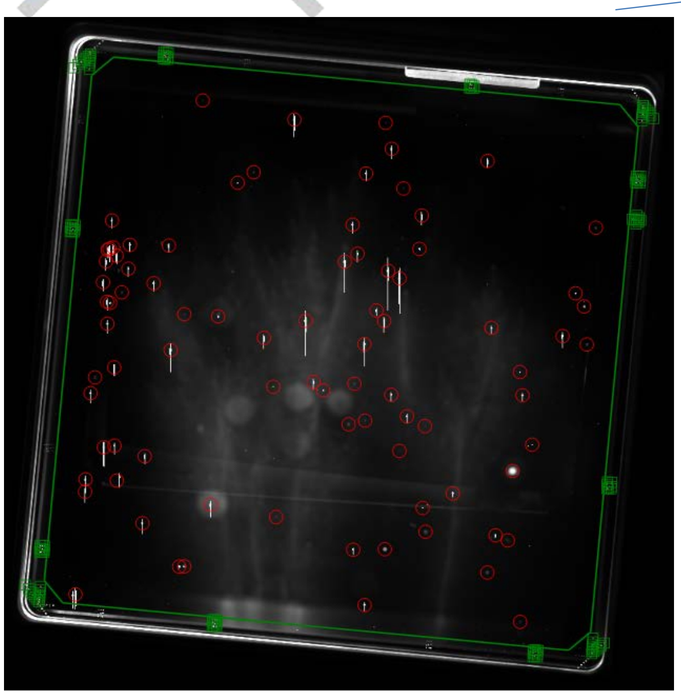

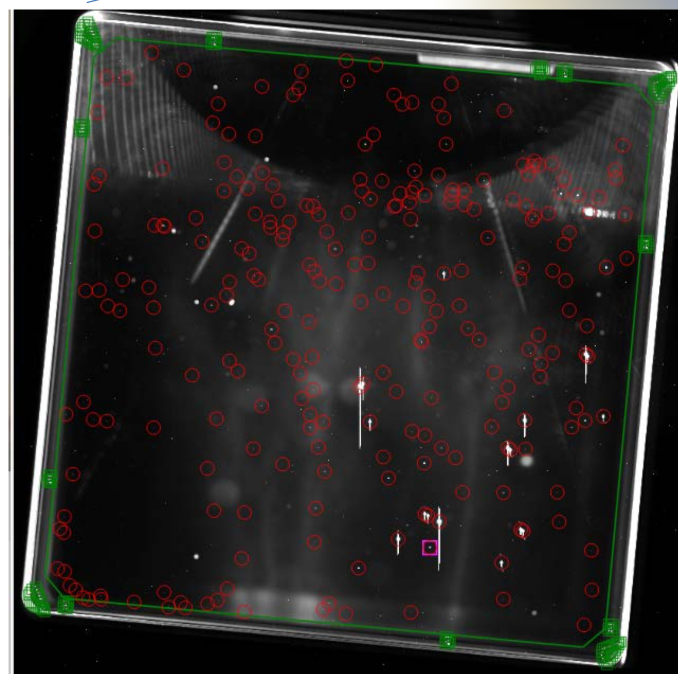

B438-N180103-003-000_180104_052624 WFLC GDS

Some detectable sites are damage/pits (actionable), while others are hardware reflections, stray/out-of-focus light, previously-repaired damage or camera flaws (not actionable)

We monitor growth of relevant sites and any approaching its optic-specific size limit is “blocked” until it can be removed and repaired

					Flaw Information									Blocker Information					
Quad	Beam	Optic Type	Optic SN	Taxon	Flaw Id	Observed Size (mm)	Dwg X (mm)	Dwg Y (mm)	Flaw Classification	First Seen	Last Seen	Growing?	Has Grown?	Available Blockers	NCR Number	I-Stat	Image	Growth Chart	Generate Blocker NCR 5 (5)
Q24B	B246	WFLF	272049	Mod4_A2_RAM	10507953	0.219	206.5	109.1		FODI	N120422-001-999_120423_112222	1	1		NULL			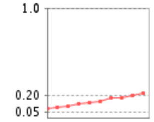	☐
Q24B	B247	WFLF	271167	Mod4_A2RAM_RAM	3326371	0.250	381.1	125.7	OMF SCRATCH, UNDAMAGED	ELV	N120422-001-999_120423_111632	1	1		456633	4.1		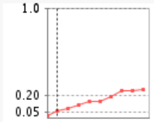	☑
Q24B	B247	WFLF	271167	Mod4_A2RAM_RAM	11151729	0.266	367.5	109.3	MITIGATED - RAM2000-E_4	FODI	N120422-001-999_120423_111632	1	1		457388	4.1		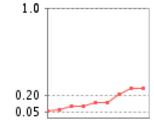	☑
Q24B	B247	WFLF	271167	Mod4_A2RAM_RAM	11115471	0.256	347.6	114.7	MITIGATED - RAM2000-E_4	FODI	N120422-001-999_120423_111632	1	1		458003	4.1		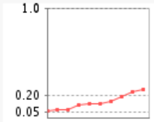	☑
Q24B	B247	WFLF	271167	Mod4_A2RAM_RAM	3305757	0.332	349.2	55.8	MITIGATED - RAM2000-E_4	OPL	N120422-001-999_120423_111632	1	1		456733	4.1		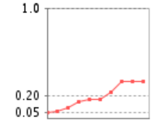	☑
Q24B	B247	WFLF	271167	Mod4_A2RAM_RAM	11036678	0.302	313.3	186.4	MITIGATED - RAM2000-E_4	FODI	N120422-001-999_120423_111632	1	1		457263	4.1		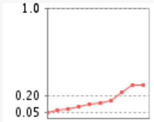	☑
Q24B	B247	WFLF	271167	Mod4_A2RAM_RAM	11058819	0.222	242.2	263.4		FODI	N120422-001-999_120423_111632	1	1		NULL			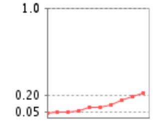	☐

Optics are removed from NIF, repaired and then re-used on NIF

Optics are removed from NIF and brought to this Optics Mitigation Facility to repair each relevant damage site

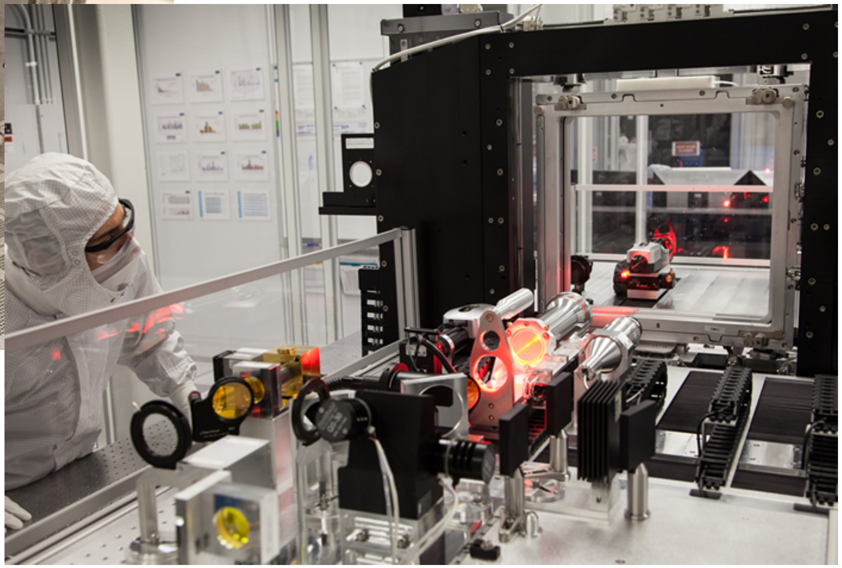

NIF recycles optics by finding, tracking and repairing sub-millimeter damage sites

- OMF repairs sub-mm damage found on NIF optics by etching a small cone over the damage site.
- This will effectively "erase" the damage from the view of NIF's pulsed laser light allowing it to disperse evenly.

Slide by Nathan Mundhenk

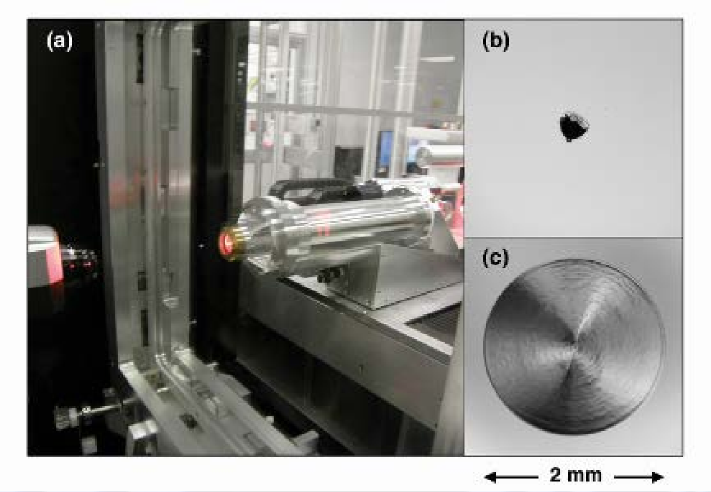

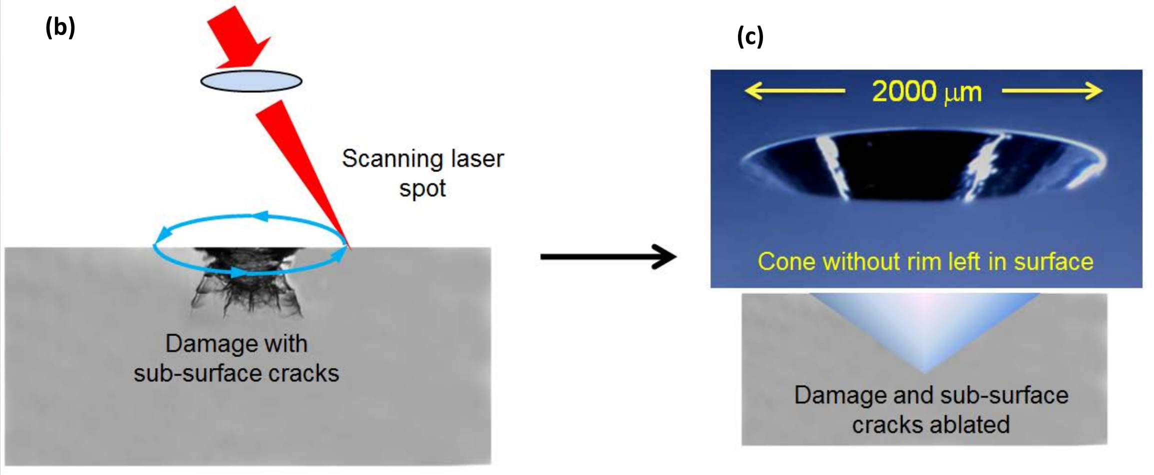

Remnant damage was difficult-to-detect even for expert operators

- Expert operators watched a video of the entire repair process (0.3 – 5 minutes) before an image of the final repair is captured.
- The final, still image can be nearly inconclusive to the human operator without context from the video.
- In 2016 we compared machine learning methods using only the final, still image. [TN Mundhenk, LM Kegelmeyer, SK Trummer]

Slide by Nathan Mundhenk

Remnant damage (Incomplete Repair)

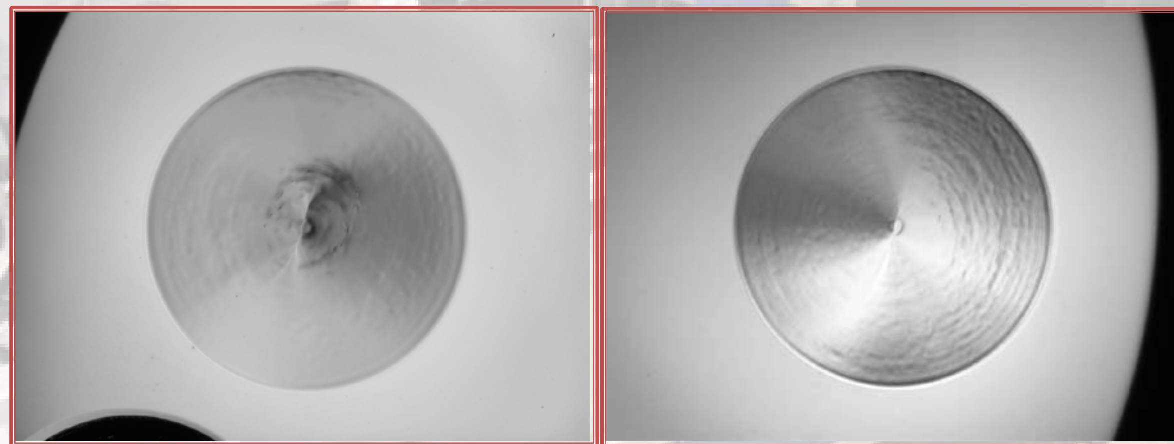

No remnant damage (Complete Repair)

Deep Learning is a subset of Machine Learning that can automatically determine which features to use for solving the problem at hand

Machine Learning

Labeled "truth" data.

100's

Feature Discovery & Extraction

Custom, Hand-crafted.

Training

Feature values

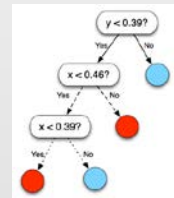

Runtime Decision Making

Sub-second

Deep Learning

1000's

Machine-learned (black box)

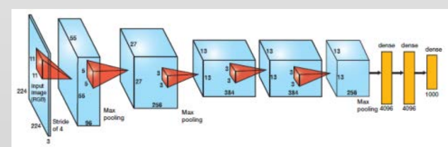

Entire image. Many layers. Huge number of iterations. GPU-optimized.

Sub-second with GPU

Transfer Learning: take a Convolutional Neural Net (CNN) model trained on a different (huge) dataset and re-tune it to work with the image dataset at hand

- ImageNet (Database): Millions of images from Google with labels Cat/Dog/Truck/Car …
- AlexNet (Large CNN): Trained to find what distinguishes one image type from another
- DamageNet (LLNL): Modified (tuned) AlexNet to distinguish our High-Res microscopy sites

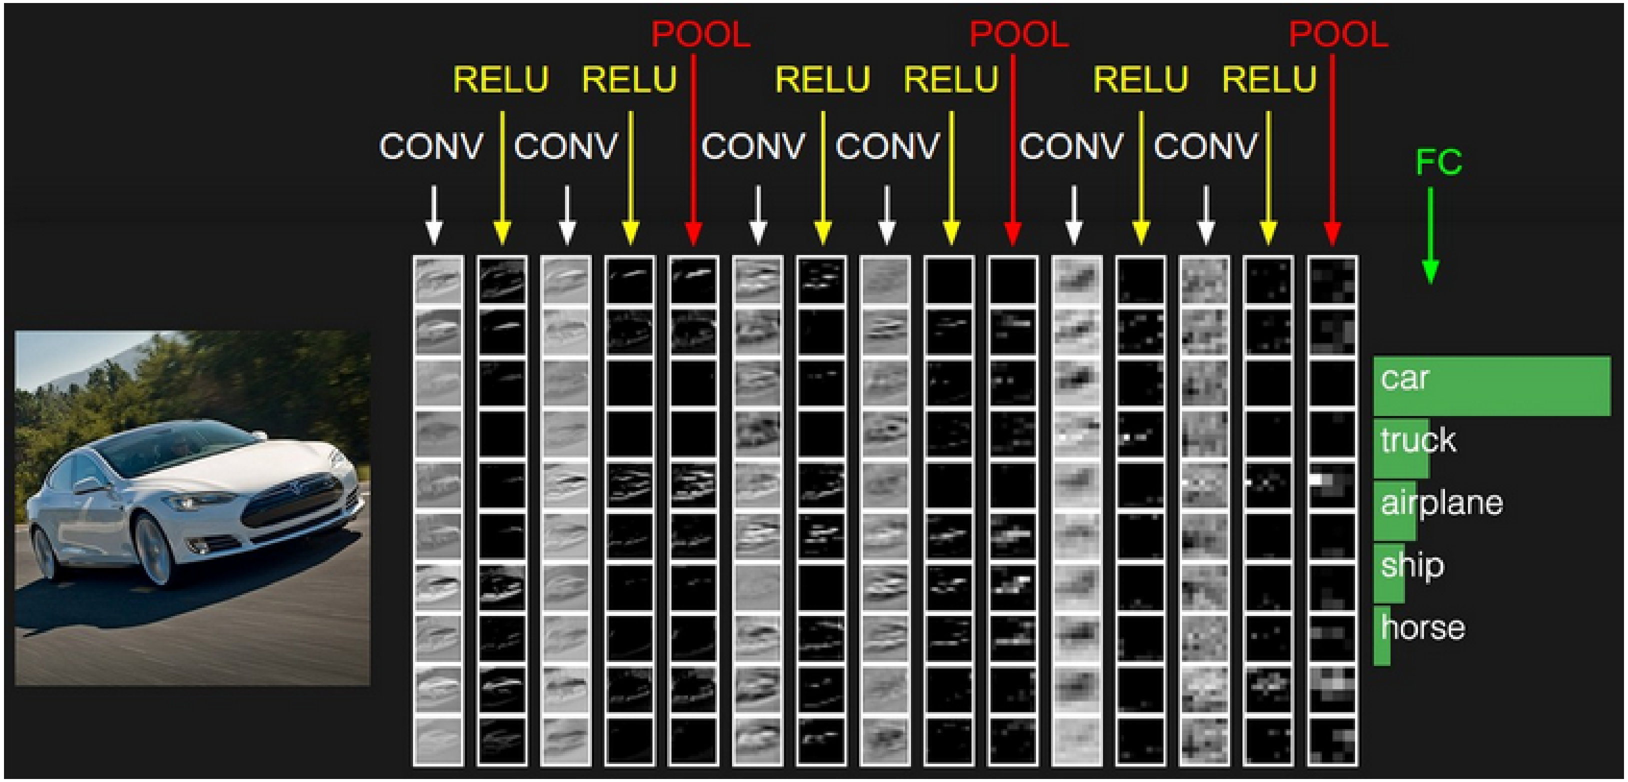

Image from Stanford Class Slides: http://cs231n.github.io/convolutional-networks/#fc

The automated method must handle the subtleties of still images, as well as various repair sizes, configurations, and illuminations

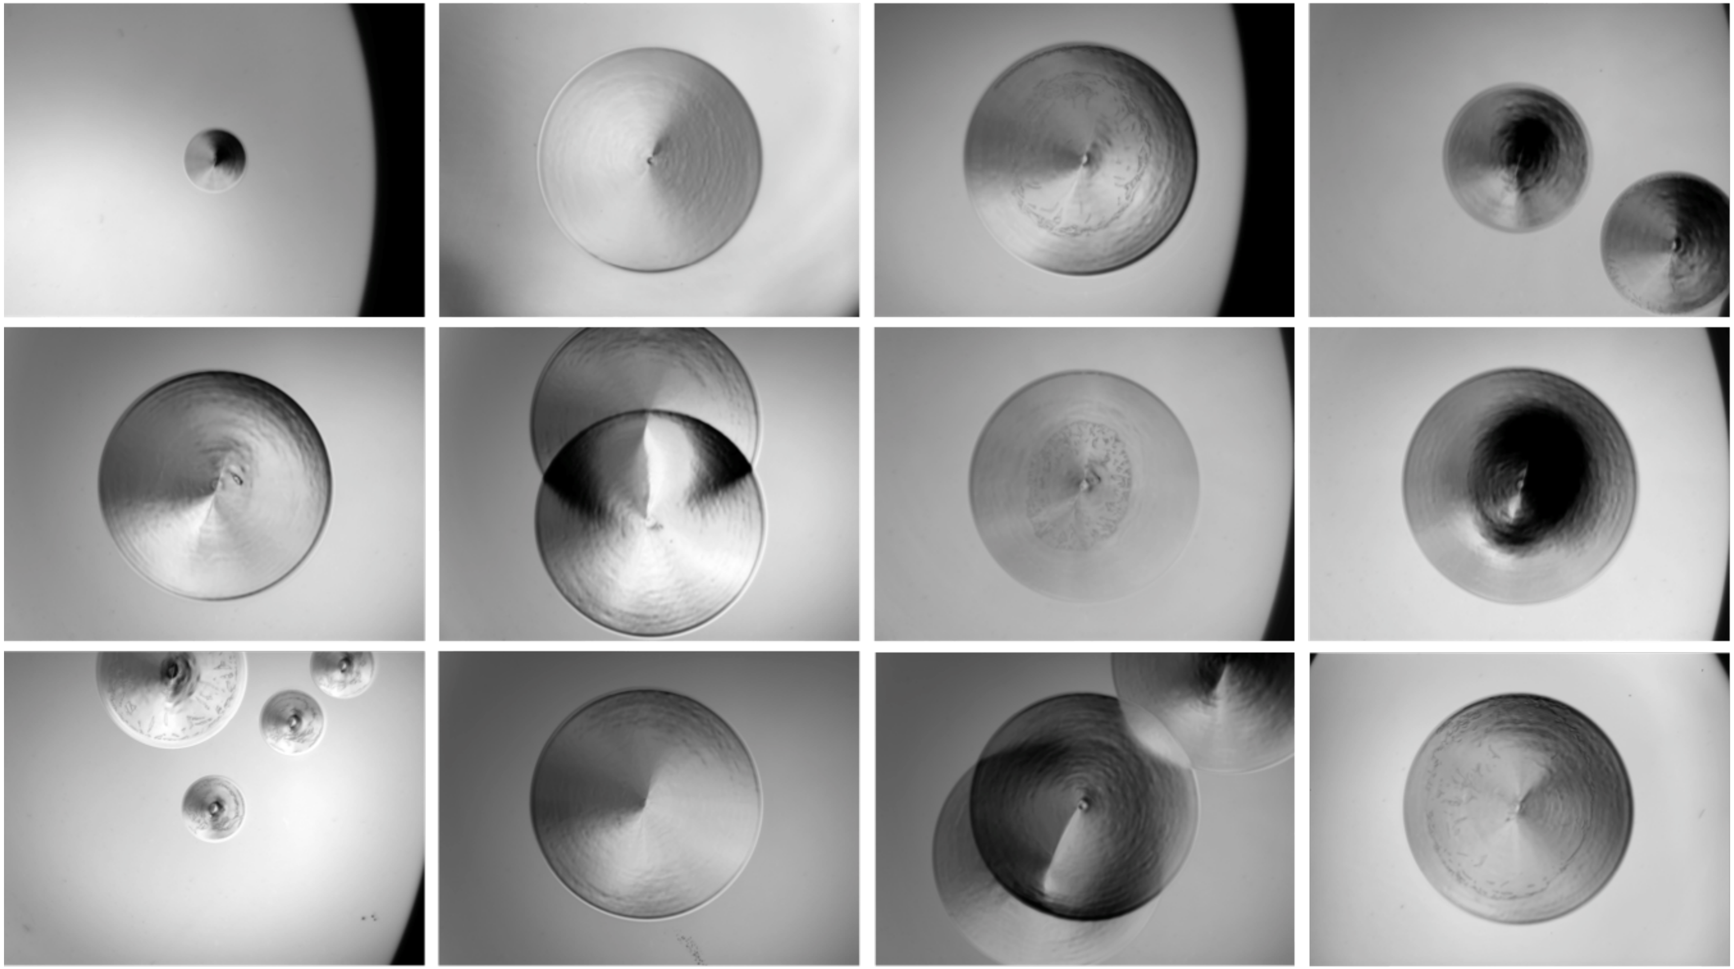

Slide by Nathan Mundhenk

We evaluated various supervised ML techniques and found likely improvement over human accuracy (estimated at ~91% worst-case)

Method	Accuracy
Decision Trees	93.55%
AlexNet	96.86%
ResCeption	97.52%
BN-GoogLeNet	97.65%
ResNet-152	97.91%
Inception-v2	98.17%

How to be sure the Machine Learning techniques aren't "cheating"? Use visual feedback.

Visualization of results using unsupervised feedback helped evaluate Deep Learning results

- Main idea: Rather than propagate the loss backwards through the network, we propagate the actual network output backwards.
  - This projects the output backwards to the places on the image **most responsible** for the result.
  - This is similar to how information is pushed backwards using Deep Dreams.
  - It is easy to do in *Caffe*. Just switch out the loss values in the “diff” layer with the outputs in the “data” layer and call the backwards phase function of the network.

Slide by Nathan Mundhenk

The network projection looks like its training data

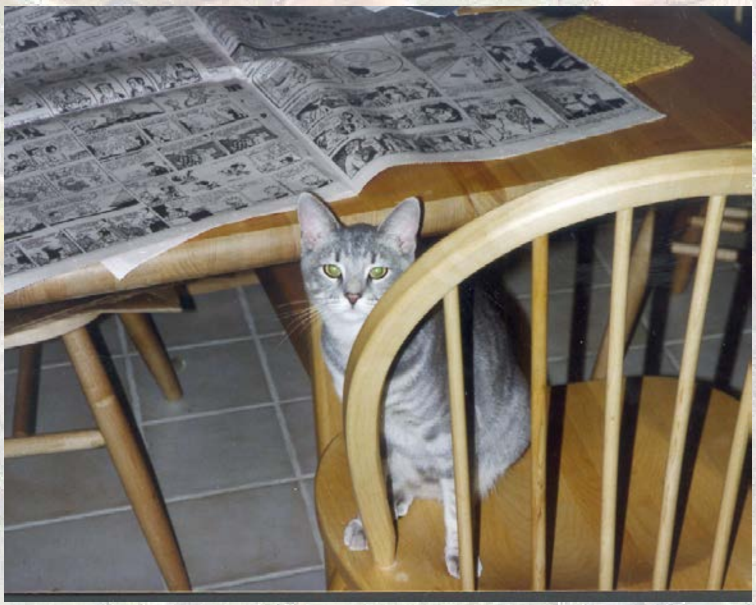

Input Image

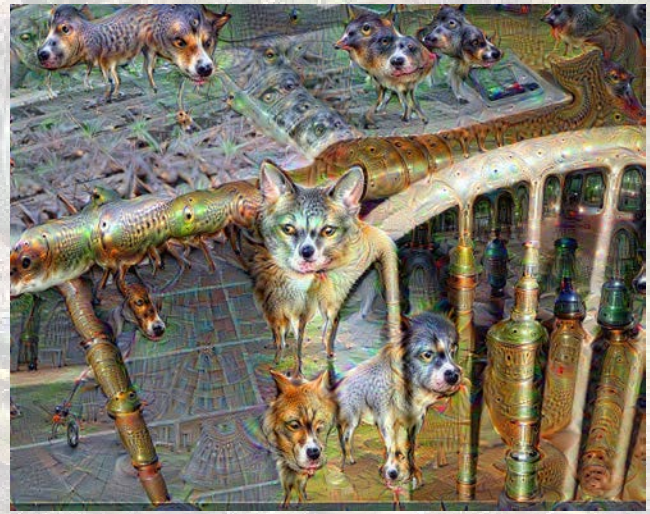

Network trained on *ImageNet*

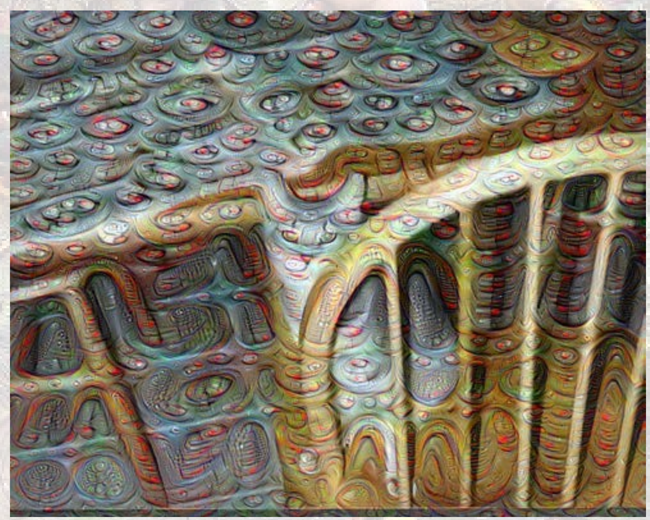

Network trained on *CompCars* (cars)

Slide by Nathan Mundhenk

Unsupervised Results – backwards projection creates a “heat map” showing areas of focus for the neural net

- Three images selected at random with detected remnants.
- Yellow ellipse is the ground truth provided by operator.

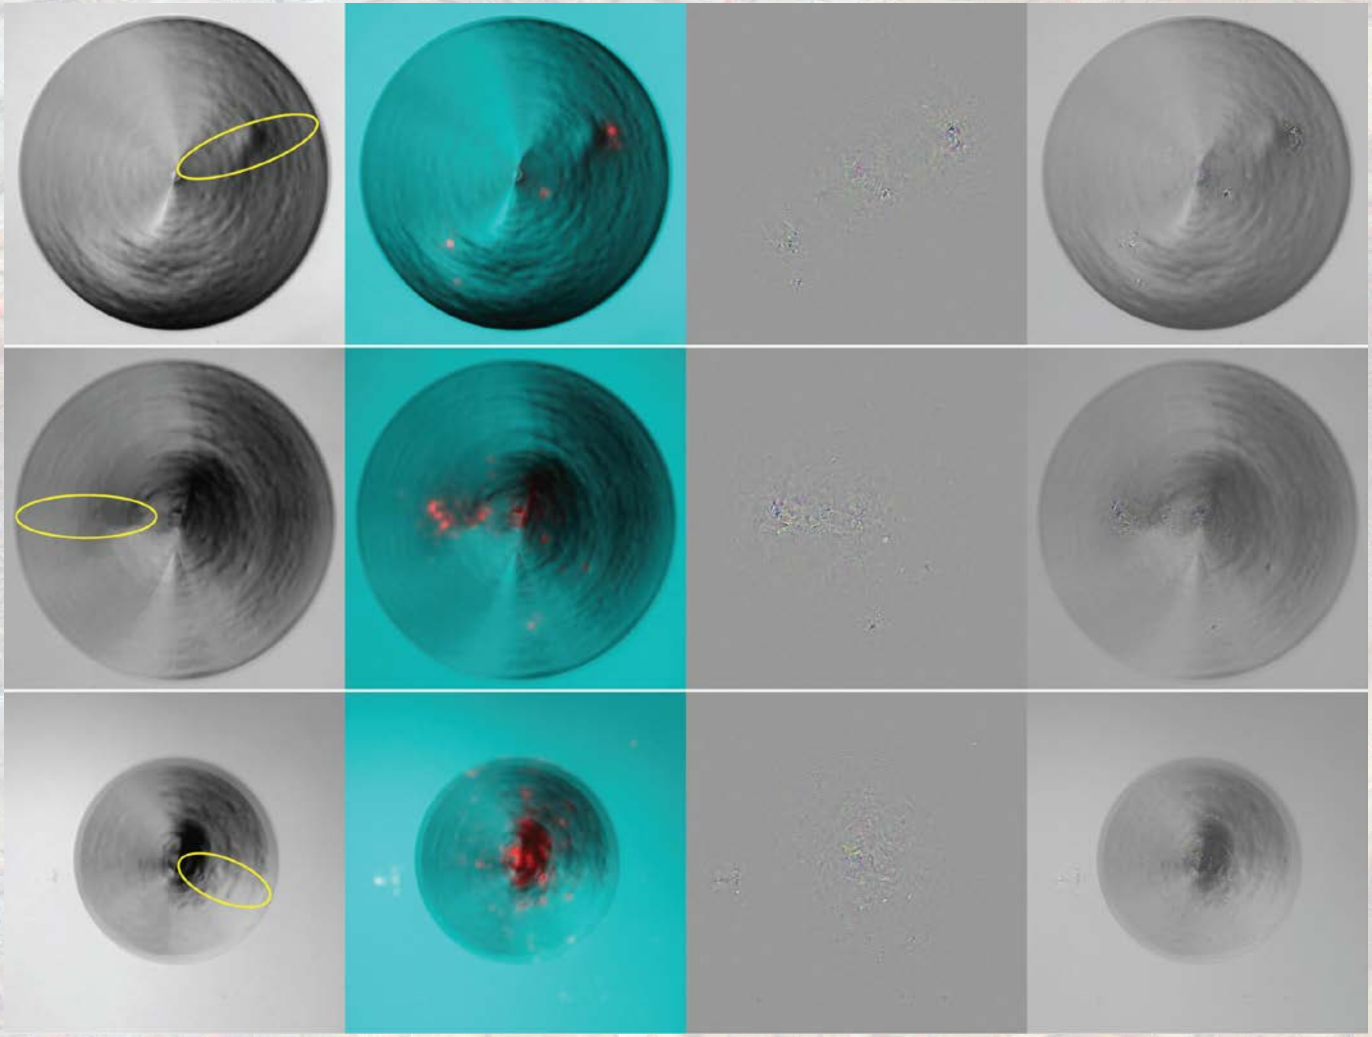

Model	Accuracy
AlexNet	96.86%
Inceptionv2	98.80%

(5x cross-validation)

Slide by Nathan Mundhenk

Summer 2017: Use transfer learning to automatically classify 12 types of damage morphology from (View/Nikon) scanning microscope

- Only a small fraction of tiny sites need to be repaired.
- Automatic classification makes it feasible to repair only these.

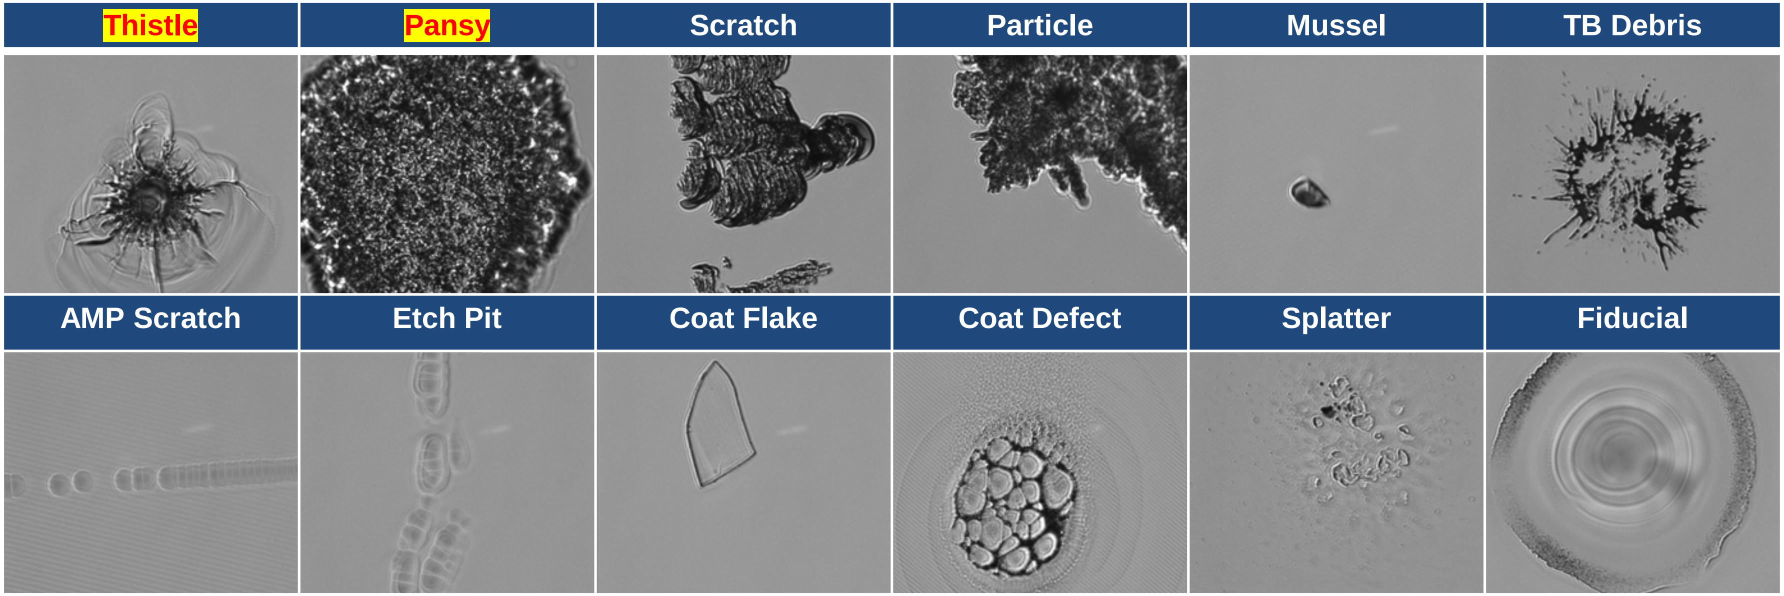

Human experts use two illuminations to classify sites. We input different modalities for our transfer learning by taking advantage of the fact that ImageNet consisted of color images.

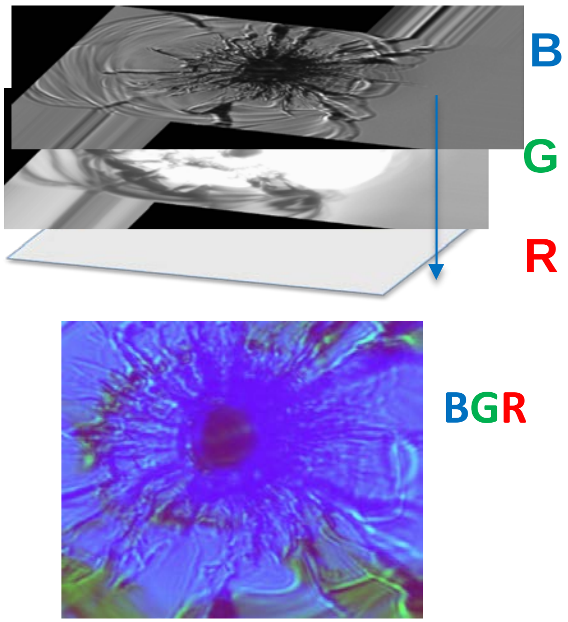

Backlight (B) and Coaxial (G) illumination images of damage sites were concatenated into a color image (BGR) to prepare them for the CNN

We improved the already-high accuracy results of two established deep learners by replacing the final decision-making layer of AlexNet with an ensemble of decision trees

[Connor Amorin (UMass Amherst), Laura Kegelmeyer, Philip Kegelmeyer (Sandia)]

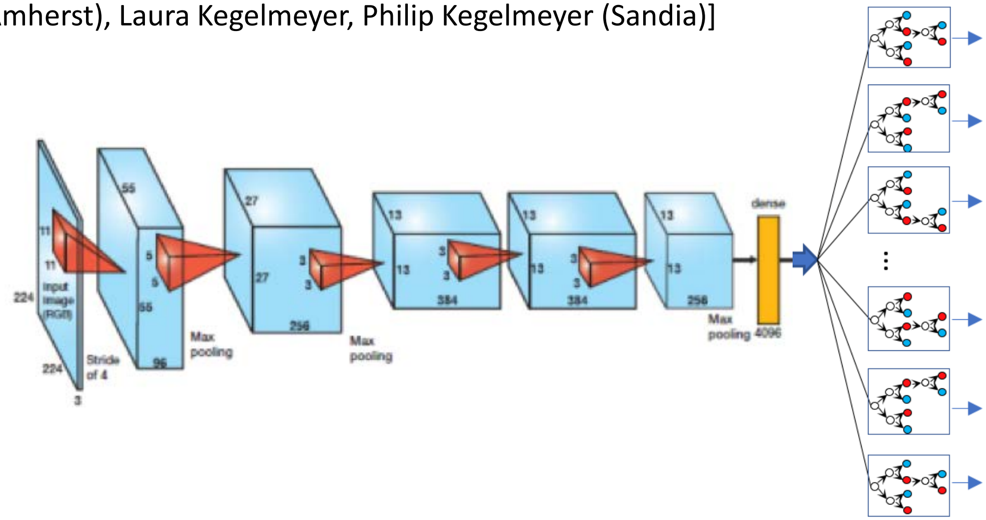

Model	Image Size	Test Accuracy *5-fold cross-validation	
		12-class Damage Dataset	2-class Remnant Dataset [1]
AlexNet	352x352	*97.40%	96.86%
Inceptionv2	352x352	*98.11%	*98.80%
AlexNet + Ensemble Decision Trees	352x352	*99.17%	*99.28%

Image analysis and machine learning for NIF Optics Inspection helps automate tedious processes and enables an efficient optics recycle loop for 1.8 MJ shots…. and beyond

- Since ~2007, machine learning, has been used to improve analysis accuracy, automation and quality control to inform and enable the NIF Optics Recycle Loop.

- Several new projects are in progress using Deep Learning and a Dual-Network solution to improve accuracy for inspection and process automation, and we have more projects in the pipeline.

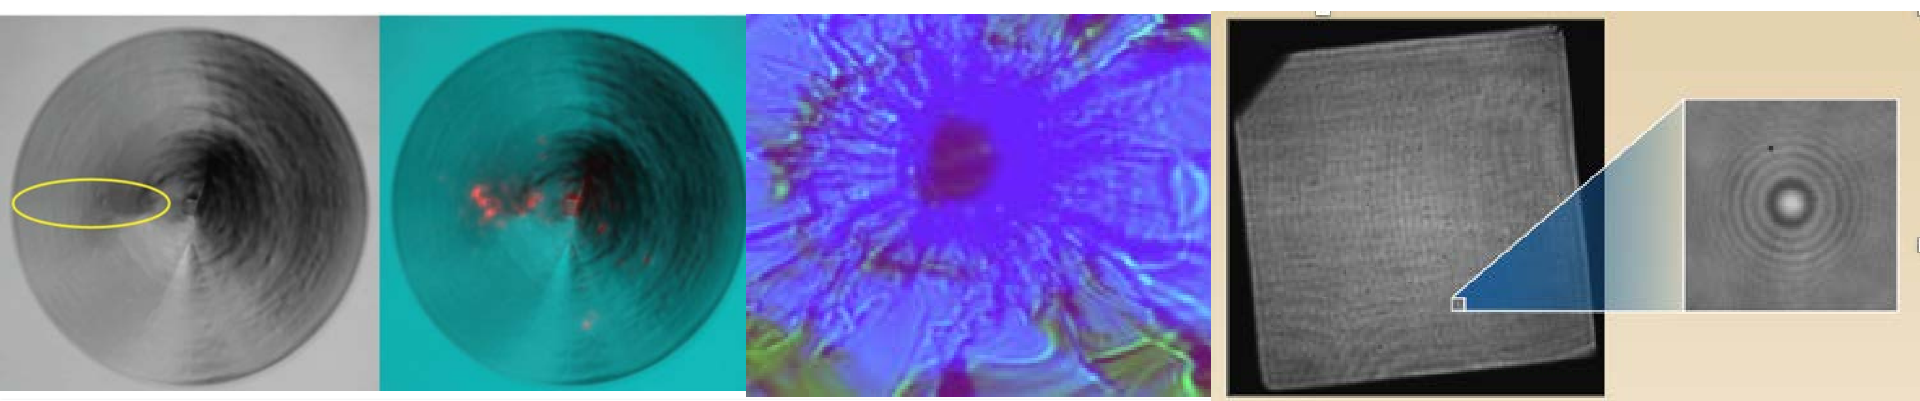

Lawrence Livermore
National Laboratory