Deep Learning for Explainable Image Classification of Chemical Laboratory Apparatus

Dr. Wesam Sakla[1], Capt. Tyler Hough[2]
[1]Lawrence Livermore National Laboratory (LLNL)
[2]Defense Threat Reduction Agency (DTRA)

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

Overview

- We want to classify images containing apparatus found in chemical laboratories.
- Of particular interest are apparatus associated with the production of aerosols.
- We use deep neural networks trained on GPUs.

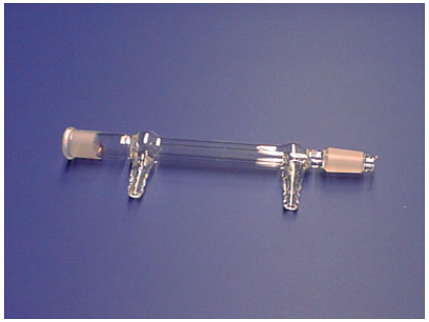

Data Acquisition

- Wrote Python script to scour *Google Images* for keywords pertaining to objects of interest
- Opens a web browser, forms the search URL, parses the resulting HTML
- Retrieves URLS corresponding to search results and automatically downloads images
- Downloaded ~ 9100 images across 15 object categories
- Had someone manually inspect the images, quickly removing those which are completely irrelevant
- Curated dataset: 5789 images

CNN Image Classifier

- Partitioned 5789 curated images into training/testing sets:
  - 4746 for training (82%)
  - 1043 for testing (18%)

- Trained four CNN architectures:
  - ALEXNet
  - GoogLeNet
  - ResNet-50
  - VGG-16

- All models are fine-tuned from models pre-trained on ImageNet

ImageNet
Training Data

Train with 1.2 million
labeled images

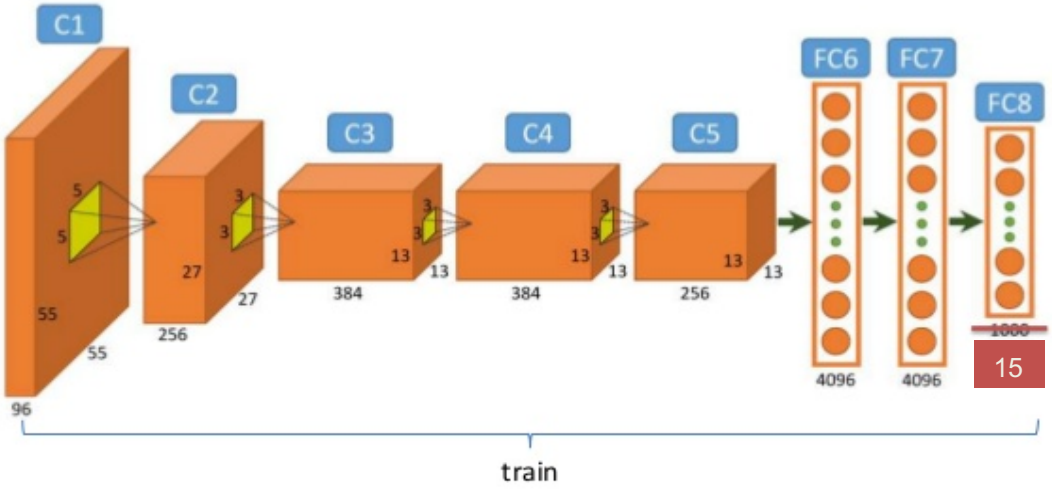

CNN Image Classifier

Training Images

CNN Model Tuning

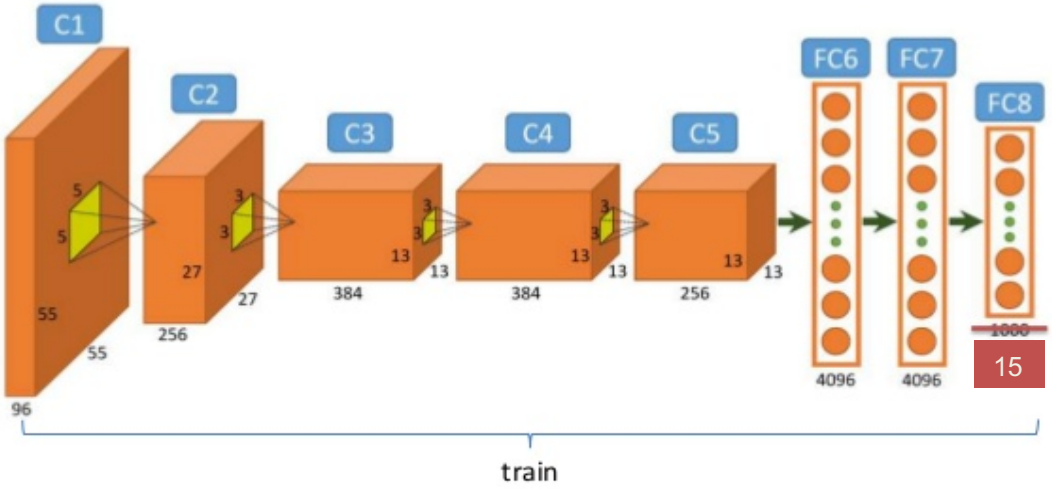

CNN Architecture	Accuracy
$ALEXNet$	84%
$GoogLeNet$	90%
$ResNet-50$	90%
$VGG-16$	90%

VGG-16 Confusion Matrix

Object	Class	# Test Images	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Class Acc. (%)
lypholizer/dryer	1	86	**74**	1	0	4	0	4	0	0	0	1	0	0	2	0	0	**86.05**
nebulizer kit	2	77	0	**73**	0	0	1	0	1	0	0	1	0	0	1	0	0	**94.81**
ultracentrifuge	3	66	1	0	**57**	0	0	3	0	0	0	1	0	0	0	2	2	**86.36**
milling machine	4	116	0	0	0	**111**	0	0	0	0	0	0	0	0	1	2	2	**95.69**
condenser	5	51	0	1	0	0	**46**	1	0	0	2	0	0	0	0	1	0	**90.20**
rotary evaporator	6	65	2	0	1	1	5	**54**	0	0	0	0	1	0	0	1	0	**83.08**
biosafety suit	7	63	0	0	0	0	0	0	**58**	0	0	0	2	0	2	1	0	**92.06**
3-neck flask	8	50	0	1	0	0	2	0	0	**45**	2	0	0	0	0	0	0	**90.00**
add/sep funnel	9	42	0	2	0	0	3	0	2	1	**34**	0	0	0	0	0	0	**80.95**
magnetic stir plate	10	86	0	1	0	1	0	0	0	0	1	**80**	0	0	1	1	1	**93.02**
respirator	11	76	1	0	0	1	0	0	3	0	0	0	**71**	0	0	0	0	**93.42**
graduated cylinder	12	75	0	0	0	1	1	0	0	1	0	0	0	**71**	0	1	0	**94.67**
biosafety cabinet	13	76	1	0	0	1	0	0	1	0	0	0	0	0	**71**	1	1	**93.42**
bioreactor/fermenter	14	77	5	1	0	0	0	1	0	0	0	1	0	2	2	**64**	1	**83.12**
DNA/RNA synthesizer	15	37	0	0	2	3	0	1	0	0	0	0	0	0	0	3	**28**	**75.68**
TOTAL		1043																**90**

Interpretability of CNNs

- CNNs excel at object image classification, detection, segmentation, etc.
- Large parameter space makes intuitive interpretation difficult.

Meaningful integration of AI requires transparent models that can explain why they predict what they predict

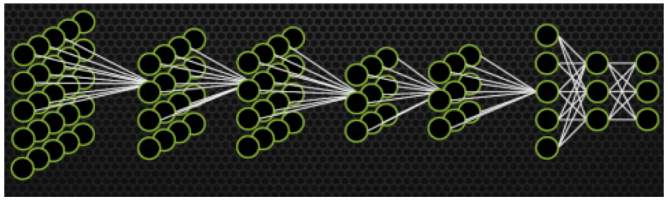

"ultracentrifuge": 0.92

Gradient-based Class Activation Mapping

- Visual explanation technique that does not require modification/re-training of network
- Can be applied to any CNN-based task (i.e., image captioning, visual question answering)
- The last convolutional layers in a CNN learn *semantic, class-specific* information (i.e., object parts)

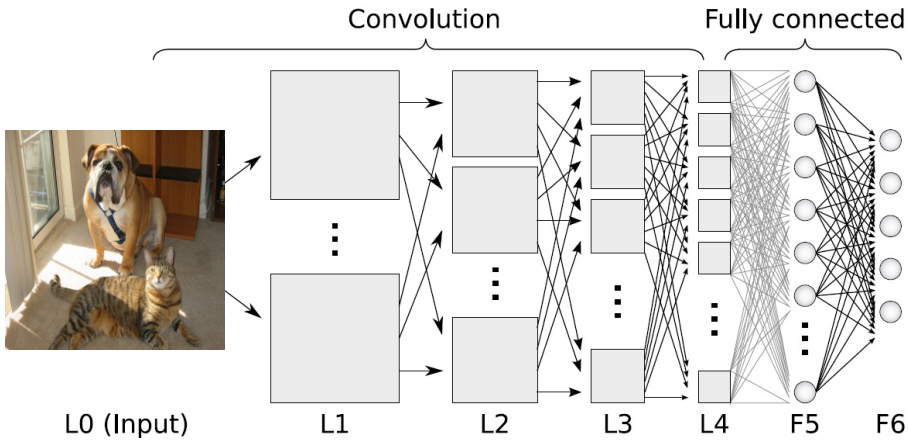

Use the gradient flowing into last convolutional layer of CNN to understand importance of each neuron for a prediction (classification)

Selvaraju *et al.*, Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization, arXiv:1610.02391, 2016.

Guided Grad-CAM

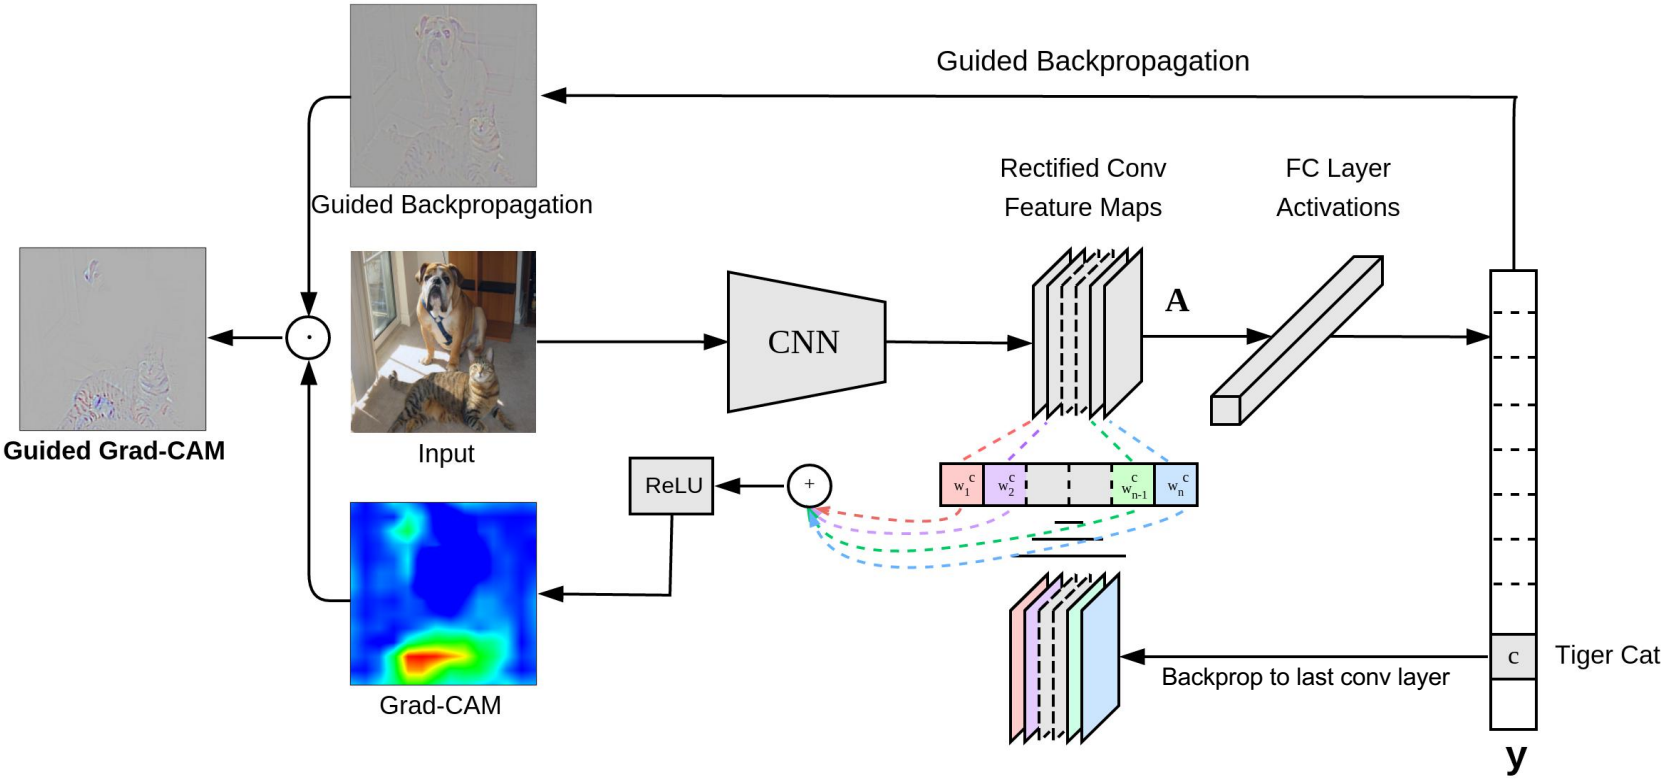

$$L^c_{\text{Grad-CAM}} = ReLU\underbrace{\left(\sum_k \alpha_k^c A^k\right)}_{\text{linear combination}}$$

$$\alpha_k^c = \overbrace{\frac{1}{Z}\sum_i\sum_j}^{\text{global average pooling}} \underbrace{\frac{\partial y^c}{\partial A_{ij}^k}}_{\text{gradients via backprop}}$$

Selvaraju *et al.*, Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization, arXiv:1610.02391, 2016.

Visualizations

GB | Guided Grad-Cam | Grad-Cam

ALEXNET
✓ *grad. cylinder: 100%*

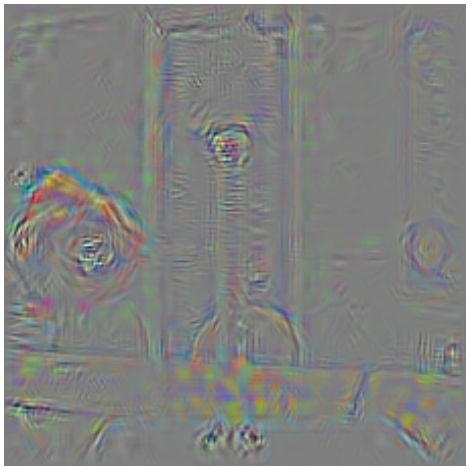

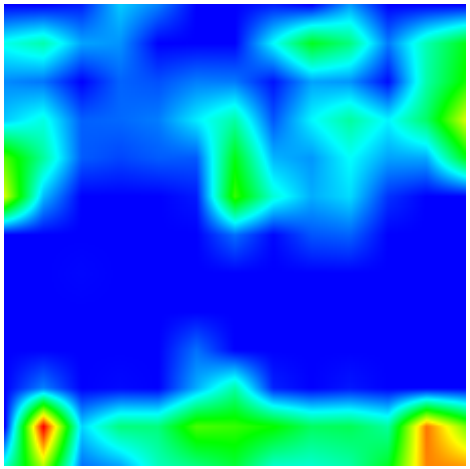

VGG-16
✓ *grad. cylinder: 100%*

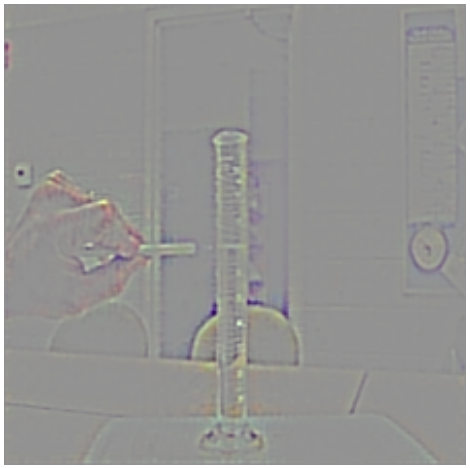

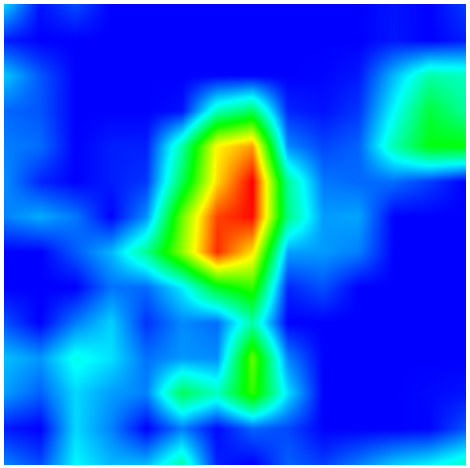

Visualizations

GB | Guided Grad-Cam | Grad-Cam

ALEXNET
✓ *synthesizer: 100%*

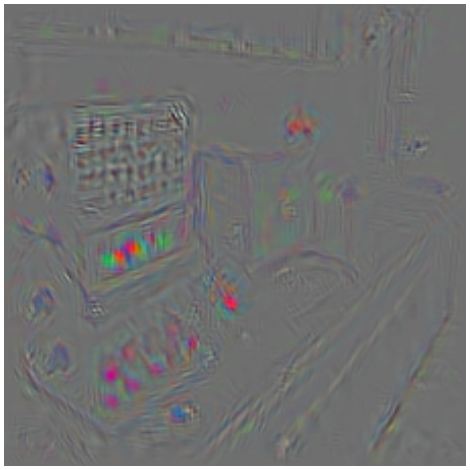

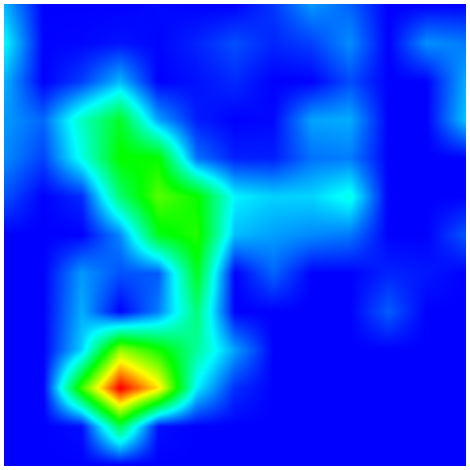

VGG-16
✓ *synthesizer: 100%*

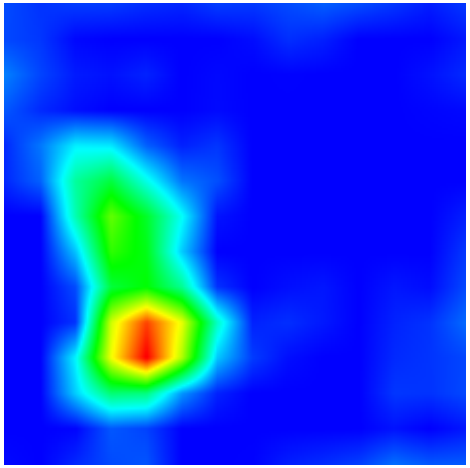

Visualizations

GB | Guided Grad-Cam | Grad-Cam

ALEXNET
✓ *ultracentrifuge: 100%*

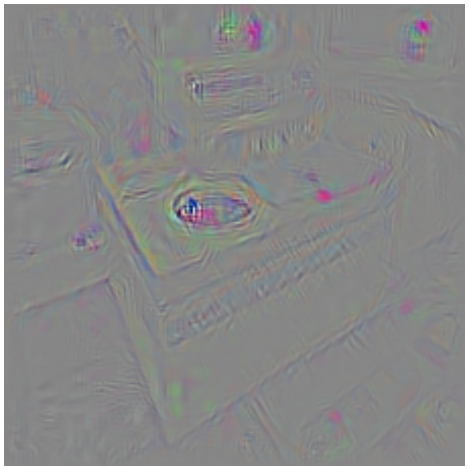

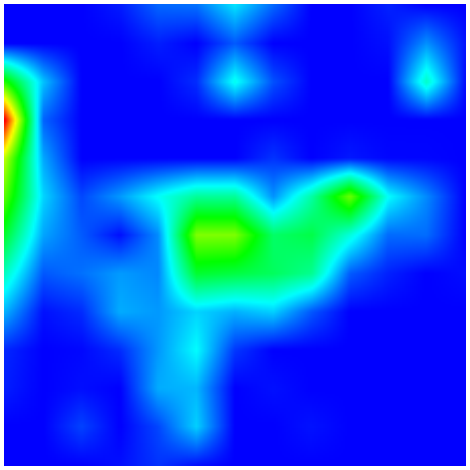

VGG-16
✓ *ultracentrifuge: 100%*

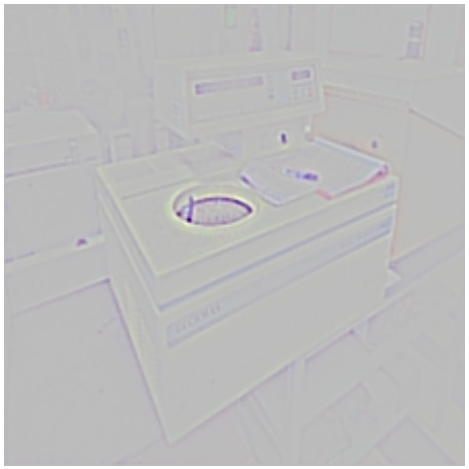

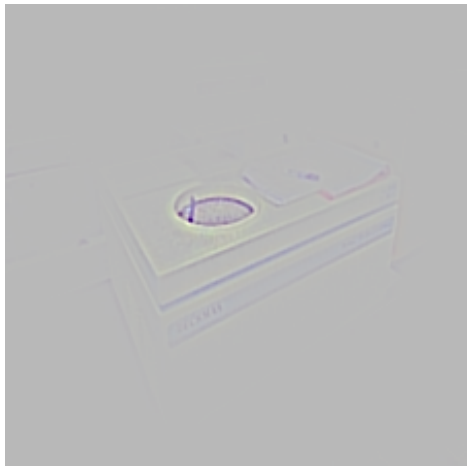

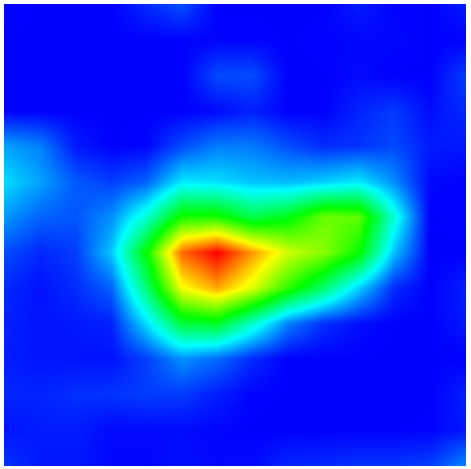

Visualizations

GB

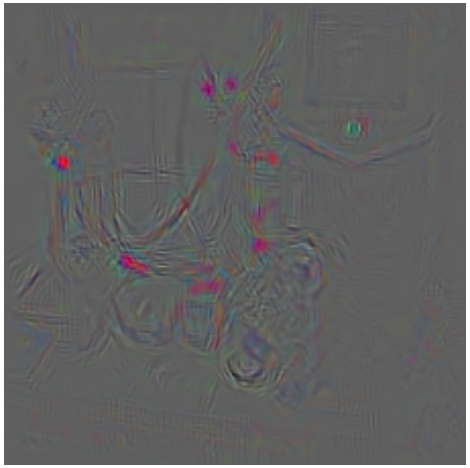

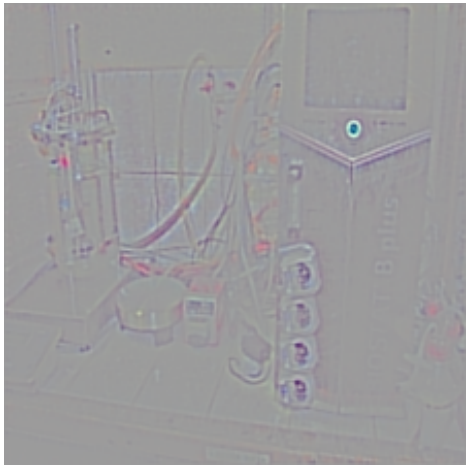

Guided Grad-Cam

Prediction

ALEXNET
× *Rotary evaporator*
Score: 79.95%

VGG-16
✓ *Bioreactor/fermenter*
Score: 94.30%

Conclusion

- Trained CNN classifiers of chemical apparatus using open-source, web-scraped imagery,
- Obtained 90% classification rates on held-out test set containing ~1000 images.
- Used recent attention mechanism techniques for visualizing and interpreting model predictions.
- Future work to focus on more expansive, “wild” image sets and training detection models.

Lawrence Livermore
National Laboratory