Machine Learning at LCLS
Locating Detector Features

David Schneider, davidsch@slac.stanford.edu

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

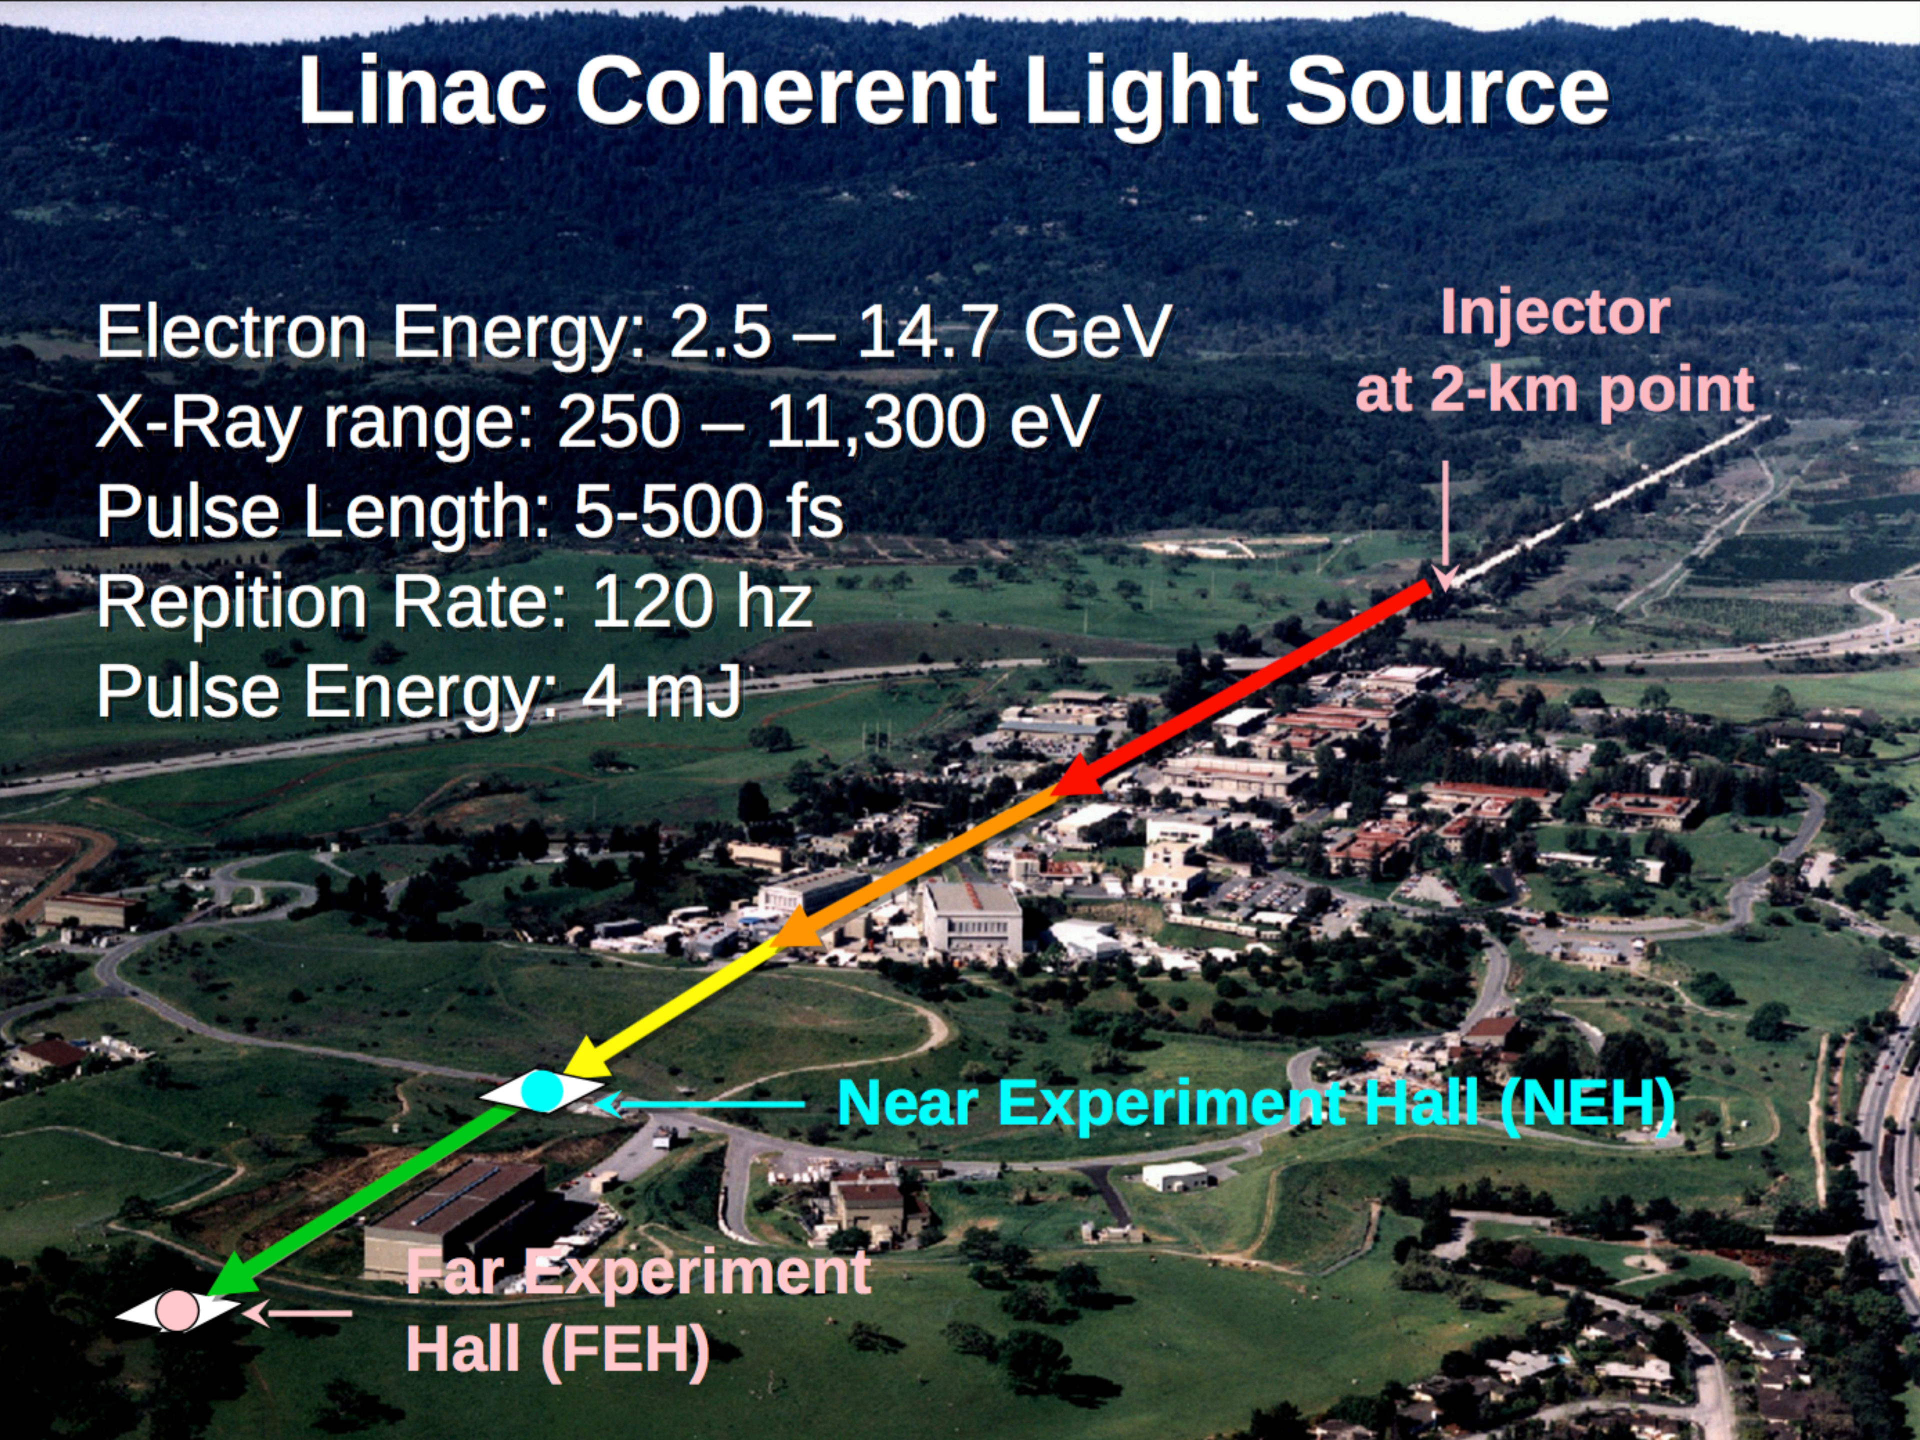
Linac Coherent Light Source
Electron Energy: 2.5 – 14.7 GeV
X-Ray range: 250 – 11,300 eV
Pulse Length: 5-500 fs
Repition Rate: 120 hz
Pulse Energy: 4 mJ
Injector
at 2-km point
Near Experiment Hall (NEH)
Far Experiment
Hall (FEH)

LCLS Experimental Floor

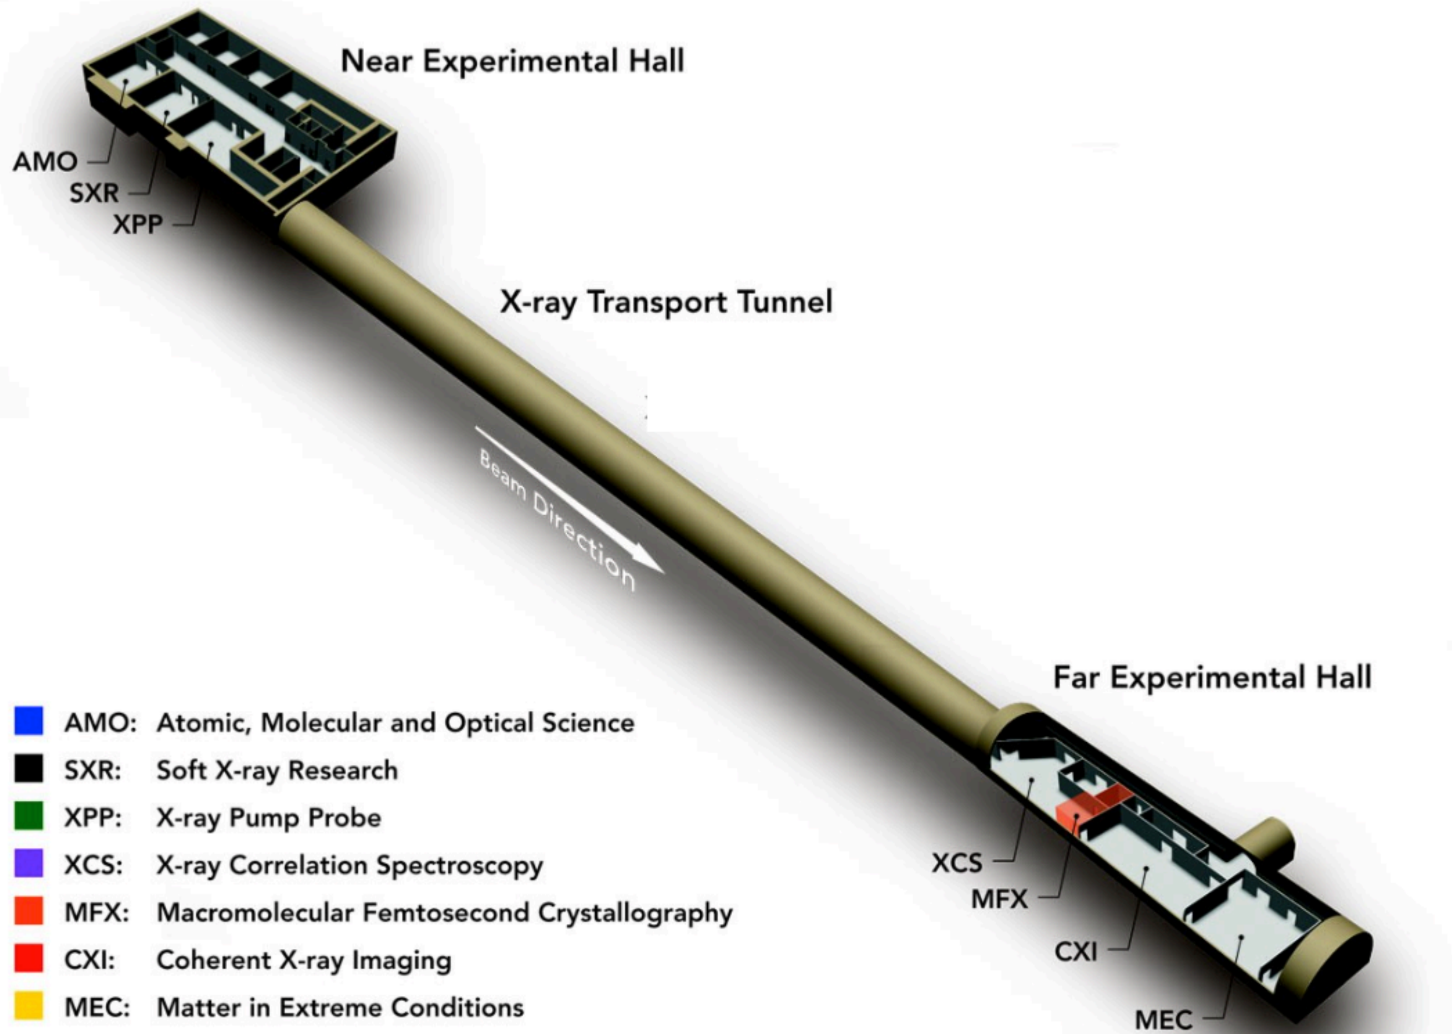

- AMO: Atomic, Molecular and Optical Science
- SXR: Soft X-ray Research
- XPP: X-ray Pump Probe
- XCS: X-ray Correlation Spectroscopy
- MFX: Macromolecular Femtosecond Crystallography
- CXI: Coherent X-ray Imaging
- MEC: Matter in Extreme Conditions

7 hutches = wide variety of data/experiments

Detector Images

Sample/Shot Interaction

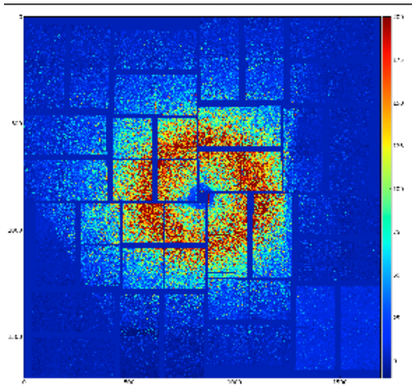

Find Bragg's peaks in a Diffraction shot at CXI

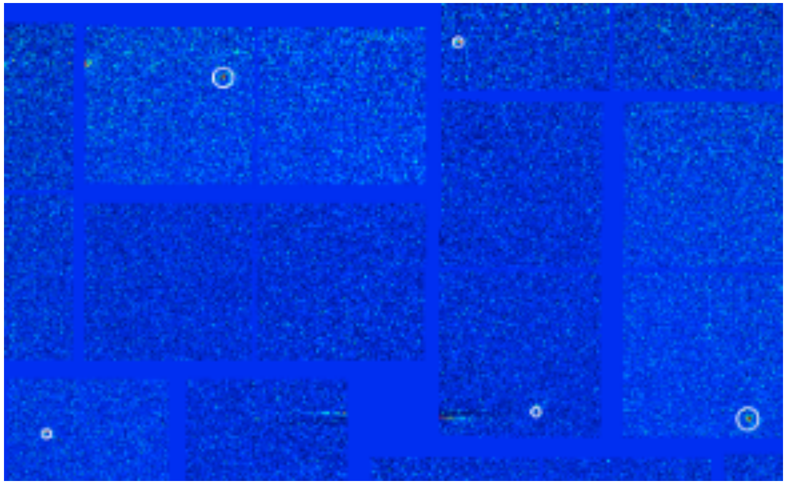

Shot Characterization

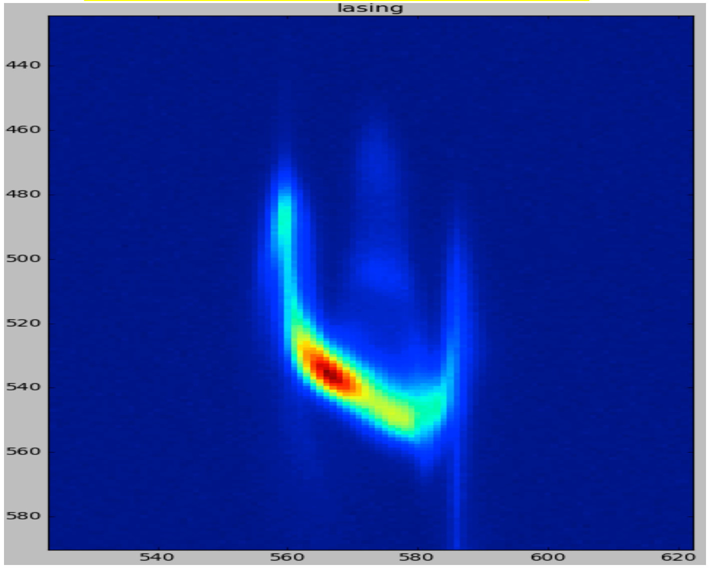

Find Time resolved power during a shot at AMO

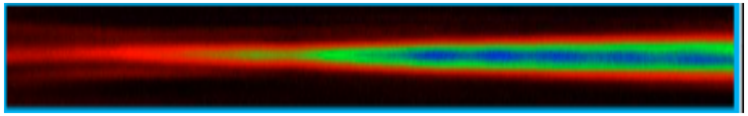

Find delay between pump and probe during shot at XPP

Detector Features: Classification and Regression

Shot characterization requires detailed ***Image Processing***

LCLS provides many detector features ***(peak finders, GUI for tuning peaks, Time Tool)***

Problem: Experiments change.

New image processing is time consuming to develop

Under study: determining when ***Machine learning*** is more effective at shot characterization

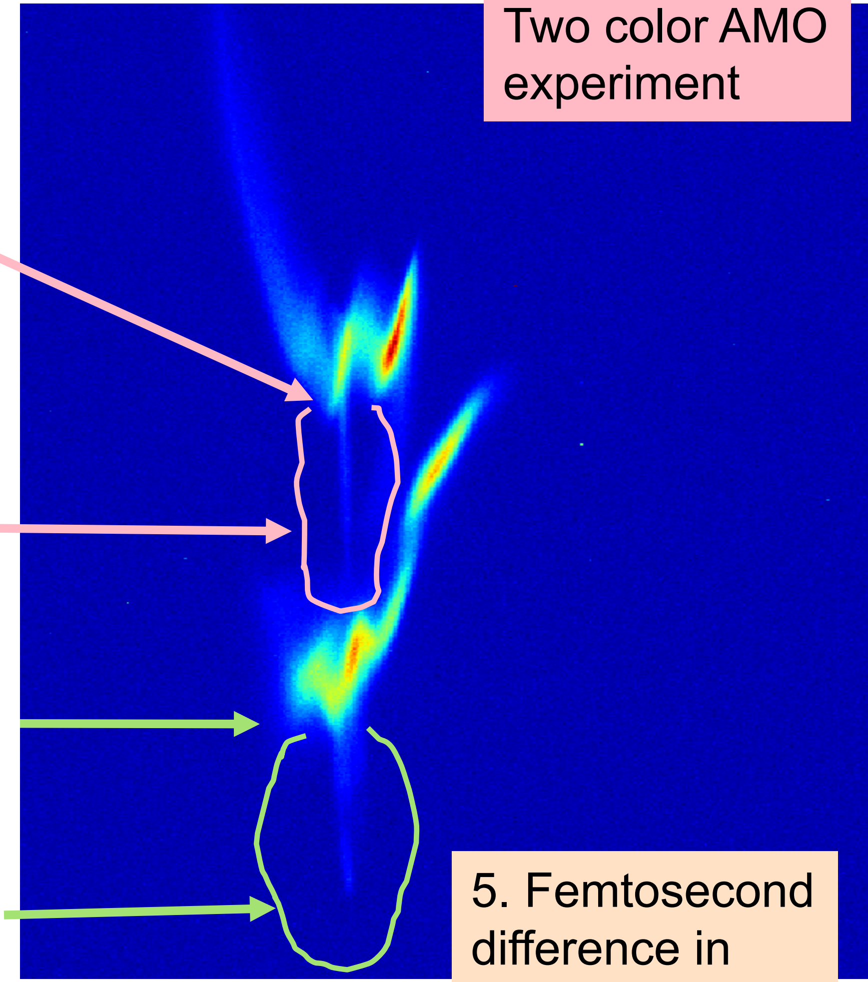

Challenges with Machine Learning

SLAC

- Labels for Training data
  - LCLS data - no simulator
    - 100k shots for a typical experiment
  - In contrast
    - HEP - standard model
    - rich set of simulated data for supervised learning
- Processing LCLS Data
  - Large datasets - megapixel images
  - Data format – like binary XML
    - expensive to load event
      - detector corrections
      - Parse all the machine data
    - For ML, presently, make a copy to hdf5

Two Color AMO Experiment (R. Coffee et all)

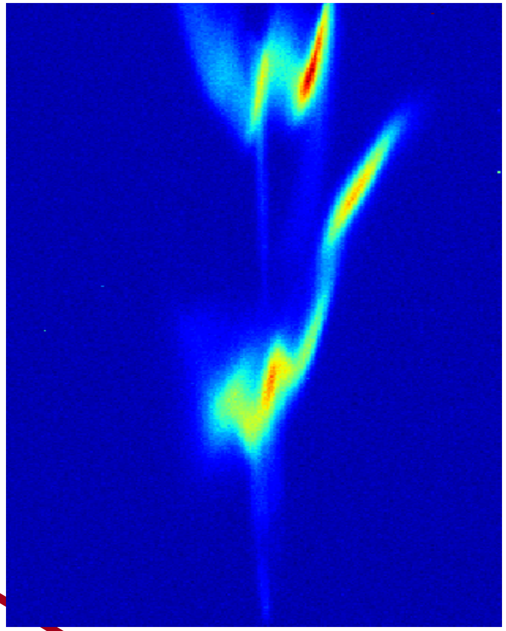

Replaces Image processing with 1D Peak Finding

Special detector:
- high/low energy lasing
- early/late peaks on spectrometer
- energy level – peak position

Issue - same detector needed for
- sample/shot interaction

Experiment:
- Training runs
  - Detector used for shot characterization
  - classification: 4 labels from peaks
    (neither, low, high, both)
  - Regression: 2 energy values

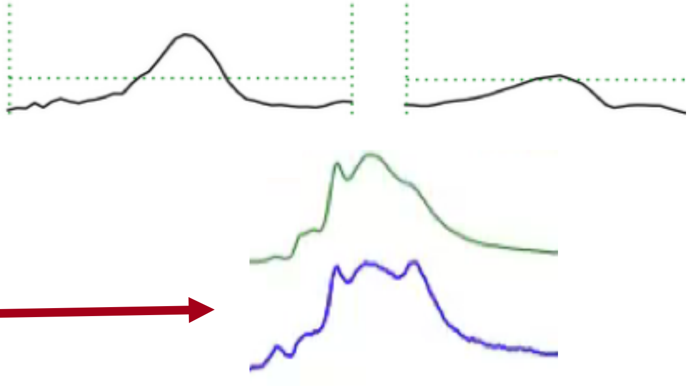

- Molecular runs
  - Predict shot (0,1,2,3, low/high energy)
  - sort down to see structure

M. Mongia, R. Coffee, C.O' Grady, D. Schneider

Input:
Image + machine data stable between training and prediction

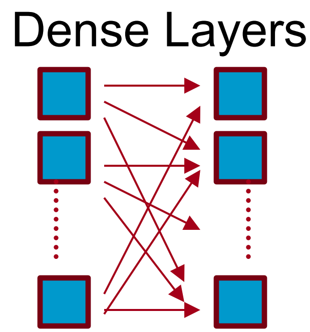

Concatenate machine data with flattened CNN output before final dense layers

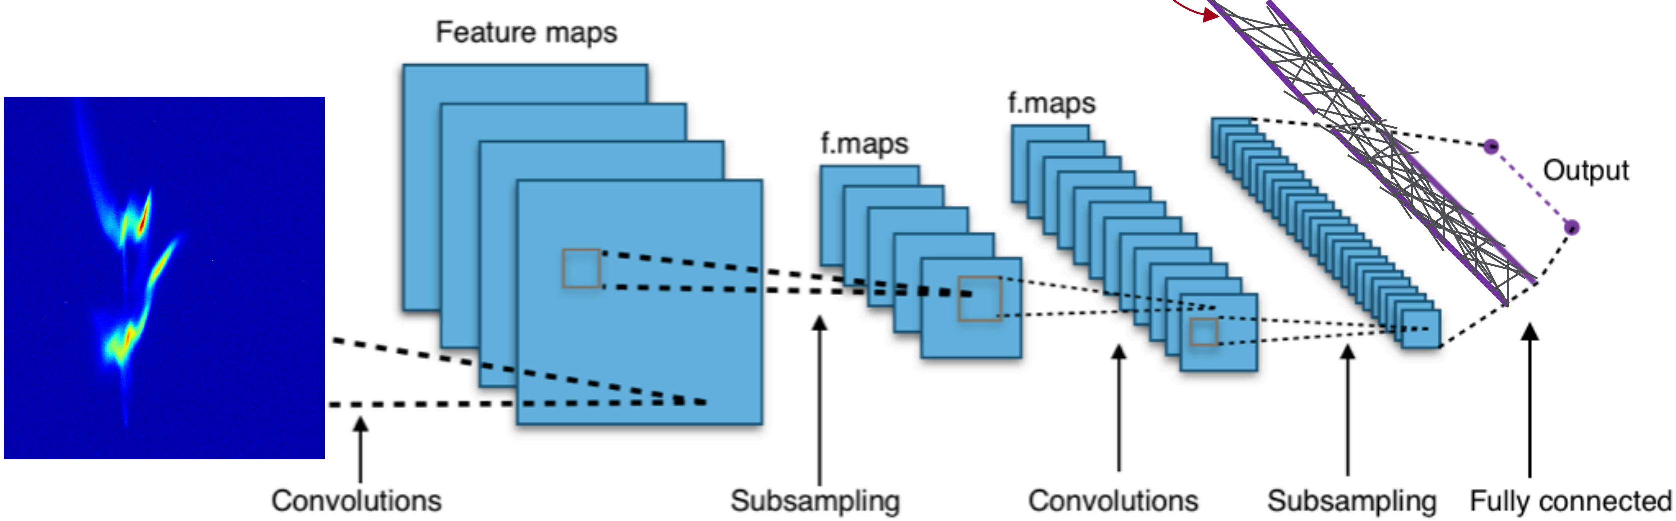

Guided Back Propagation for Feature Finding

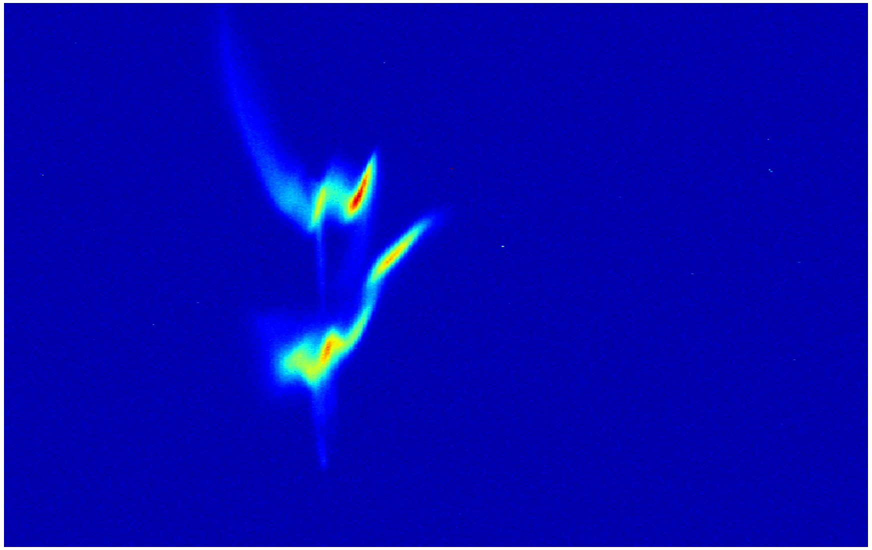

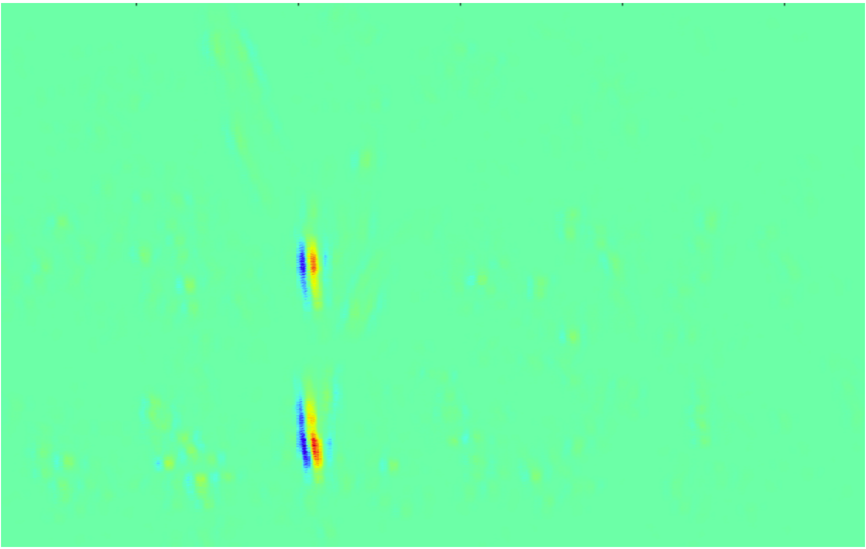

Question: why did the Neural Network make it's choice? Will scientists trust these models?

Guided back propagation – intuition is you ask the network how to ***change*** the ***image*** to improve the prediction

Below – technique highlights exactly what is important

Locate Lasing from G. Backprop? Still need femtosecond delay between lasing, not on spectrometer. Can we robustly get it from guided backprop?

Machine Learning from Reference vs. Signal?

Special detector gave us 4 labels for machine learning

May not generalize to other experiments

However – many experiments record:

reference runs to compare to ***signal runs***

General XTCAV Analysis:

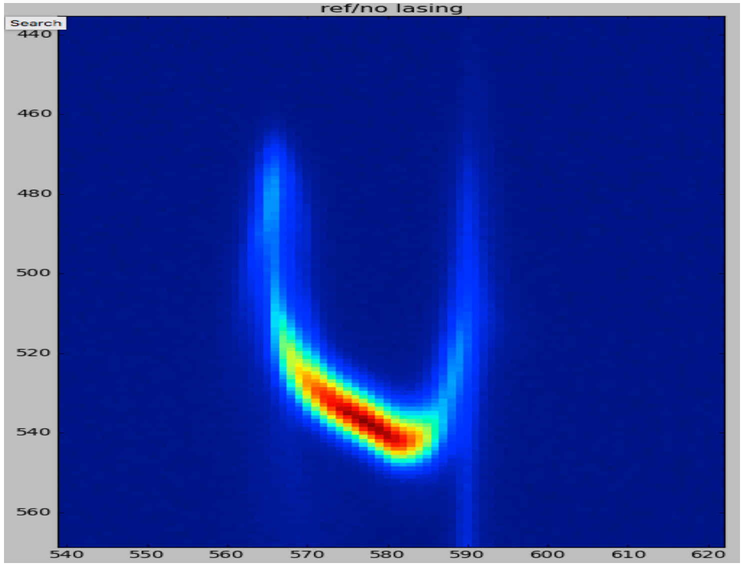

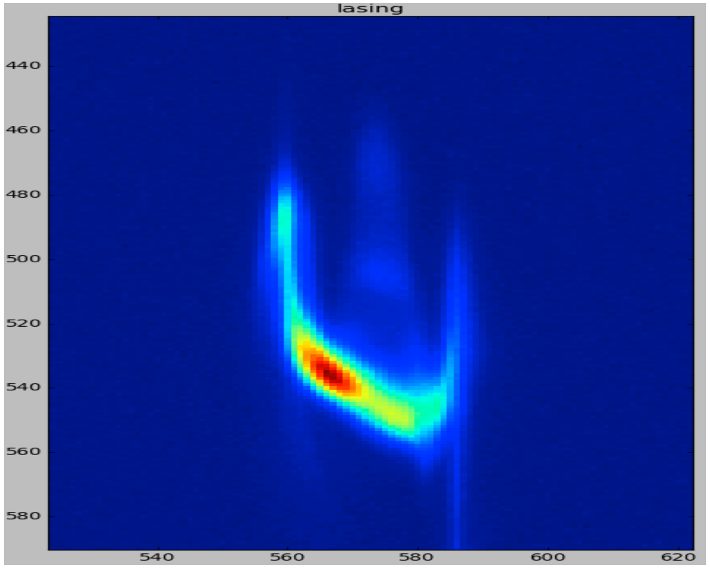

Same pattern for Time Tool, Calibrating Detectors, Crystallography

Reference vs. Signal, Saliency style Algorithms

0/1 Classifier for Two color AMO experiment

1 (high), 2 (low), 3 (both) -> 1 (signal)

but old 1 (high) weaker signal (dataset not balanced)

Guided Backprop: no robust result with it yet

More recent, similar:

Relevance Propagation

(A. Binder, et all) 2016

Formulas tricky to implement for CNN's

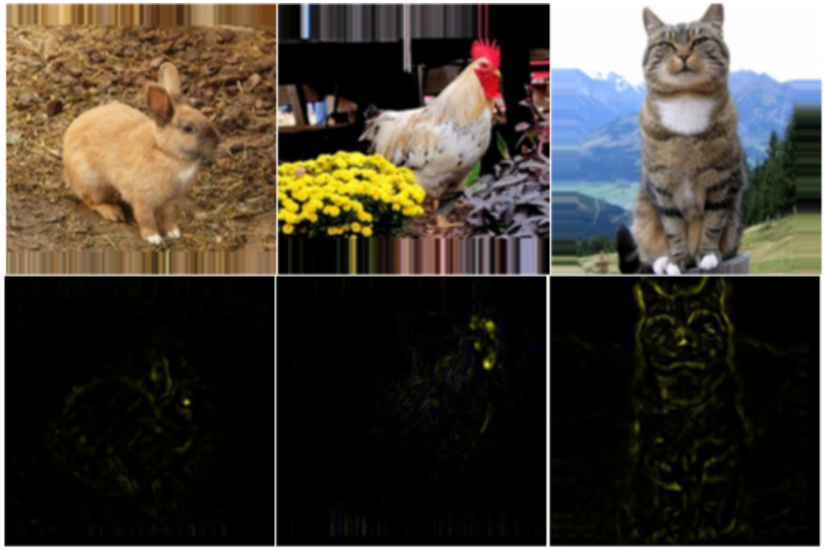

Others: LIME, Grad-CAM

Can we generate the right reference from a signal shot?

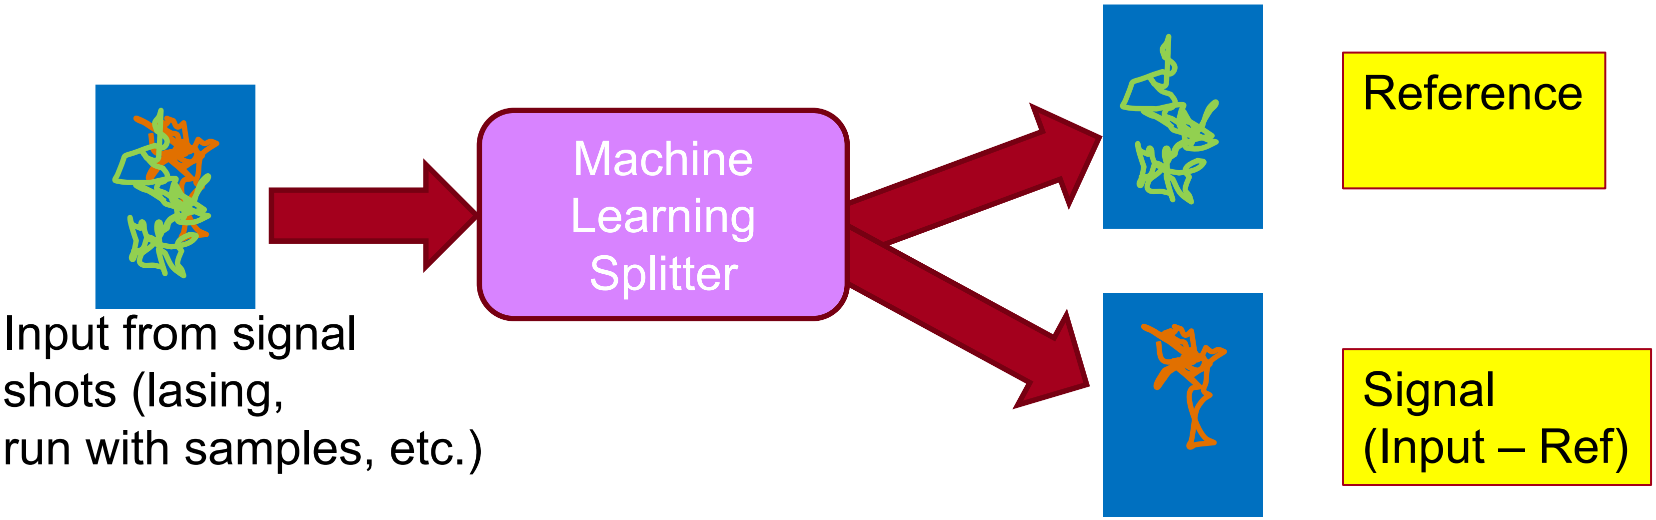

May tie into machine learning area of *Generative Models*
recent work such as *Adversarial Networks* or *Variational Auto Encoders*

Generative Models

Training data: faces

Z
Low dimensional representation

Sample from Z to generate New faces

Radford et all, ICLR 2016
Images via CS231n slides

Generative Models: Vector Math

Training data: faces

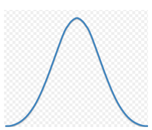

samples

Radford et all,
ICLR 2016
Images via
CS231n slides

Averages in Z space
math in Z space

Generate Faces
From these Z's

Lasing Vector Math?

L=
Lasing average
In latent Z space

N=
No Lasing average
In latent Z space

V = N – L   could this remove lasing?

X =
Particular
Lasing shot
in latent Z
space

X+ V = magically,
the right no lasing
reference?

Numerous issues,
But idea is what applications
of the representation in Z
space are there?

Variational Auto Encoder (VAE)

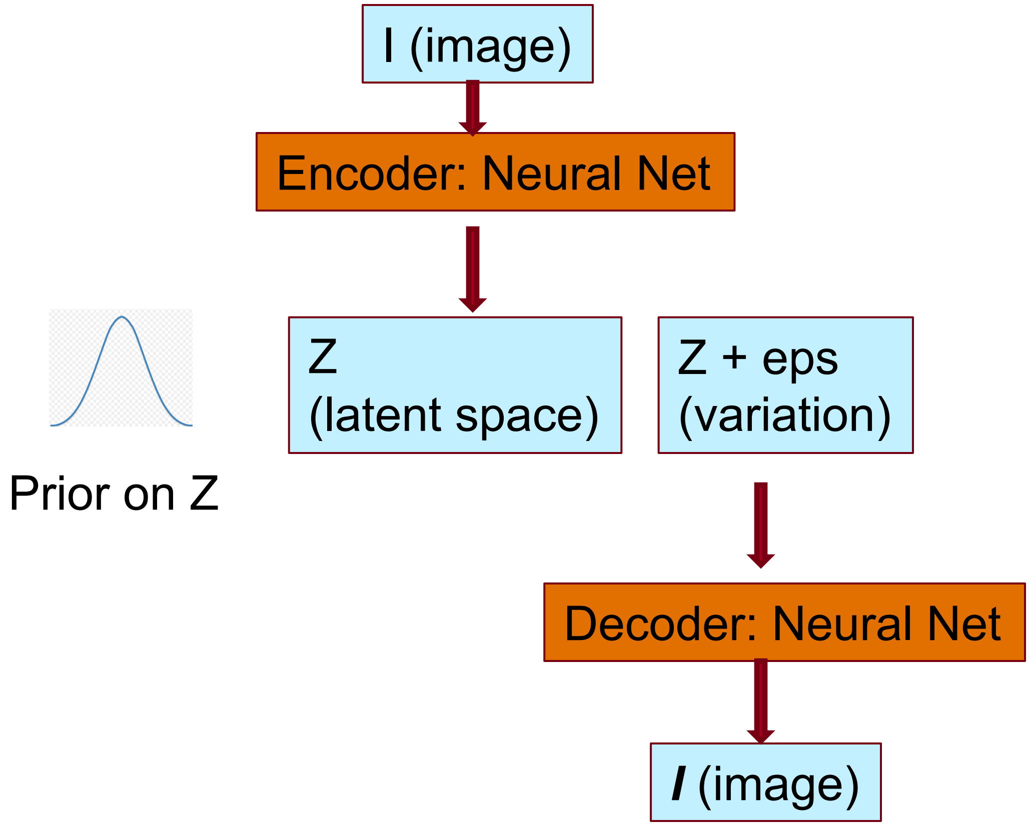

Loss: two terms: image recovered: (I – ***I***)^2 -> 0

latent space looks like prior: (KL divergence term)

VAE on MNIST (from Keras Blog)

I: 784 flattened pixels, values in [0,1]

Encoder: One Dense Layer
Relu Activation

Z: 2 dim

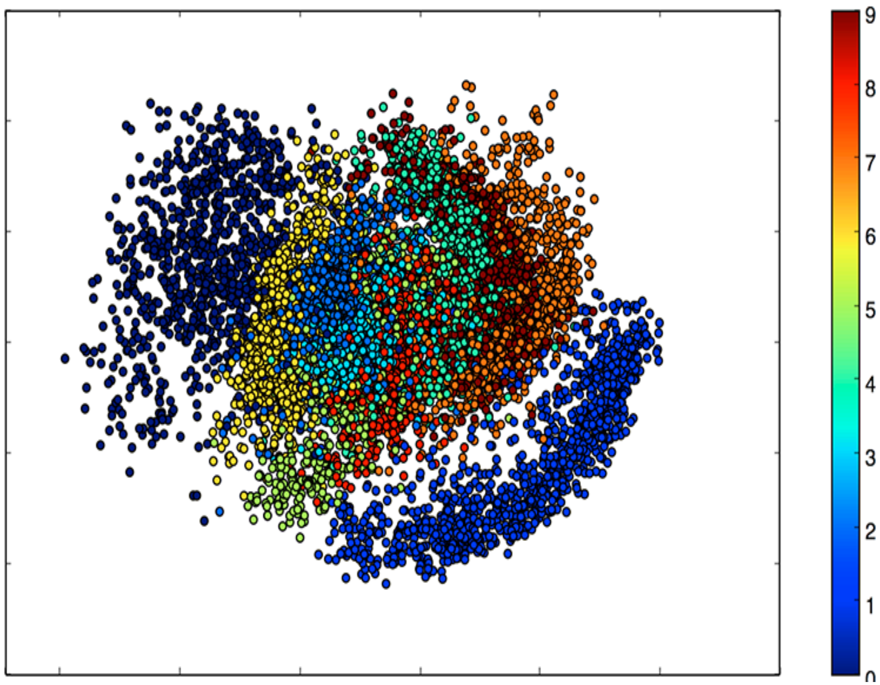

Z - interpolation

Decoder: Dense 2 –> 256, Relu
Dense 256 -> 784, Sigmoid

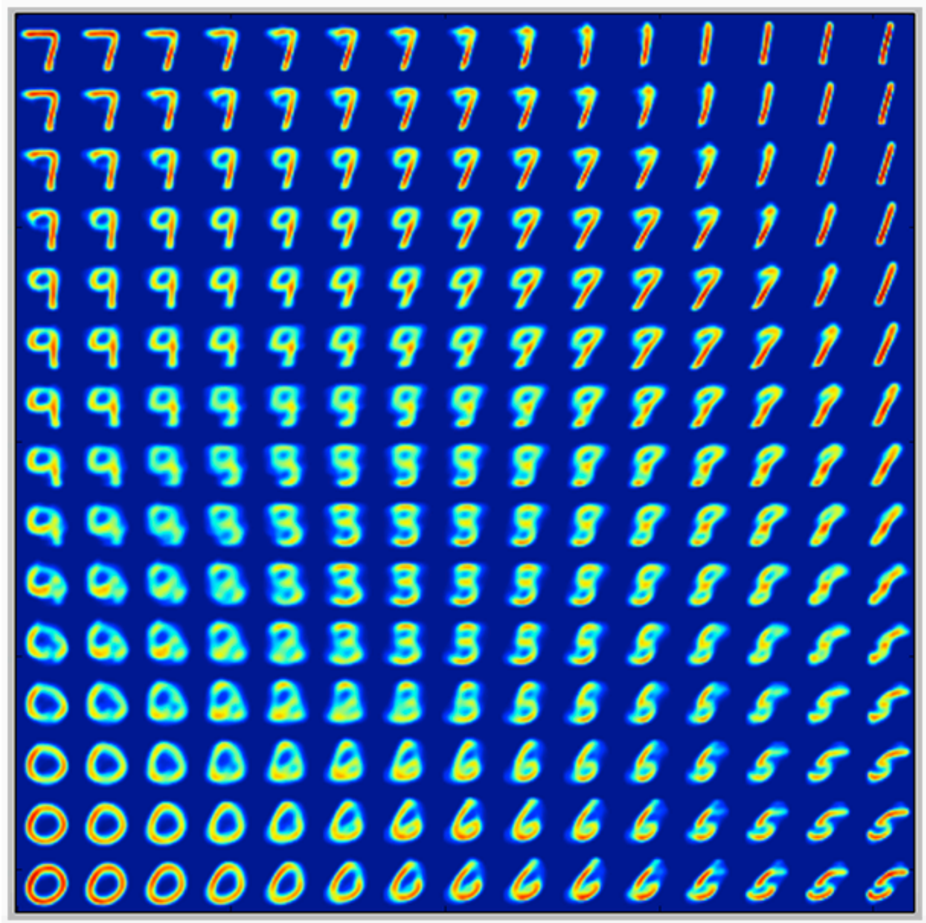

VAE on Two Color AMO Data – Latent Space

Reduce Images: 5000 pixels
728 x 568 -> 100 x 50
Log/Thresh at 300 ADU
center on horiz ROI

Encoder: two dense layers
5000 -> 256 –> 2
Decoder: two dense layers
2 -> 256 -> 5000

Latent Space
Z: 2dim – 6000 lasing (red)
1000 no lasing (blue)

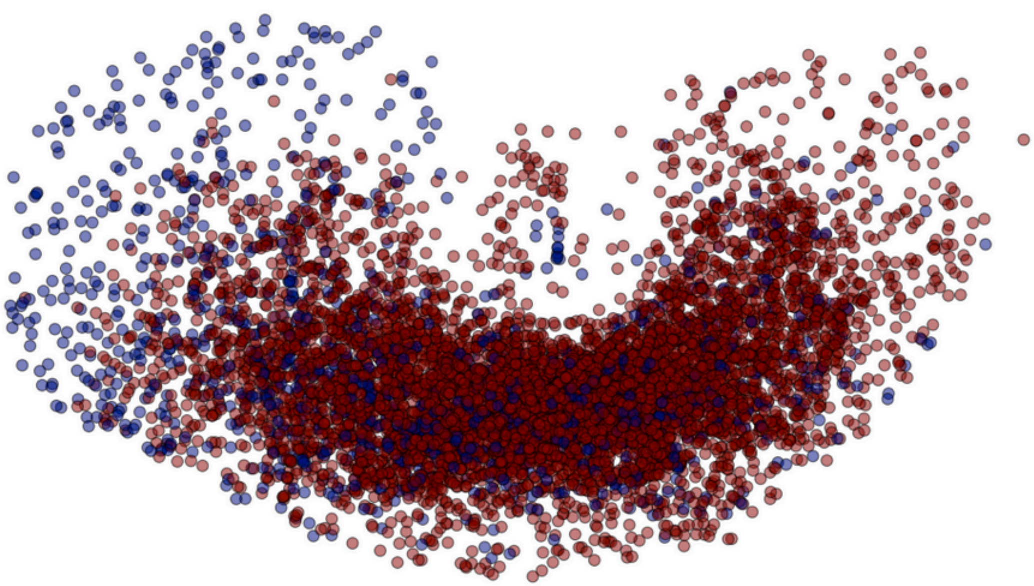

VAE – Lasing Reconstruction

Orig

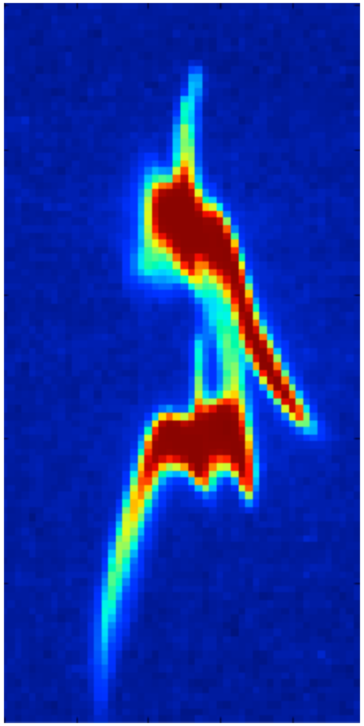

Reconstruction

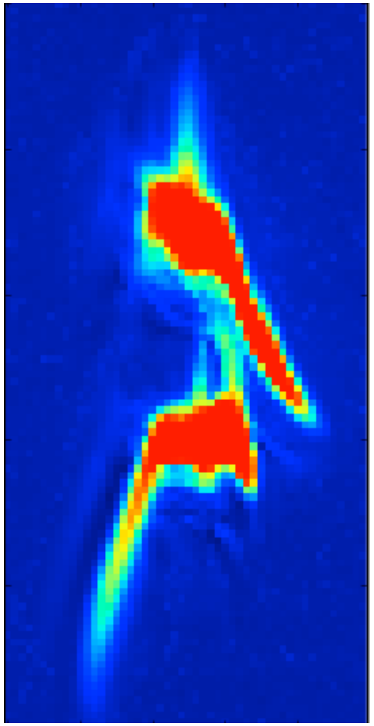

Orig - Reconstruction

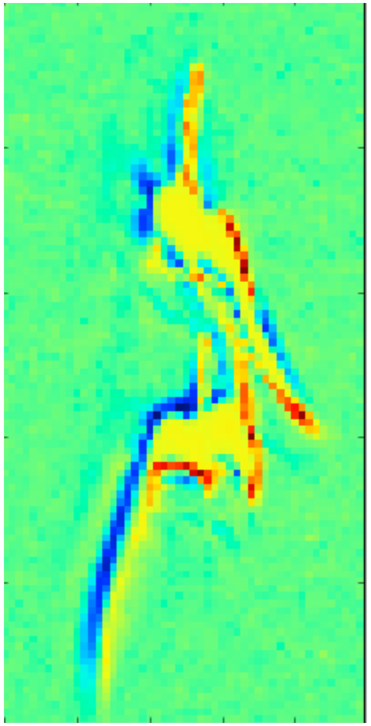

VAE – Bad No Lasing Reconstruction

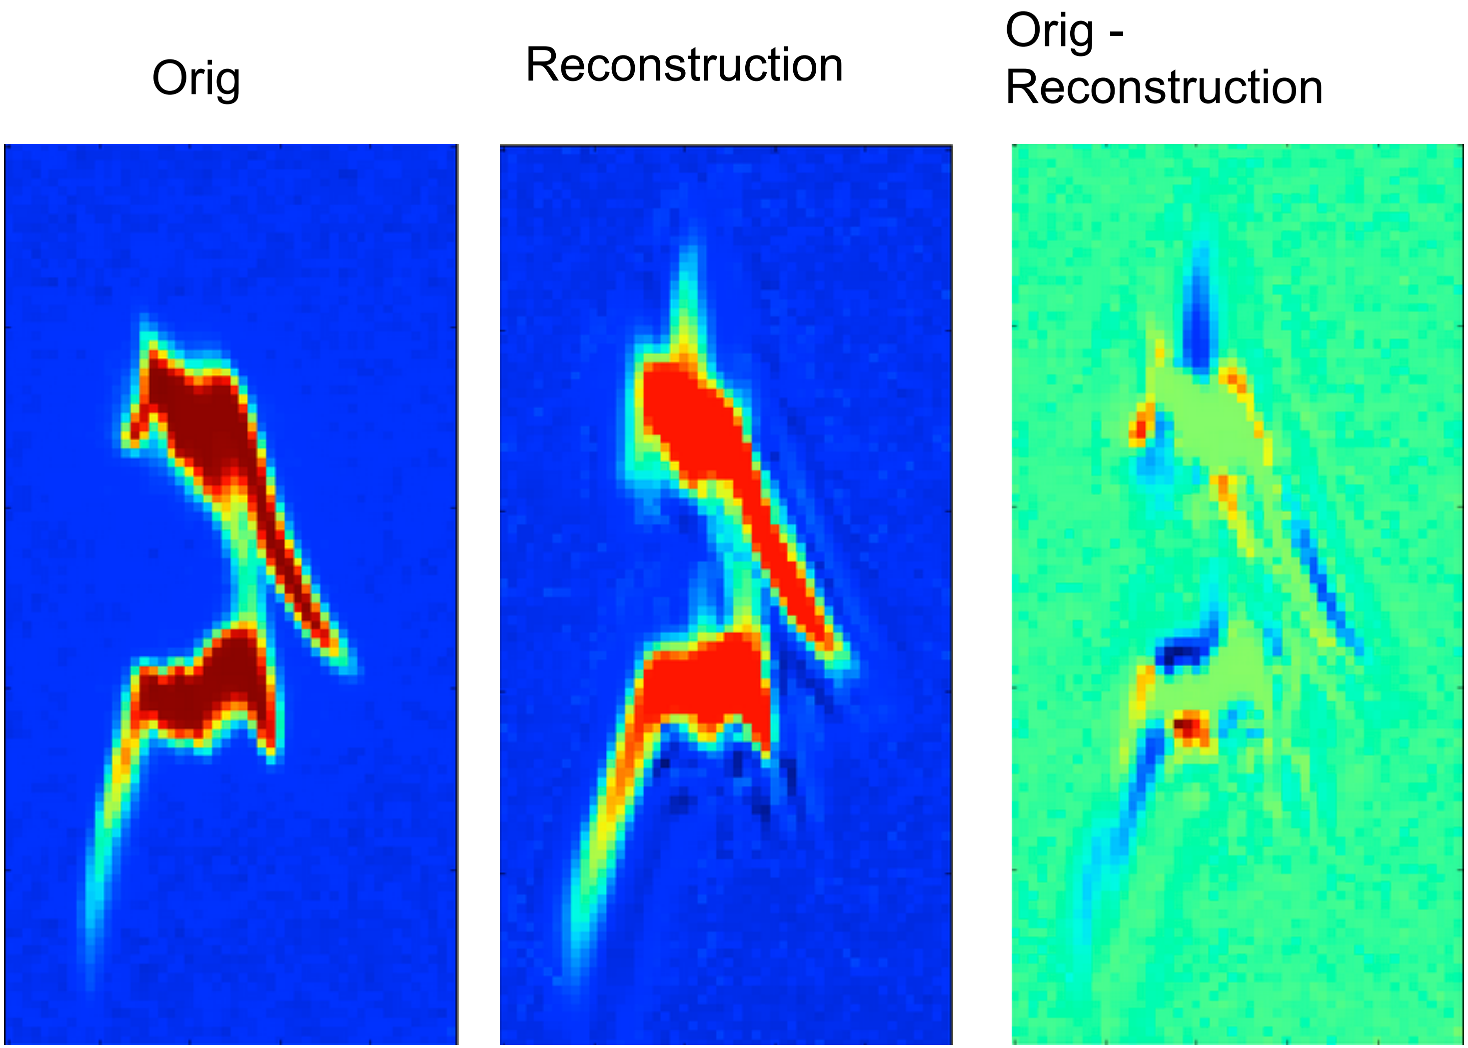

VAE Interpolation

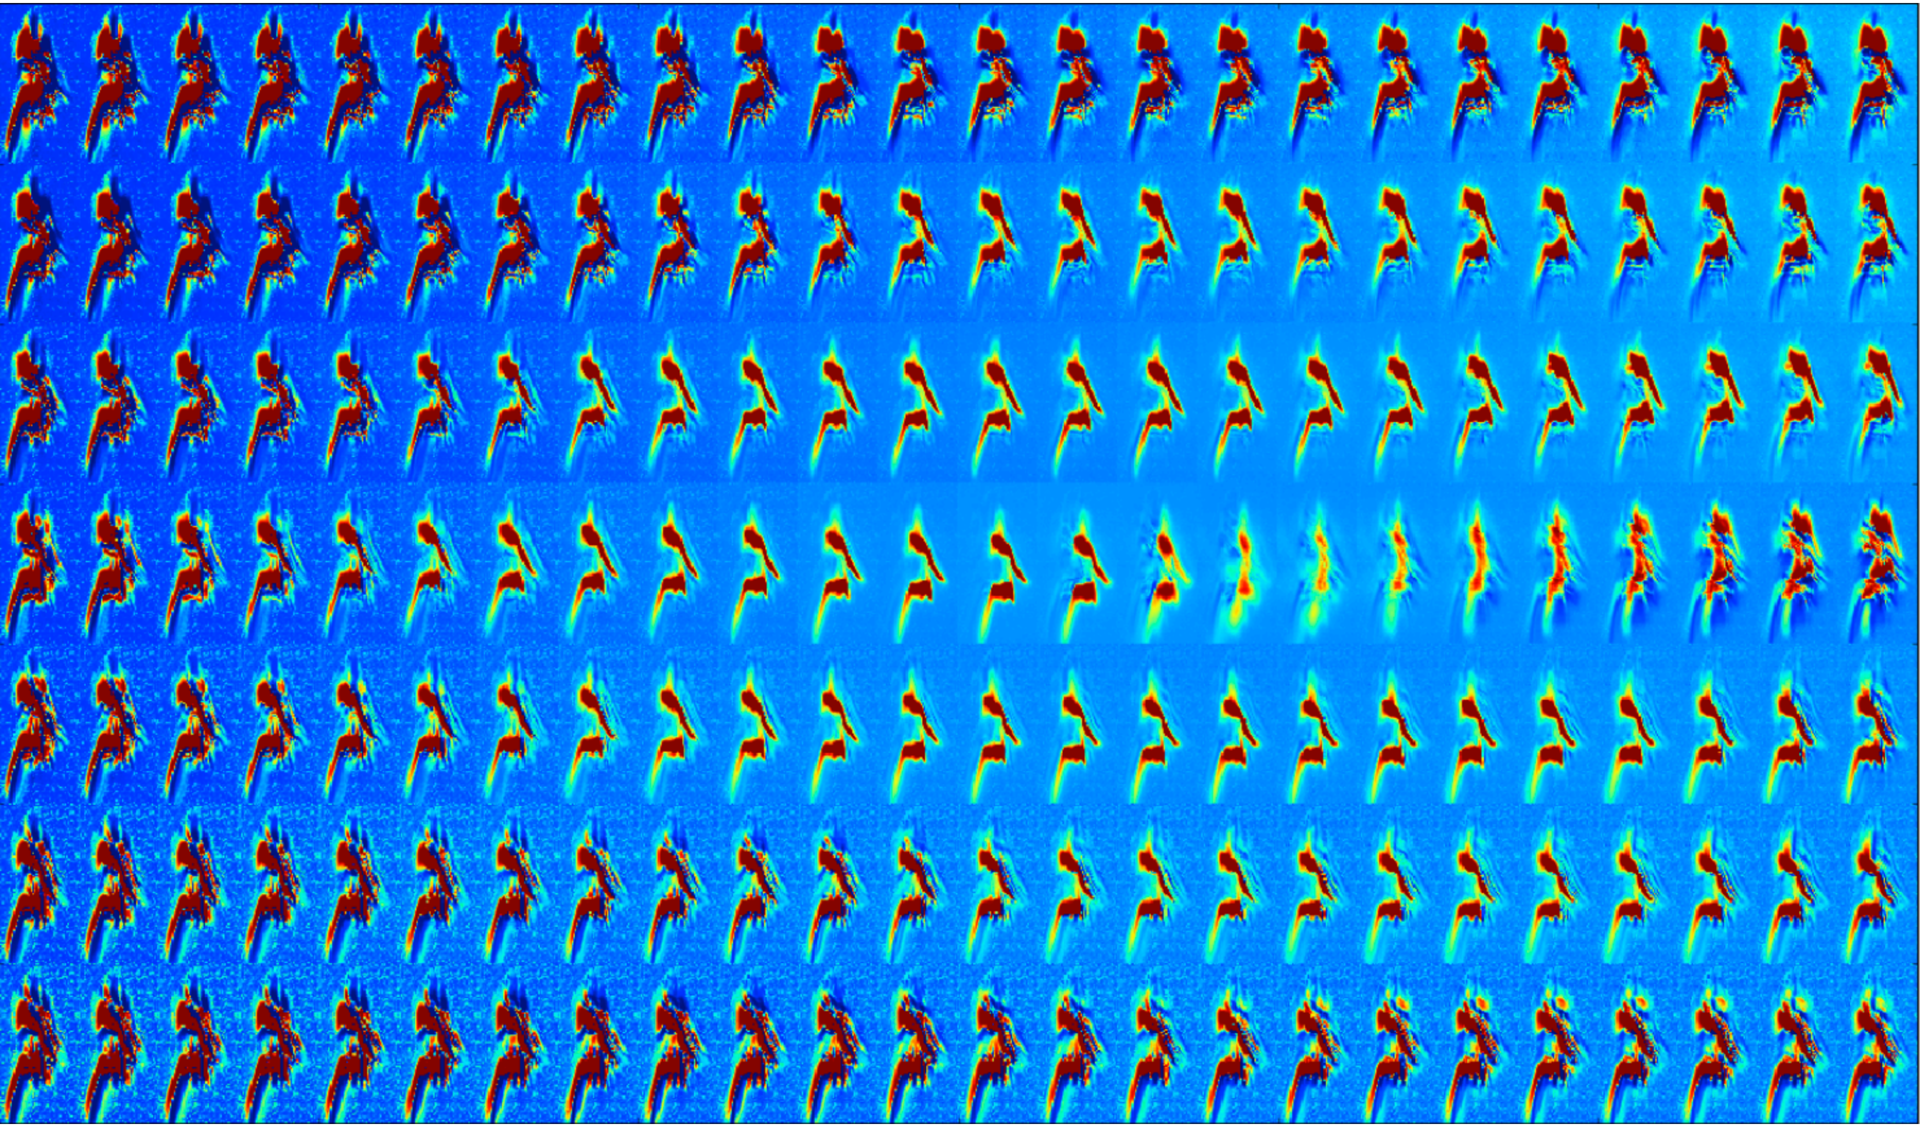

Deep Learning Splitter Idea – train like a GAN

I_0 (no lasing)

I_1 (lasing)

F (Nnet): I -> I
Generates correct reference from signal – removes Lasing -
(like the GAN generator, but no noise input)

F_0

F_1

D: (separate Nnet): discriminator:
4 way, I_0 , I_1 , F_0 , F_1

Loss Terms:
D 4-way
D 2-way 0=(I_0, F_0, **F_1**) 1= I_1

$I_0 = F_0$ (don't alter no-lasing)
$I_1 = F_1$ minimize changes to lasing to make it no-lasing
but better, model physical process, conversation law? Etc.

Transfer Learning

Reuse award winning Network trained on ImageNet

Millions of images

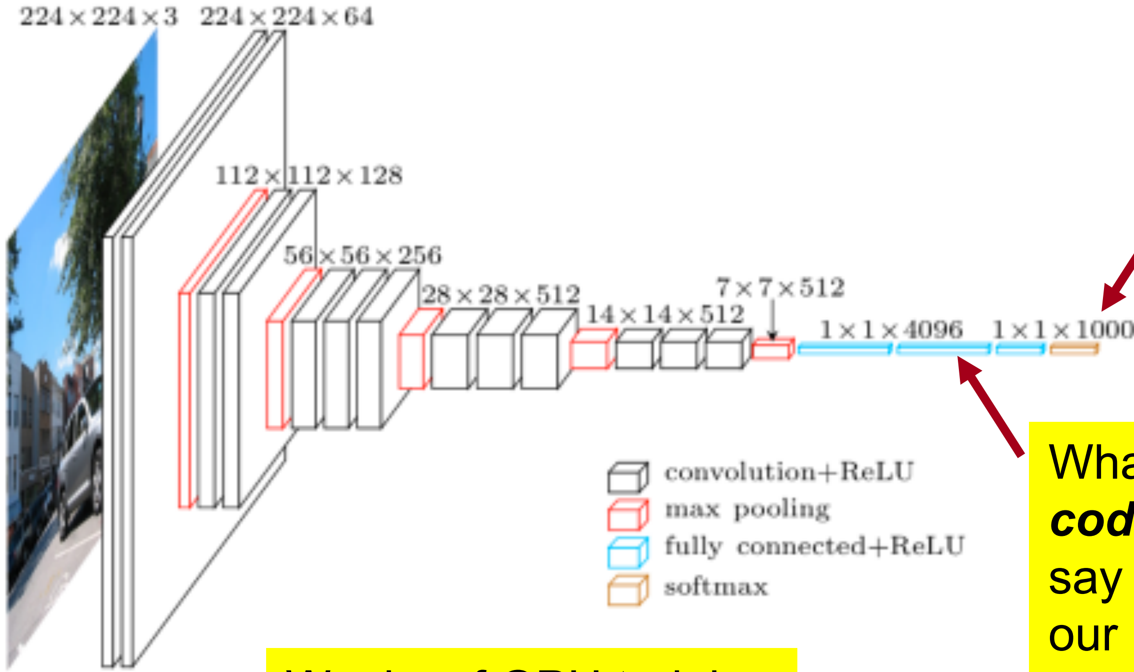

1000 Categories (cat, dog, Bus, ect)

Weeks of GPU training

What do ***codewords*** say about our Detector images?

Vgg16 – from http://www.cs.toronto.edu/~frossard/post/vgg16/

Transfer VGG16 to XTCAV

- Take *reference* and *signal* runs of XTCAV
- Run *Images* of each through vgg16 to get *codewords*
- Top rows: *reference*, bottom: *signal*

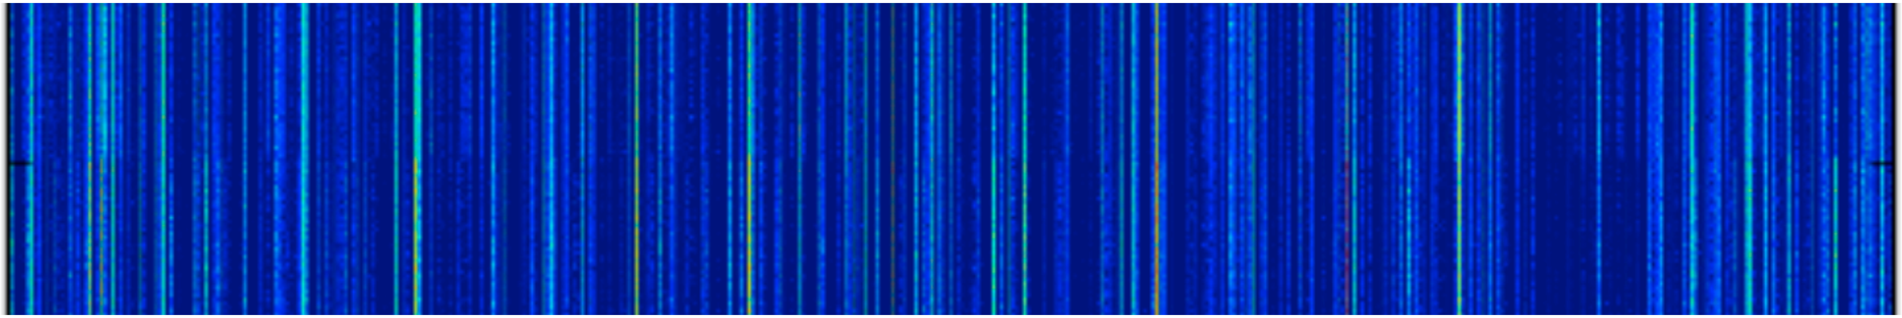

Features that distinguishes images (lasing electrons) are not spatially localized

But in codewords – we see vector components that do

Possible Application: Look up Reference

XTCAV analysis – given *signal (lasing)* shot, find the corresponding *reference (no lasing)* shot (first step to isolate lasing).

Transfer Learning approach: minimize Euclidean distance in codeword space

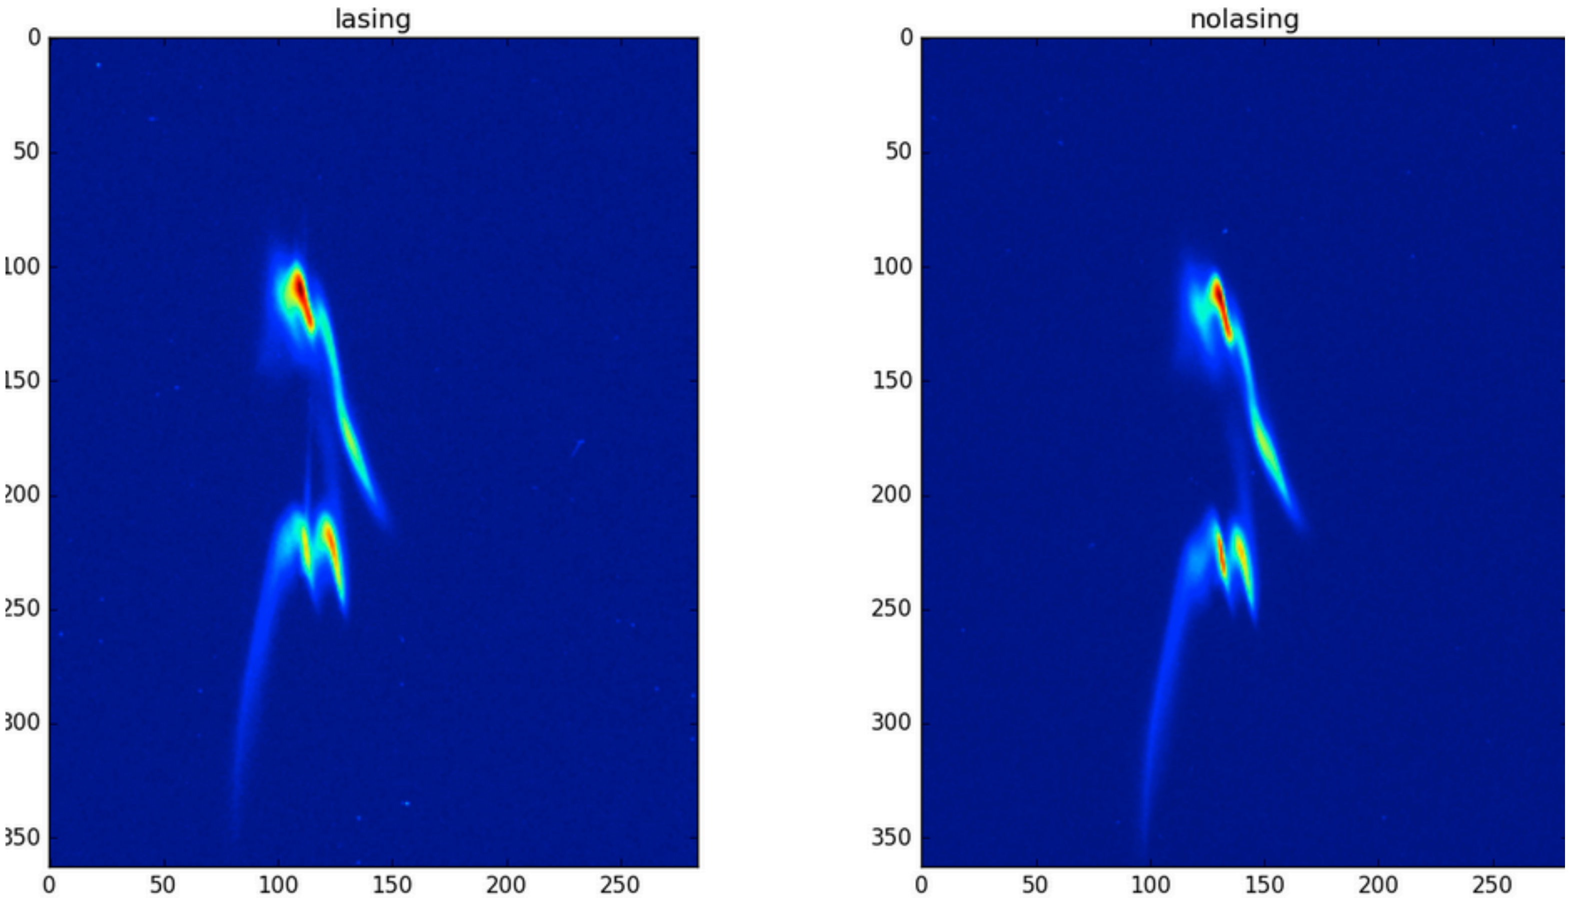

First Pass: Doesn’t work well

subtract *reference* from *signal* in Image space, but they don’t line up

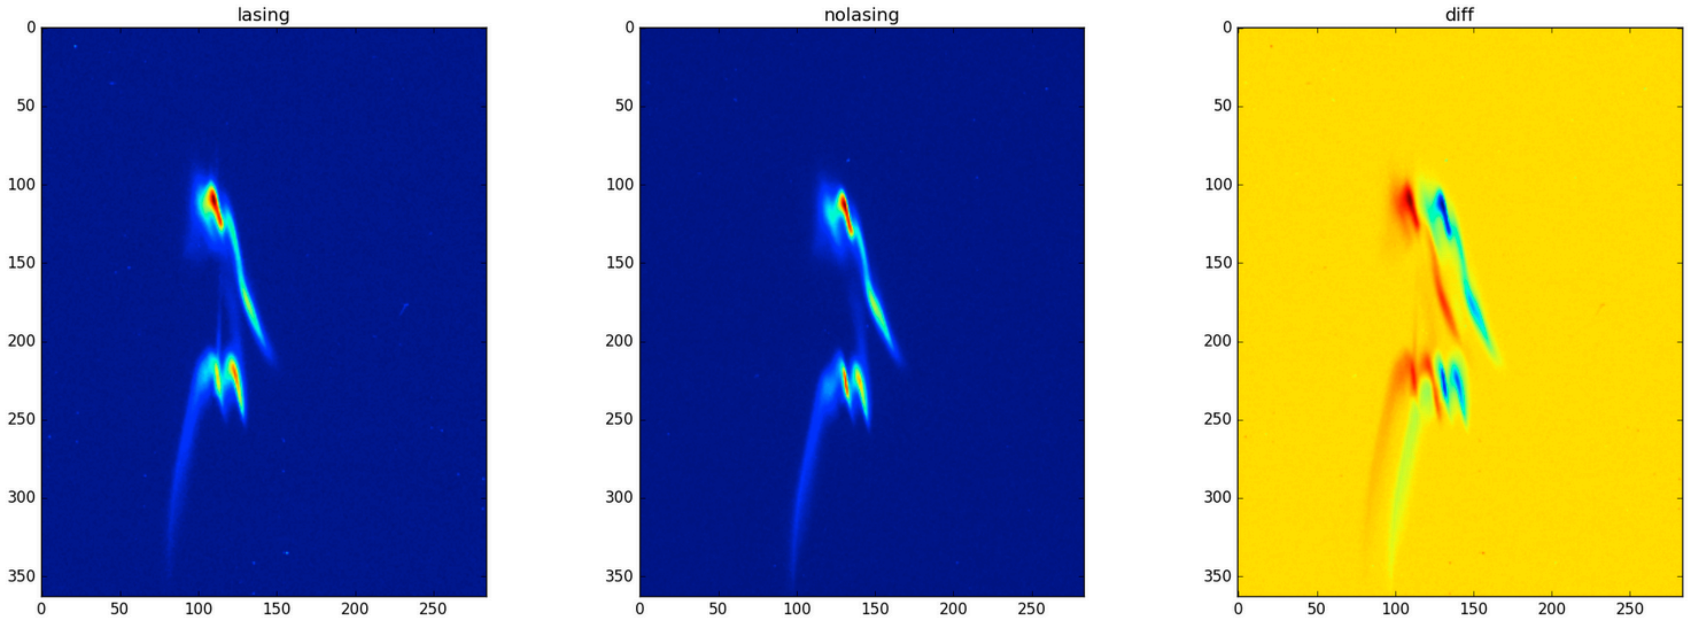

Can we invert a codeword to get the right reference?

- Start with a reference codeword
  - close to signal
  - may not be for existing reference shot
- Optimization: invert through vgg layers
  - Has statistics of references for that layer
  - Is ‘close’ to the signal shot output for that layer

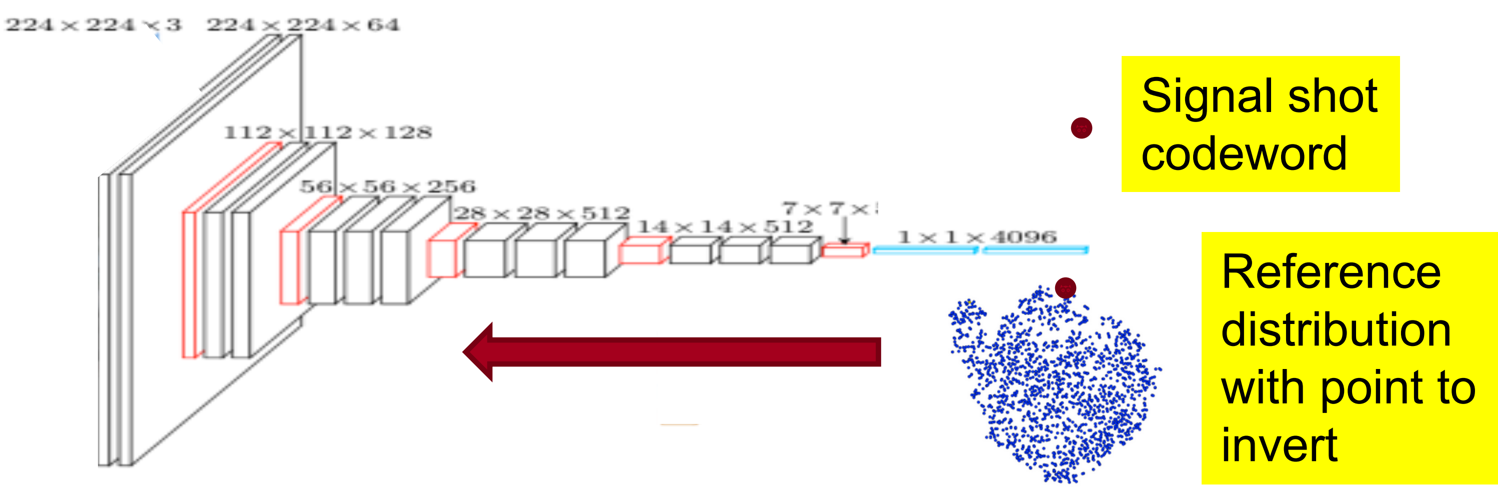

Occlusion – first pass – not working well

Train classifier to discriminate *reference* from *signal*

Occlude different spots of *signal* shot, classify each

Want lowest score for *signal* when feature occluded

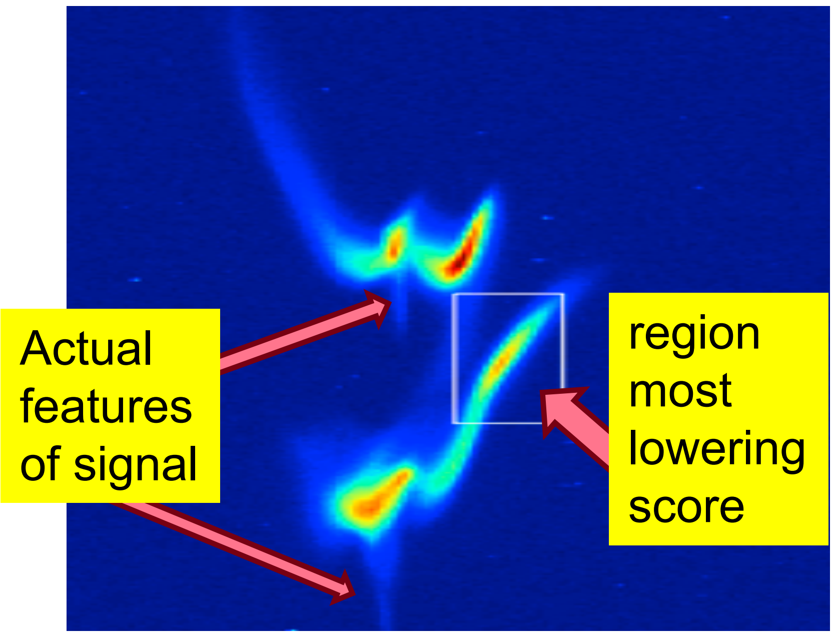

Signal shot before occlusion

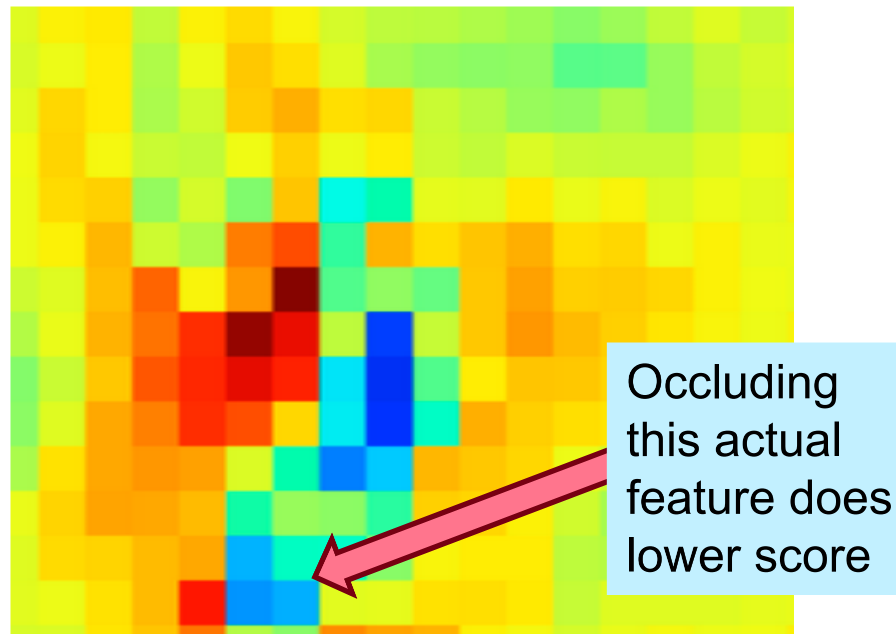

Heat map: *signal score* as function of occluded region

Occlusion – should it work?

Classifier

- trained on *reference* and *signal*
- Has never seen occluded images before
- Why would *signal* score for occluded 'do the right thing'?
- We will look at t-sne plots to evaluate
  - t-sne – one of several techniques to embed high dimensional data in a low dimensional space

T-sne for XTCAV Images

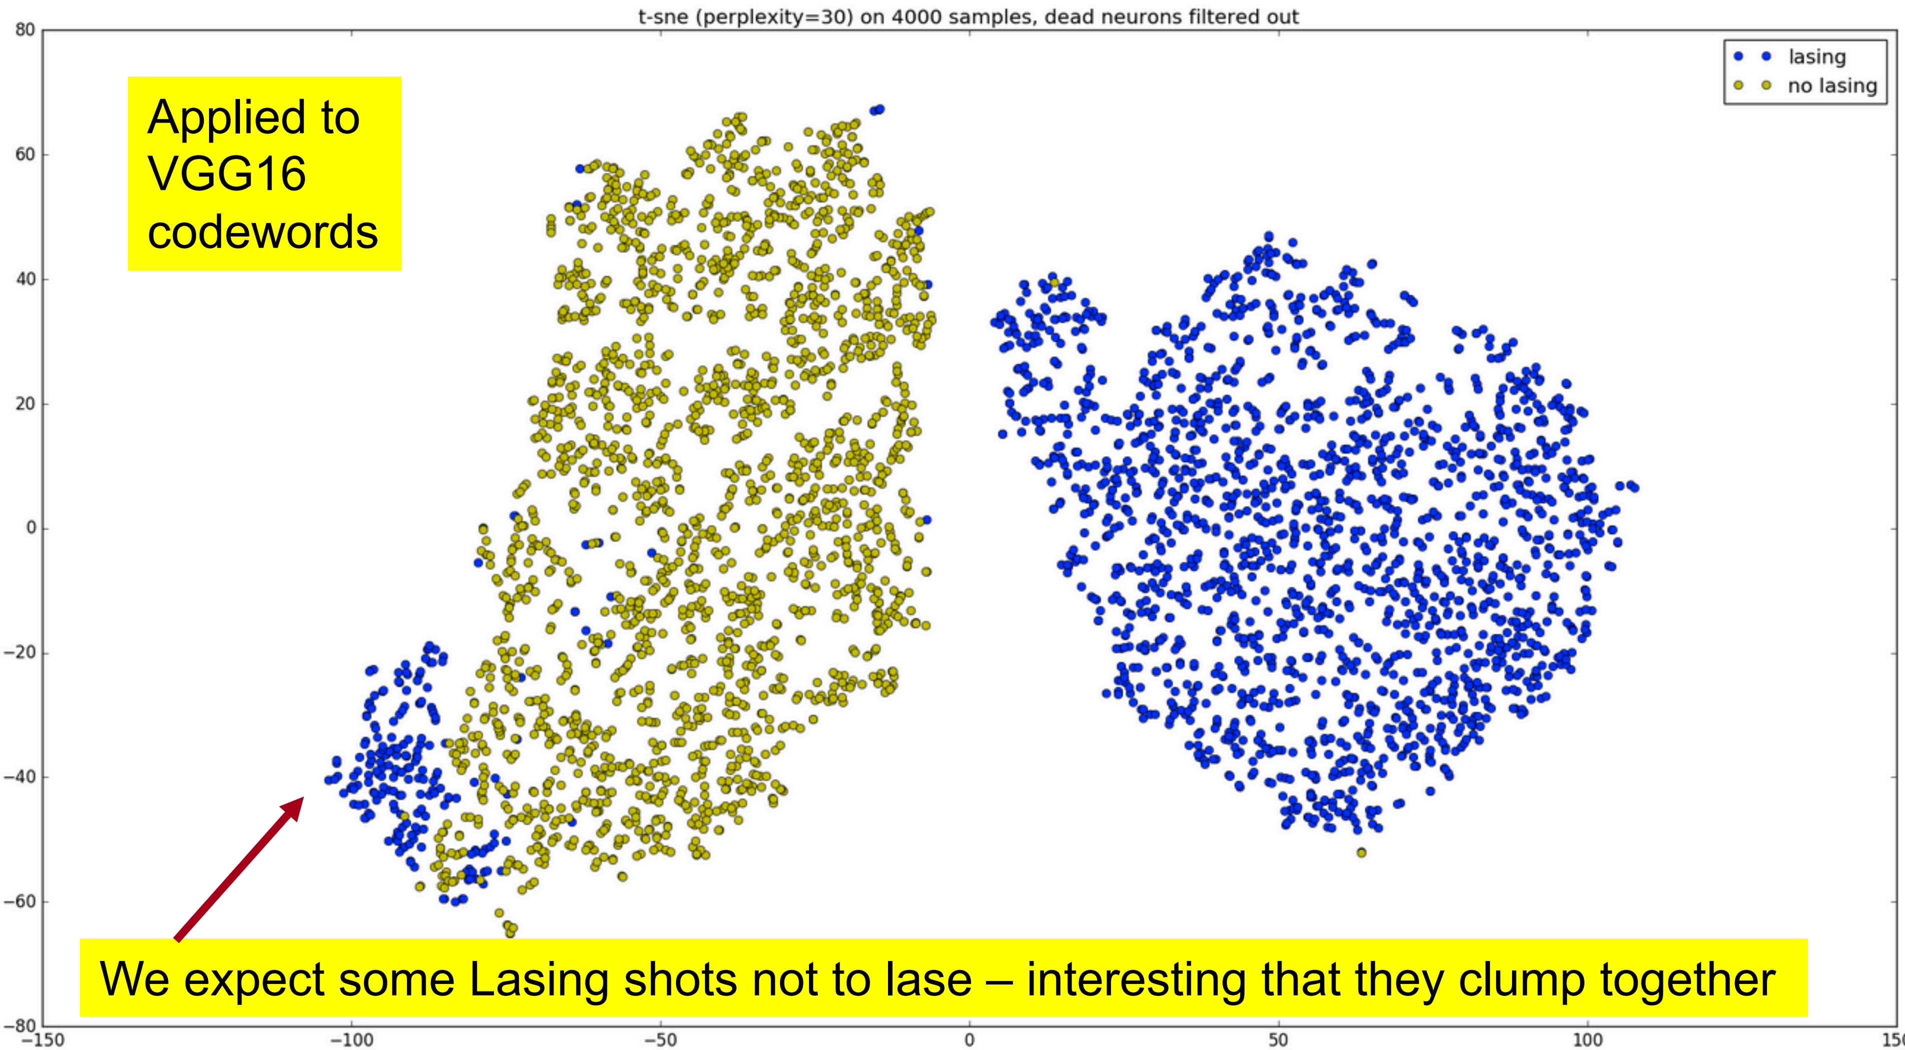

Depends on how good classifier is

t-sne plot (unsupervised clustering) of convnet codewords

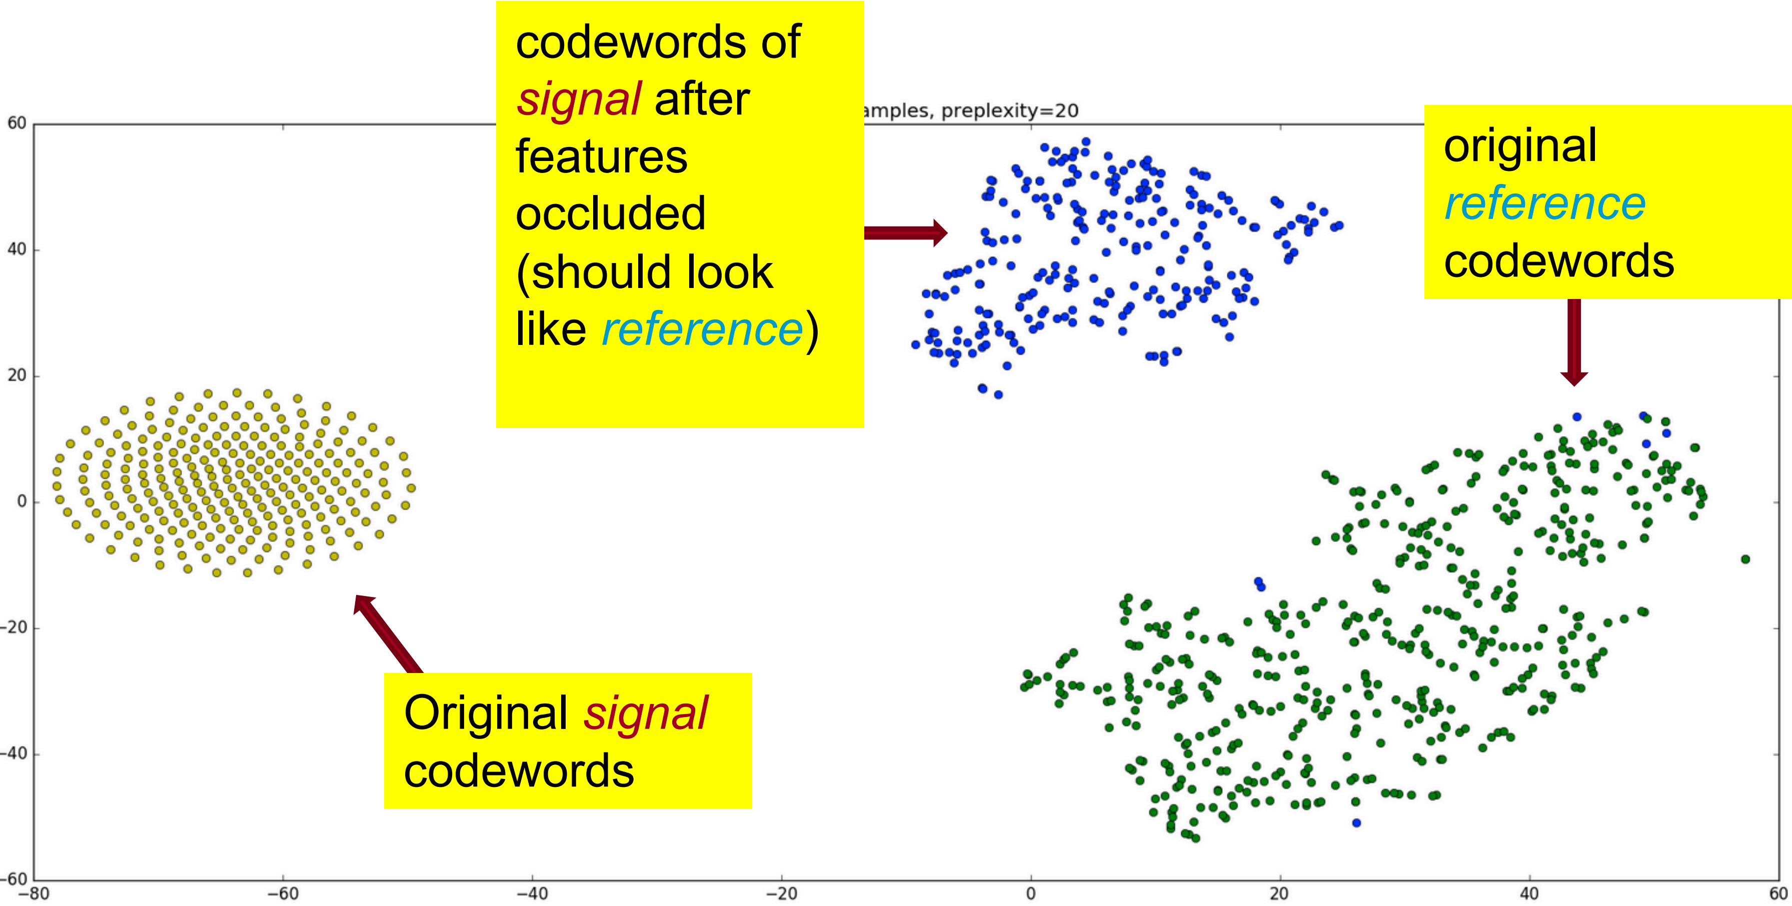

Depends on how good classifier is

Ideal Signal vs. Reference dividing line – occluded looks like reference
Maybe SVM would give us this

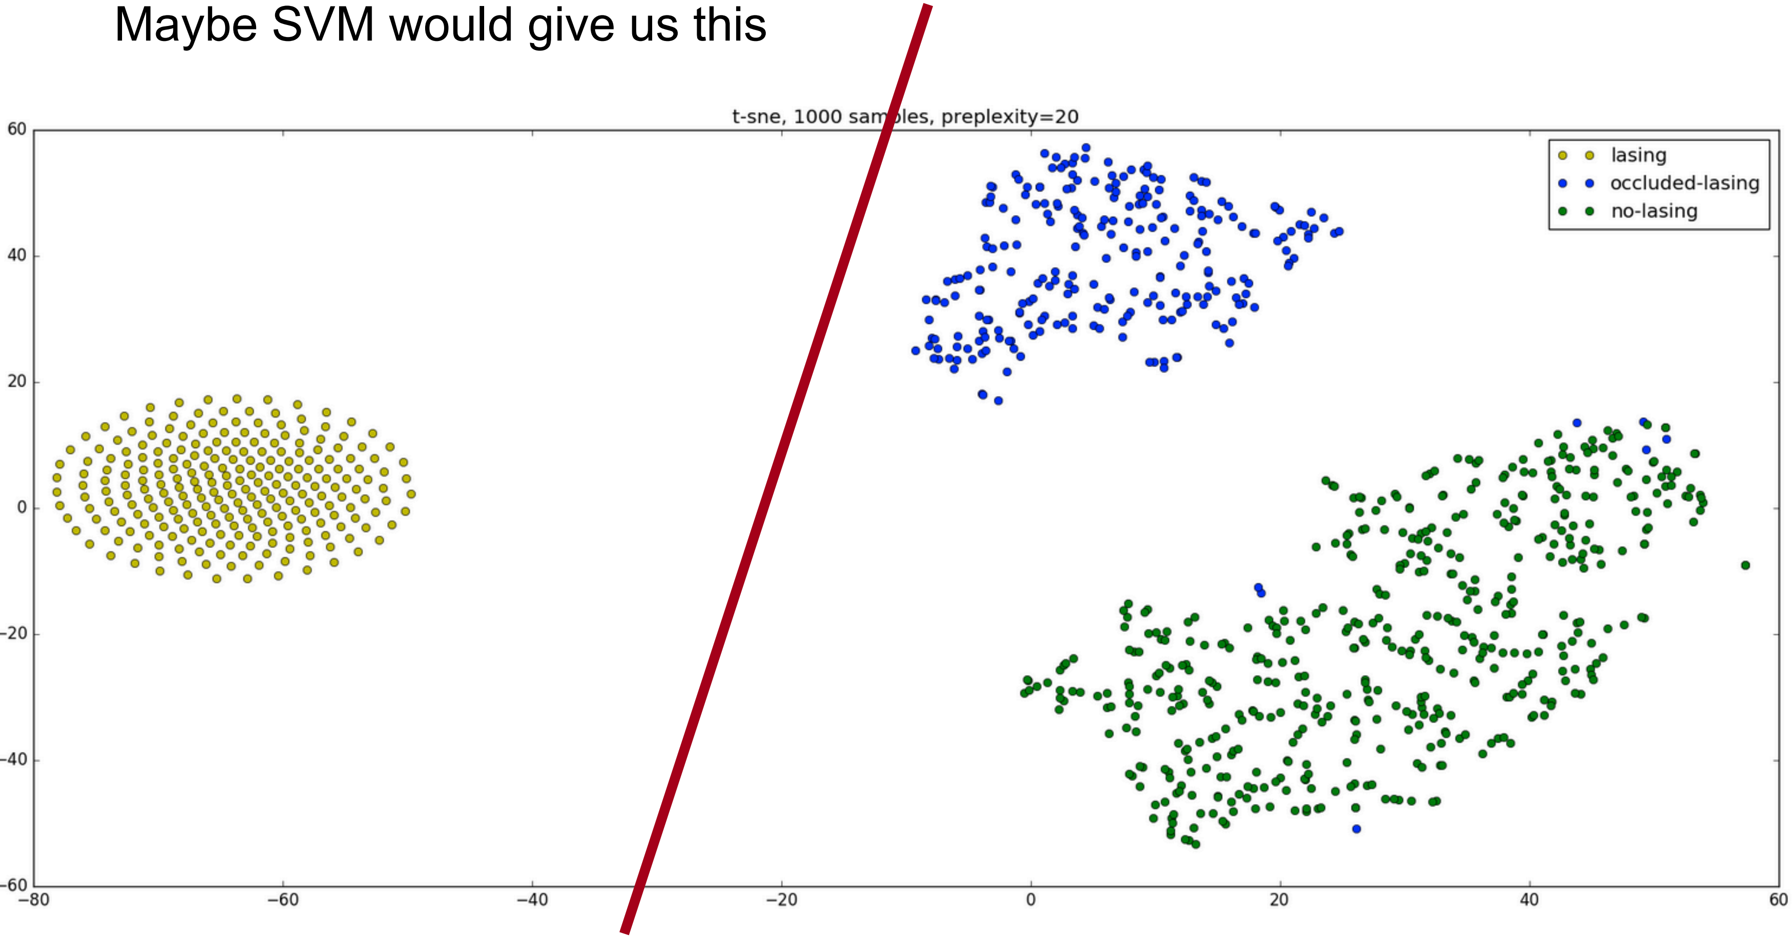

Depends on how good classifier is

Poor division for occlusion (still perfect for signal/reference classifier)

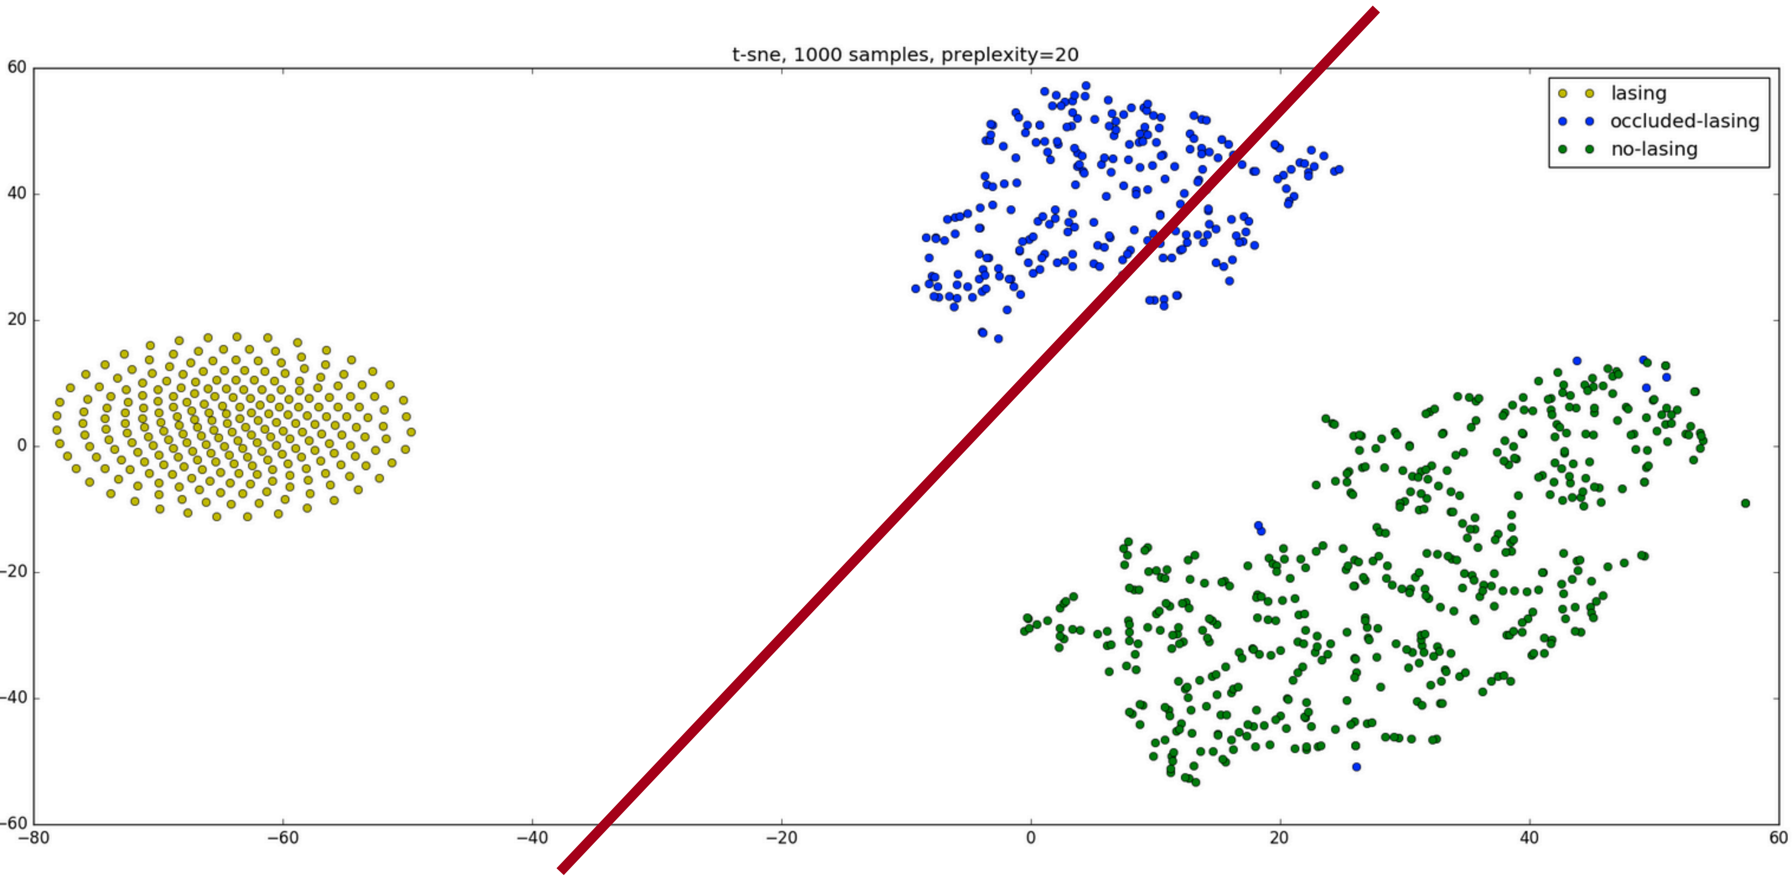

Label Boxes and Train Regression from Codewords

SLAC

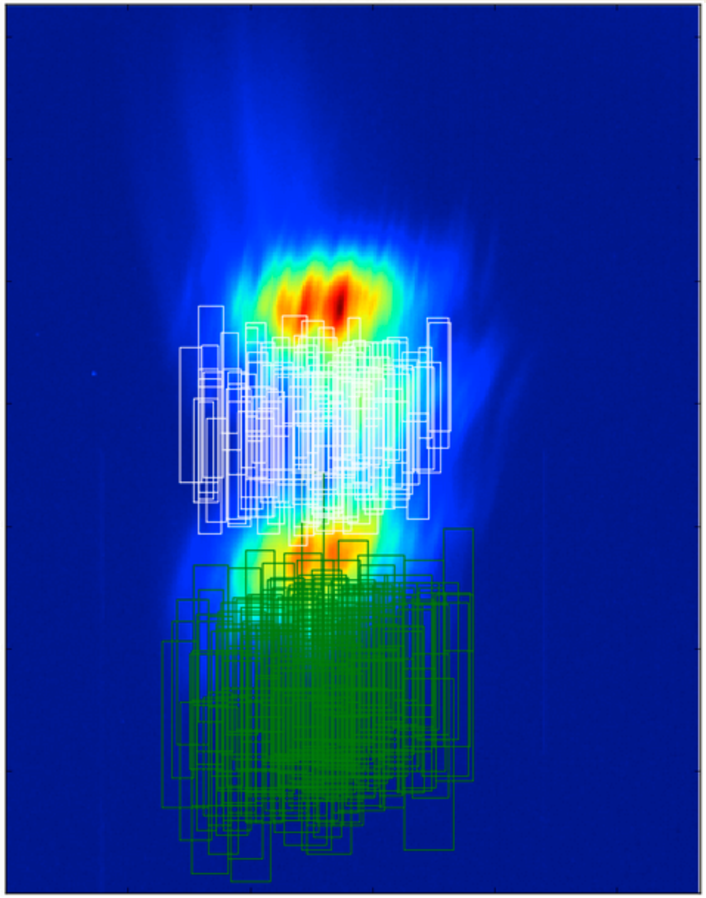

Label 250 images
plot shows variation in labels
White: feature 1
Green: feature 2

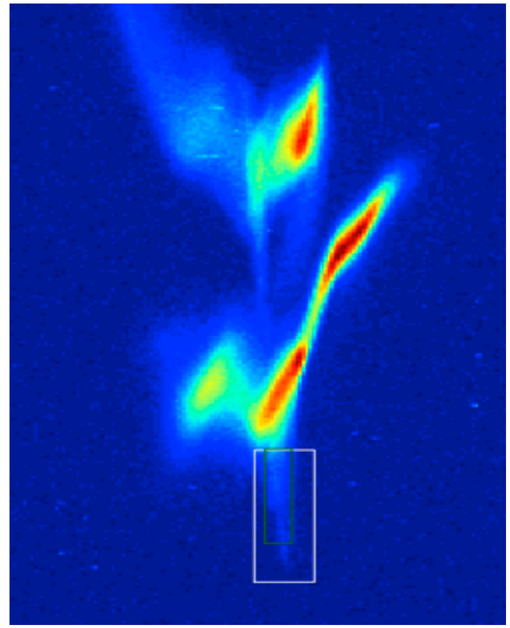

First Regression:
codewords to
box for feature 2

Green: Truth
White: predicted

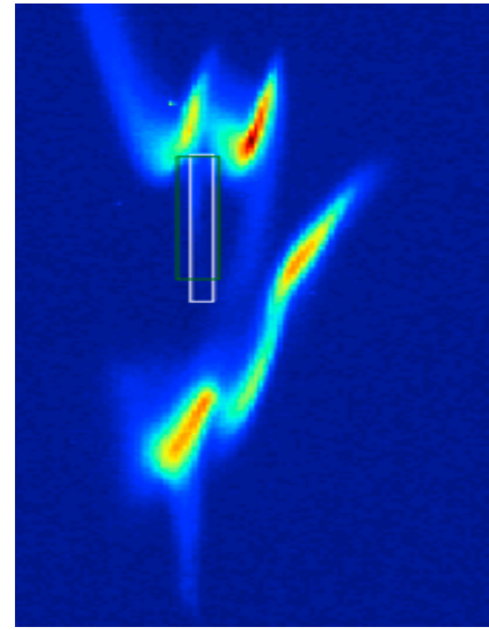

Second:
codewords to
box for feature 1

Green: Truth
White: predicted

ImageNet Convolutional Networks can Find Beams on Accelerator Screens

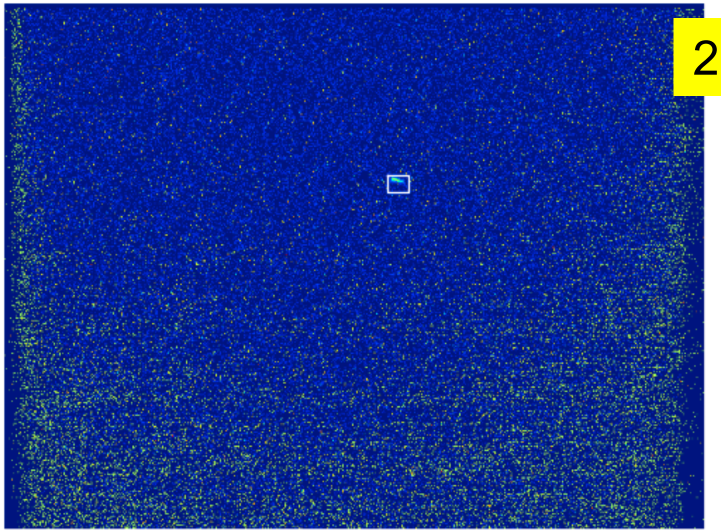

250 labeled boxes

VCC screen – high variation in beam

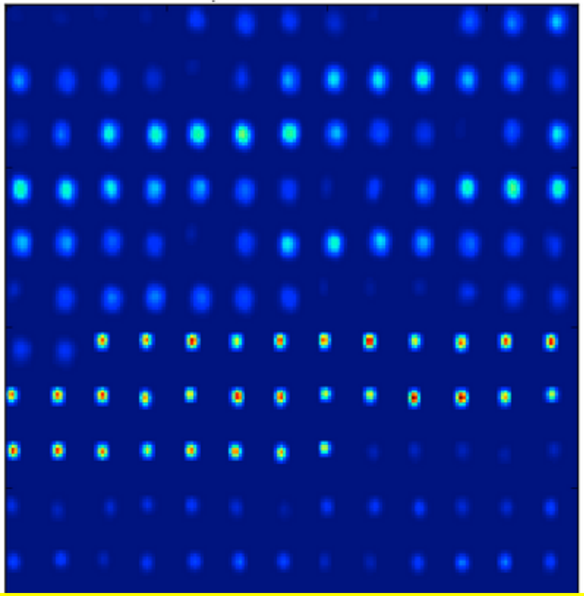

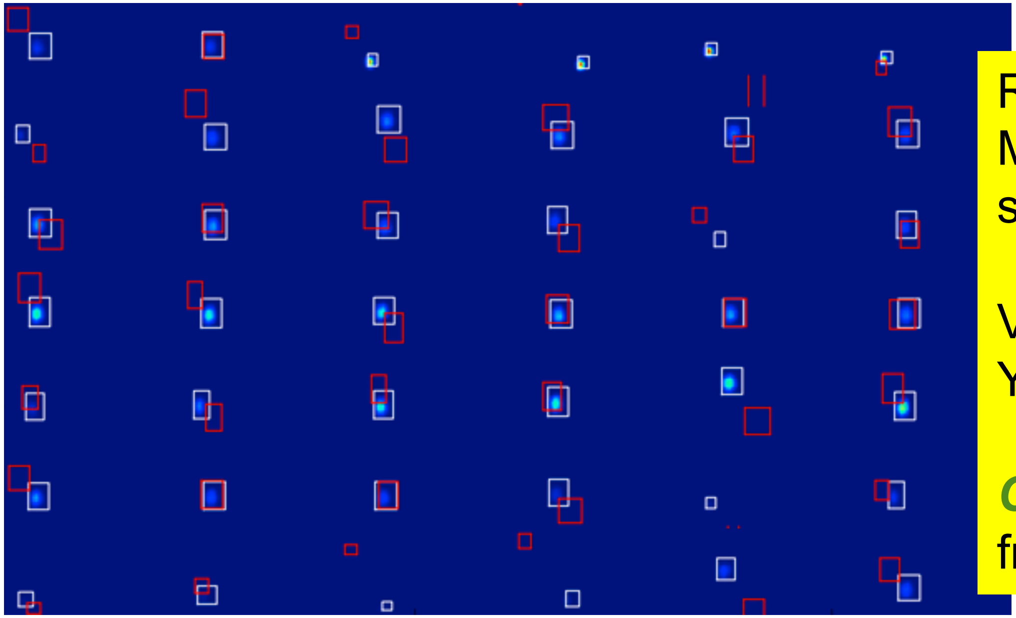

Regression Results:
Measurement of accuracy: for how many shots is area of intersection/union > 0.5?

VCC Screen – only ***0.09*** (left)
YAG Screen – ***0.89*** (not shown)

Critical step: de-noising preprocessing from A. Rashed

“DetectorNet” + “ImageNet”

Transfer learning lesson: the neural Network trained on millions of ImageNet is effective for Detector Data

How much more effective if also trained on millions of detector images, from 1000’s of categories like Diffraction of Molecule A, B, Beam Diagnostic C, etc?

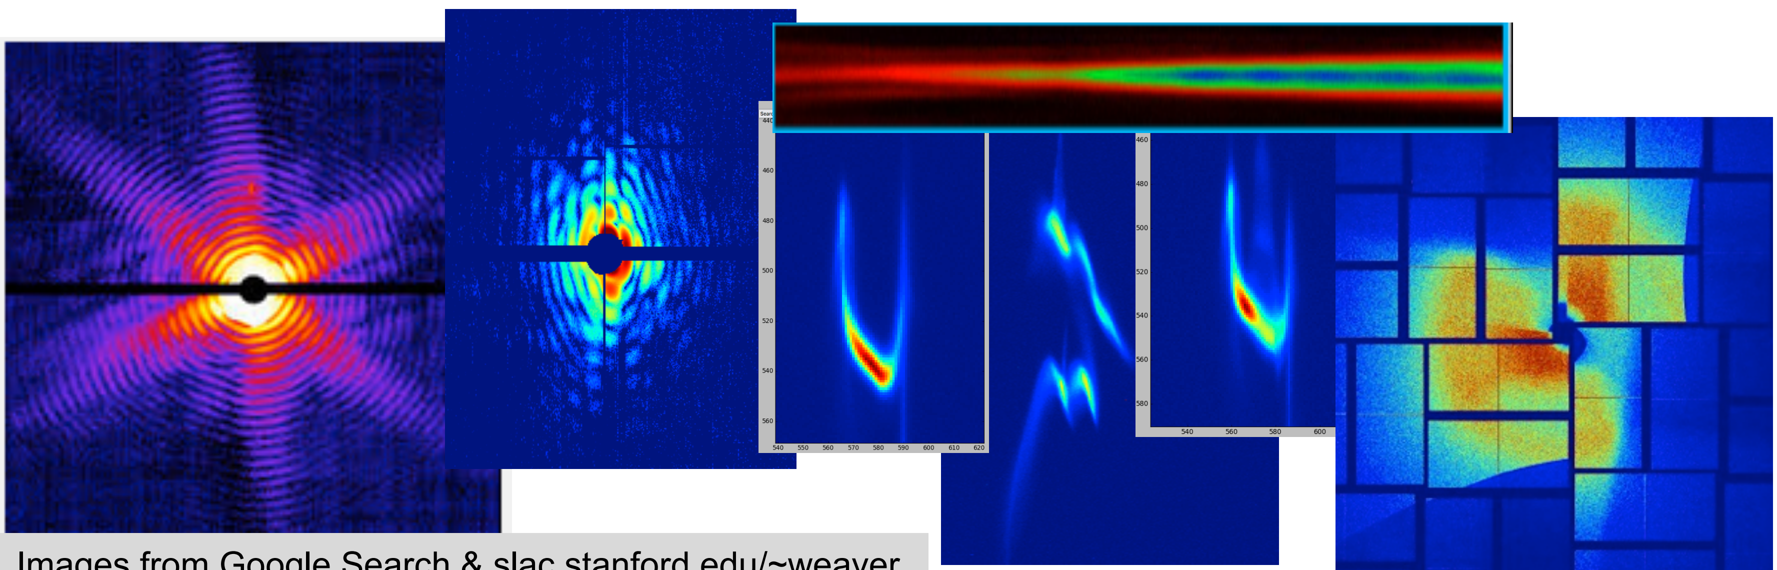

Images from Google Search & slac.stanford.edu/~weaver

Image Captioning for Detector Shots?

If we train a *Recurrent Neural Network* on detector images with captions, will the machine be able to write them in the future?

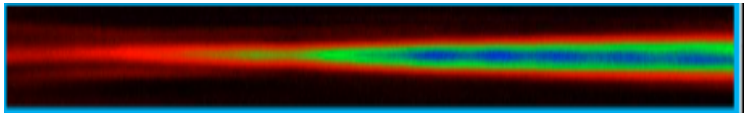

Laser/Beam interaction for Time Tool

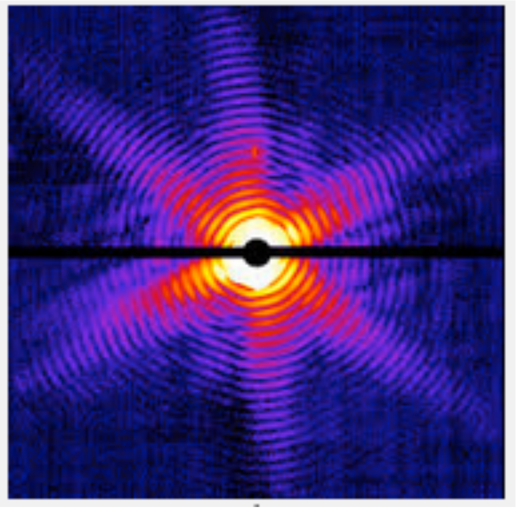

Great shot of
mimi virus diffraction!

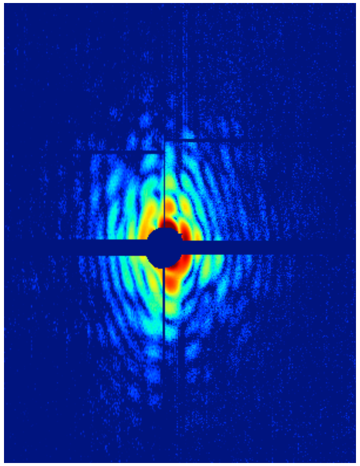

X-ray diffraction pattern
produced by a cyanobacterium
at the LCLS

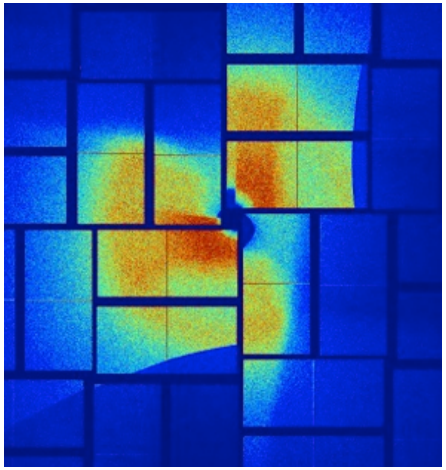

Strange effect on
a CSPAD

Images from Google Search (LCLS public, DESY) & slac.stanford.edu/~weaver