Online Data Quality Monitoring and Anomaly Detection

Thomas Britton
David Lawrence
Kishansingh Rajput

Online DQM

- **Who cares?**
  - Accelerator operation is very expensive
  - Quality of the data = Quality of the physics results
  - Human monitoring the plots as they are generated is Labor intensive, Mundane
  - Human error are likely to occur

- **Things that Matter the most**
  - Amount of False Positives (False Alarms) Vs True positives
  - Inference time

From DAQ to ANA

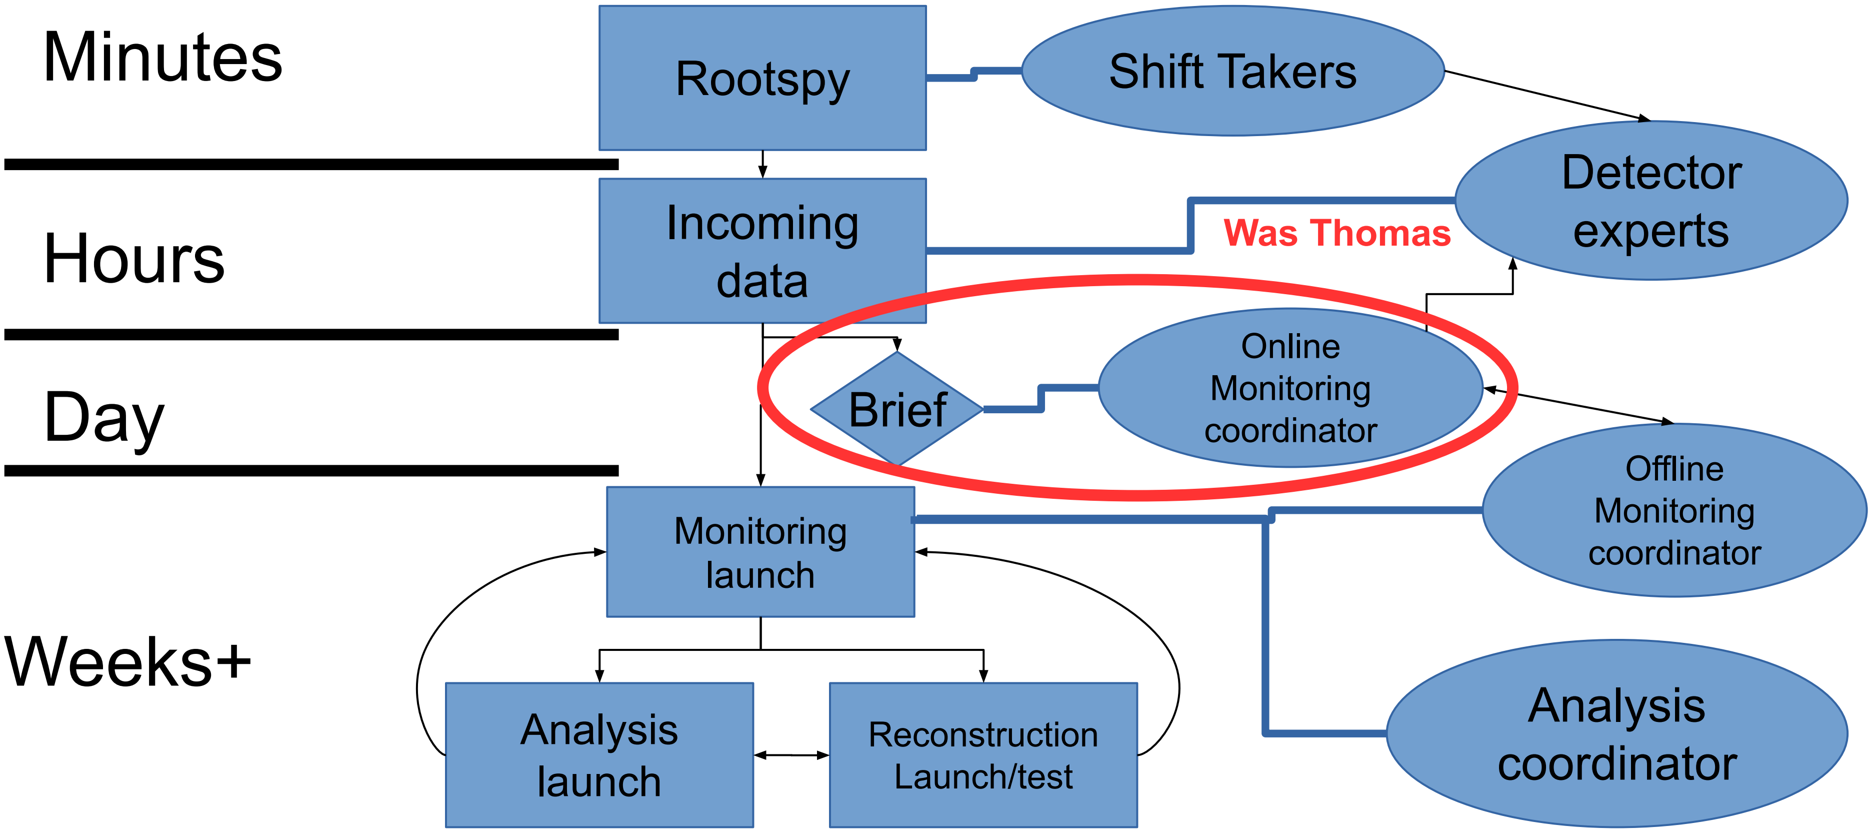

The Challenge

- **Every run** produces an initial **22 plots**. More thorough monitoring is performed offline and produces **109 plots**. With a run lasting **~3 hours** every day there are **between ~175 and 875** plots to look at.
  - A shift taker might miss a few plots, but there is no reason a machine couldn't aid in looking at all of them…
- Often times a single plot being "off" is not an indication of problems. Need to look at all the plots to determine cause and severity
  - Trigger studies: Often look like big problems but are not. Can be hard to catch when shift logs have scant details

LHC Also Using A.I. for Data Quality Monitoring

CHEP IOP Publishing
IOP Conf. Series: Journal of Physics: Conf. Series **898** (2017) 092027 doi:10.1088/1742-6596/898/9/092027

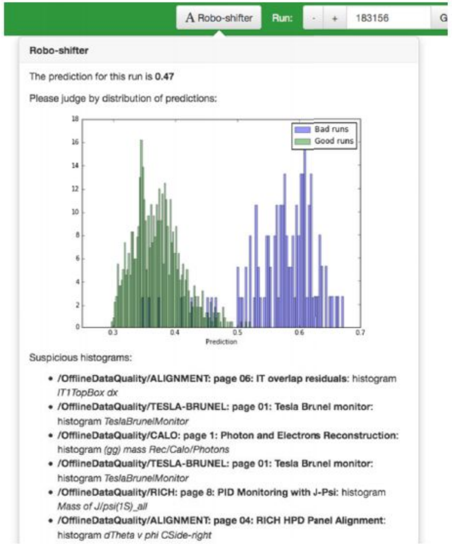

Figure 3. Robo-shifter interface

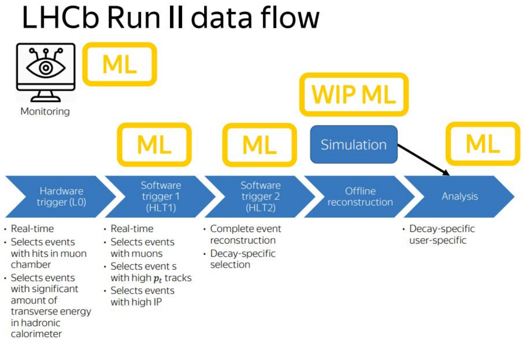

[Real-time physics, alignment, and reconstruction in the LHCb trigger] 3

CERN openlab Summer Student Report 2016

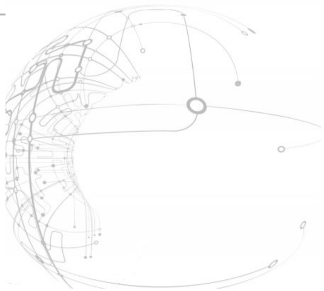

Data Quality Monitoring at CMS with Machine Learning

July-August 2016

Author:
Aytaj Aghabayli

Supervisors:
Jean-Roch Vlimant
Maurizio Pierini

CERN openlab Summer Student Report 2016

Introducing Hydra

- Hydra aims to be an extensible framework for training and managing A.I. for near real time monitoring
  - If you need it to tell a dog from cat I can have hydra do that, without system modification, now

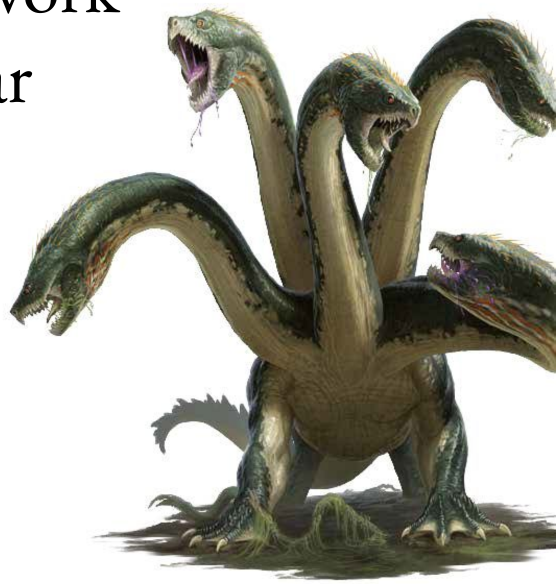

Koboldpress.com

From DAQ to ANA

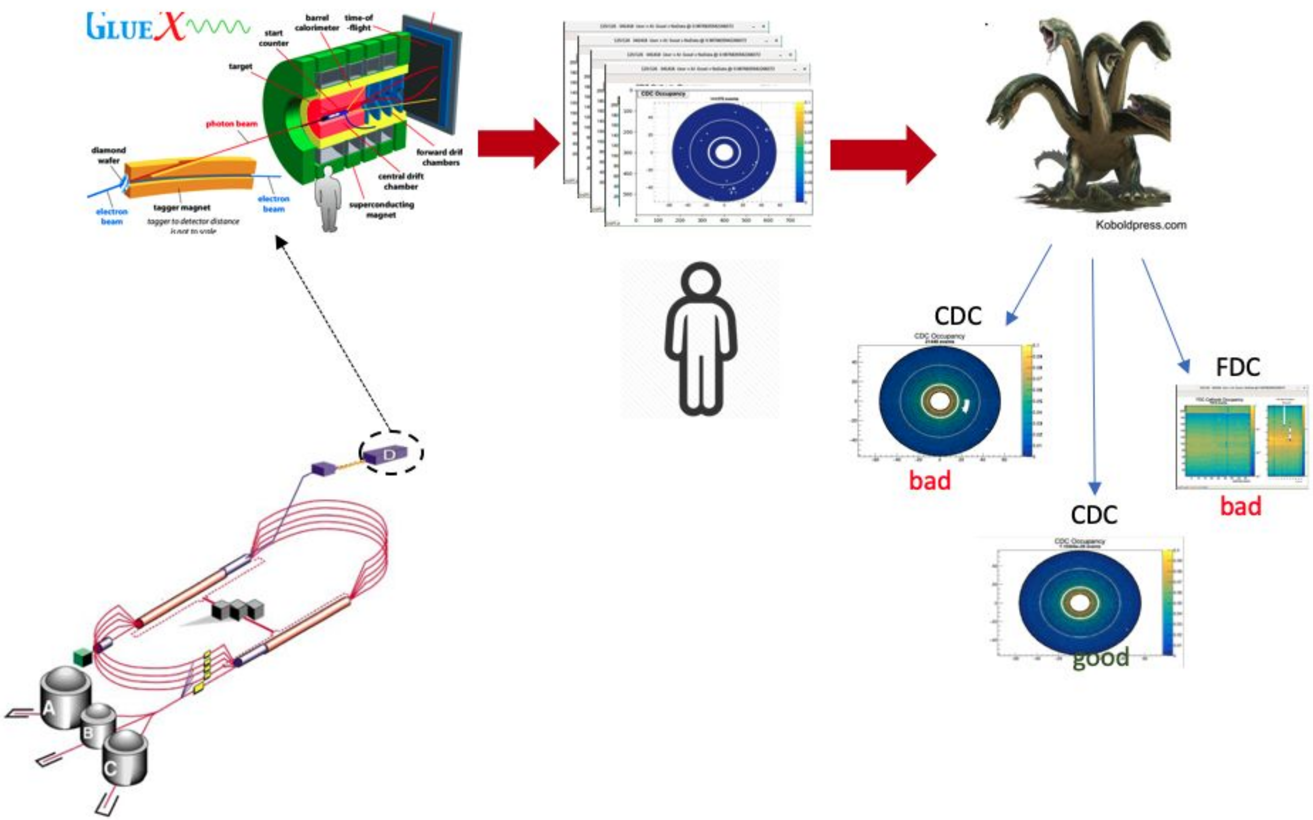

Getting the Data

The Data is Already There!
(no additional processing)

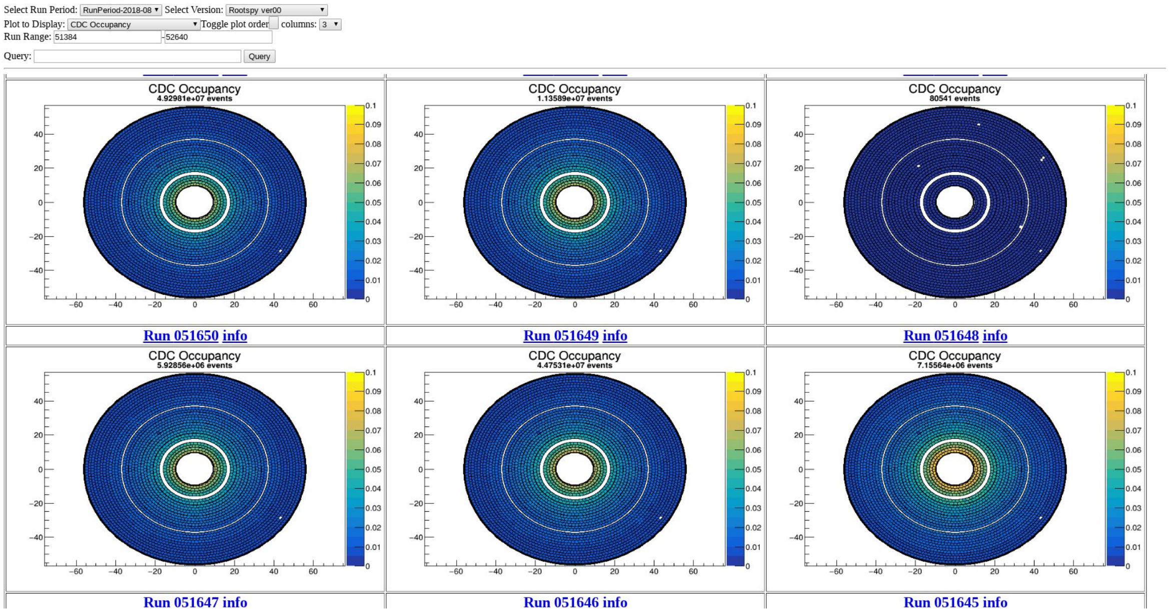

Labeling that Data

- Webpage:

~A few hours to label **all** of the plots of a given type

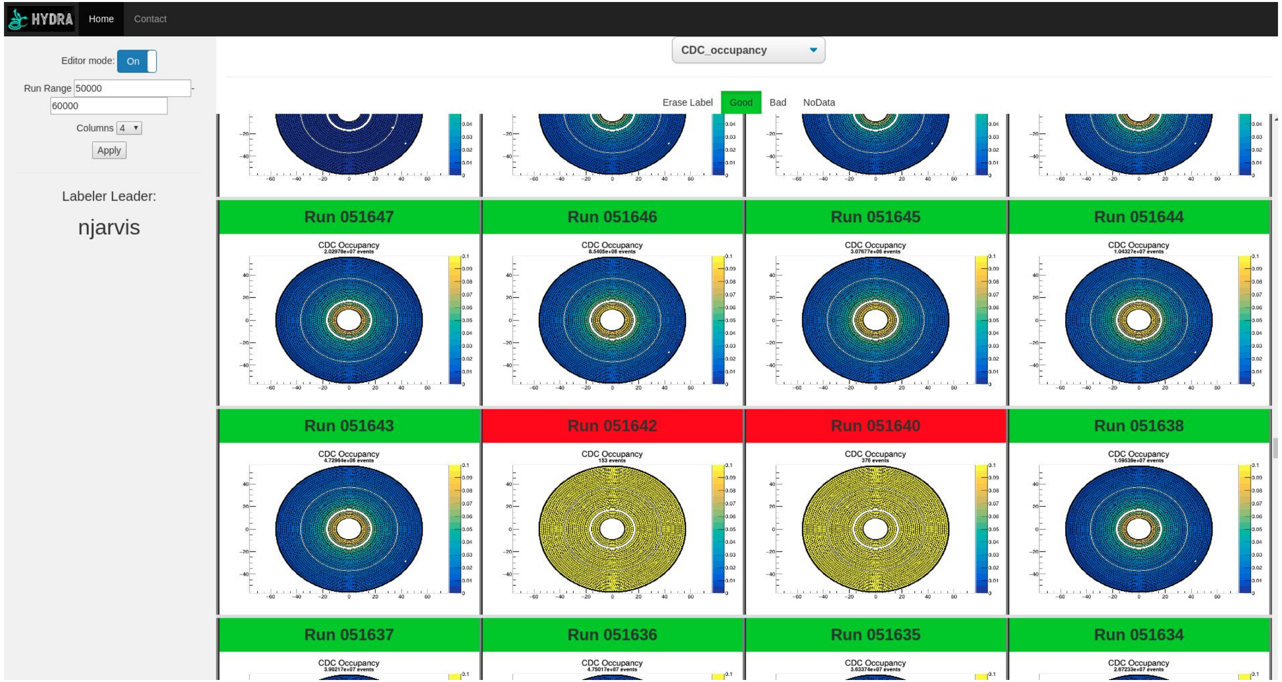

The Backend

- Supported by a database
  - **All** plots
  - **All** user defined labels
  - **All** models
  - **All** models' classifications with confidence*

    *Only saved plots

- Training is virtually push button to allow for automated retraining as needed

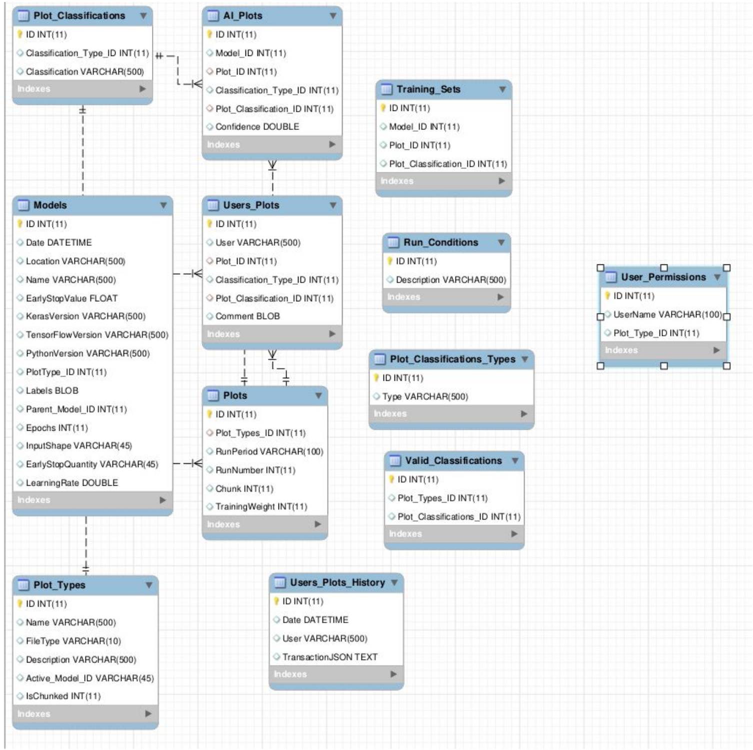

The Inception v3 Network

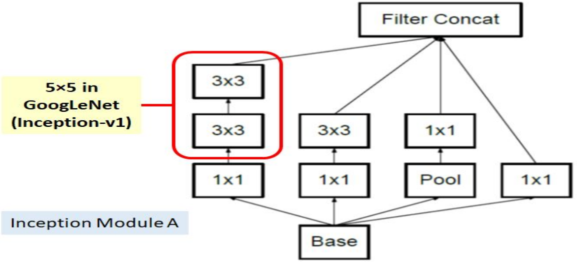

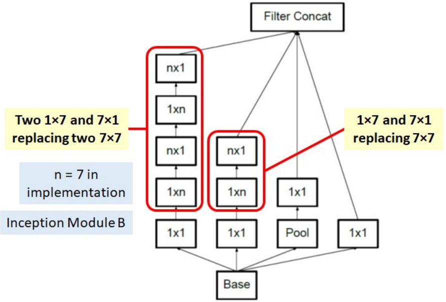

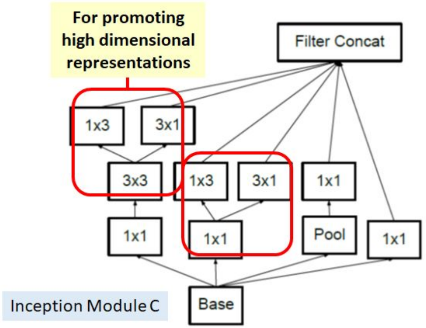

Already in Keras and Tensorflow

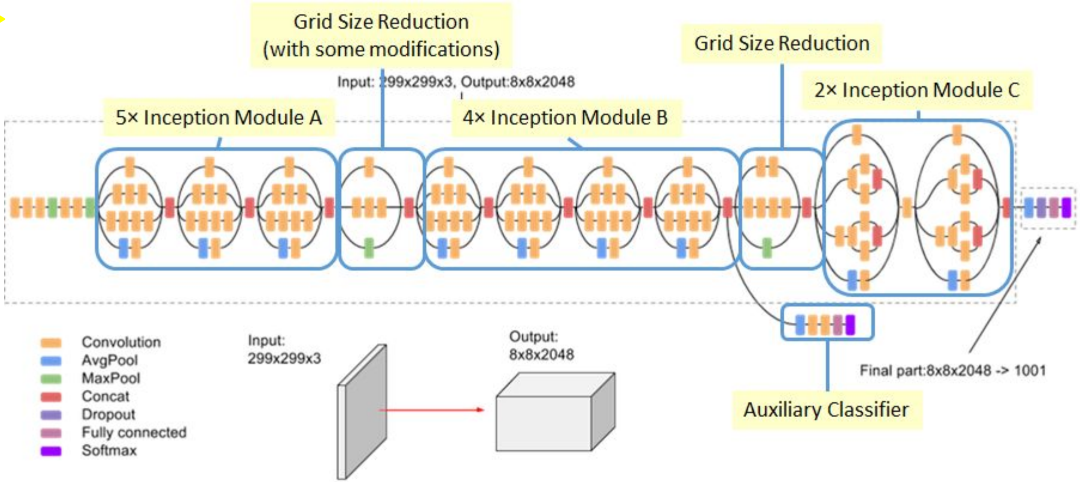

Knocked-out Data

- David Lawrence produced a plugin to "simulate" data as needed. We can train it to **recognize issues that may not have happened** yet.
  - Can also use it to test for "over-fitting"

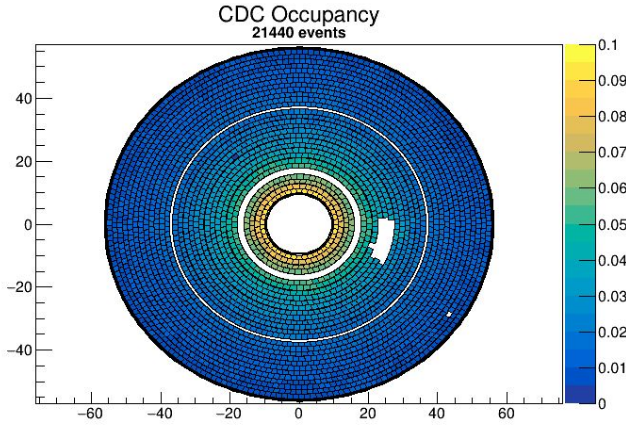

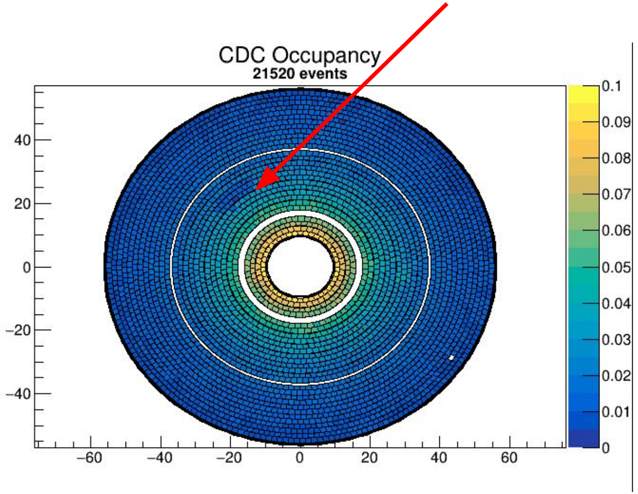

Inferences

- Prediction performed by a script call
  - **3 modes**:
    - **Datum** (full path to image)
    - **Directory** (non-blocking attempts to analyze all data it finds)
    - **Crawl** (over all saved data in the DB)

- Writes and publishes (0MQ) **JSON report** including each label the model knows about and it's confidence in the label for each piece of data analyzed
  - Can be used in downstream processes
- Automatically captures examples for future training/analysis
  - "Bad" examples
  - Disagreements
  - Every Nth example

Dashboard

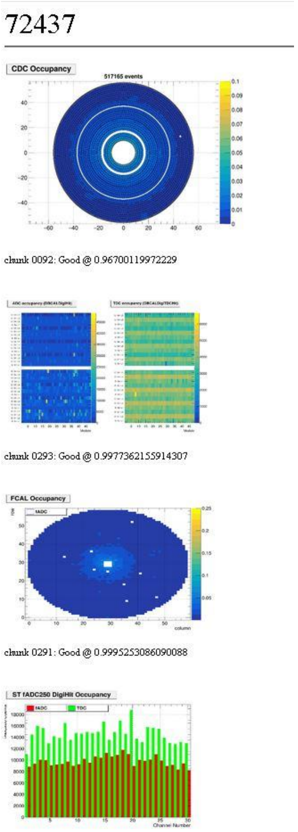

- **Real-time dashboard** viewable from anywhere shows the last plot analyzed as well as the class and confidence
- Can identify beam trips within a minute
- Able to detect hot channels in some detectors for later calibration. Can detect these problems early indicating hardware that may soon need replacing
- Future variations will give a go/no-go indicator for the plots. When something seems off shift crews can see the plot and **focus on the plots that matter**
    - Notify experts, sound an alarm, take corrective action

CDC Results

At false positive rate of 0.005
True positive rate for Anomaly is 0.96

An Anecdote

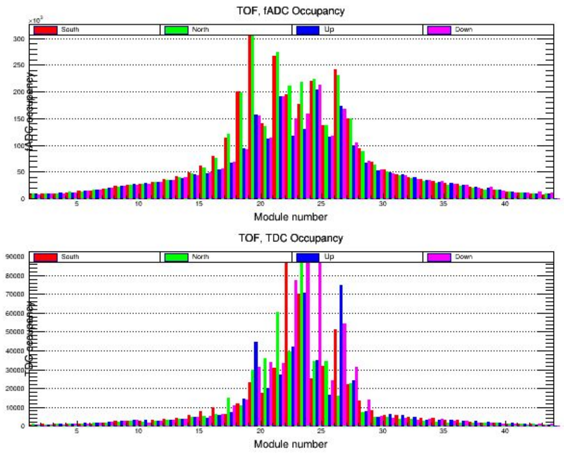

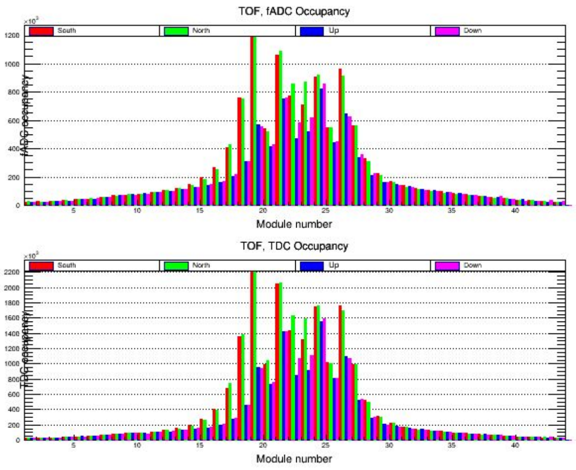

- Both of these look “good” at first glance (both initially labeled good)
  - The one on the left is actually bad (the A.I. caught it)
- A.I. seems to be able to look at subtle differences in shape

Another anecdote

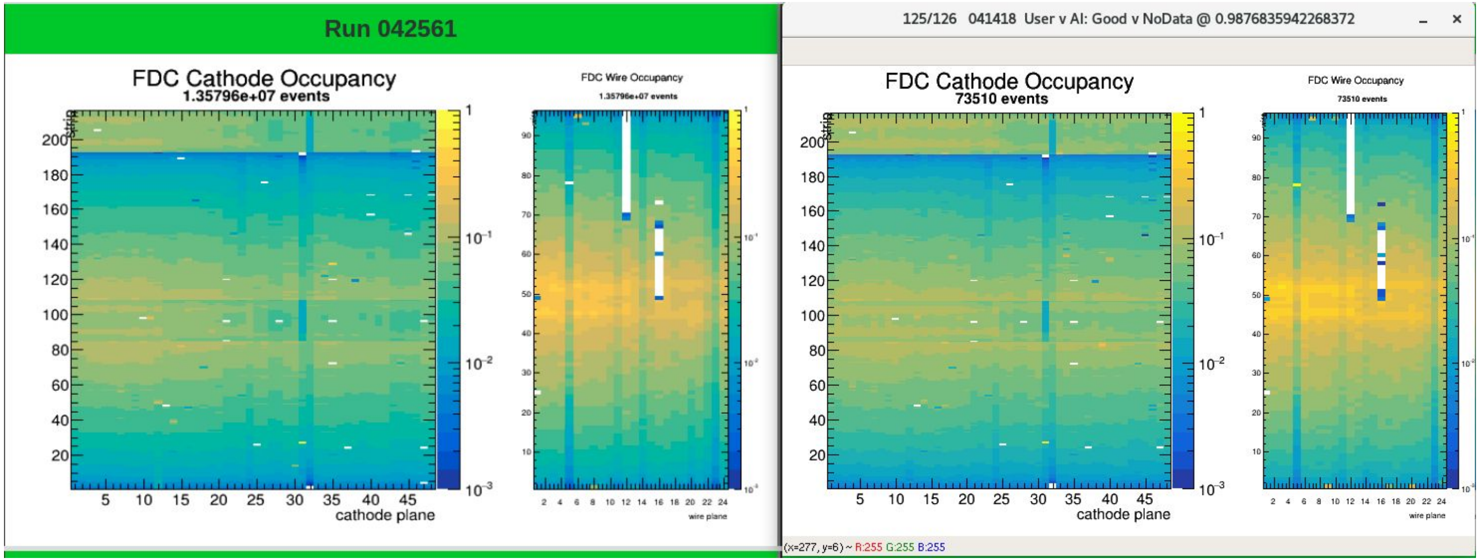

- The labeler was instructed by the detector expert to label any plot containing **fewer than 100k events as "NoData"**. This is one example of several in which the labeler labeled as "Good" and the **A.I. predicted "NoData"**…the true label given the number of events

Some things noticed on the way

- Label balancing is important
  - Detectors are designed to work!!
    - E.g. Good : No data : bad = **80:15:5**

- Strategic undersampling! For CDC, There's very low variation in the good/No data plots as compared to **bad** plots.
  - Random undersampling of the larger categories to match the smaller categories(bad) **yields similar results!**

- Breakfast

Hydra Fast Facts

- Hydra looks at a finer time scale then any higher level monitoring the shift crew performs. **Approximately every minute**
  - Because who hits reset?
- **Operates** (conservatively) at about **3-4Hz**
  - From receiving an image to action **~300ms**. Most of the time spent on model inference
  - Inference accounts for **~71%** of the total processing time and is driven primarily by model size
- Currently focused on **go/no-go decisions**
  - Doctor classifying you as sick with no diagnosis as to what you are sick with. Refinement underway

HydraRun also saw the FDC problem, which I probably would have missed inspecting it by eye.

Koboldpress.com

Layerwise Relevance Propagation

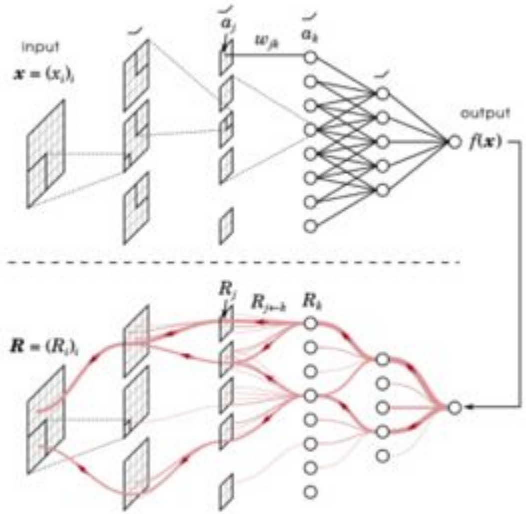

W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders and K. -R. Müller, "Explaining Deep Neural Networks and Beyond: A Review of Methods and Applications," in *Proceedings of the IEEE*, vol. 109, no. 3, pp. 247-278, March 2021, doi: 10.1109/JPROC.2021.3060483.

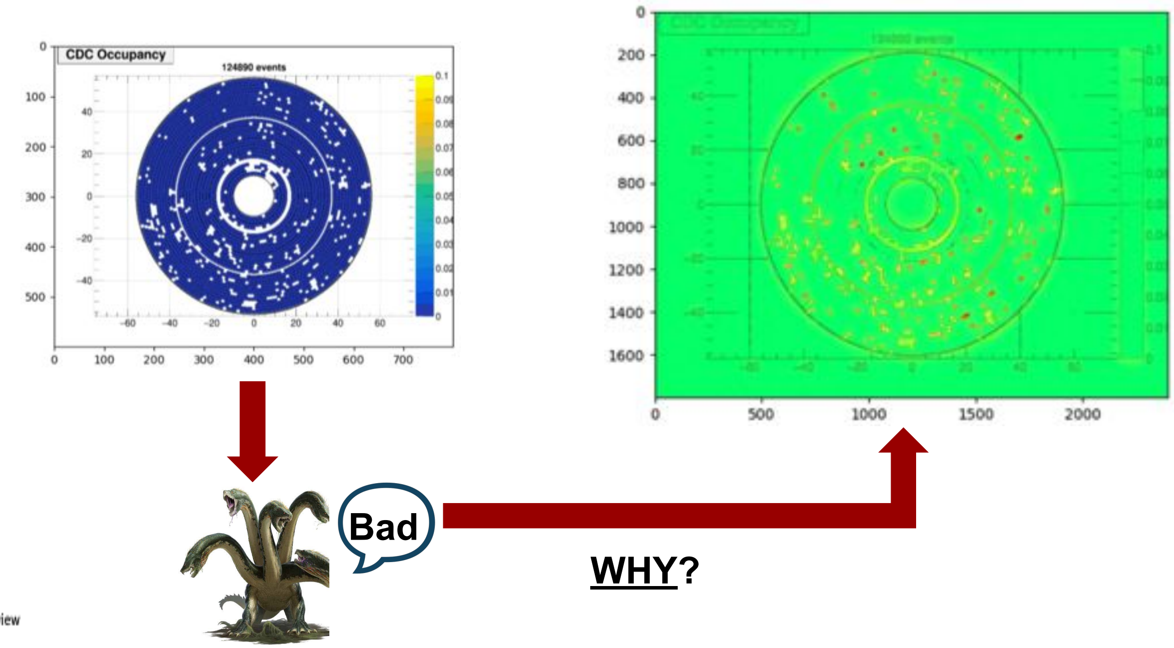

Hydra Future Development

- **Classification splitting**
  - From a doctor saying "you are sick" to actually diagnosing a condition
- **Custom, optimized models**
  - Inference accounts for ~71% of the total processing time and is driven primarily by model size
- Ability to actually **take corrective action** as needed
  - Will require trust and more data on in situ running
- **More plot types!!**
  - Data types too

Backup

Results (Start Counter)

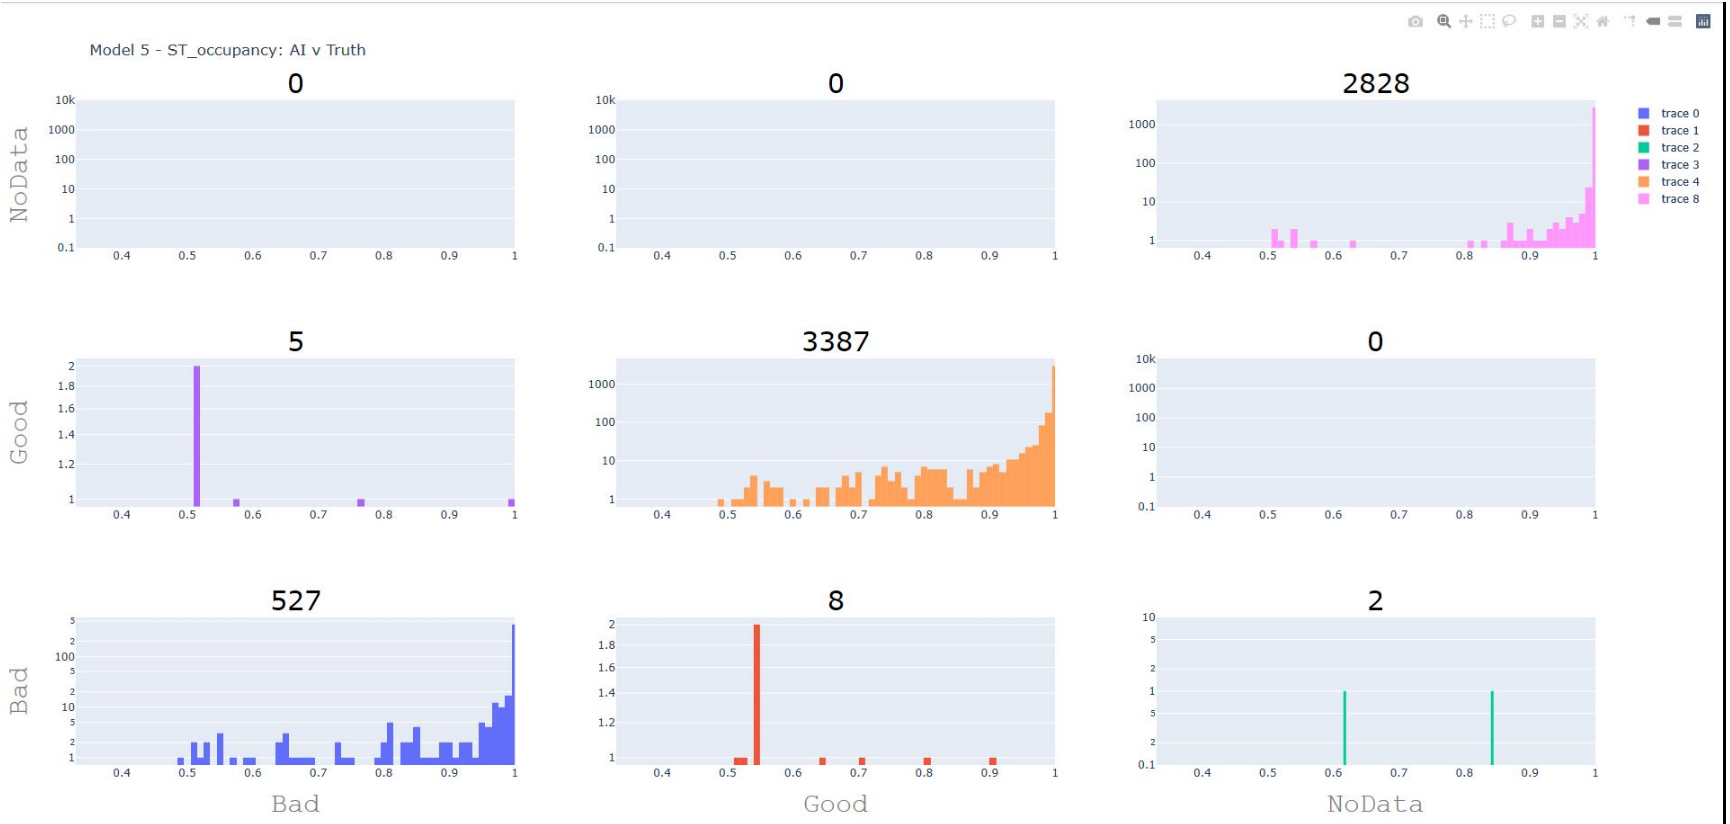

Results (FCAL)

Results (DIRC)

Results (BCAL)

Results (FDC)

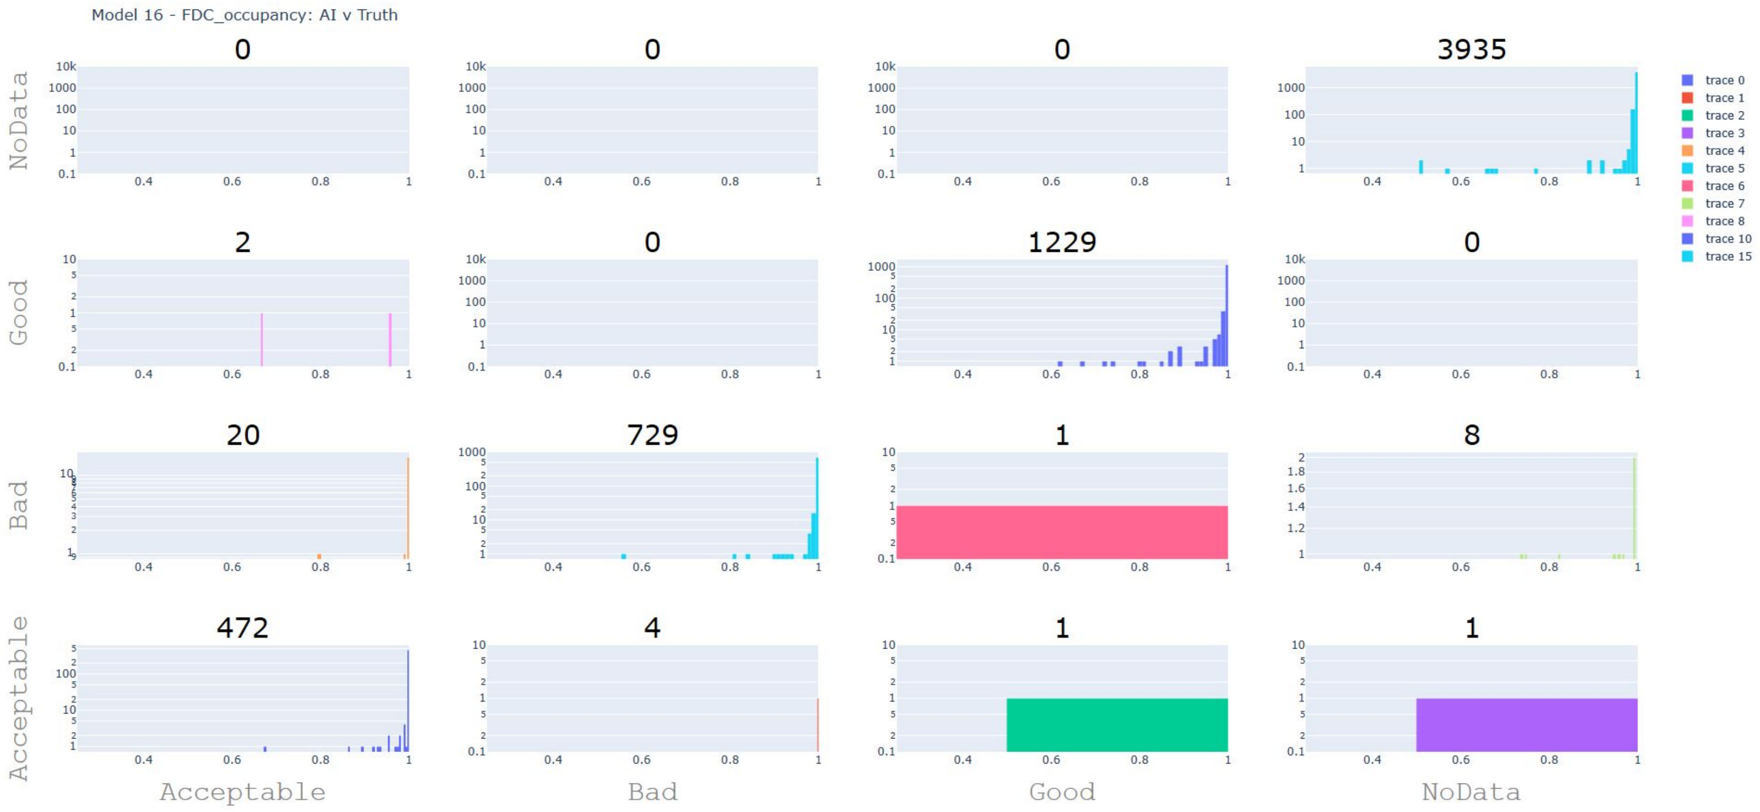
Model 16 - FDC_occupancy: AI v Truth
0
0
0
3935
2
0
1229
0
20
729
1
8
472
4
1
1
NoData
Good
Bad
Acceptable
Acceptable
Bad
Good
NoData
trace 0
trace 1
trace 2
trace 3
trace 4
trace 5
trace 6
trace 7
trace 8
trace 10
trace 15

Results (TOF)


Model 20 - TOF_occupancy: AI v Truth
NoData
Good
Cosmic
Bad
2
0
0
3115
0
0
7705
0
1
474
0
0
400
1
16
5
Bad
Cosmic
Good
NoData
trace 0
trace 1
trace 2
trace 3
trace 4
trace 5
trace 10
trace 12
trace 15