Data Quality Monitoring with AI: Hydra

Thomas Britton
David Lawrence
Kishansing Rajput
Torri Jeske

What is ML?

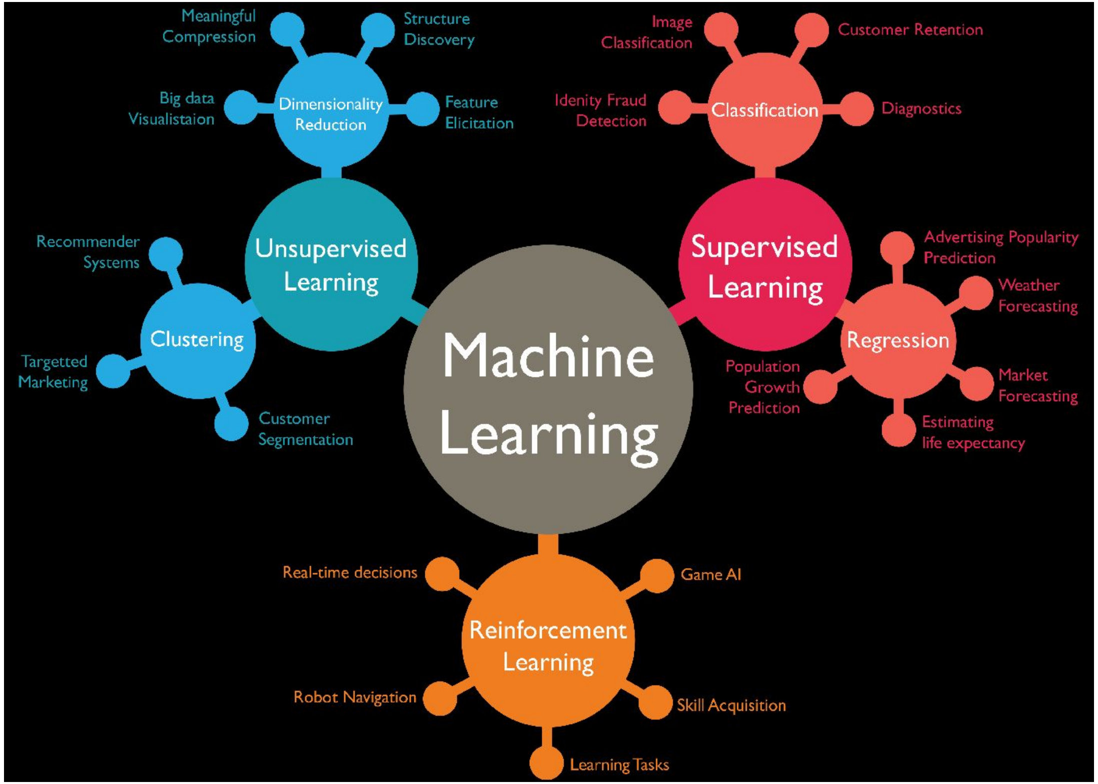

Machine Learning
Unsupervised Learning
Dimensionality Reduction
Meaningful Compression
Structure Discovery
Big data Visualistaion
Feature Elicitation
Clustering
Recommender Systems
Targetted Marketing
Customer Segmentation
Supervised Learning
Classification
Image Classification
Customer Retention
Idenity Fraud Detection
Diagnostics
Regression
Advertising Popularity Prediction
Weather Forecasting
Population Growth Prediction
Market Forecasting
Estimating life expectancy
Reinforcement Learning
Real-time decisions
Game AI
Robot Navigation
Skill Acquisition
Learning Tasks

Reinforcement Learning

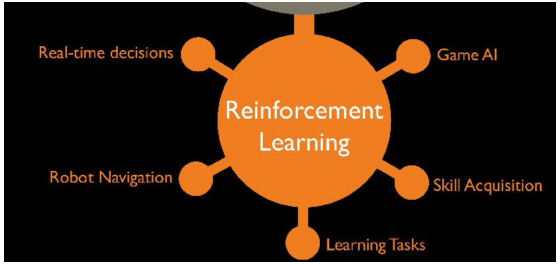

- Generally useful to solve problems with competing objectives
    - State dependent

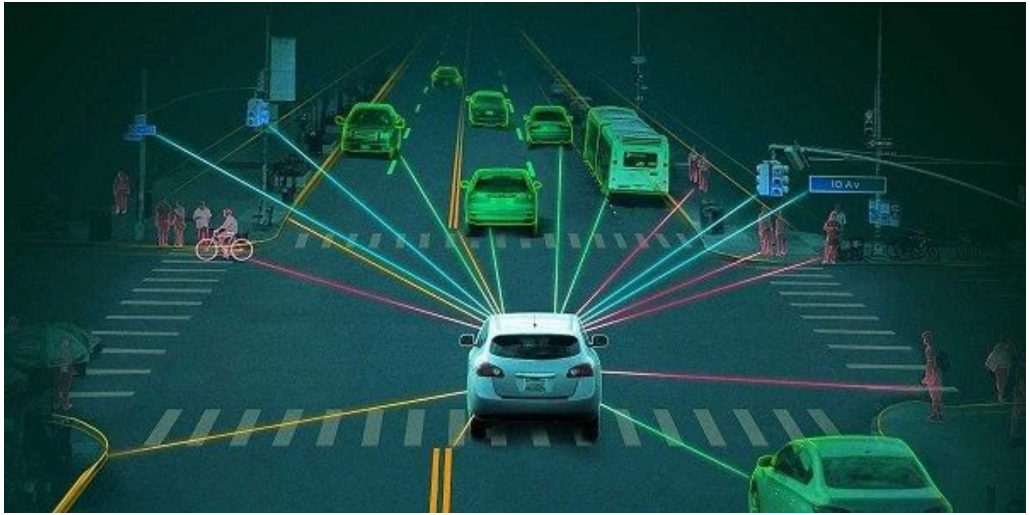

How it Works

- Can you make mistakes while learning?
  - Simulation
- Trying to learn an optimal policy
  - Given the **state** of things what **action(s)** can be performed to maximize my **reward**

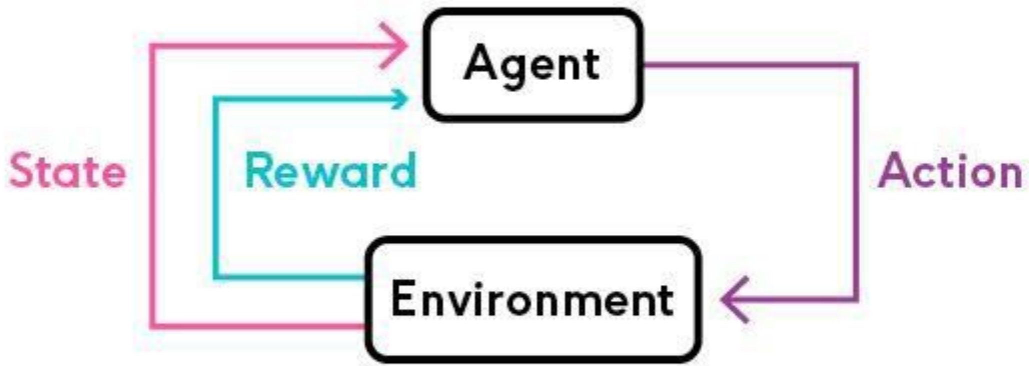

Unsupervised Learning

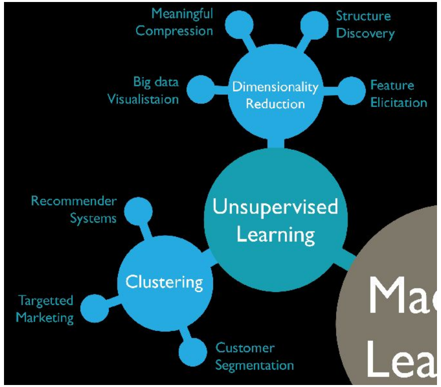

- Generally used for data which cannot be labeled or have no labels
    - Comparative in nature
    - Feature extraction (encoders)
    - Can be used to jumpstart labeling
- The reason people think big tech is actively spying
    - People are really easy to predict..

How it Works

- Think line of best fit!
  - Clustering (borders)

- Use features with a distance measure to determine the cluster it belongs to
  - Need to choose k-clusters
    - As k-clusters -> Nevents
      - You get cluster sizes of 1

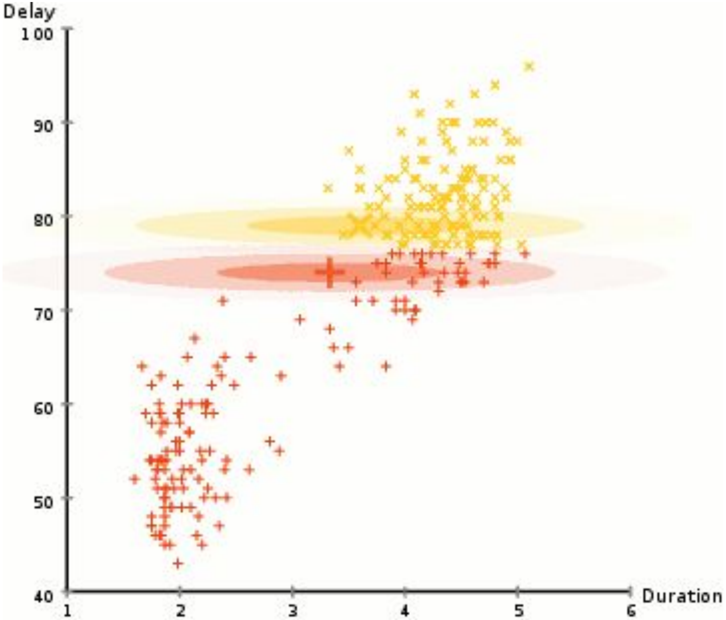

Supervised Learning

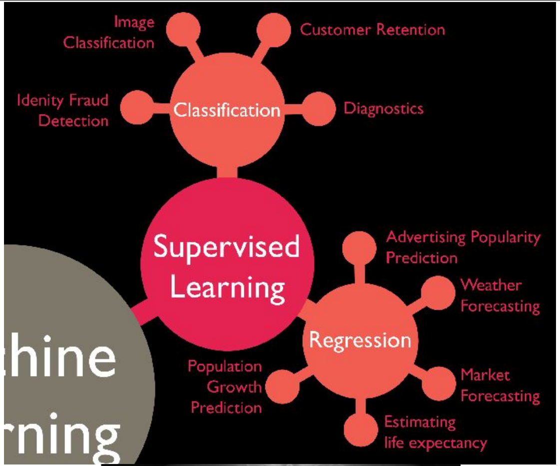

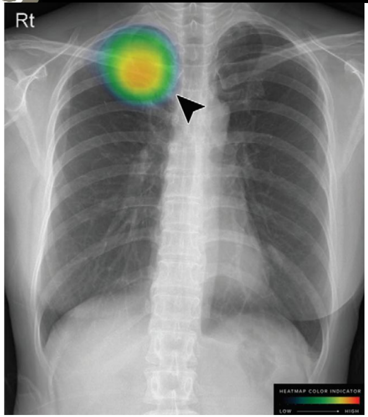

- **Expert driven systems**
  - Need labels
    - Can be expensive to obtain
  - **You've** been training them for years

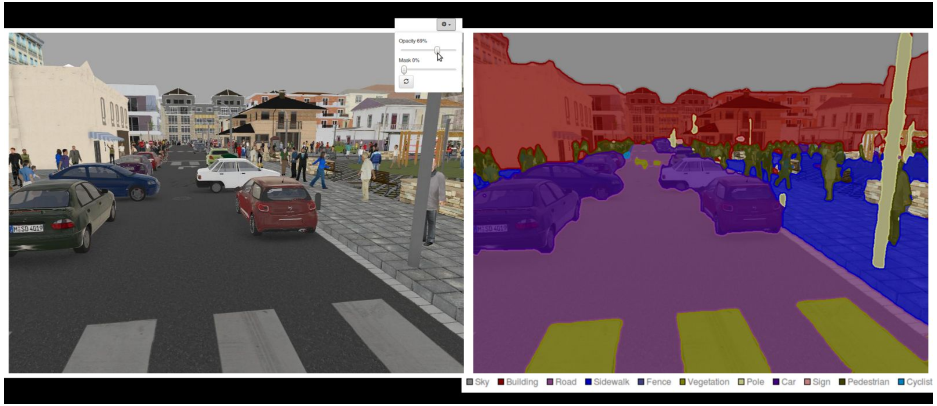

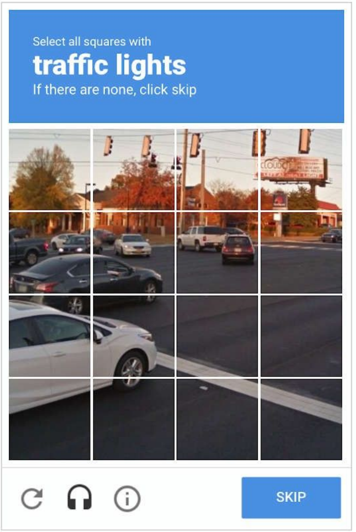

How it Works

- Two main branches
  - Classification
  - Regression
- Requires a loss function
  - How far off are you?
- Alter the trainable parameters to minimize the loss
  - Have to worry about all the typical fitting issues (local min)

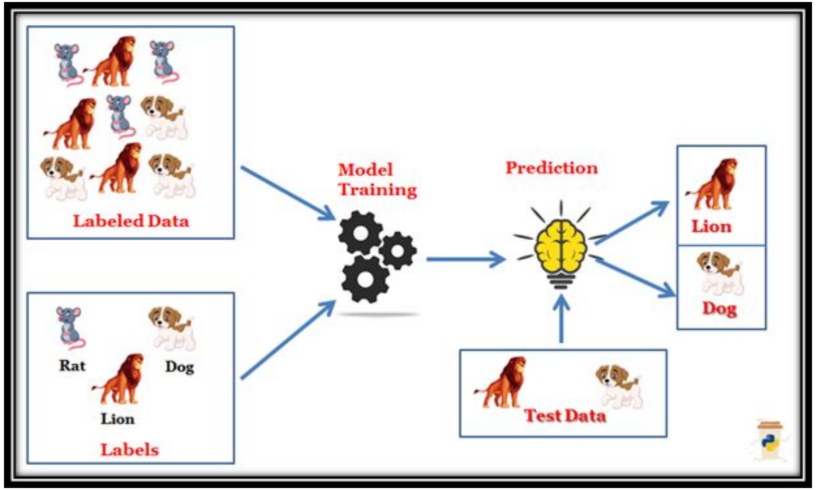

Neural Nets: Node Anatomy

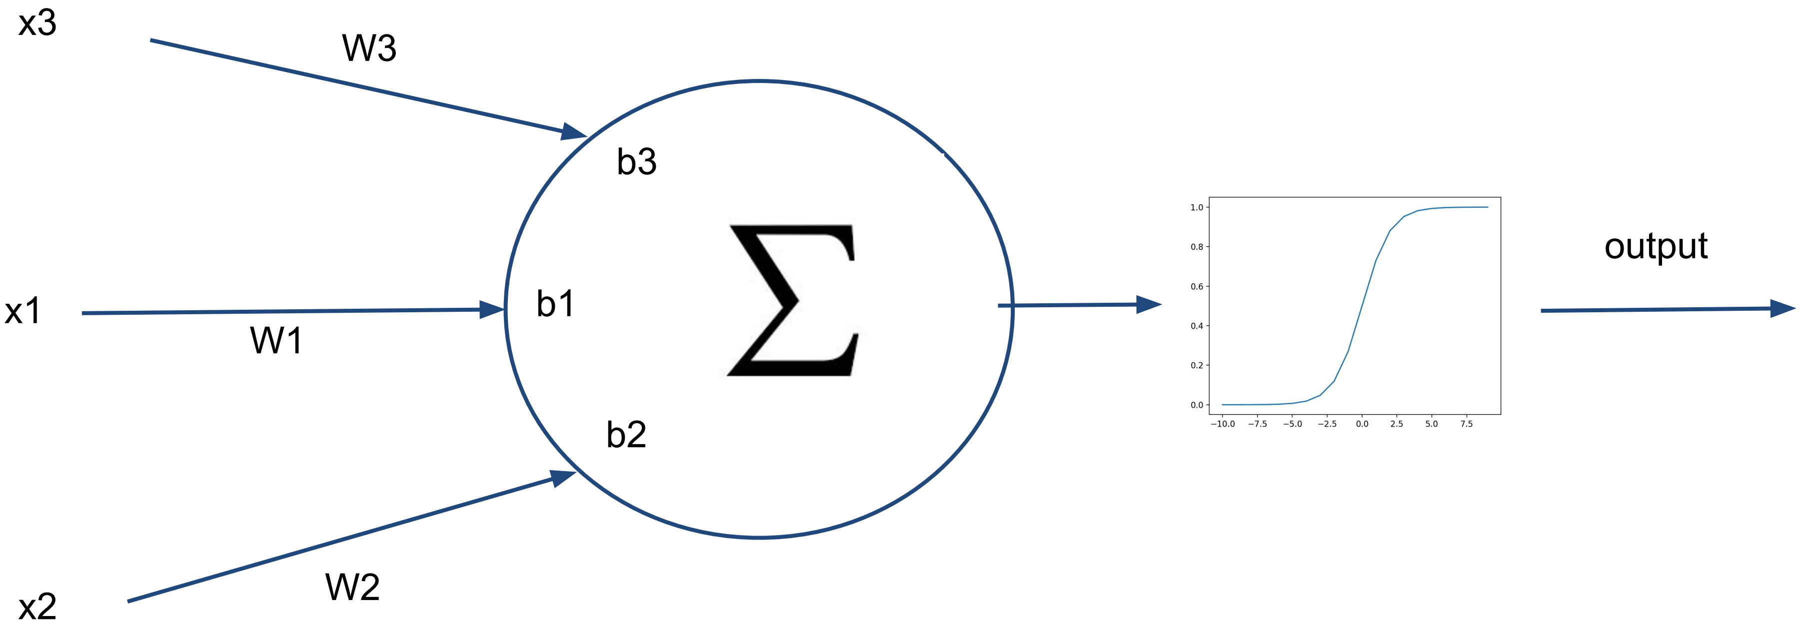

Neural Nets

Feed forward

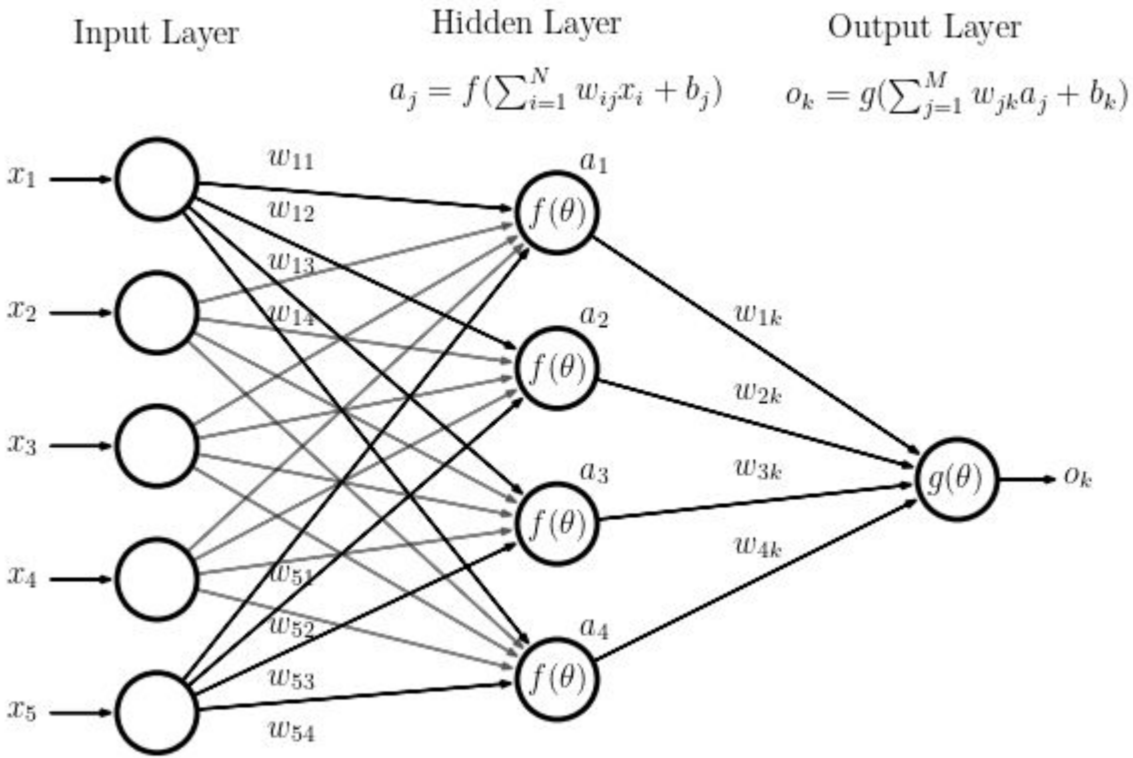

Backpropagation

$$\frac{\partial C}{\partial w_{jk}^{l}} = \frac{\partial C}{\partial z_j^l}\frac{\partial z_j^l}{\partial w_{jk}^{l}} \qquad \textit{chain rule}$$

$$z_j^l = \sum_{k=1}^{m} w_{jk}^{l} a_k^{l-1} + b_j^l \qquad \textit{by definition}$$

$$m = \textit{number of neurons in } l-1 \textit{ layer}$$

$$\frac{\partial z_j^l}{\partial w_{jk}^{l}} = a_k^{l-1} \qquad \textit{by differentiation (calculating derivative)}$$

$$\frac{\partial C}{\partial w_{jk}^{l}} = \frac{\partial C}{\partial z_j^l} a_k^{l-1} \qquad \textit{final value}$$

Equations for derivative of C in a single weight (w_jk)^l

Modify the weights (and biases) based on the loss value backwards layer by layer node by node

Look it's a gradient!

https://towardsdatascience.com/

CNN (math)

https://towardsdatascience.com/

$$G[m,n] = (f * h)[m,n] = \sum_j \sum_k h[j,k] f[m-j, n-k]$$

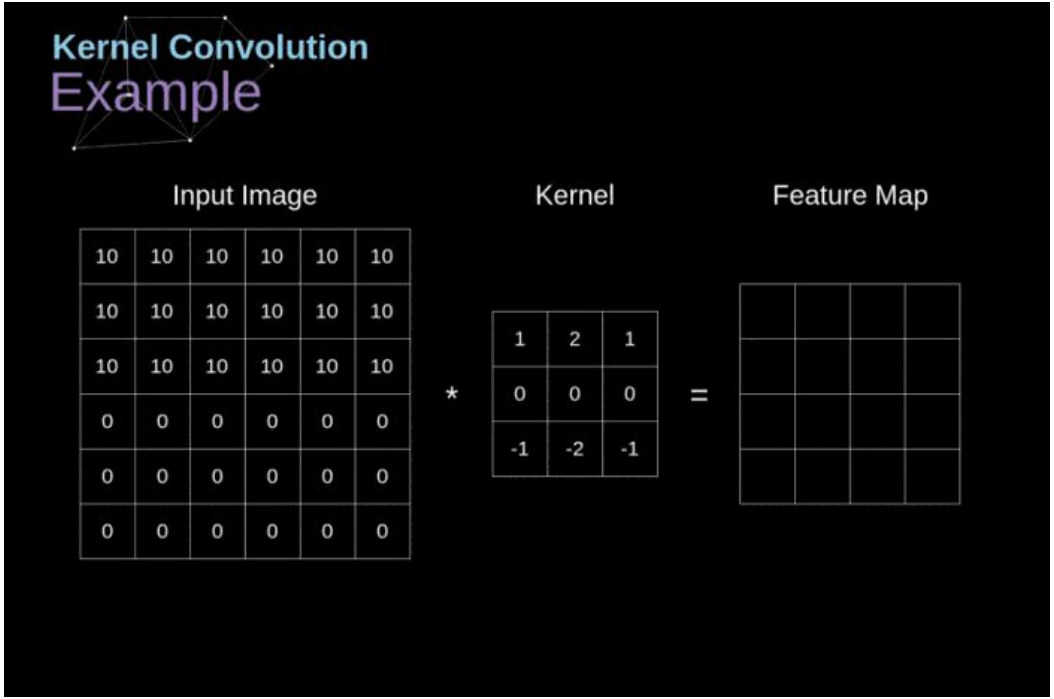

* stride of 1

CNN (kernels)

https://towardsdatascience.com/

- **Different kernels produce different outputs** (features).
  - so you are training a CNN layer to extract different features from the data

What is DQM?

- Modern Detector systems are **highly complex and interconnected**
  - Narrow operational envelope

- **Need multiple levels of checks and controls**
  - Unknown unknowns

- Not necessarily only for uncorrectable issues….

Monitoring: From Data Acquisition to Analysis

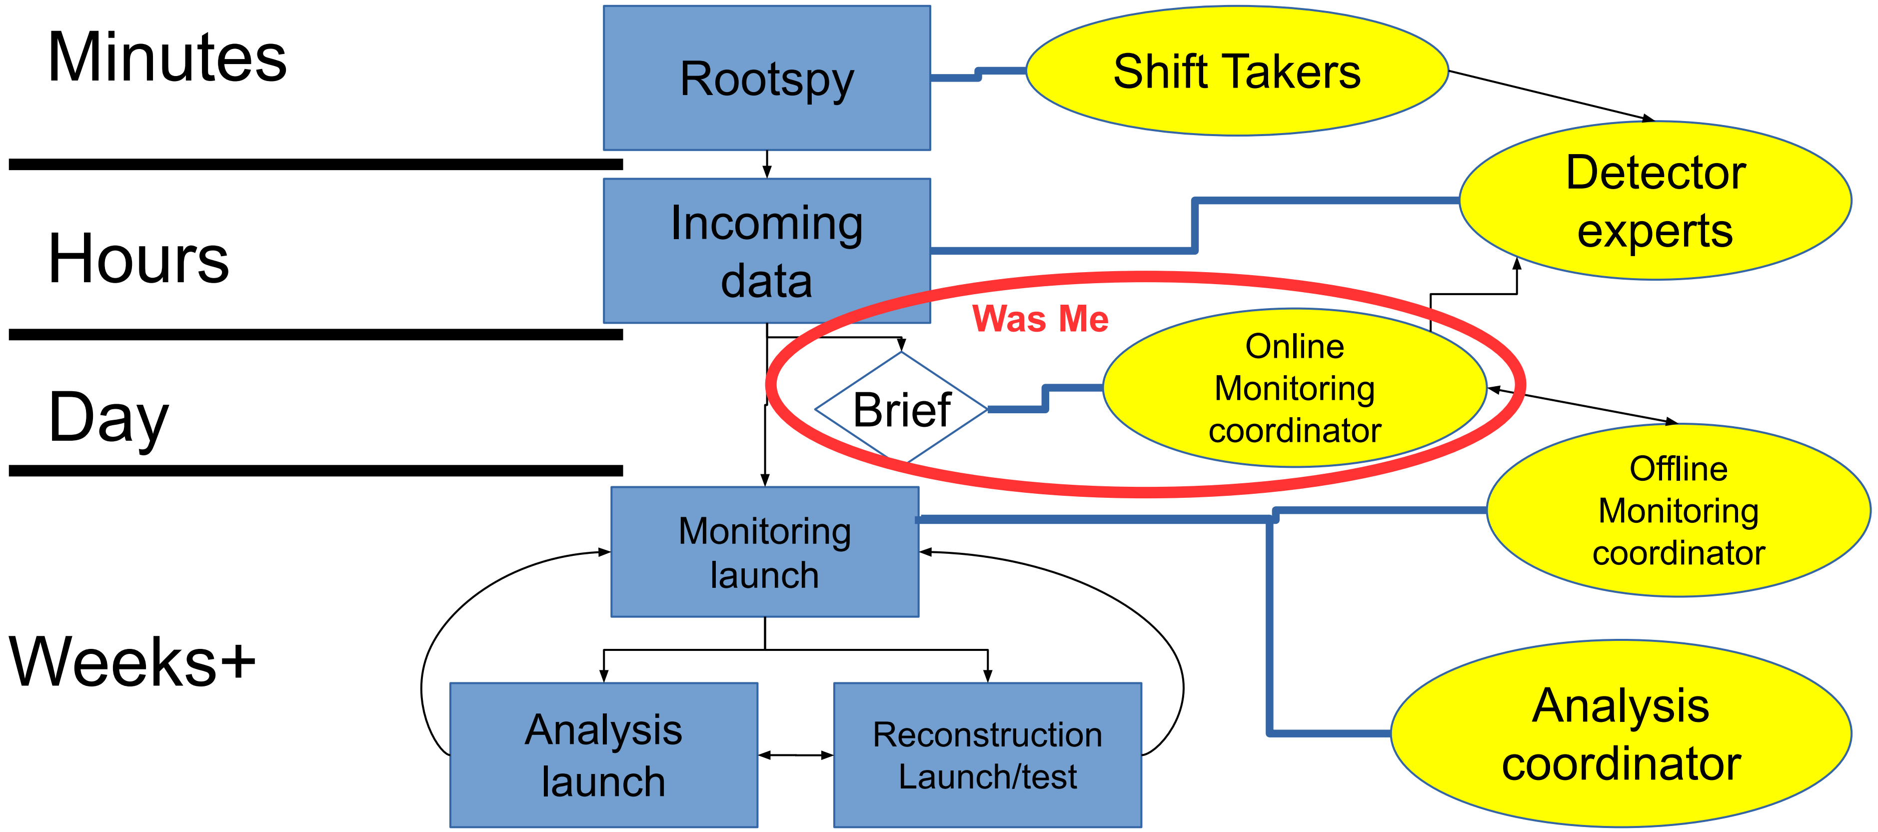

Let's Play

Ladies and gentlemen: the story you are about to hear is true. Only the names have been changed to protect the innocent.

*Logbook searches not included

Questions to keep in mind

What is the problem? Is it isolated? Is this due to running conditions or a test? Should we stop taking data/call the shift crew?

It's 5:15am the email has just come in...

Disclaimer

- What follows involves many plots from many runs at GlueX
  - **Occupancy plots** - rate at which detector sensors fire
  - **Beam properties** - Normalized energy spectrum and beam bunch frequency
  - **Physics signals** - fully reconstructed particles
  - **Trigger status** - which trigger conditions are being met and live time

- Not shown are the logbook entries and ~20 other plot types checked

Now let me guide you through an abridged 5:15am…

Let’s Play

BCAL occupancy changed!

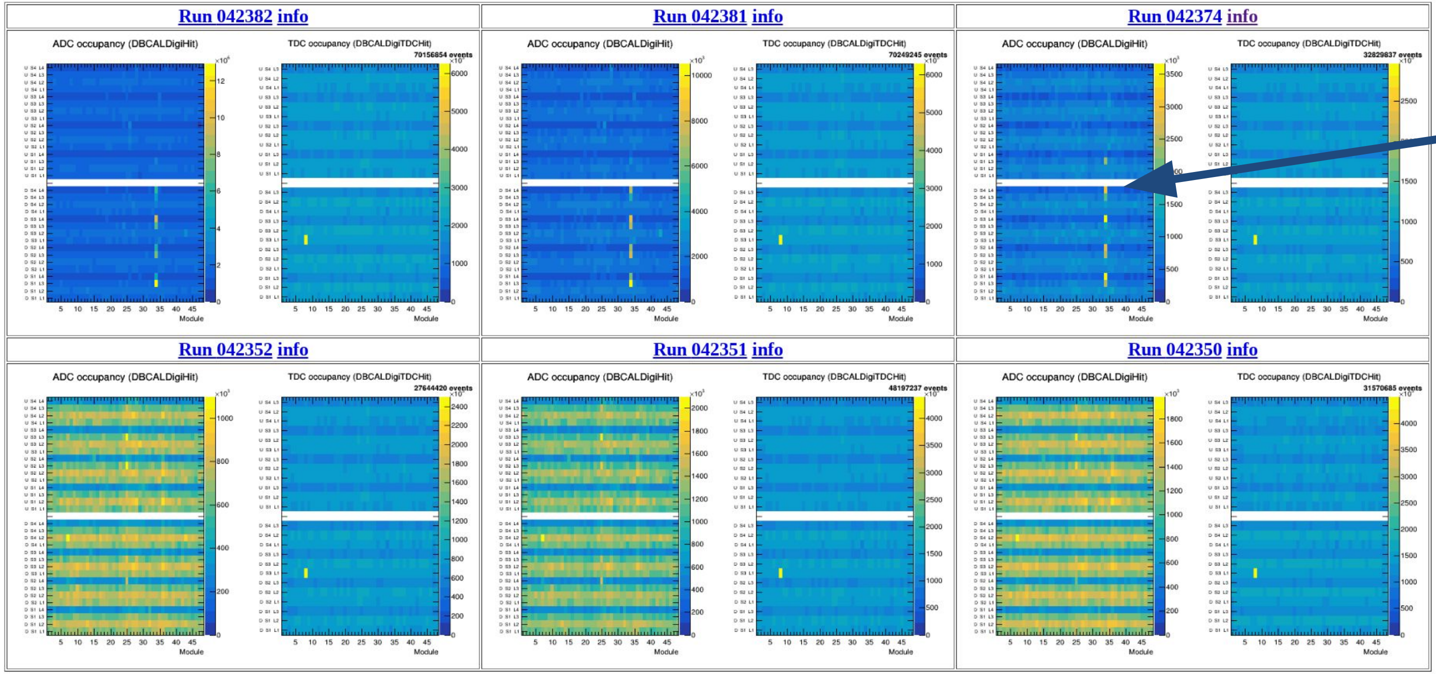

Anomaly detected!

Let's Play

Nothing looks really off with the beam

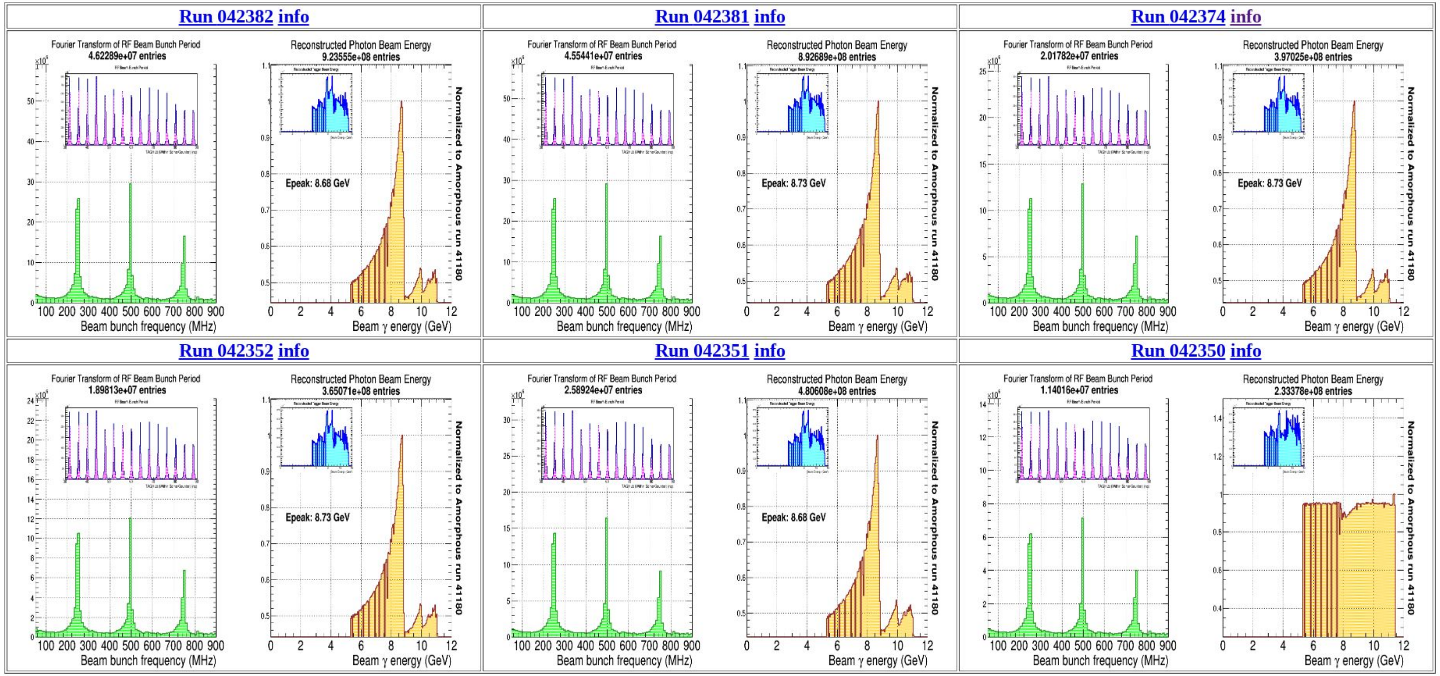

Let's Play

CDC "dims" between 42350->2351….hmmmm. Remember 42351 from the BCAL looked normal? Probably uncorrelated…..

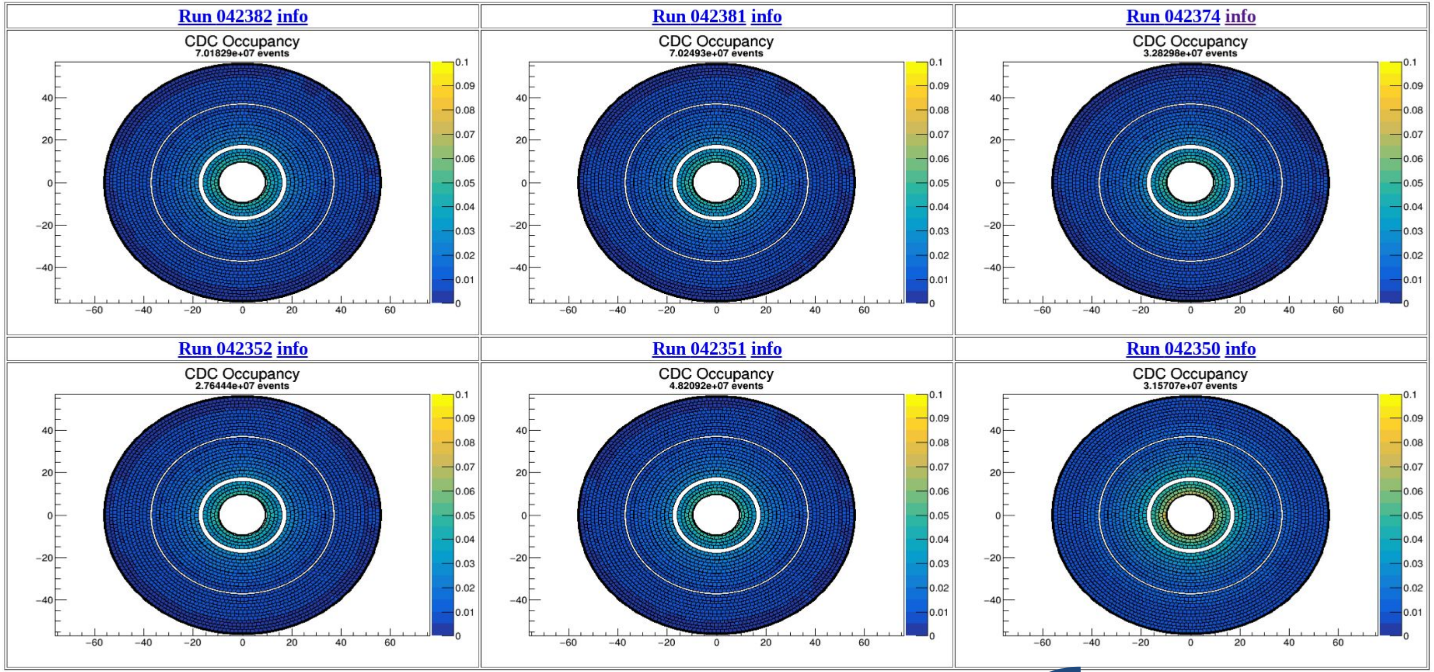

Let’s Play

FCAL looks normal. Note a few new dead modules….

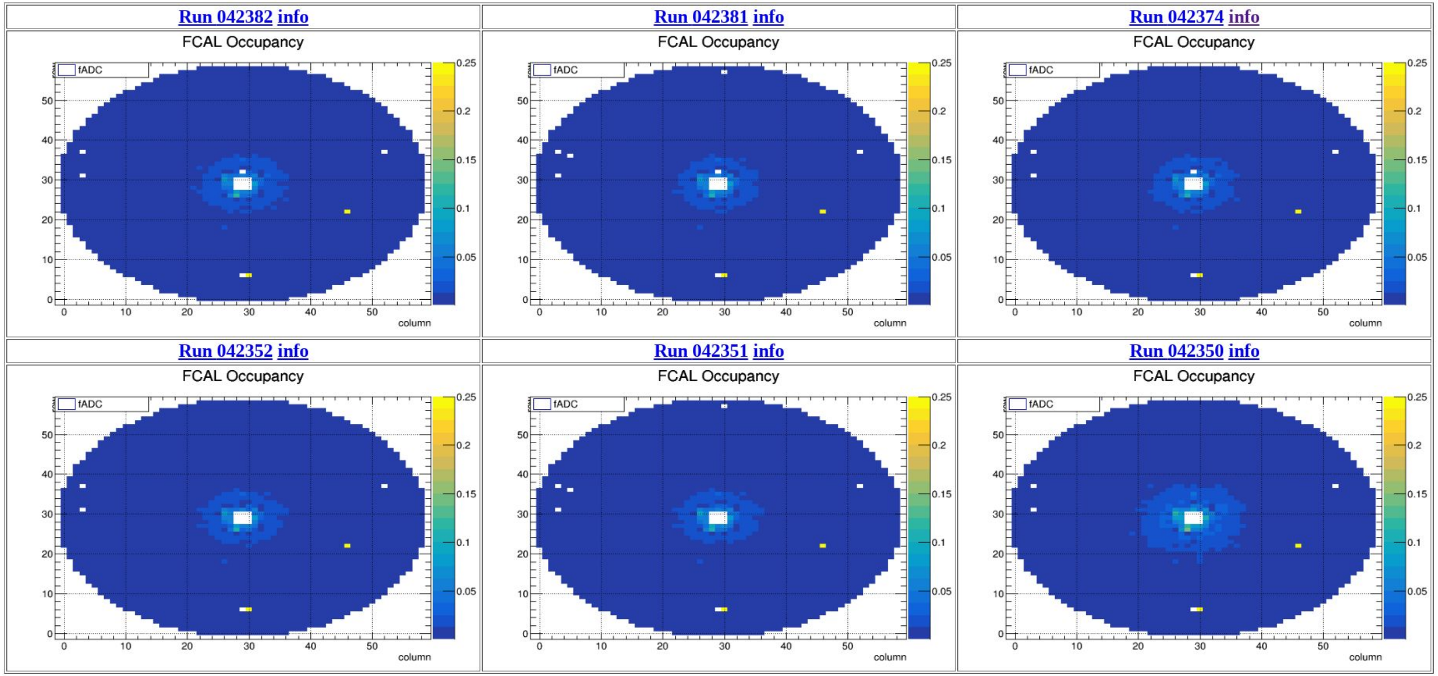

Let's Play

13 plots later….
Physics signals still in tolerance

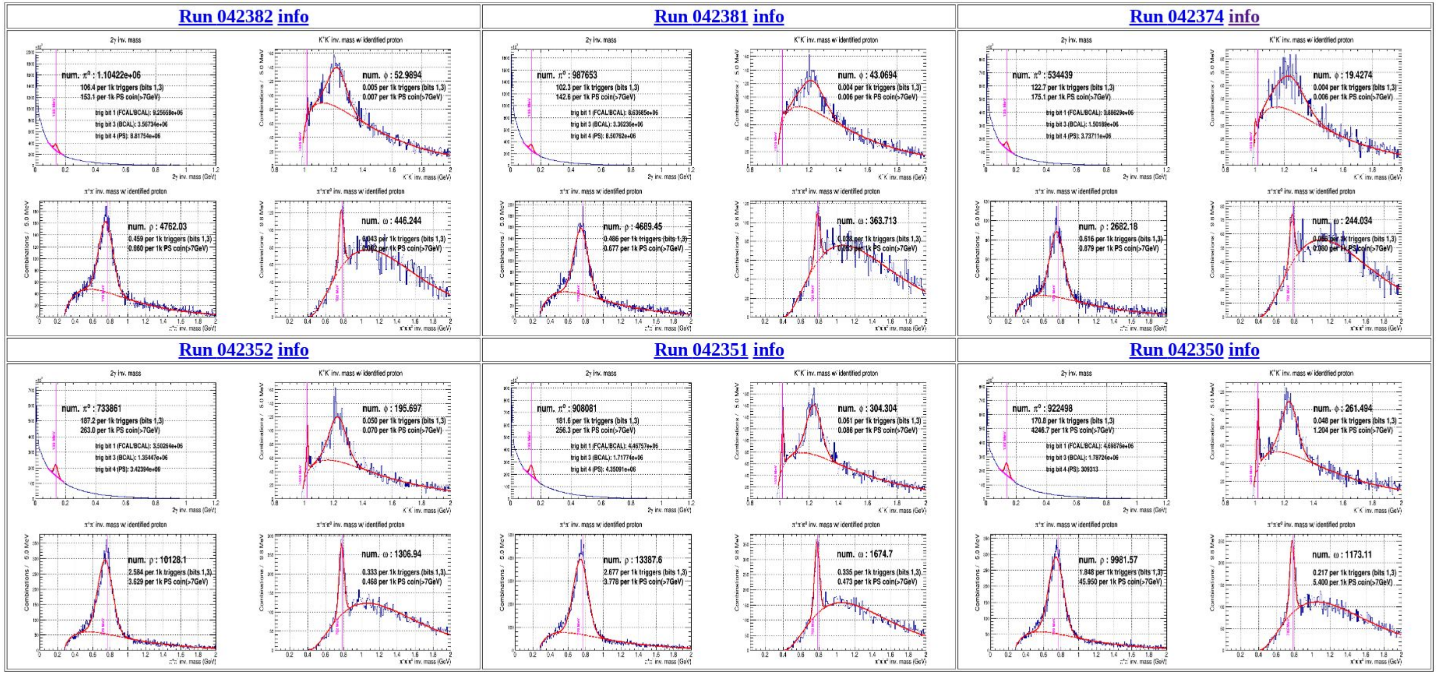

Let’s Play

5 plots later...

Increased dead time at 42351

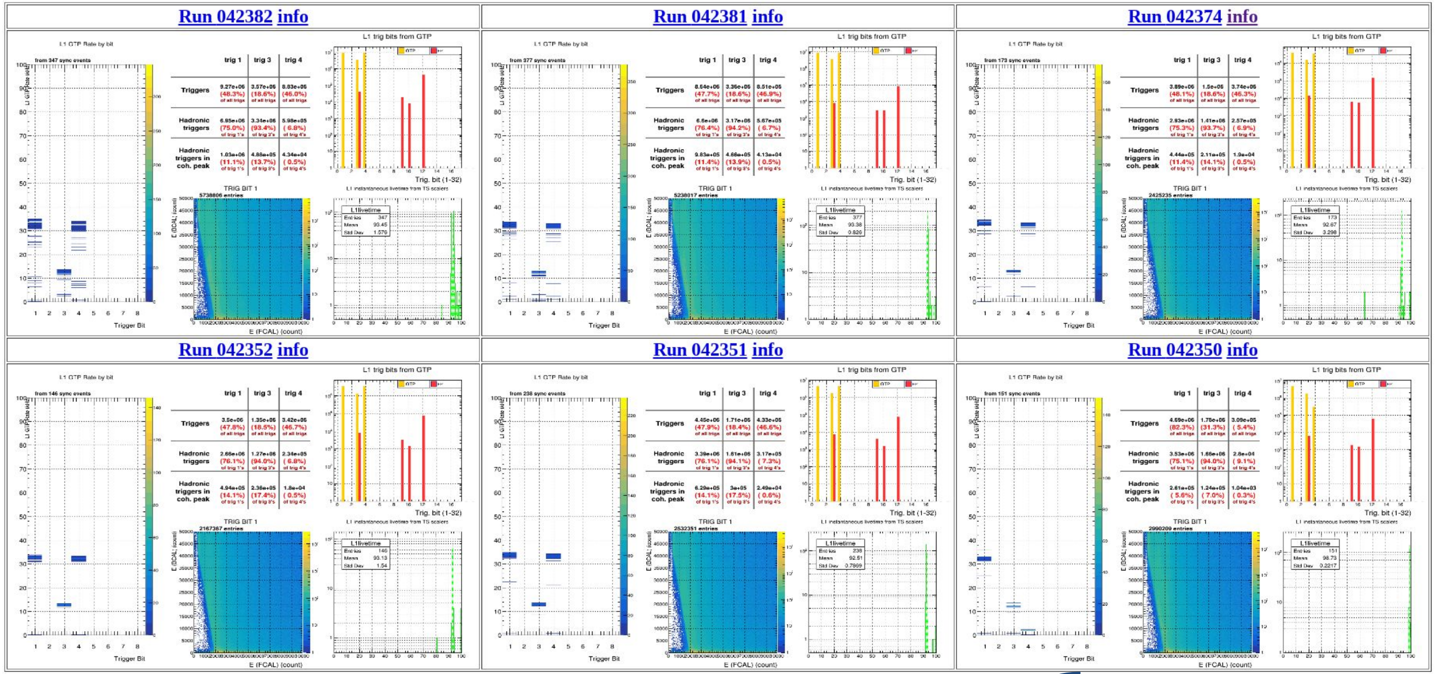

Conclusion

Likely, several BCAL modules have become noisy contributing to an increased deadtime. The reason for occurrence is unknown, but it likely happened <u>during</u> run 42351. It is not significantly impacting physics data. Detector experts should be notified and further investigation should be performed. An opportunistic access to the hall may be needed

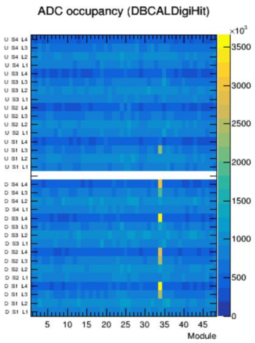

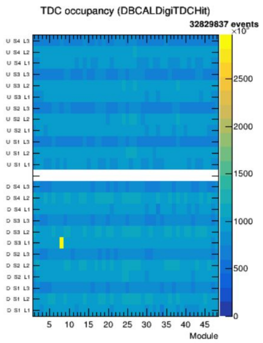

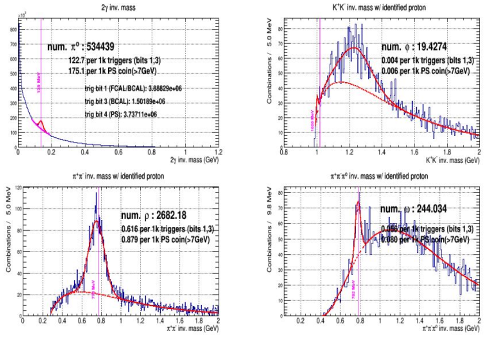

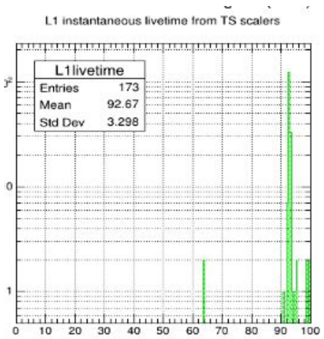

Epilogue:

In roughly 24 hours the noisy components were found and replaced. Normal operations resumed

Issues

- Plenty of issues that are not alarmable in the traditional sense but detectable

They suspect that the red light in the FCAL is left on since yesterday

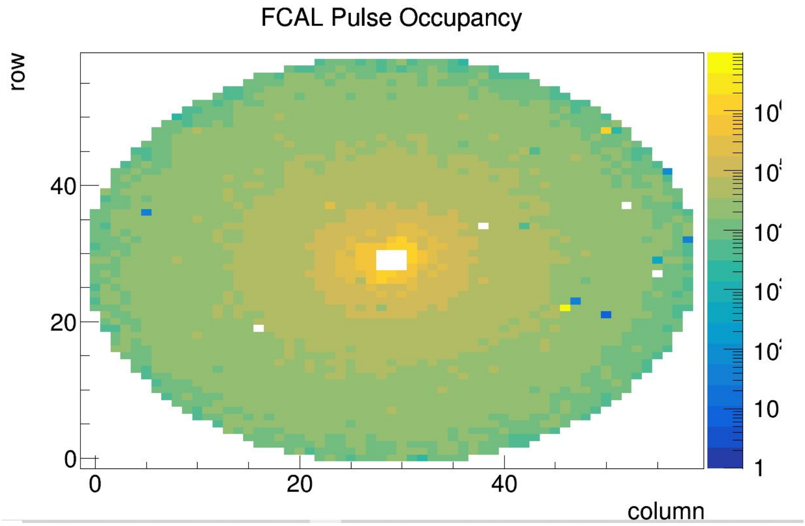

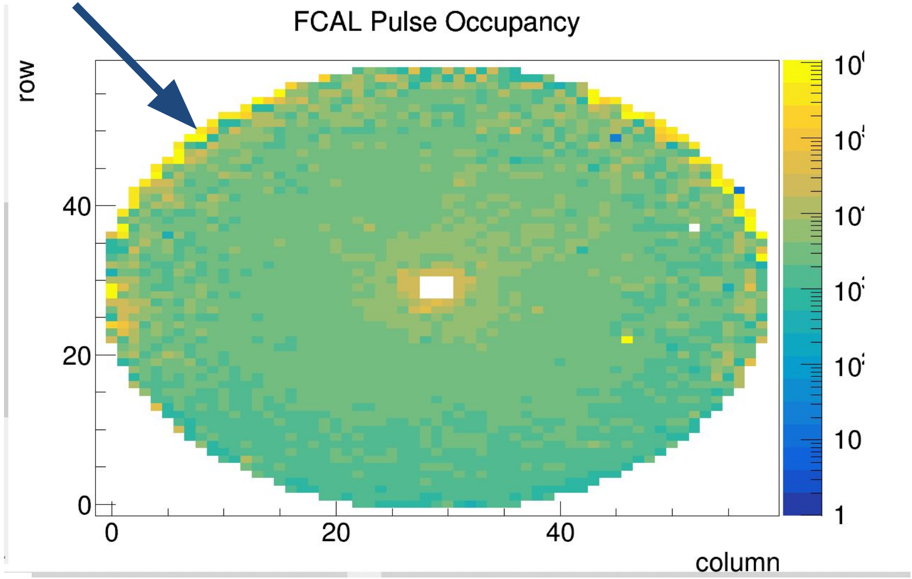

The Challenge

- **Every run** produces an initial **22 plots**. More thorough monitoring is performed offline and produces **109 plots**. With a run lasting **~2-3 hours** every day there are **between ~175 and 875** plots to look at.
  - To preserve sanity I looked at closer to 175 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
  - <u>Trigger studies</u>: Often look like big problems but are not. Can be hard to catch when shift logs have scant details

A better way?

- Advances in hardware and proliferation of frameworks have lead to rapid advancements in Machine Learning
- If AI can detect cars, people and read signs in real time then it should be possible to build a system to handle monitoring….

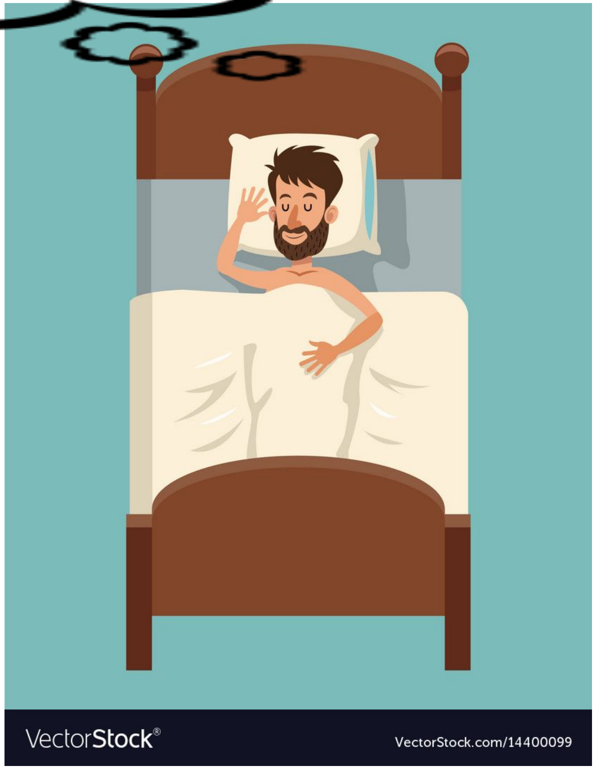

AI

- What is involved is a lot of **pattern recognition** and paying attention to the **correlations** between plots coupled with **domain knowledge**
  - When plot A looks like X plot B should look like Y but it looks like Z and by looking at plot C probably Q is wrong...
- This is precisely where **AI can shine**
  - Large, multi-dimensional datasets full of correlations

Considerations

- Many considerations in developing and deploying AI models
  - Data collection:
    - How will you get your **training data**?
    - How will you **test the model**?
  - Labeling (in the case of supervised learning):
    - How will you **label all of your data**?
    - What will that **cost**?
  - Model development:
    - What **requirements** does the model need to fulfill? (e.g. inference time, error rates etc)
  - Sociological:
    - How will **people interact** with the system?
    - Are there any **barriers to deployment** from the human side?

Considerations (cont)

- Can we leverage computer vision?
  - I had 0 experience with AI
  - Interested in computer vision
  - I performed DQM via looking at images

- Are there off the shelf technologies I can employ?
  - Side project
  - Lack of dedicated resources
  - No experience!

Data Collection

- Fortunately, I already had all the data I would need to get started
- GlueX already had monitoring plots from all the detectors from all the runs in image format

- Using these is slightly wasteful in terms of data needed for inference

- Get to use all of the developments in machine vision
    - Convolutional Neural Nets

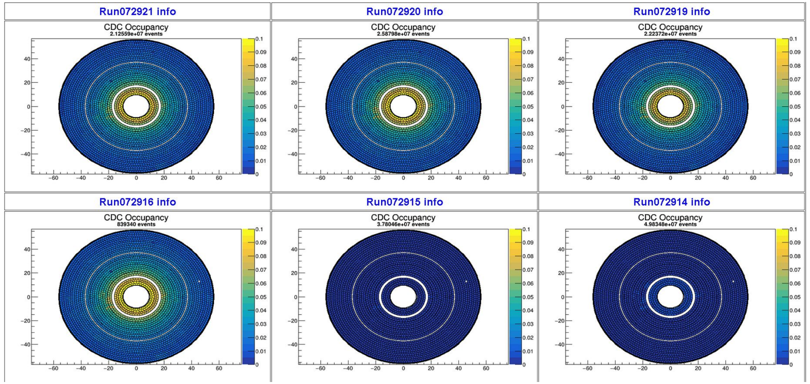

- <u>Cuts down on startup time</u>

Labeling

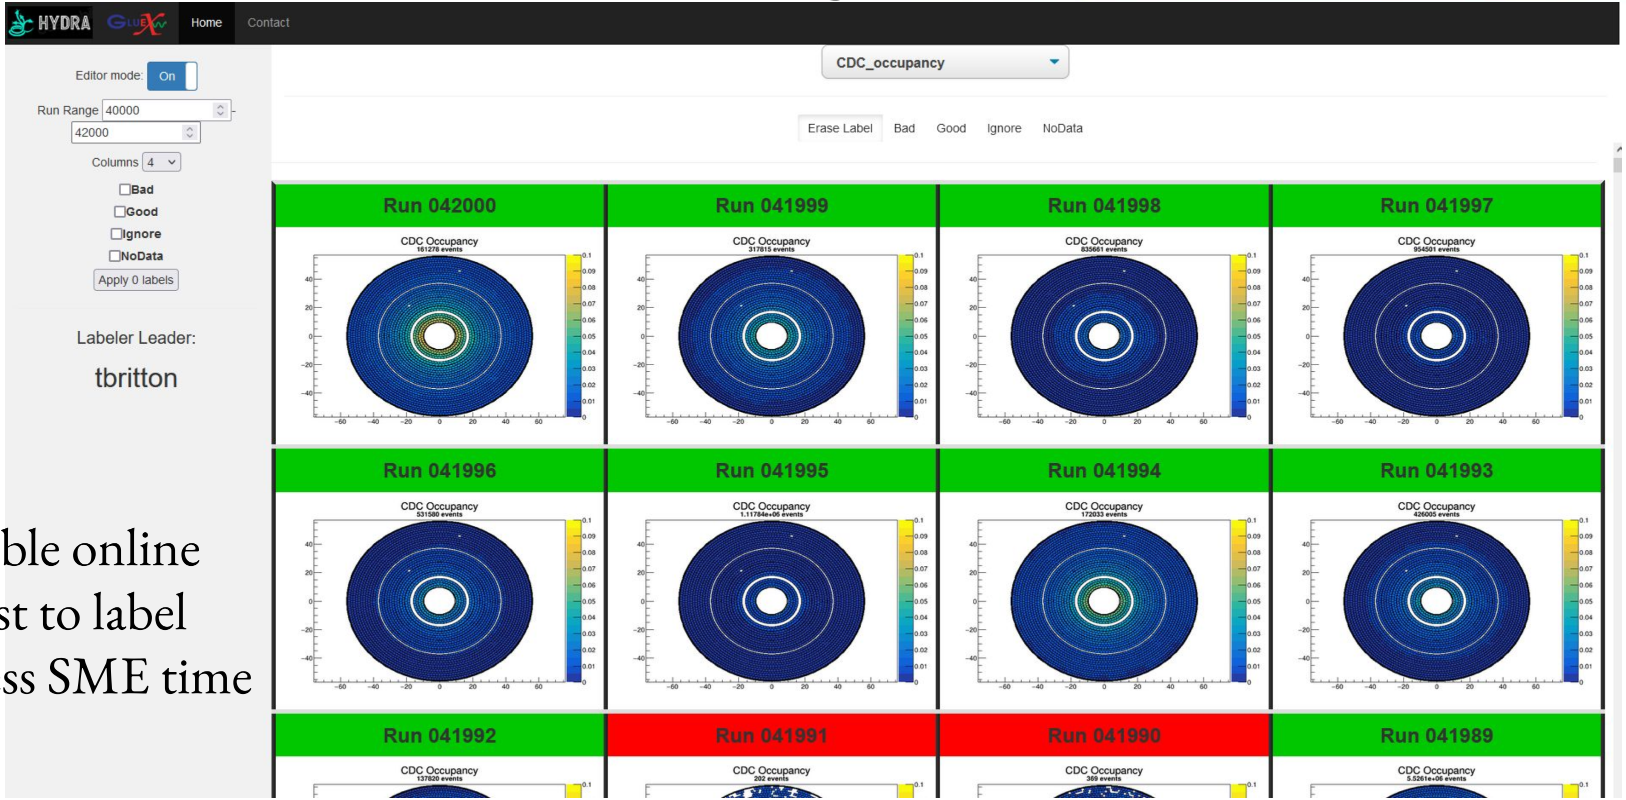

- Accessible online
- Very fast to label
  - Less SME time

Model "Development"

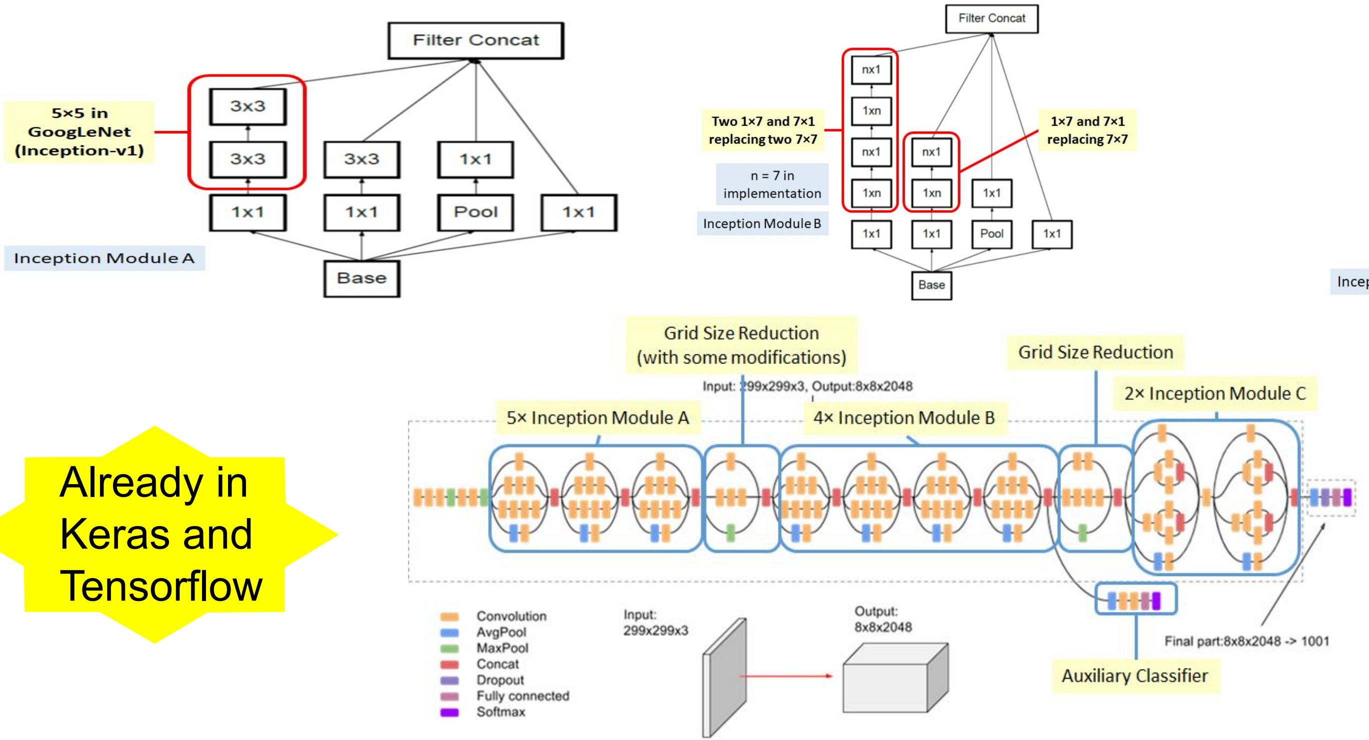

Already in Keras and Tensorflow

The Inception v3 Network

Challenges in Training

- **Unbalanced data sets**
  - Think 80% “Good”, 10% “Bad”, 10% “Other”
  - Can lead to biased inferences
    - Returning to the mean (weighted)

- **Lack of examples**
  - Especially in problematic states
  - Unknown unknowns

Unbalanced Labels/Small Datasets

- One standard method is to **super sample**
  - Repeatedly show example(s) of lesser labels to mimic having balanced labels
- **K-fold**
  - Quarantine a test set and repeatedly train a model using K blocks of samples with each block being used in K-1 training sets and exactly once in validation

- These methods increases training time, sometimes quite significantly

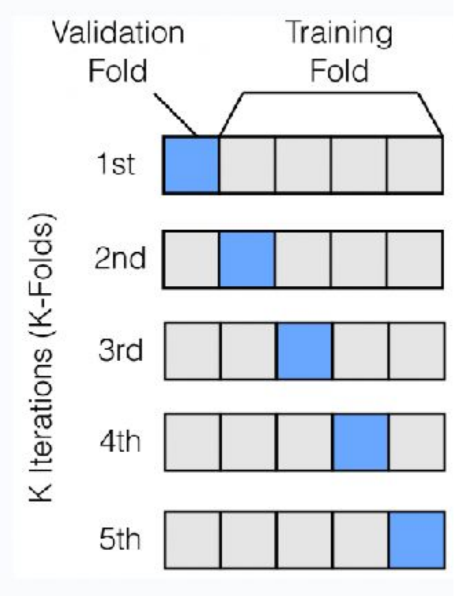

Strategic Undersampling

- Narrow operational envelope means that when you've seen one "good" plot you've seen them all
  - Ie there are **many many** ways it could be wrong, but a smaller subset of ways it will be right

- Can exploit the relative plot variance to randomly undersample the lower variance labels to equal the smaller, higher variance labels
  - Increases training speed by **~5x** with no loss in model accuracy

Lack of Examples

- Can use “simulated” issues **Useful for overfitting checks!**

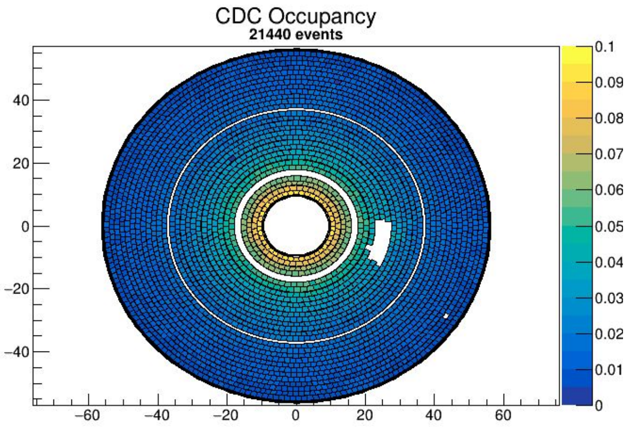

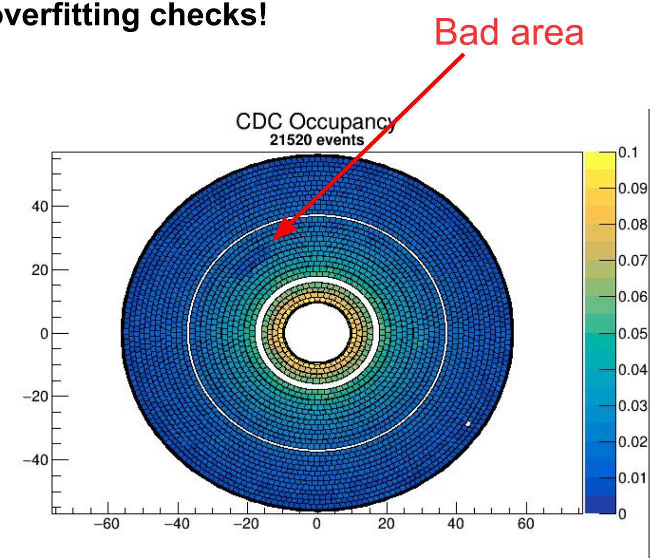

Model was **able to use a small fraction of single board issues and generalize it** to all single board and double board issues

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
- Most importantly, Hydra allows me to embrace my inner sloth:

Koboldpress.com

Where Hydra Lives

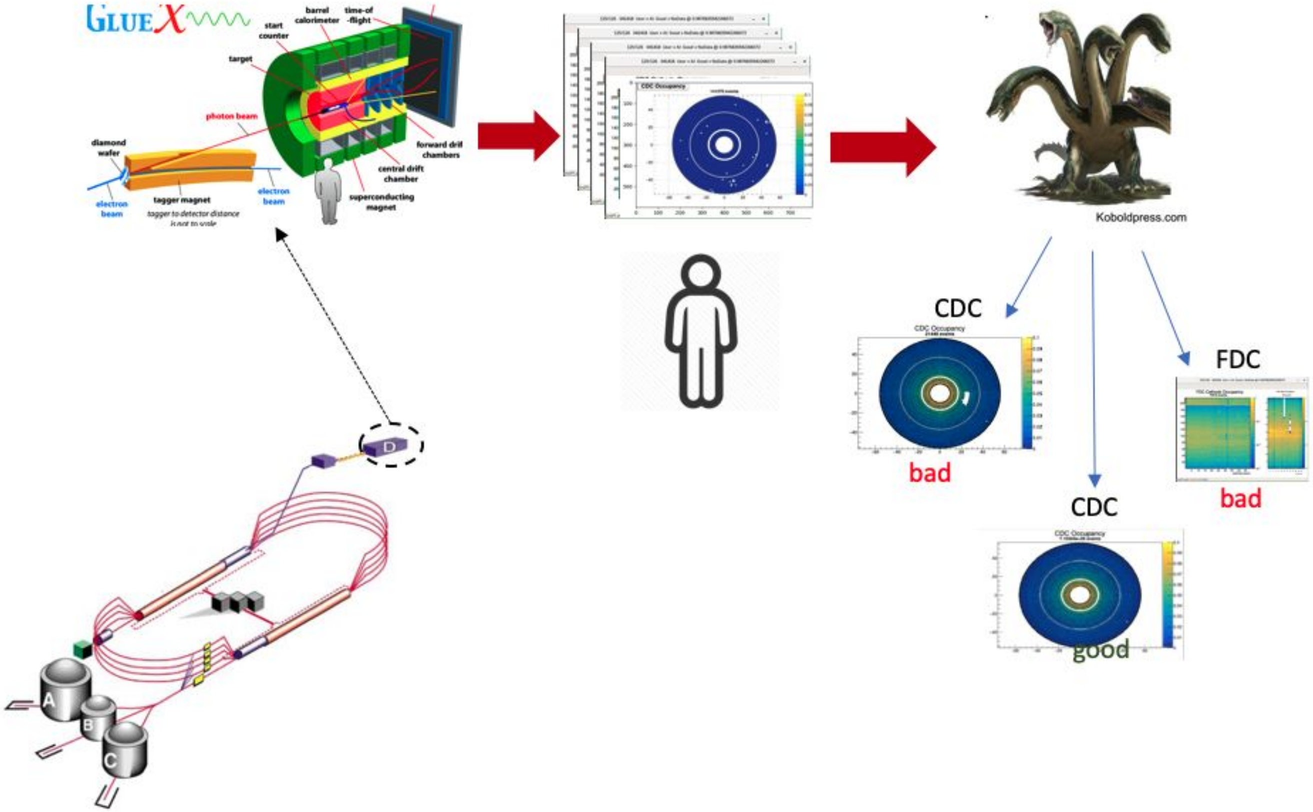

The Backend

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

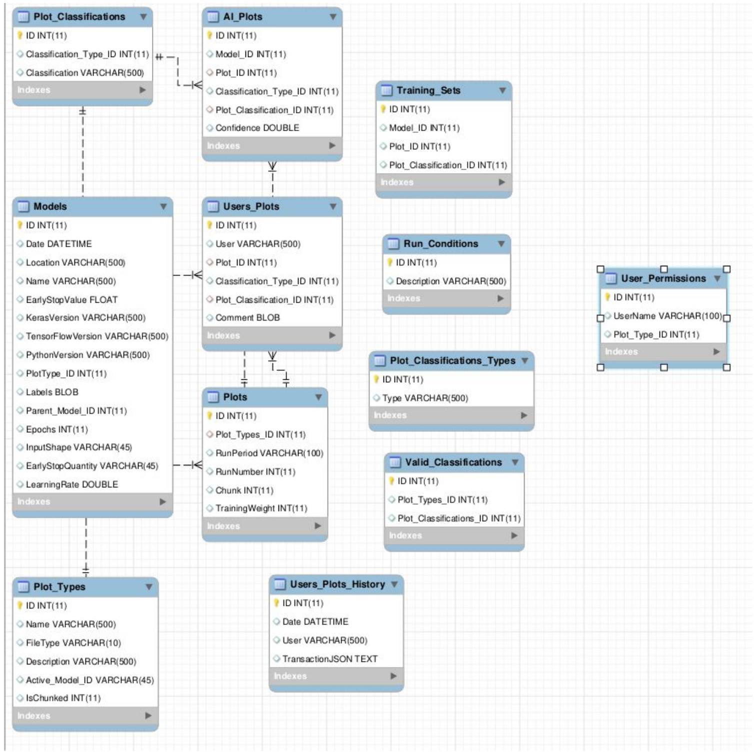

- Training is virtually push button to allow for automated retraining as needed

E.g. CDC Results

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

E.g. CDC Results

Hydra

Experts / Truth

thresholding!

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

An Anecdote

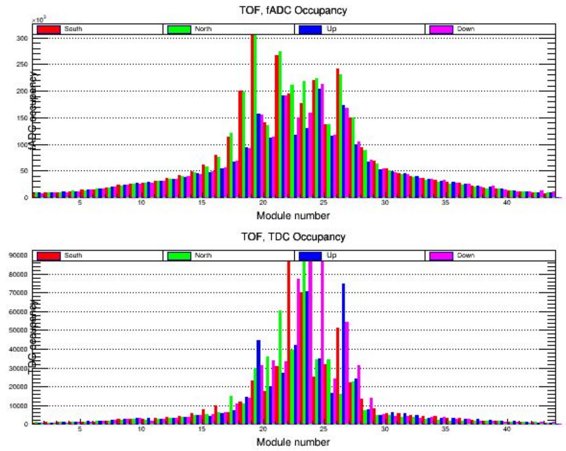

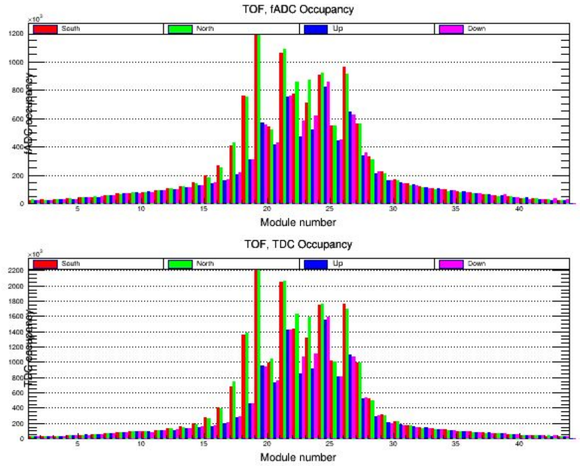

- Both of these look “good” at first glance (both initially labeled good)
  - The one on the left is actually bad (the A.I. caught it)
- A.I. seems to be able to look at subtle differences in shape and **maybe even “read” the y-axis**

Another Anecdote

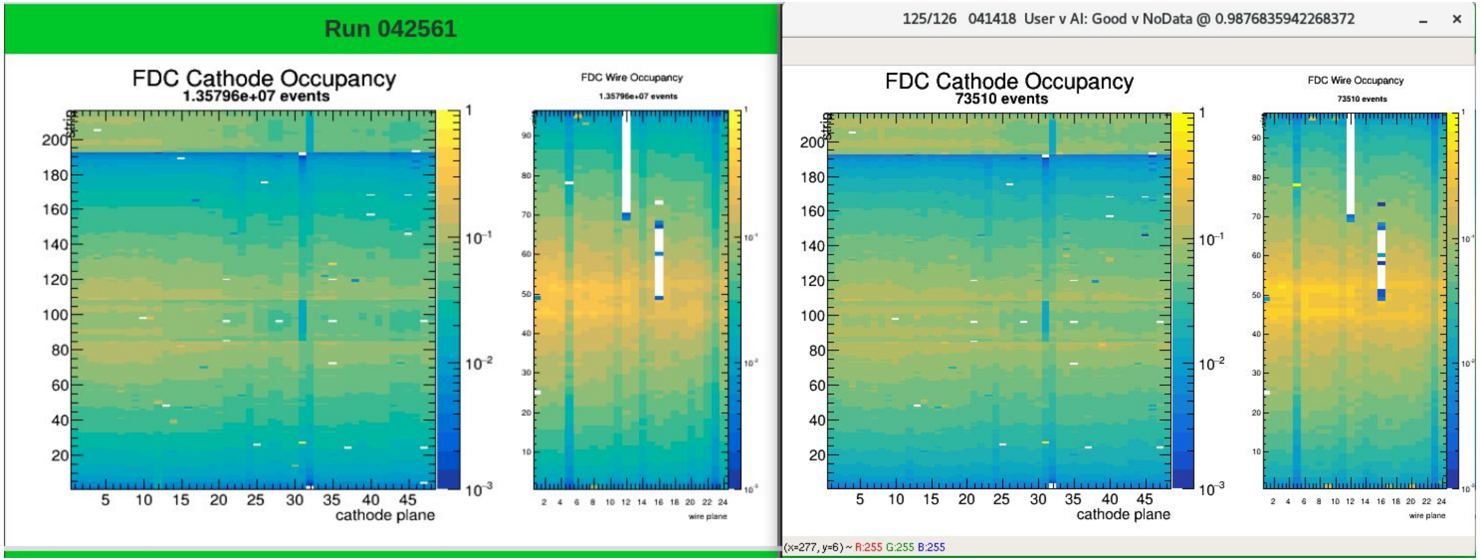

- The labeler was instructed by the detector expert to label any plot containing **fewer than 100k events as "NoData"**. This is one example of several in which the labeler labeled as "Good" and the **A.I. predicted "NoData"**…the true label given the number of events

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

Limitations

- Need broad buckets
  - Okay for now as most important are the “go no-go” decisions
  - Lower diagnostic capabilities currently

- DL requires a **LOT** of data (sometimes)
  - Many examples of every problem
  - ML-ops! (beyond the scope)

- **What can we do to develop the capabilities?**
  - Generic anomaly detection
  - Multi-modal data

A Perk for Using Images....

Gradient Class Activation Map (GradCAM)

- Enables us to identify regions of image important for classification
  - Dog vs cat:
    - The dog's face is important to the network when deciding if it is a dog or cat
  - Dog Breed:
    - Region important for pit bull (husky) classification
- For Hydra, can provide more context into network's decision making in regards to 'Good' or 'Bad' images from various detectors

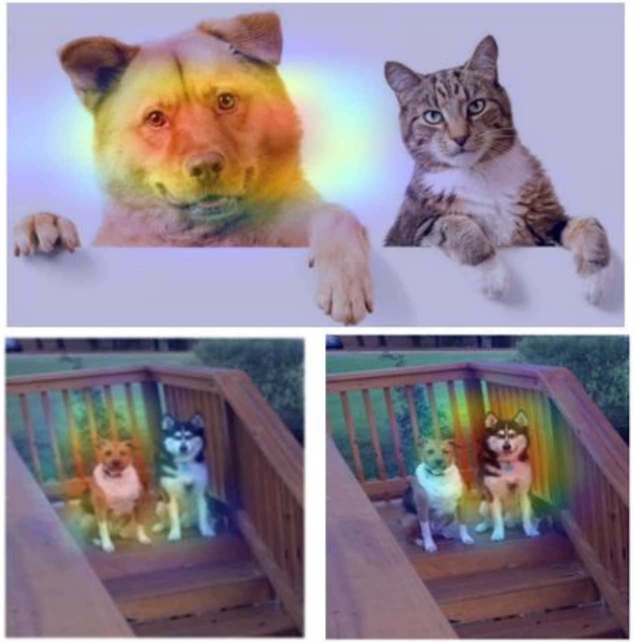

Layered Approach

- There are many **generic anomaly detection methods**
  - Usually **comparative in nature**
    - Define some "distance" metric and compare an example to an amalgamated standard
  - May just be a configuration change
    - Lacks predictive expectation..

- Can imagine using a generic anomaly detection algorithm which passes anomalies to a diagnosis
  - I feel sick....better go see a doctor
  - My condition is anomalous...what could be wrong?

Layered Approach

- This would improve the robustness of the entire system
  - Novel problems never before seen would still get flagged as anomalous
    - May be a higher false positive rate…but could take all these examples for further diagnosis training

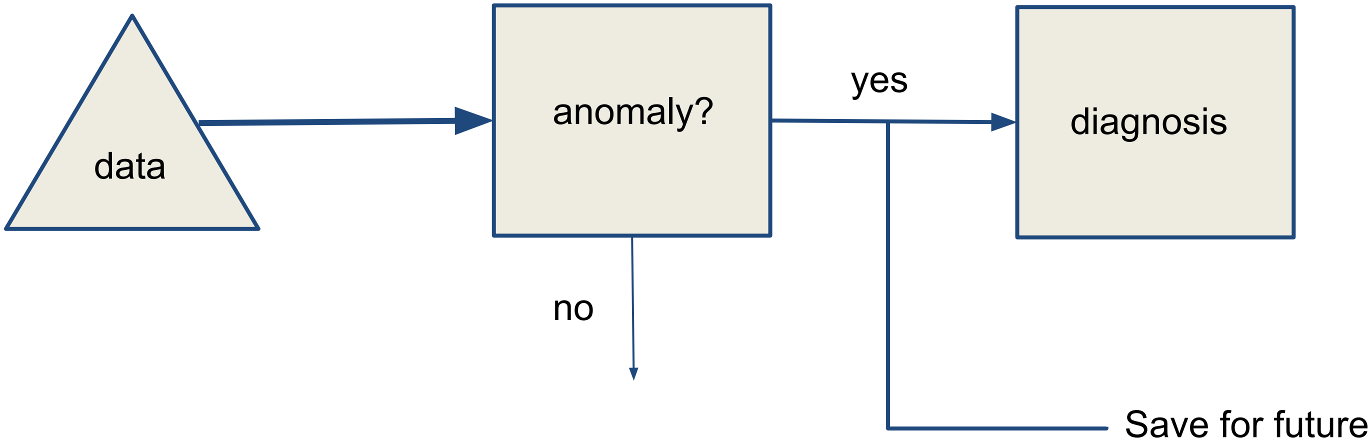

The Body of Hydra

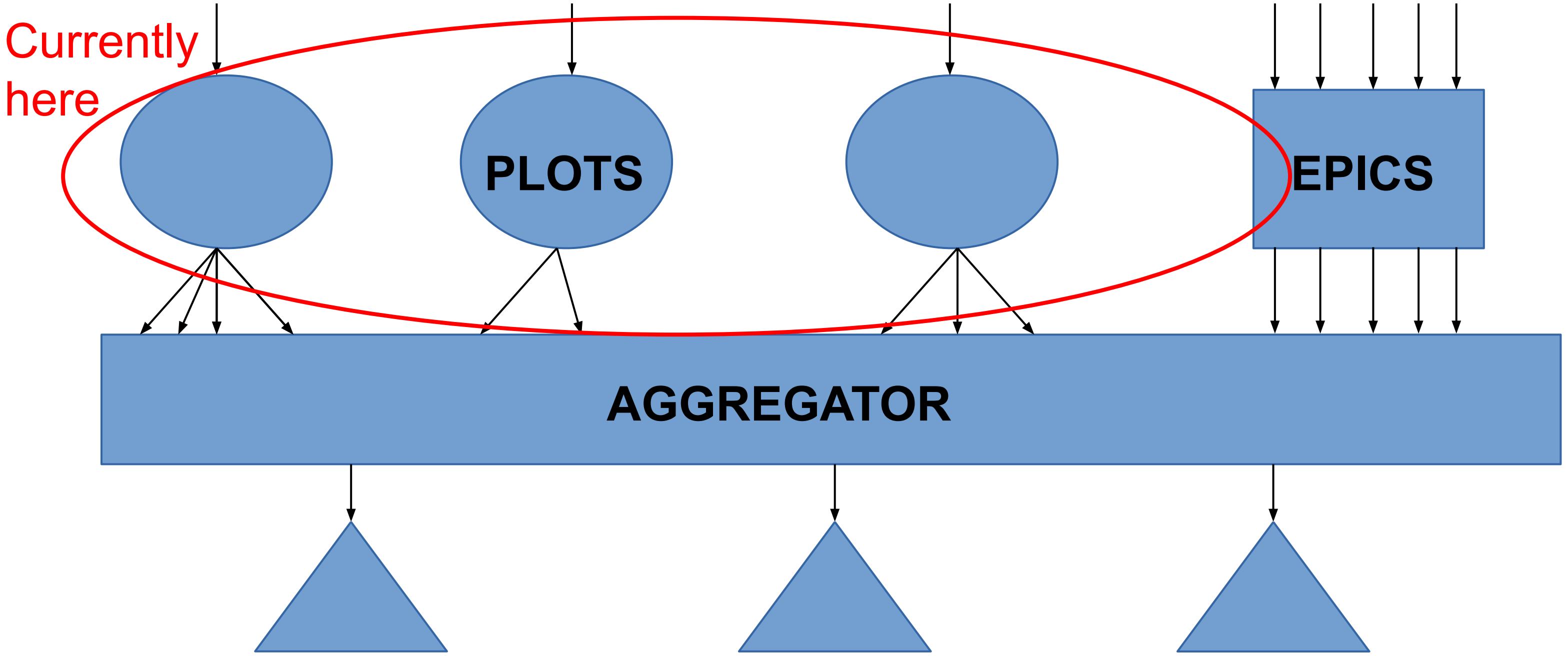

Outputs/Actions

Multi Modal

- The body of Hydra could **synthesize disparate data**
  - Think back to the example....
  - “Oh! I think there is a calibration off here because I see A looks like X and B looks like Y”

- There may be **opportunities for NLP work**
  - Read logbook entries to explain the collection of plots
  - Make autonomous logging to aid the 2am dirth
    - Configuration changes etc
  - explainability!

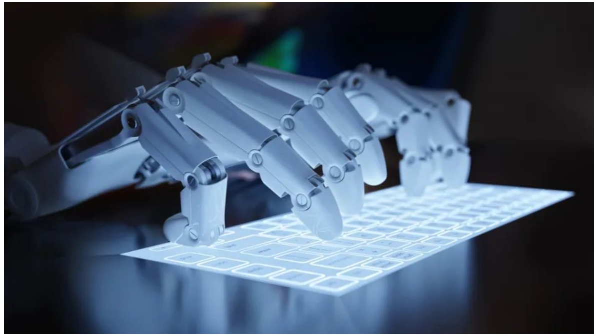

Human Interaction

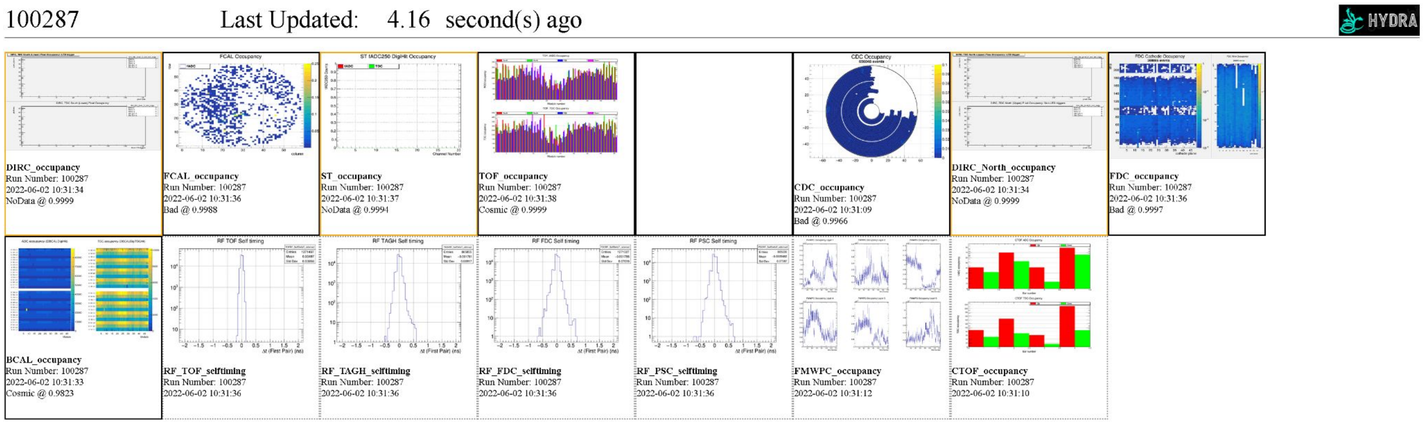

- **Real-time dashboard** viewable from anywhere shows the last plot analyzed as well as the class and confidence
- Able to detect hot channels in some detectors for later calibration. Can detect these problems early indicating hardware that may soon need replacing
- When something seems “off” shift crews can see the plot and **focus on the plots that matter**
  - Notify experts, sound an alarm, take corrective action

Human Interaction cont.

- People operate on **trust**
  - Tend to trust people

- People **Avoid discomfort**
  - Too many alarms will lead to people turning the system off

- People are **busy/lazy**
  - Need to lighten workload, not add to it

- People operate in an **app driven** world
  - UX actually a consideration

Conclusion

- Rapid advances in AI/ML tools and techniques allow even novices to being playing around with AI
  - Leverage off the shelf technology!
- Hydra has been successfully deployed in Hall-D and has already spotted problems humans have missed
  - SBS and CLAS deployments ongoing
- AI is here to stay and we need to start considering how we are going to work alongside the machines

“AI won’t replace the scientist, but scientists who use AI will replace those who don’t.”

-Adapted from a Microsoft report, “The Future Computed”