Jefferson Lab

HYDRA

computer vision for data quality monitoring

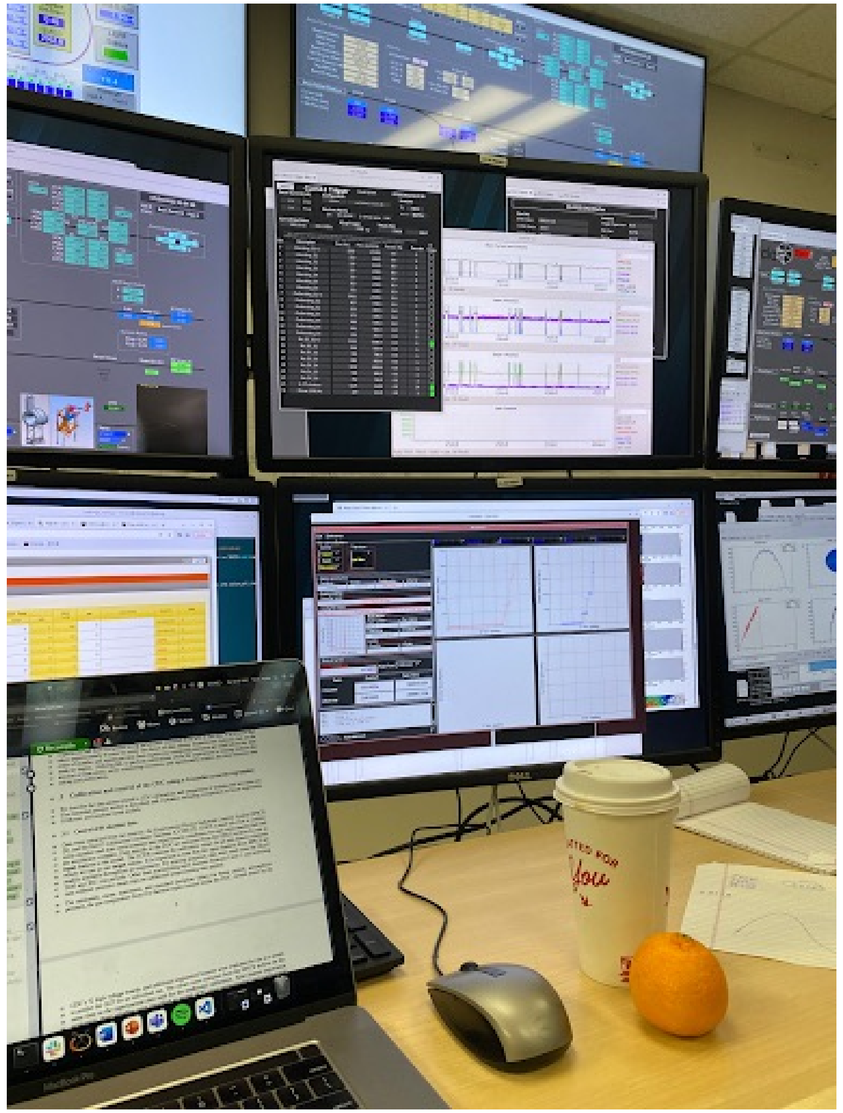

Brad Sawatzky -- on behalf of the Hydra team

brads@jlab.org

Humans cannot accurately monitor hundreds of images

Fatigue

Continuous monitoring is mentally exhausting.

Inconsistency

Human judgement can vary significantly between individuals.

Slow response time

Humans can only process so much information at a given time.

Scalability

Human-based systems do not scale efficiently as data volumes increase.

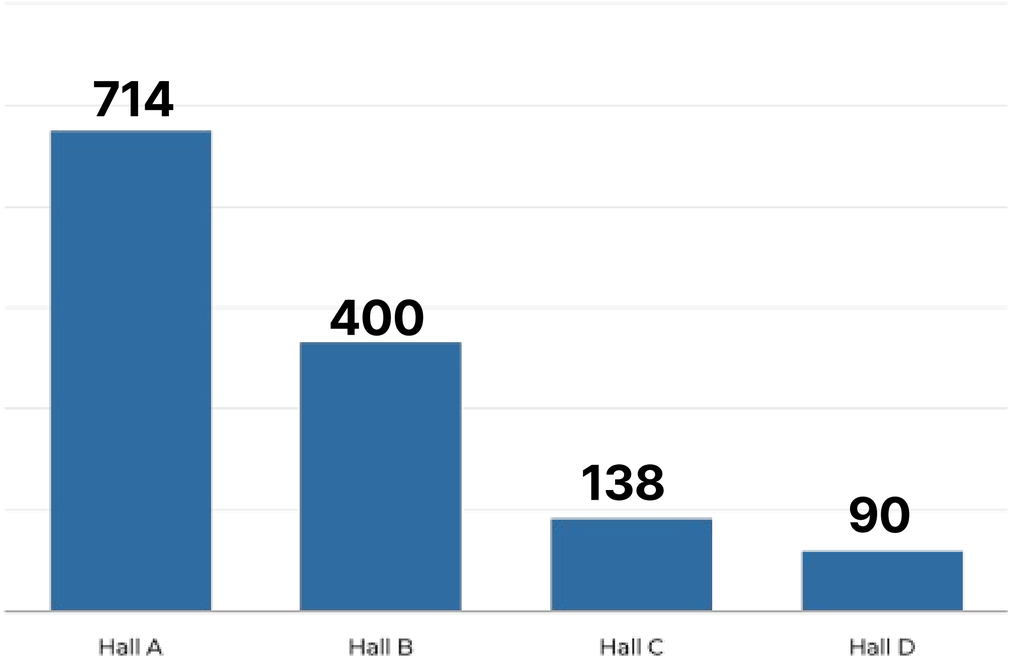

Approximate number of individual histograms per experiment per run, monitored by the shift crew

HYDRA

20524 Last Update: 6.00 (s) ago

showing 71 / 74 images Show All

RICH_s1
Run Number: 20524
2024-10-14 15:34:25
Bad

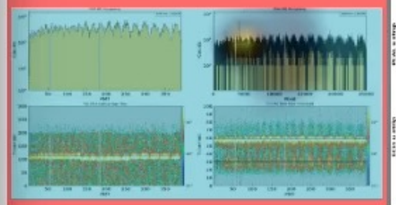

ECAL_adcTime_s3
Run Number: 20524
2024-10-14 15:34:20
Good

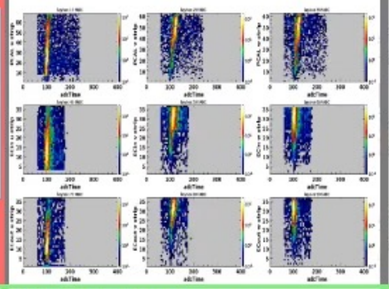

ECAL_tdc_s3
Run Number: 20506
2024-10-14 15:30:31
Good

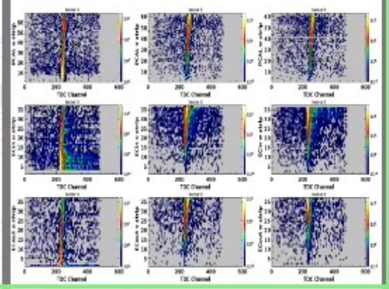

RF_adc
Run Number: 20524
2024-10-14 15:34:24
Good

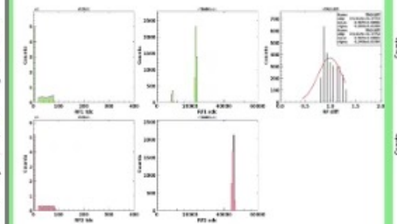

BST_multiplicity
Run Number: 20524
2024-10-14 15:34:16
Good

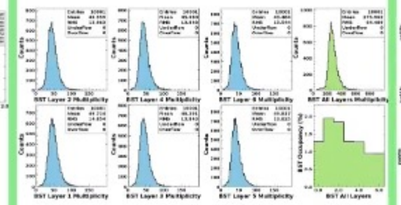

FTOF_adcOccupancy
Run Number: 20524
2024-10-14 15:34:22
Good

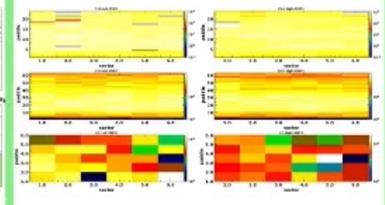

ECAL_adcEnergy_s3
Run Number: 20524
2024-10-14 15:34:19
Good

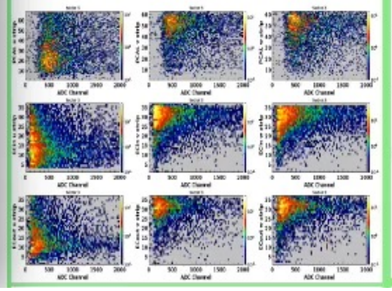

FTOF_tdcOccupancy
Run Number: 20506
2024-10-14 15:30:36
Good

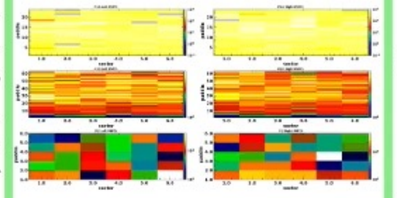

HEL_board
Run Number: 20524
2024-10-14 15:34:22
Good

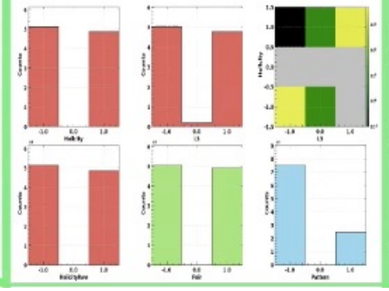

ECAL_tdcOccupancy
Run Number: 20506
2024-10-14 15:30:30
Good

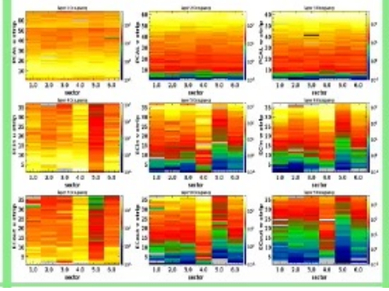

FTOF_adcTime_s6
Run Number: 20524
2024-10-14 15:34:22
Good

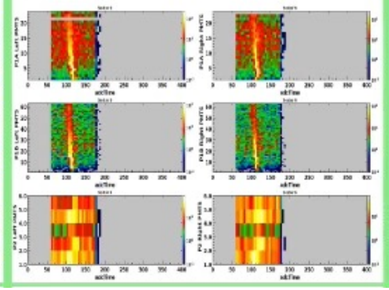

FTOF_adcTime_s4
Run Number: 20524
2024-10-14 15:34:22
Good

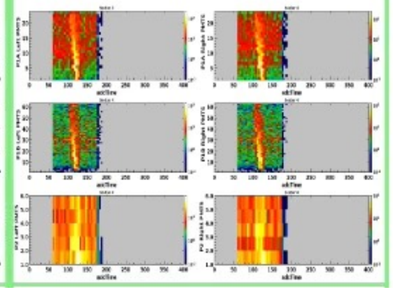

Hydra performs data quality monitoring alongside the shift crew, non stop.

Hydra at a glance

Monitoring workflow

Python modules for image management, inference, and action

Back end

MySQL database facilitates system-front end communication and storage

Front end

Web app offers actionable insights to users from anywhere without needing technical system or AI knowledge.

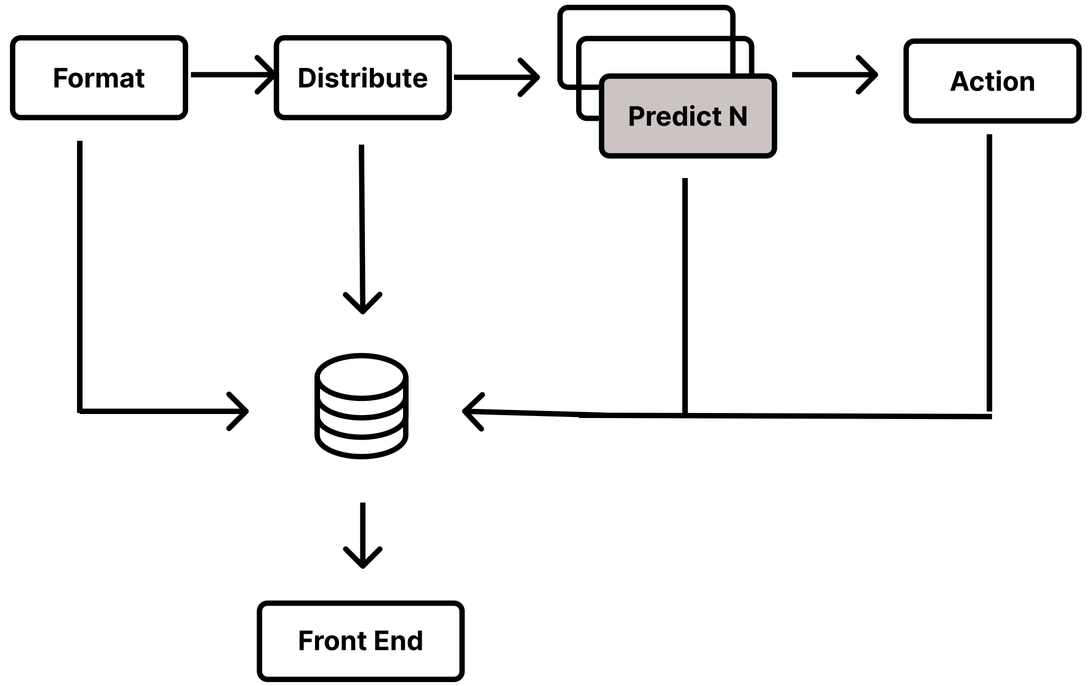

Use cases

Full

Label images, train models, run inference and interact with full FE

Online monitoring

Front end + Database

View labeled and unlabeled images on web interface

Offline monitoring, calibration

Workflow + Database

Train models, run inference, and store results in database

Offline processes

Summary Statistics

	Hall D	Hall B	Hall A	Hall C	miniHydra
Year deployed	2019	2022	2024	2024	2024
Labeled Images	139,900	236,422	5,438	302	229,872
Plot Types	15	73	42	16	46
Active Models	6	41	0	0	-
Batch / frequency	15 / min	73 / 3-5min	-	-	-
Runs with 1 bad image	1,035	1,019	-	-	-
Total analyzed	1,101,160*	1,199,771	-	-	-

Full FE+DB

Over 600k labeled images, over 2 million classified images across all deployments.

*since record keeping began

Monitoring Workflow

Model updates

Hydra can incorporate any computer vision model (e.g., VGG16) or your own custom built model.

The current default model Hydra uses is InceptionV3 with CBAM added.

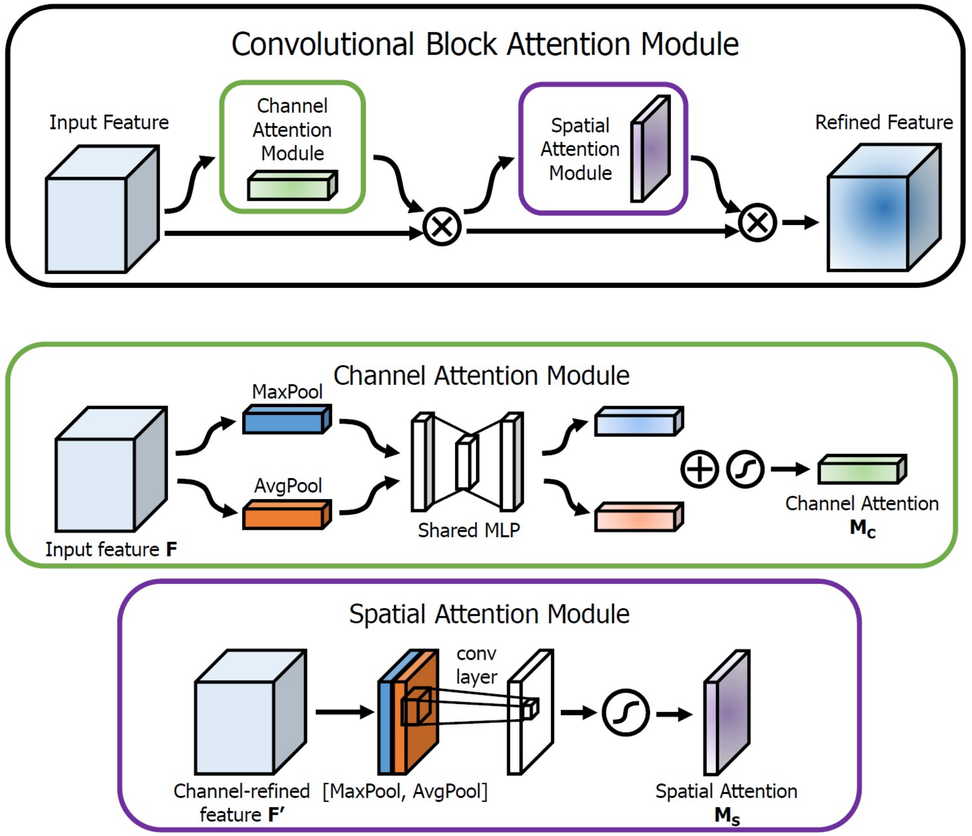

Interpretability

GradCAM

Gradient-weighted Class Activation Maps provide visual explanations for the model's classification by highlighting important regions of the image

this is a normal image

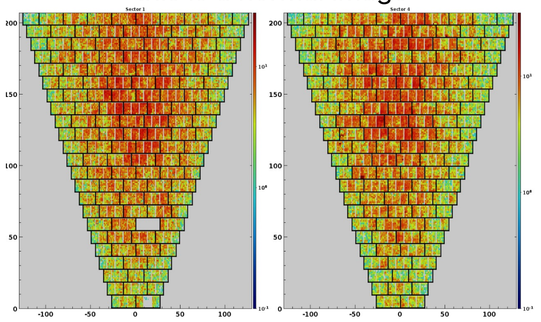

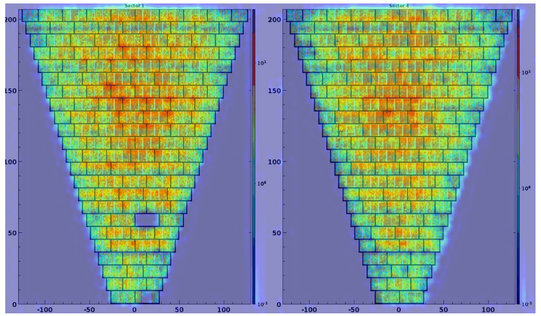

this is a bad image

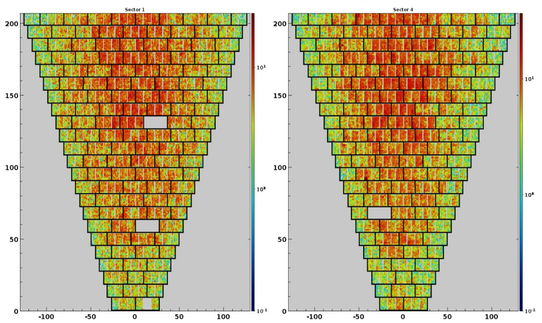

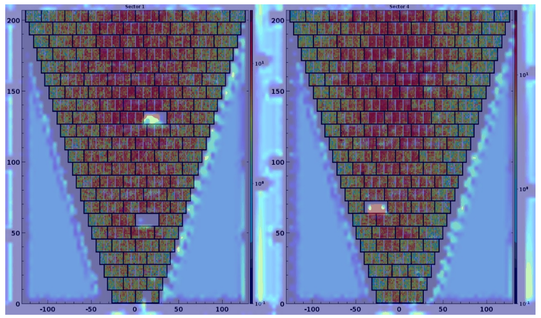

Heat maps are produced from mixed layers in InceptionV3+CBAM

Front end updates

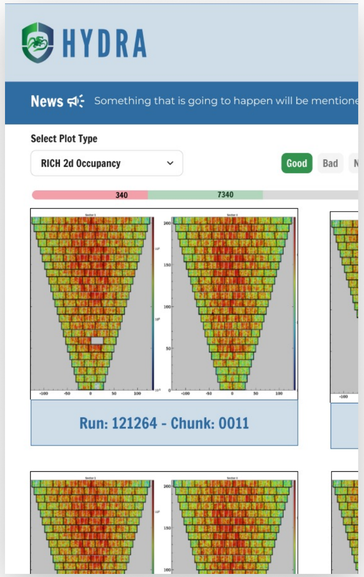

Image Labeler

Efficiently label thousands of images used for training a model.

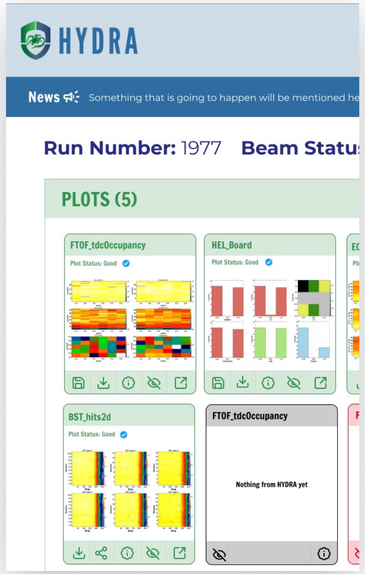

Run

See predictions in near real-time. This page continuously updates with new images during an experiment.

Library

Contains information useful to evaluate a given model's training and performance.

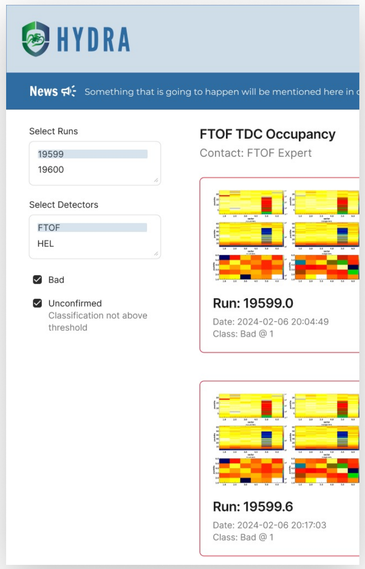

Log

Displays problematic and potentially problematic images from a trailing 24 hour window.

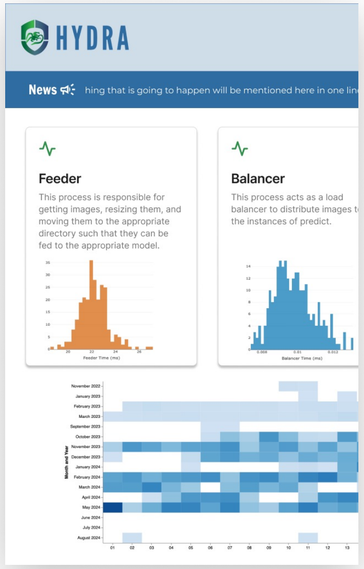

Status

Primarily used by administrators to monitor system performance.

Grafana

Displays all predictions over time. Trend analysis on predictions can indicate when it is time to retrain a model.

Towards a containerized Hydra

Lowers operational overhead

Easier to deploy and manage

Automated CI/CD pipelines

Faster development, early detection of software issues

Modularized system processes

Centralized processing, dynamically scale monitoring workflow processes as needed

Containerized Hydra

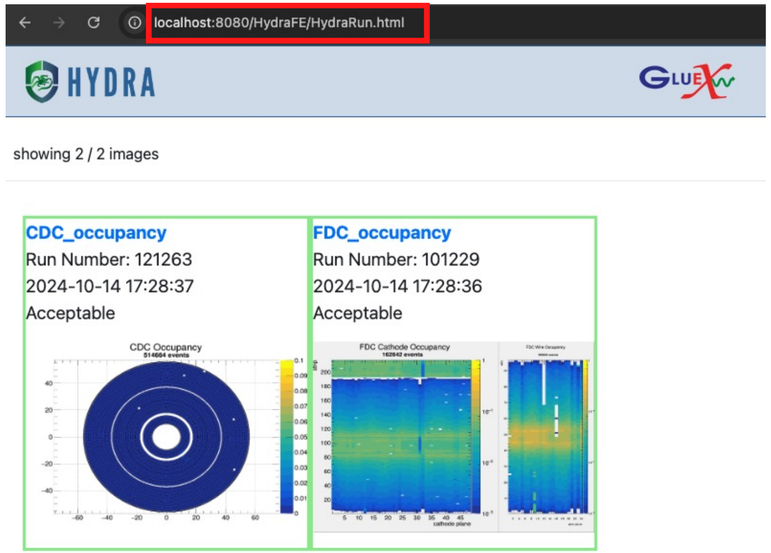

Demo container

Users can utilize pre-trained models to run inference on a sample image set with Hydra. An internal web server provides the same user interface used in production.

Development container

Currently hosted on JLab's GitLab instance, it utilizes Podman secrets to securely store GitHub SSH key that enables cloning at runtime. A python script is used to set up the container environment.

Documentation + Support

User Guide

Hydra has a complete User's Guide included for both demo and development containers and on Read The Docs.

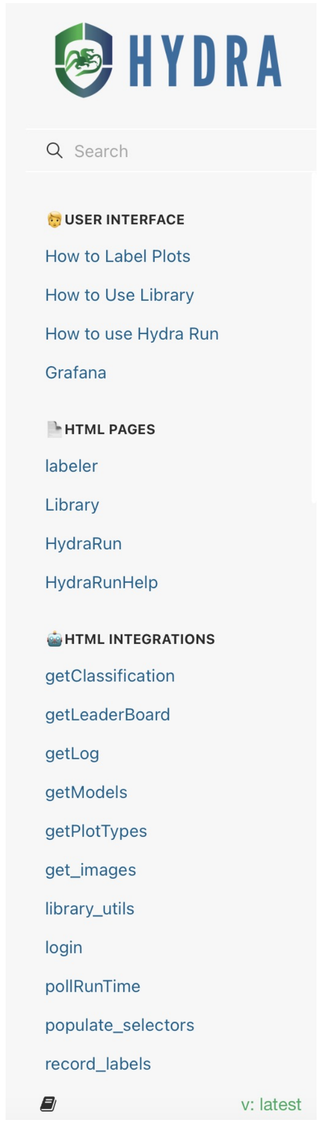

Welcome to Hydra's documentation! **Hydra** is an extensible framework for training, managing and deploying machine learning models for real time data quality monitoring.

Check out the Github for further information.

User Interface

How to interact with Hydra's web-based pages.

HTML Pages

The backend of Hydra's webpages.

HTML Integrations

The link between Hydra's backend and frontend.

HTML Utils

Background scripts for the frontend of Hydra.

Hydra Libraries

Background scripts for the backend of Hydra.

Hydra Scripts

The main scripts that run Hydra.

Hydra Utils

Different dependencies, configurations, and connections that Hydra uses.

Development Roadmap

Q1 FY25

Migrate existing web application to React, UI/UX testing

Q2 FY25

Migrate existing deployments to containerized Hydra

Q3 2025

Automated CI/CD pipelines

Q4 FY25 and beyond

Implement infrastructure for unsupervised learning

Multimodal Hydra, Deployed on Kubernetes

Meet the team

David Lawrence

EPSCI Lead

Thomas Britton

EPSCI
Creator/Lead Developer

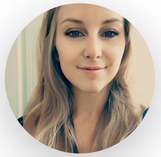

Torri Jeske

EPSCI
Operations, Full Stack Development

Nataliia Matsiuk

Information + Records
Containerization, QA

Raiqa Rasool

EPSCI
Full Stack Development

User's Guide

Download demo container

hydrateam@jlab.org

Jefferson Lab

EPSCI

extras

Anecdote:
Total response time: 7 hours 44 minutes

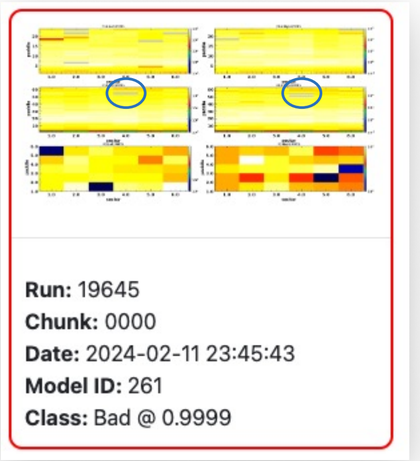

2/12/2024 00:02:00

Shift crew submits monitoring images to logbook. Problems with detector indicated.

2/11/2024 23:45:43

Hydra alarms for problem with Forward Time of Flight detector

2/12/2024 07:30

Shift crew first becomes aware of problem with detector.

Labeler

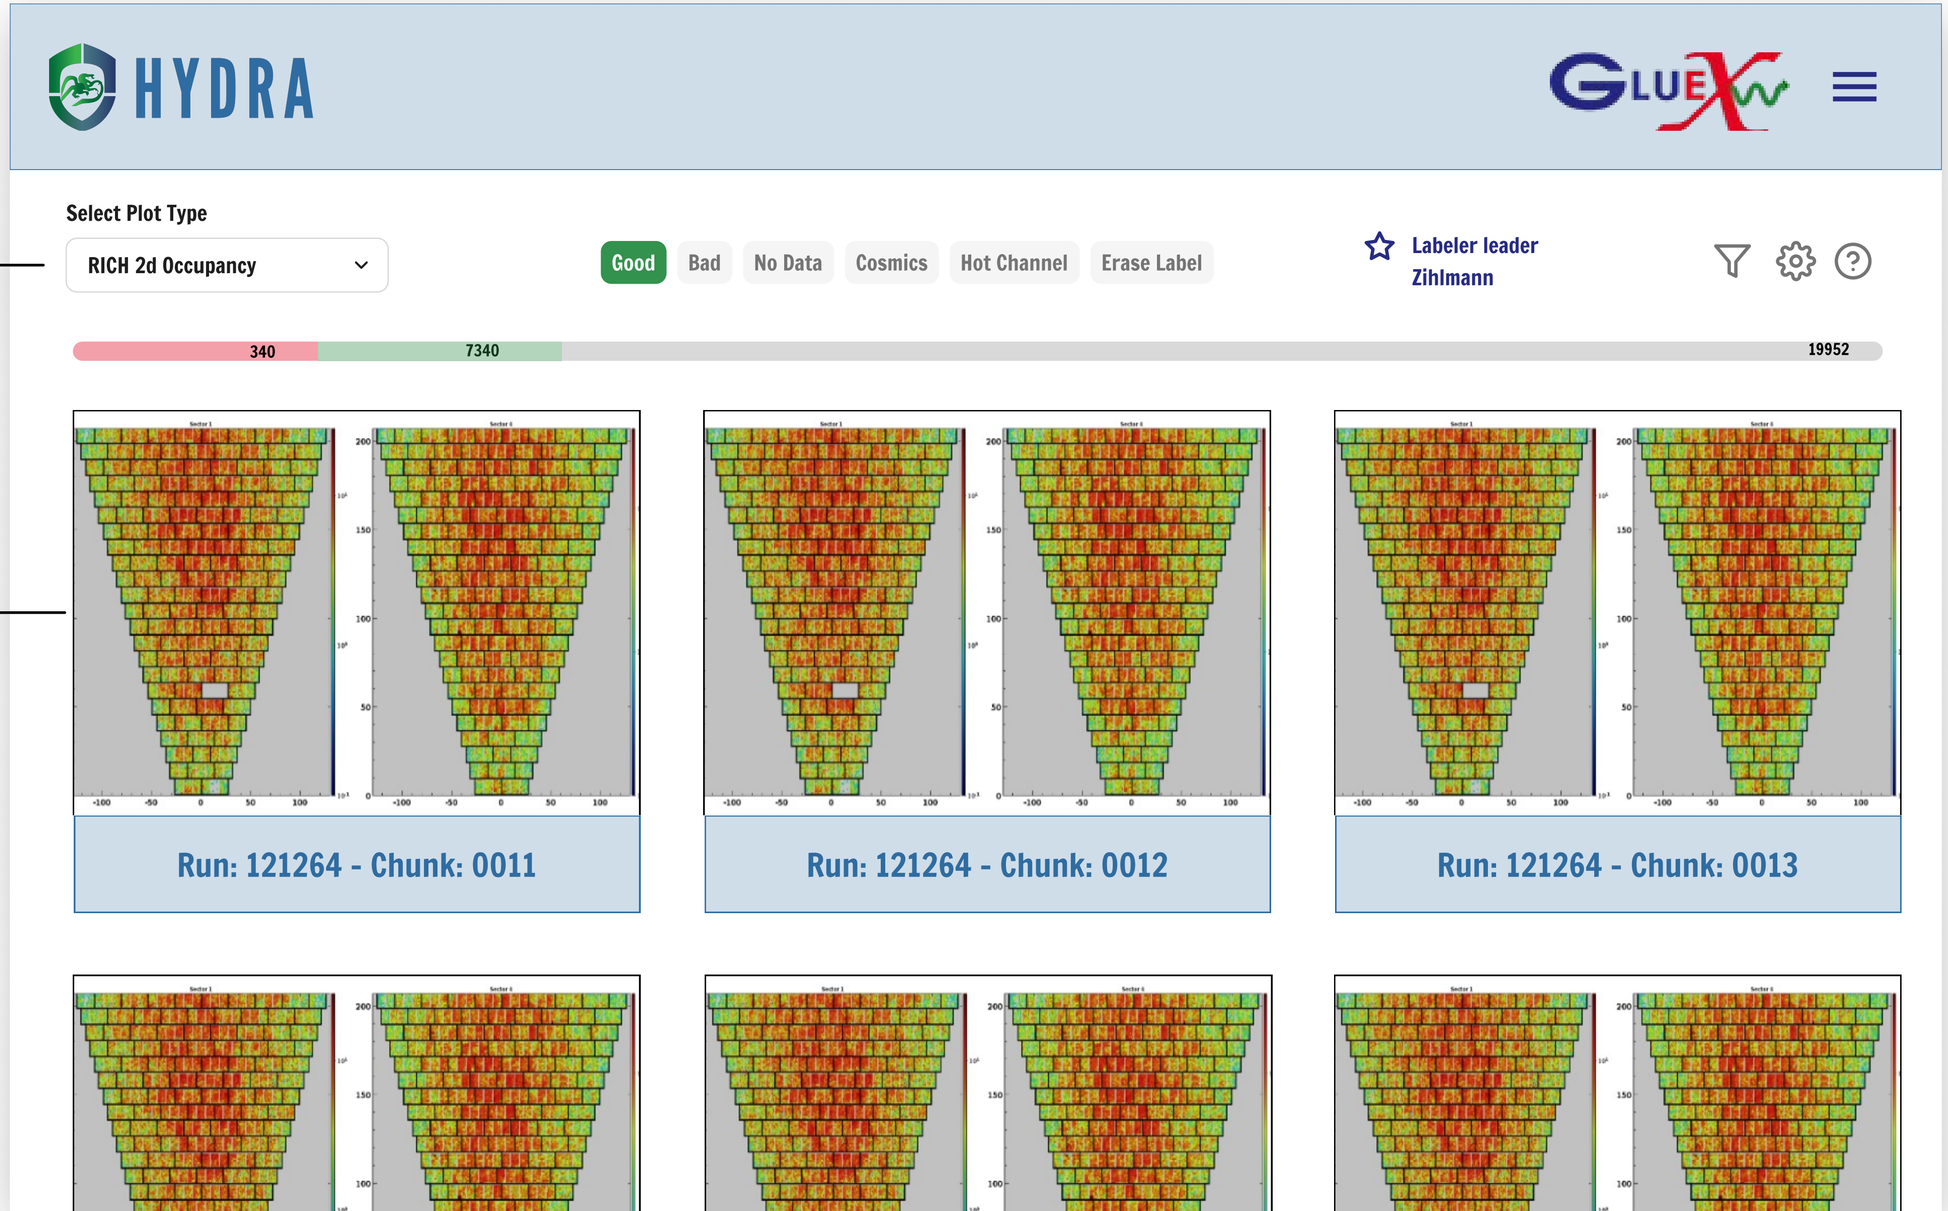

Plot Type Selector

Images to label

Run

HYDRA

clas
CEBAF Large Acceptance Spectrometer

News Something that is going to happen will be mentioned here in one line

status indicators

Run Number: 1977 **Beam Status:** ON **Last refresh:** 1.95 s ago

is Hydra broken?

Main image gallery

PLOTS (5)

Hidden Plots | Saved Plots | Filter Plots: Show All Plots

FTOF_tdcOccupancy
Plot Status: Good

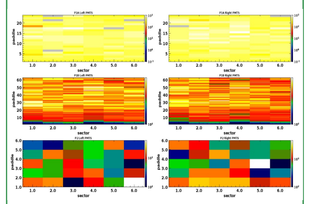

HEL_Board
Plot Status: Good

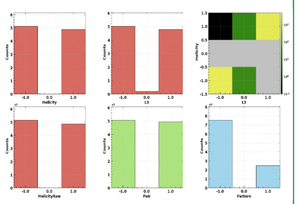

ECAL_adcOccupancy
Plot Status: Good

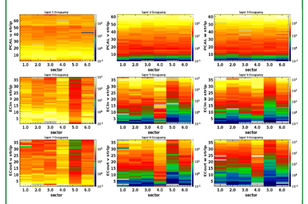

DC_OccupancyNorm
Plot Status: Good

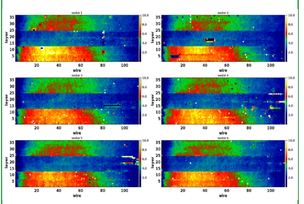

BST_hits2d
Plot Status: Good

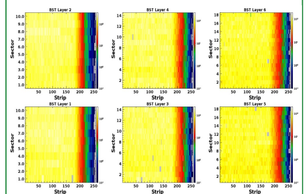

FTOF_tdcOccupancy
Nothing from HYDRA yet

FTOF_adcOccupancy
This plot type has a bad plot.
Click here to find it in bad plots section

RICH_occupancy2d
Plot Status: Good

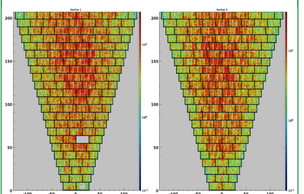

BAD PLOTS (2)

reserved for Bad plots

FTOF_adcOccupancy
Plot Status: Bad

Heatmap Opacity: 0.5

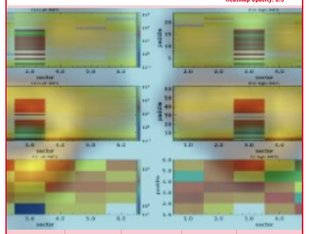

FTOF_adcOccupancy
Plot Status: Bad

Heatmap Opacity: 0.5

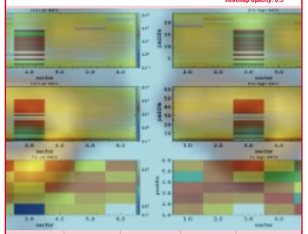

Grafana

Visualize output weights of model vs time for each plot type

Library

HYDRA

GLUEX

News Something that is going to happen will be mentioned here in one line

Select Plot Group: RICH

Select Plot Type: RICH 2D Occupancy

Select Model: (330) RICH 2D Occupancy

Enhanced Confusion Matrix | Loss | Accuracy

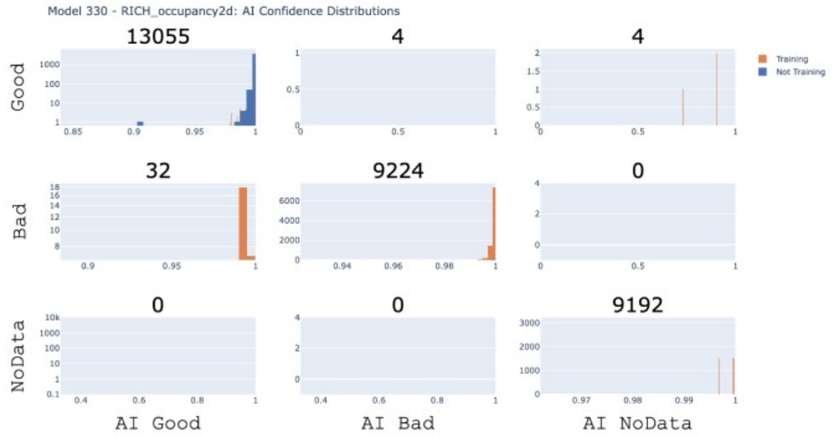

	Model ID	Date	Overall Accuracy	Thresholds	Trained By	
☐	330	2024-08-09 10:19:33	99%	Good: 0.54, Bad 0.57, NoData 0.99	roark	⚙
☐	332	2024-02-07 08:02:52	95%	Good: 0.53, Bad 0.70, NoData 0.99	roark	⚙
☐	337	2024-01-26 11:40:12	99%	Good: 0.98, Bad 0.99, NoData 0.99	roark	⚙

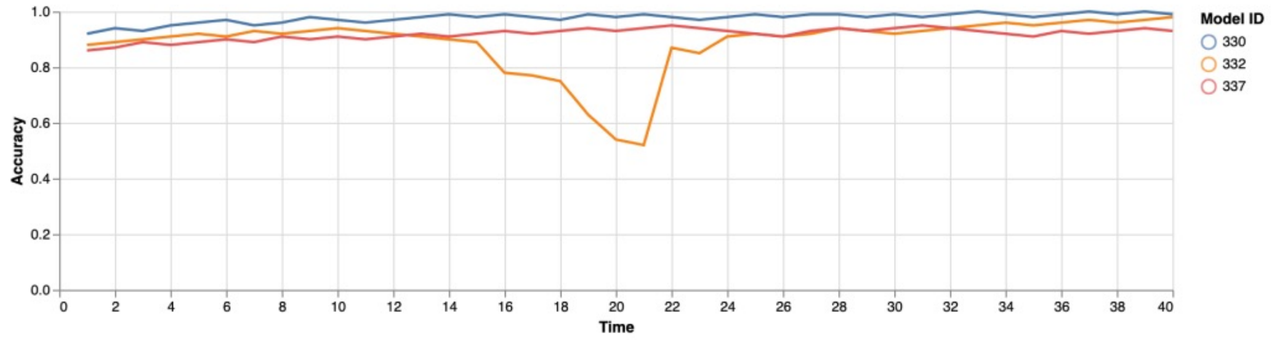

select models individually or by groups

view info from Models table, adjust thresholds

Running accuracy vs time ***

this is only reliable with frequent labeling

Log

HYDRA

GlueX

News Something that is going to happen will be mentioned here in one line

Select Runs

19599
19600

Select Detectors

FTOF
HEL

- [x] Bad
- [x] Unconfirmed
  Classification not above threshold

FTOF TDC Occupancy

Contact: FTOF Expert

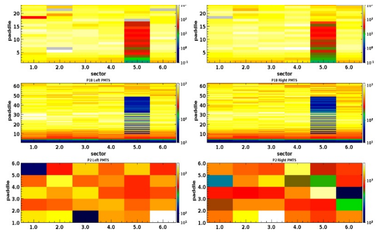

Run: 19599.0

Date: 2024-02-06 20:04:49
Class: Bad @ 1

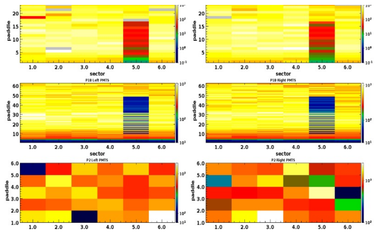

Run: 19599.2

Date: 2024-02-06 20:09:55
Class: Bad @ 1

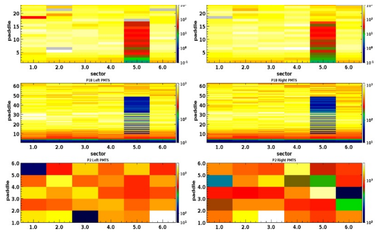

Run: 19599.4

Date: 2024-02-06 20:13:59
Class: Bad @ 0.99

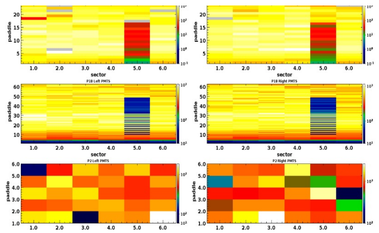

Run: 19599.6

Date: 2024-02-06 20:17:03

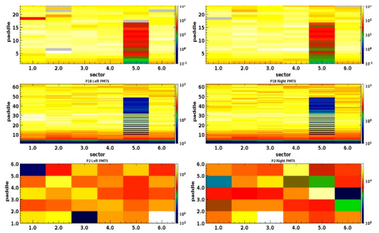

Run: 19599.8

Date: 2024-02-06 20:20:48

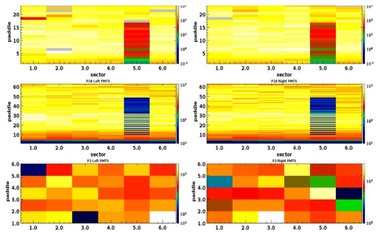

Run: 19599.10

Date: 2024-02-06 20:24:11

Not shown, selectors for Date/ Time, Run Period, Run Range