



# Monitoring Plot Anomaly

## Detection With Hydra

Thomas Britton ([tbritton@jlab.org](mailto:tbritton@jlab.org))

David Lawrence ([davidl@jlab.org](mailto:davidl@jlab.org))

Kishansingh Rajput ([kishan@jlab.org](mailto:kishan@jlab.org))

# 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.
  - To preserve sanity I looked at closer to 175 plots a day, 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

# 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

# Hydra Fast Facts

- Hydra looks at a finer time scale than 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

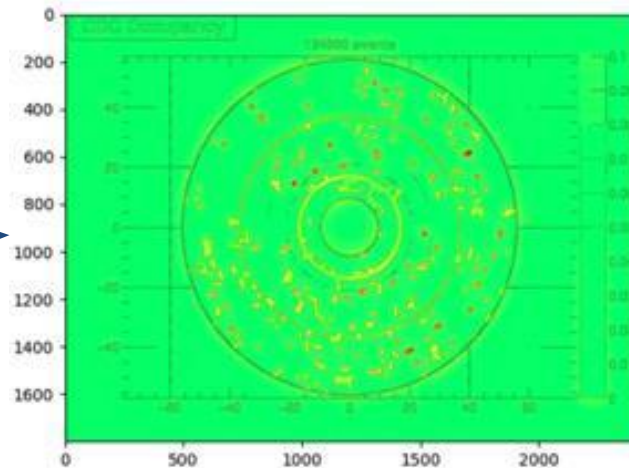
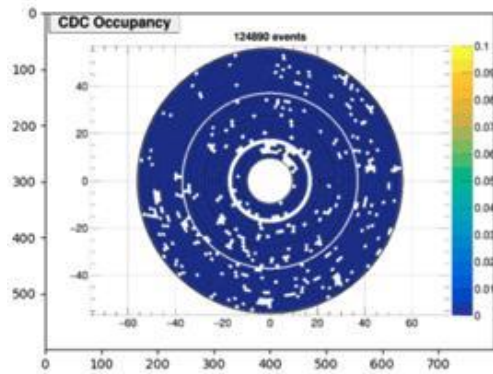


Koboldpress.com

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

# Layer-wise Relevance Propagation

- **Better than knowing that an image is “Bad” is knowing why it is bad**
  - Layer-wise Relevance Propagation (**LRP**) takes the output vector and ‘back-propagates’ it through a model to the input pixels’ weights
  - Whether a given contributes to (+) or detracts from (-) from a label
  - Has the negative of being channel dependent (e.g. is it the redness of a pixel or...?)



Blue  
channel  
only

# Hydra Future Development

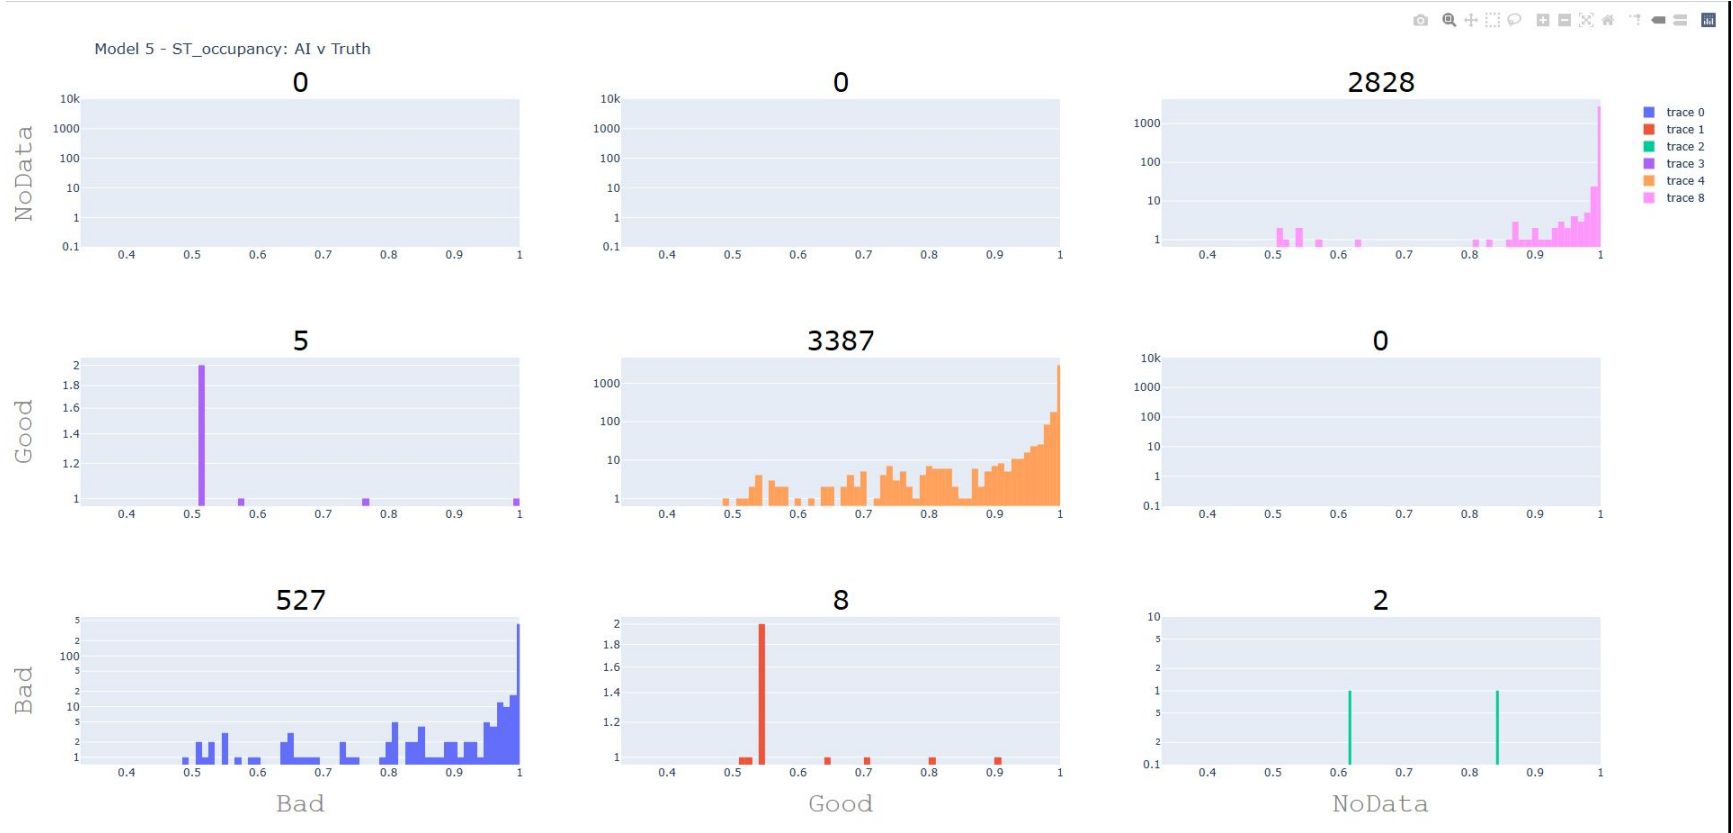
- **Classification splitting**
  - From a doctor saying “you are sick” to actually diagnosing a condition
- **Custom, optimized models**
  - Inference accounts for ~71% of the total processing time and is driven primarily by model size
- Ability to actually **take corrective action** as needed
  - Will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



# Backup Slides



# Preliminary Results (Start Counter)





# Preliminary Results (FCAL)



Model 6 - FCAL\_occupancy: AI v Truth



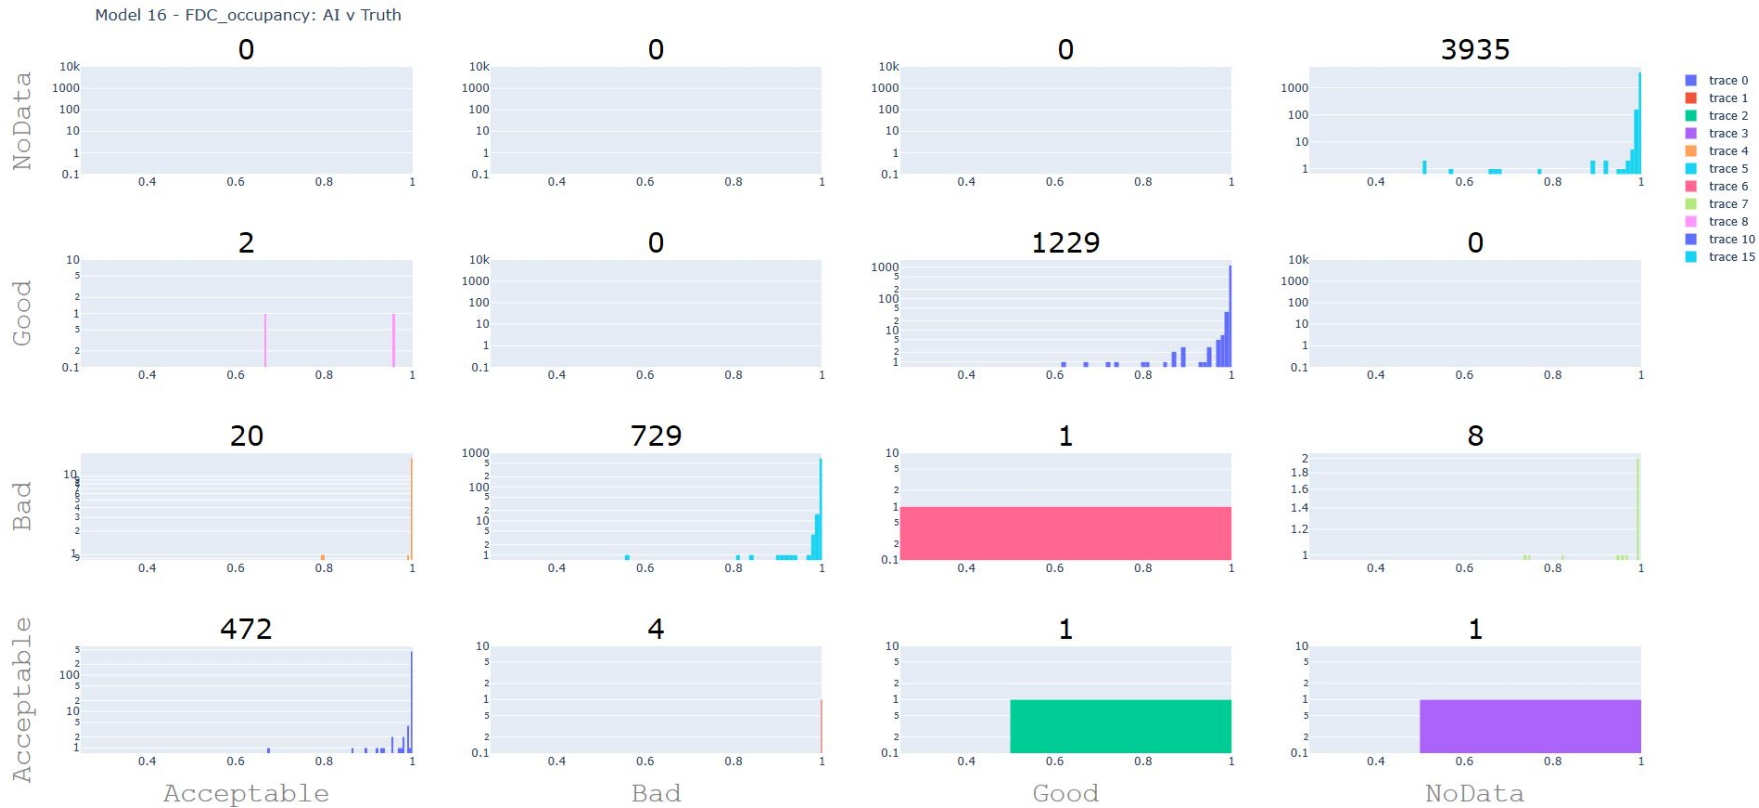
# Preliminary Results (DIRC)



# Preliminary Results (BCAL)



# Preliminary Results (FDC)



# Preliminary Results (TOF)

Model 20 - TOF\_occupancy: AI v Truth

