



U.S. DEPARTMENT OF
ENERGY



Problem of the month May 2019

Who what where why how?

- Who?
 - Open to all Jlab staff and users
 - Solo or small groups
- What?
 - A problem or puzzle to be solved with machine learning each month
- Where?
 - Ummm here....kinda
- Why?
 - Because it is fun!
 - Also prizes
- How?
 - Turn in a model file via email: tbritton@jlab.org or via common location (eg work disk) or via flashdrive to Thomas Britton



Some Fine Print

- The model must be trained from the provided data. There should be **no pre-processing** done as the models will be placed in a test harness I build
 - Translations of data inside the model itself is allowed (e.g. use a flatten layer to move from images to one dimensional arrays)
 - If you know how the data is formatted and can generate additional sets for additional training. Then go for it.
- The model's **output** should be as specified with **no additional calculations** required.
- Testing will be done on a GTX 1070 GPU and i7 7700 CPU and must fit in memory (8 Gb)
- Points will be awarded based on the metric being judged and tie-breaking will be done as the problem creator/ judge sees fit.
- Neither the judge nor problem creator will be eligible for prizes (though they may enter for the honor)


Some technical details

- The top tier teams will be invited to give a short presentation to the lunch series on what they did and why (even if it was grasping in the dark).
- Prizes will also be given out
 - Not yet determined





The Problem

- A simple tracking detector with 6 modules comprised of 6 planes each consisting of 100 wires with no magnetic field and 100% hit efficiency. Timing information is given via gray scale
- θ and Z real numbers and can be negative



The Problem

- Example images



The Problem

- The problem: given a set of hits in our detector (an image) produce the vertex (z position) as well as the angle of the track from the horizon (θ)
- Metrics (per model parameter):
 - Relative Accuracy and precision of both phi and z (smallest sum of errors squared and best resolution) to a benchmark model (David L)
 - Fastest inference
 - CPU
 - GPU

The Problem

- Complete write up and materials found here:
 - https://halldweb.jlab.org/talks/ML_lunch/May2019/
- Questions and comments can be directed to myself via slack (direct or in #ml) or email
- You have until Jun 5th
- **HAVE FUN!**

