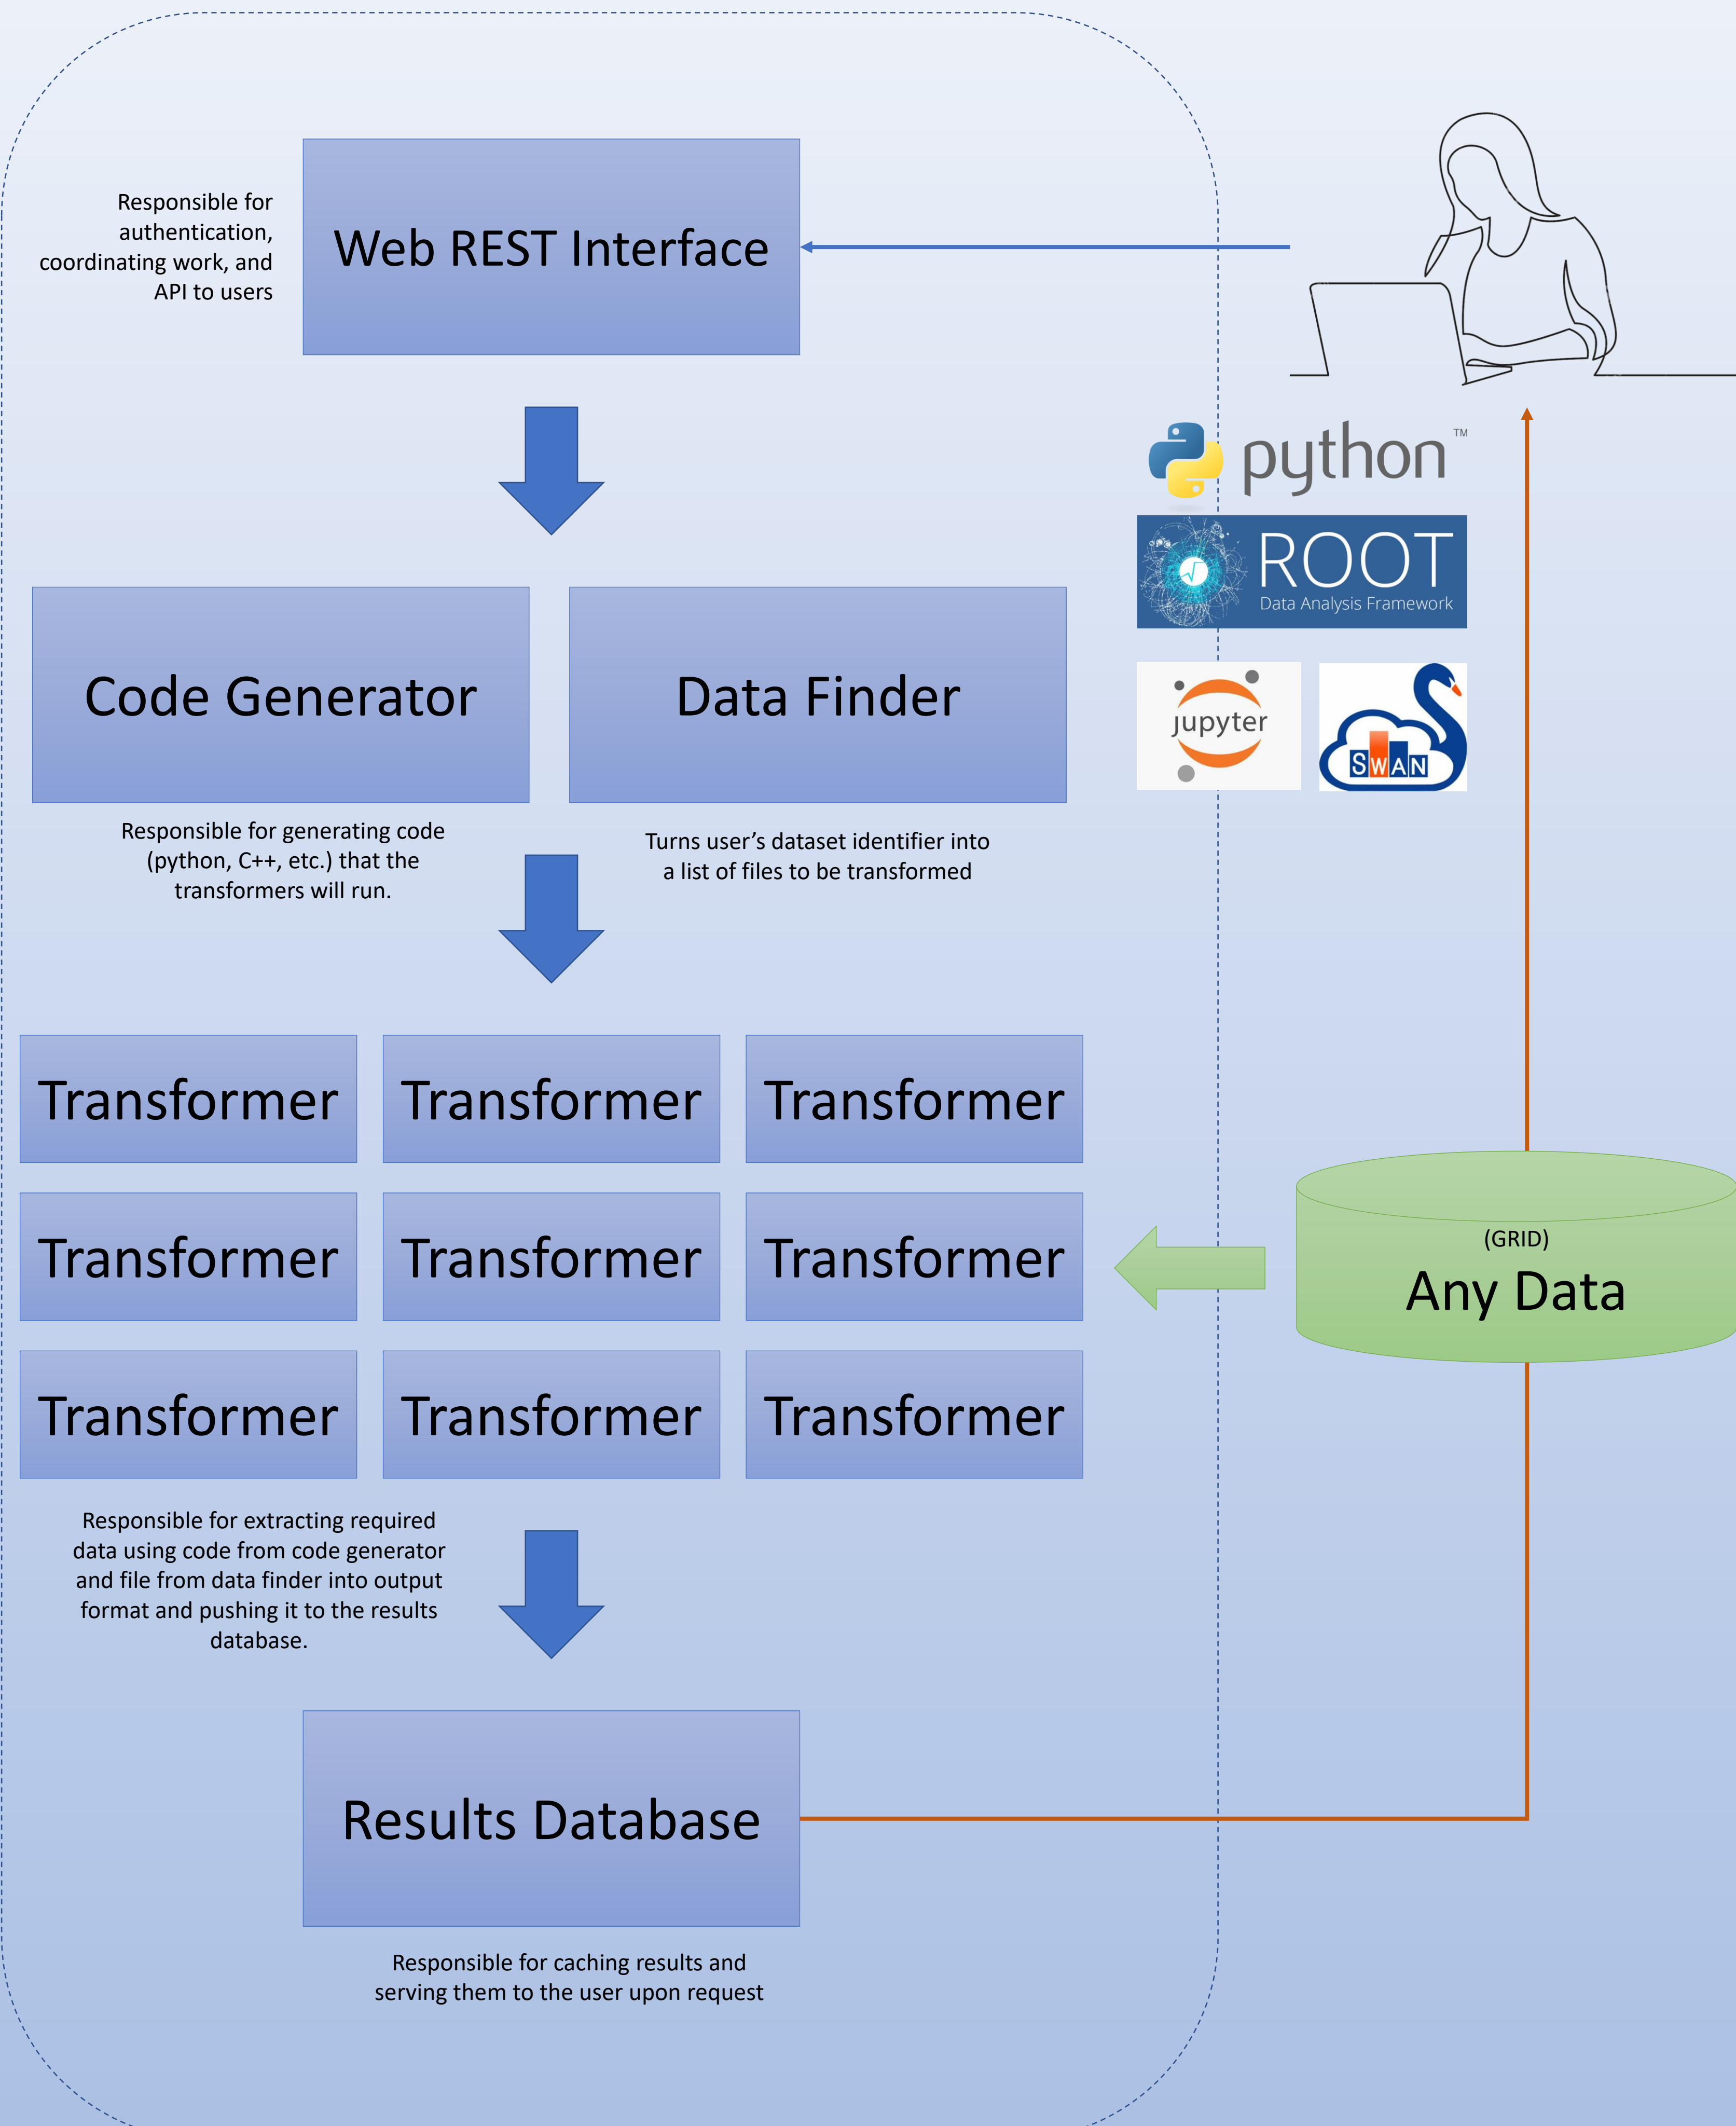


ServiceX: Making all Data Columnar

Presented by Gordon Watts (University of Washington/Seattle) for the IRIS-HEP ServiceX Team



Columnar Data – Why?

Our field, many other science fields, and industry have settled on a columnar representation of data. PAW first introduced in our field, and ROOT continues the tradition. In industry they have settled around Arrow, an in-memory implementation of columnar data. Physics data, especially older data, is not written in this format. ServiceX provides a common, highly parallel way of converting the data.

- Many modern tools can be used to analysis many types of data.
- Accessing **non-summary data** from a modern experiment
 - **Data Preservation**: an old experiment's data can be converted to be access by modern tools
 - **Combining Data From Multiple Experiments**: Data from multiple experiments can be accessed by the same analysis script when it makes sense

Why Filter The Data?

In our fields data can be 100's of TB's in size, or even PB's as we move towards the HL-LHC. We rarely need all that data. ServiceX is designed to filter the data at the source, shipping only the data the analyzer needs. This reduces the load on the network and intermediate caches and will improve the responsiveness of the system. Data can be streamed as it is available.

Cloud Native

Based on micro-services built in containers, and message passing for control, and a simple web-based REST control. It can take advantage of many industry standard tools for implementation, with us having to provide only workflow and data implementation.

CMS SH \rightarrow ZZ \rightarrow $\ell\ell\ell\ell$ Public Outreach Example

CMS has released its Higgs discovery dataset as public data. The structure of this is very similar to the ATLAS example with some important differences:

- CMS data is not a flat ROOT tree, but a CMS Run 1 AOD.

Outline

1. Use `ServiceX` for general query and object selection
2. Use `awkward` and `coffea` to do multi-object event-wide selection and plots
3. Produce the plot for running on a single MC file
4. Repeat all the MC and Data files

Why aren't we running in real time?

- The total data and MC is 70 TB. We are running on 43 TB for this notebook
- The `ServiceX` data returned locally is 10 GB
- These datasets have 1000's of files. This notebook is running in single-threaded mode, and running on a laptop at the end of a residential ADS line
- The EOS source for the data is the CERN OpenData Portal, which is not built to stream data at the rate `ServiceX` can ingest.

This talk represents nearly all the MC, with part of the 2011 and 2012 data.

Defining the base query

- We are using the CMS Run 1 AOD file type here
- No global event cuts

Filter ~35 TB of data and MC down to ~ 10 GB

```

as = ServiceXSource('cmsrun1aod')
As = as.select('name == "2011_10_02"')

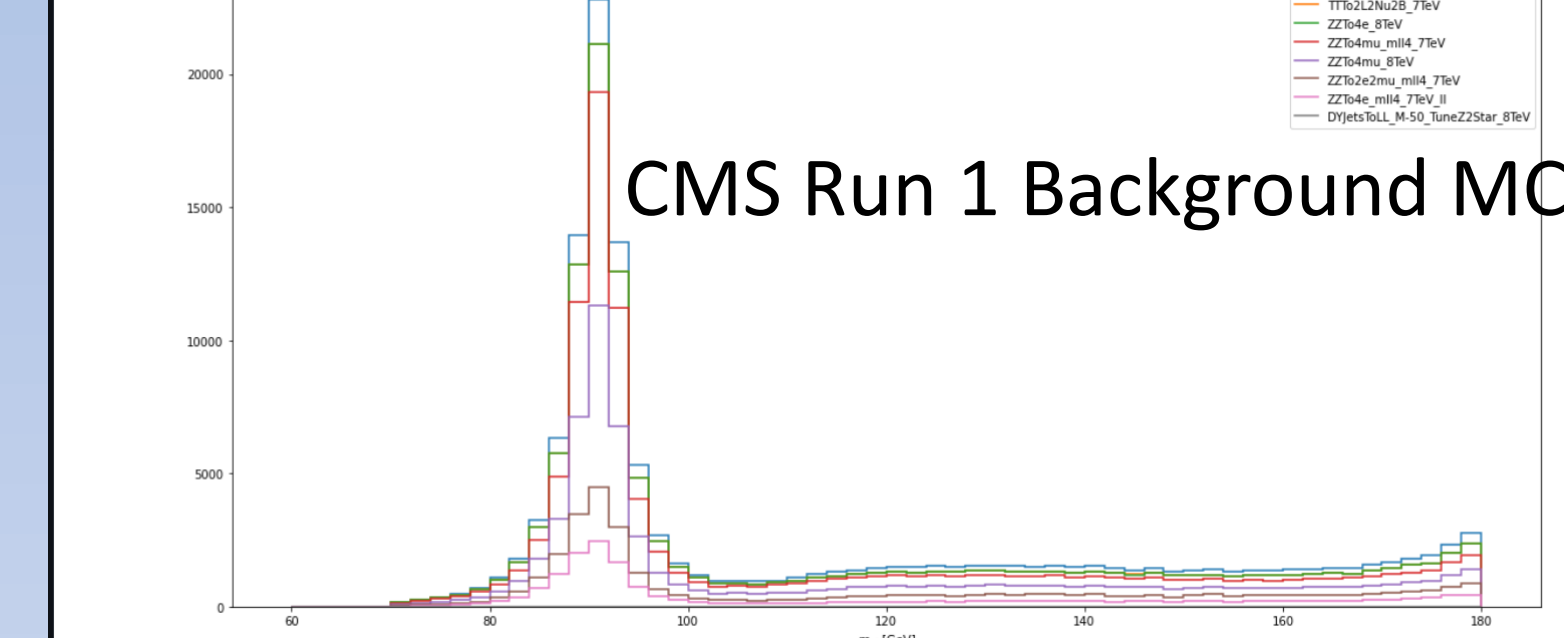
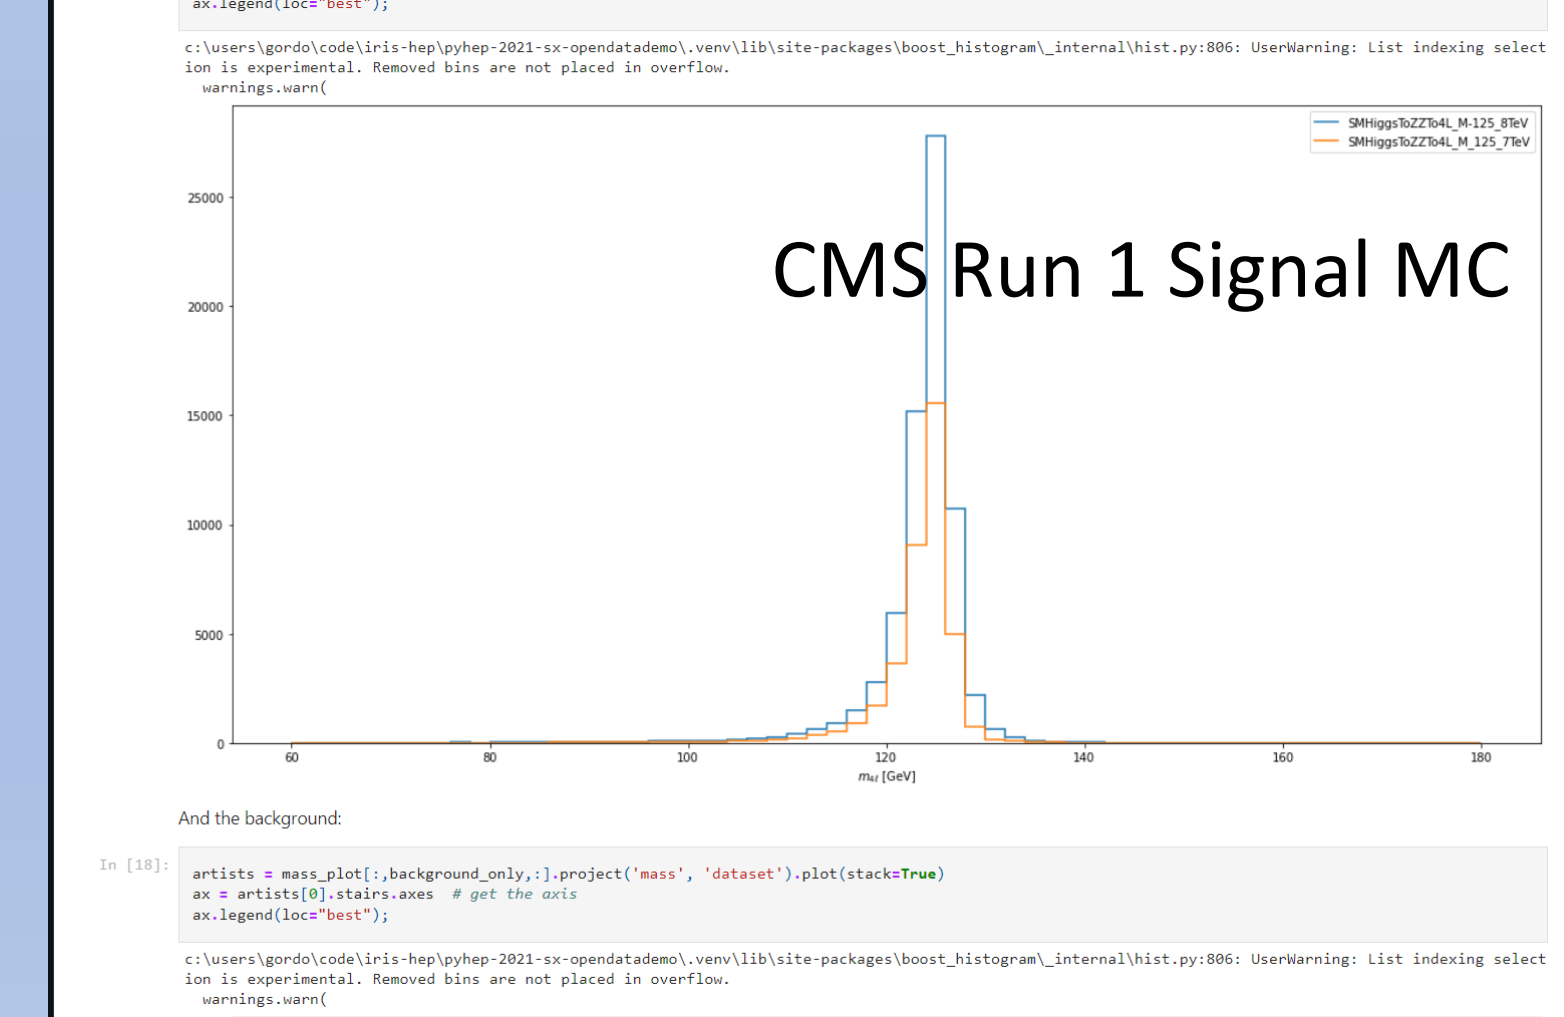
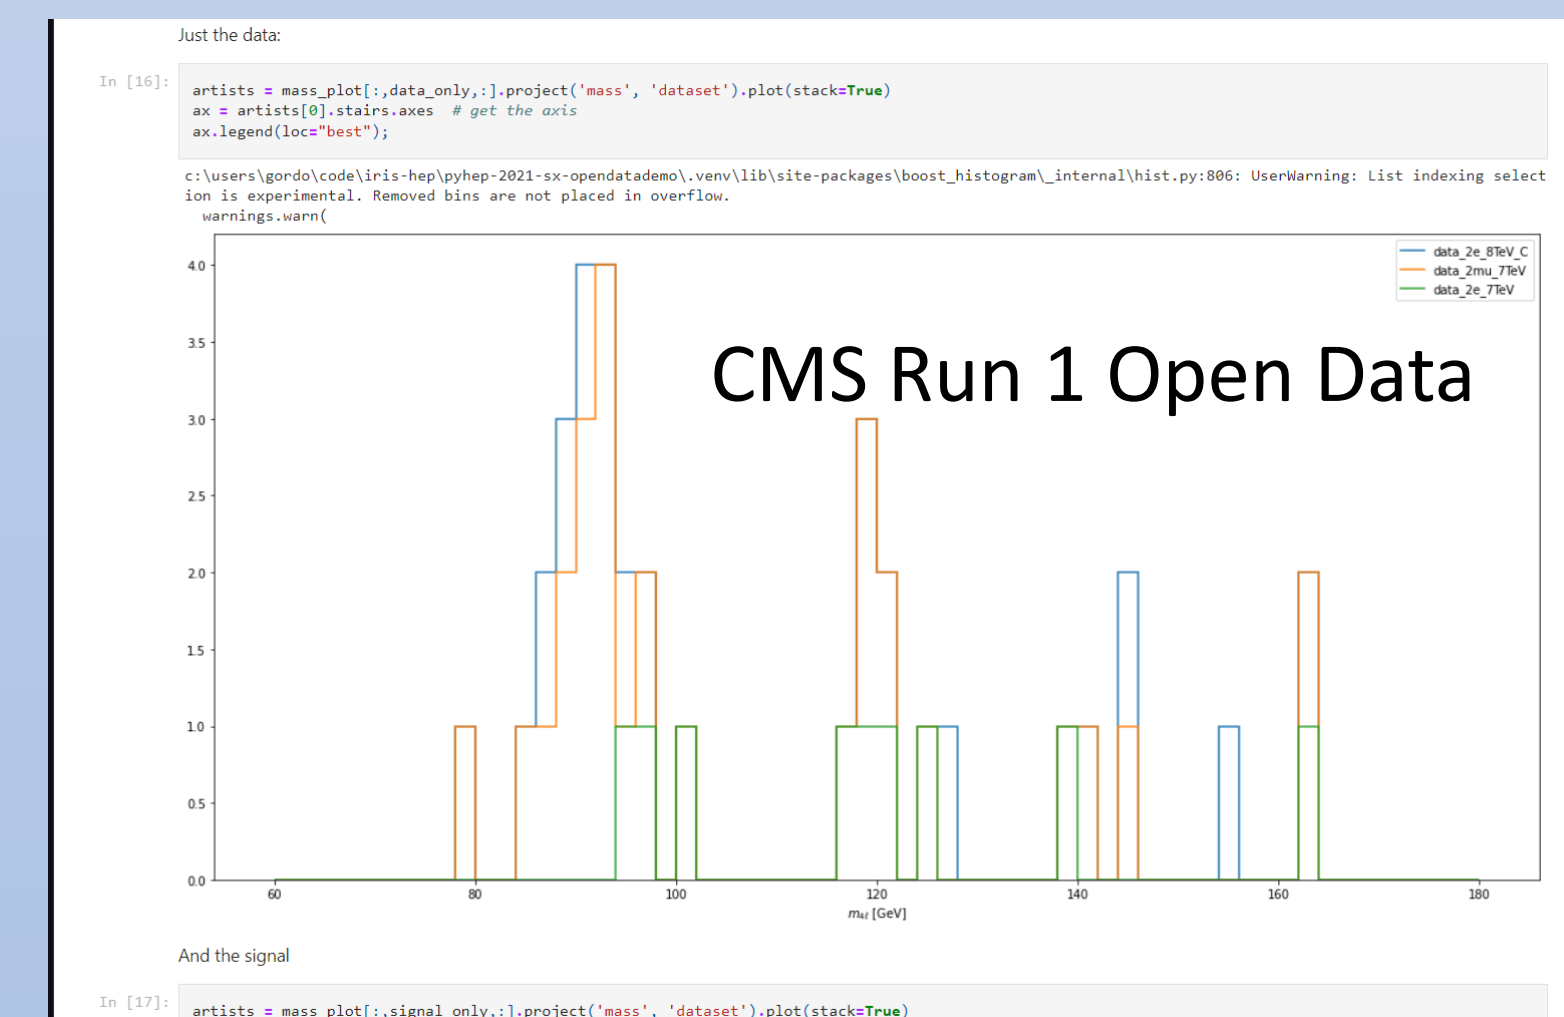
# Selecting Event Objects
lepton_source = {
    'lepton_source': [
        {'name': 'Electron'},
        {'name': 'Muon'},
        {'name': 'Tau'}
    ]
}

data_selection = {
    'lepton_source': lepton_source
}

And pull out the columns we need for selection in the code
    
```

Performing the Analysis

Analysis in Awkward Array and Coffea



Types of Transformers Available now:

- Plane Old ROOT TTree's
- ATLAS Release 21 xAOD Files
- CMS Run 1 AOD Files

In progress:

- A transformer with filtering based on the volumetric access library, yt, and data loading based on the Katai library (a binary file format specification library and I/O routine generator) (both open source community libraries)
- Snowmass produced data

Where Can Data Come From Currently?

- Anywhere accessible by rucio
- Any xrootd served file
- Any http served file
- CERN Open Data Archive (slow!)

In Progress:

- CDMS and other direct dark matter experiments

Installations

- Nebraska's CMS coffee-casa analysis facility
- U Chicago coffee-casa analysis facility
- IRIS-HEP's SSL
- And of course, any local docker on your machine as long as it is running Kubernetes!

