Likelihood Fitting and Light Curves

Fermi Summer School Analysis Tutorial Week 4

Review of Week 3 Likelihood Exercise

- Writing down likelihood
  - Model hypothesis
    - What is the model used in the exercise you worked on? Are you predicting a spatial distribution of counts? 1 or 2 dimensions? Are you assuming any background?
    - What makes a good model?
  - Write down a probability for the data given the model. To do that you need a probability distribution function.
    - What probability formulation will you use? Poisson distribution? Gaussian distribution?
  - Vary the model parameters and calculate probability to find the best fit. What values for the parameters are most probable?

We wrote down probability and calculated likelihood. Now what?

- Is the model hypothesis better than the baseline assumption about the data?
  - Remember Daniela's blue smarties? Our prior might be that the number of blue smarties in each bag is drawn randomly from a big bin and that bin is filled with an equal number of each color.
  - Priors for a gamma-ray observation?
    - A very simple assumption: data is dominated by background distributed uniformly over the sky.
    - A more informed prior: data should have diffuse emission from the Galaxy predicted by other observations, diffuse emission from residual background and unmodeled sources, and 4FGL sources.
- Calculate the likelihood for baseline and alternate model hypothesis. This is a product (sum in log space) of the likelihood for each independent bin of the data.

Likelihood ratio test and Wilks' Theorem

$$TS = 2 \ln \frac{\text{Likelihood for alternate hypothesis}}{\text{Likelihood for null (baseline) hypothesis}}$$

- Wilks' Theorem: in limit sample size n approaches ∞, test statistic for nested models is distributed like χ^2 for the degrees of freedom different between the models
  - TS = 2 Δ Ln L ~ χ^2 (N_dof)
  - TS = σ^2 in Gaussian limit
- Examples
  - Source TS: probability the source is present in the model as compared to not in the null model
  - TS_ext: compare a spatially extended source model hypothesis to a point source (radius=0)
  - TS_var: compare a model where flux is allowed to change to a model with constant flux

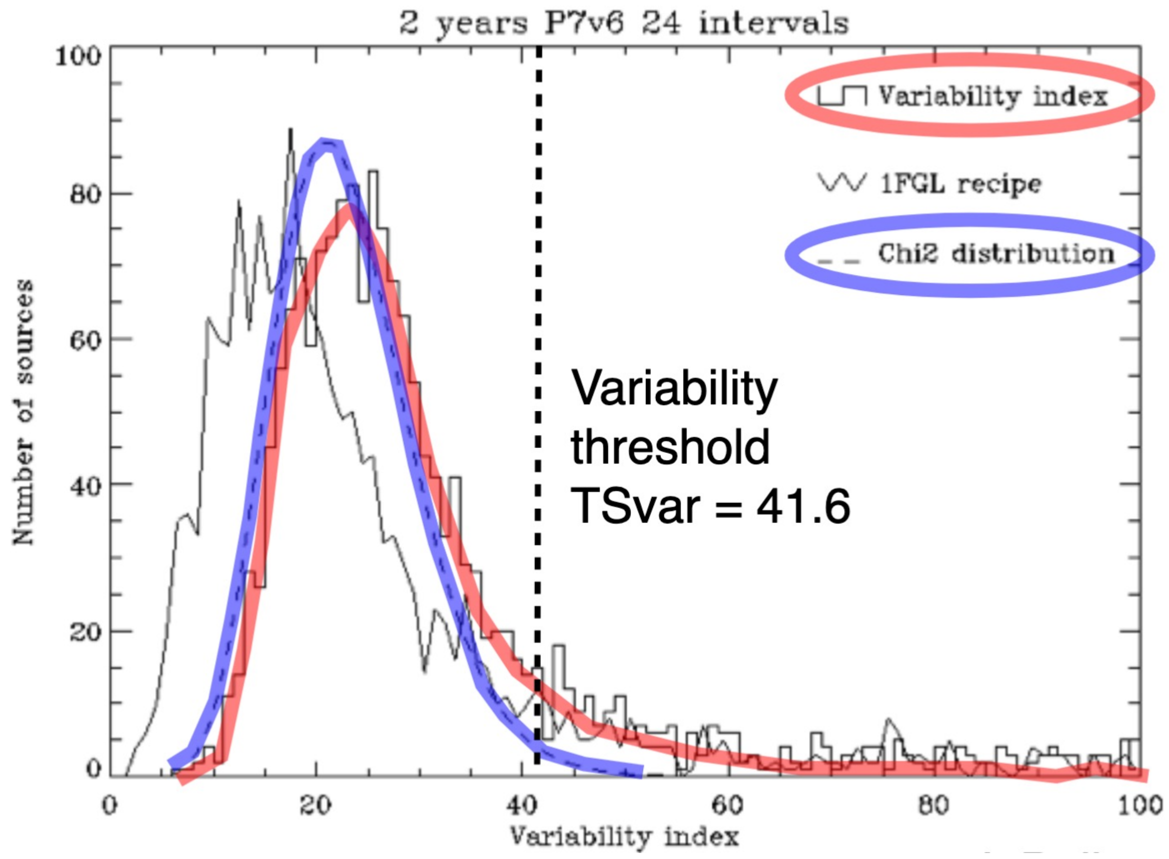

Comparing 2FGL Likelihood variability test to Chi2 distribution

Variability!

- We know the brightness of this type of source changes with time
- To answer science questions, we need to know how the source is changing with time. Let's measure the flux in a sequence of time intervals. Easy, right?
  - Note that to do Likelihood analysis in LAT, we have always been running a large number of Likelihood calculations (nbinsx*nbinsy*nbinsE). The likelihood that comes out of the LAT analysis is the sum for all of the spatial and energy bins, but we have been fitting the model parameters simultaneously for all of the bins together.
  - We will add even more bins in time, but examine the best-fit model parameters separately in those time bins

The Details: setting up your light curve

- Is the source variable?
- What time bins should you use?
- Is the source spectrum changing?
- How does your model of the region change when fitting short time intervals?
- Is your source near another variable source or a very bright steady source?
- Are there systematics that affect your analysis in important ways?

Variability Tests

- Obvious by eye variability: flux measurements differ >> than statistical errors + systematics
- Subtler variability
  - Chi-squared test of flux – consistent with a single value? (e.g. 1FGL)
  - Likelihood test for constant flux (e.g. 2FGL)
  - Tests of excess variance, example for ASCA data in [Nanda+ 1995]
- More complex questions
  - E.g., what time interval defines a period of constant rate? Bayesian blocks.

Variable sources from EGRET.
Can you guess which is a blazar? A pulsar?

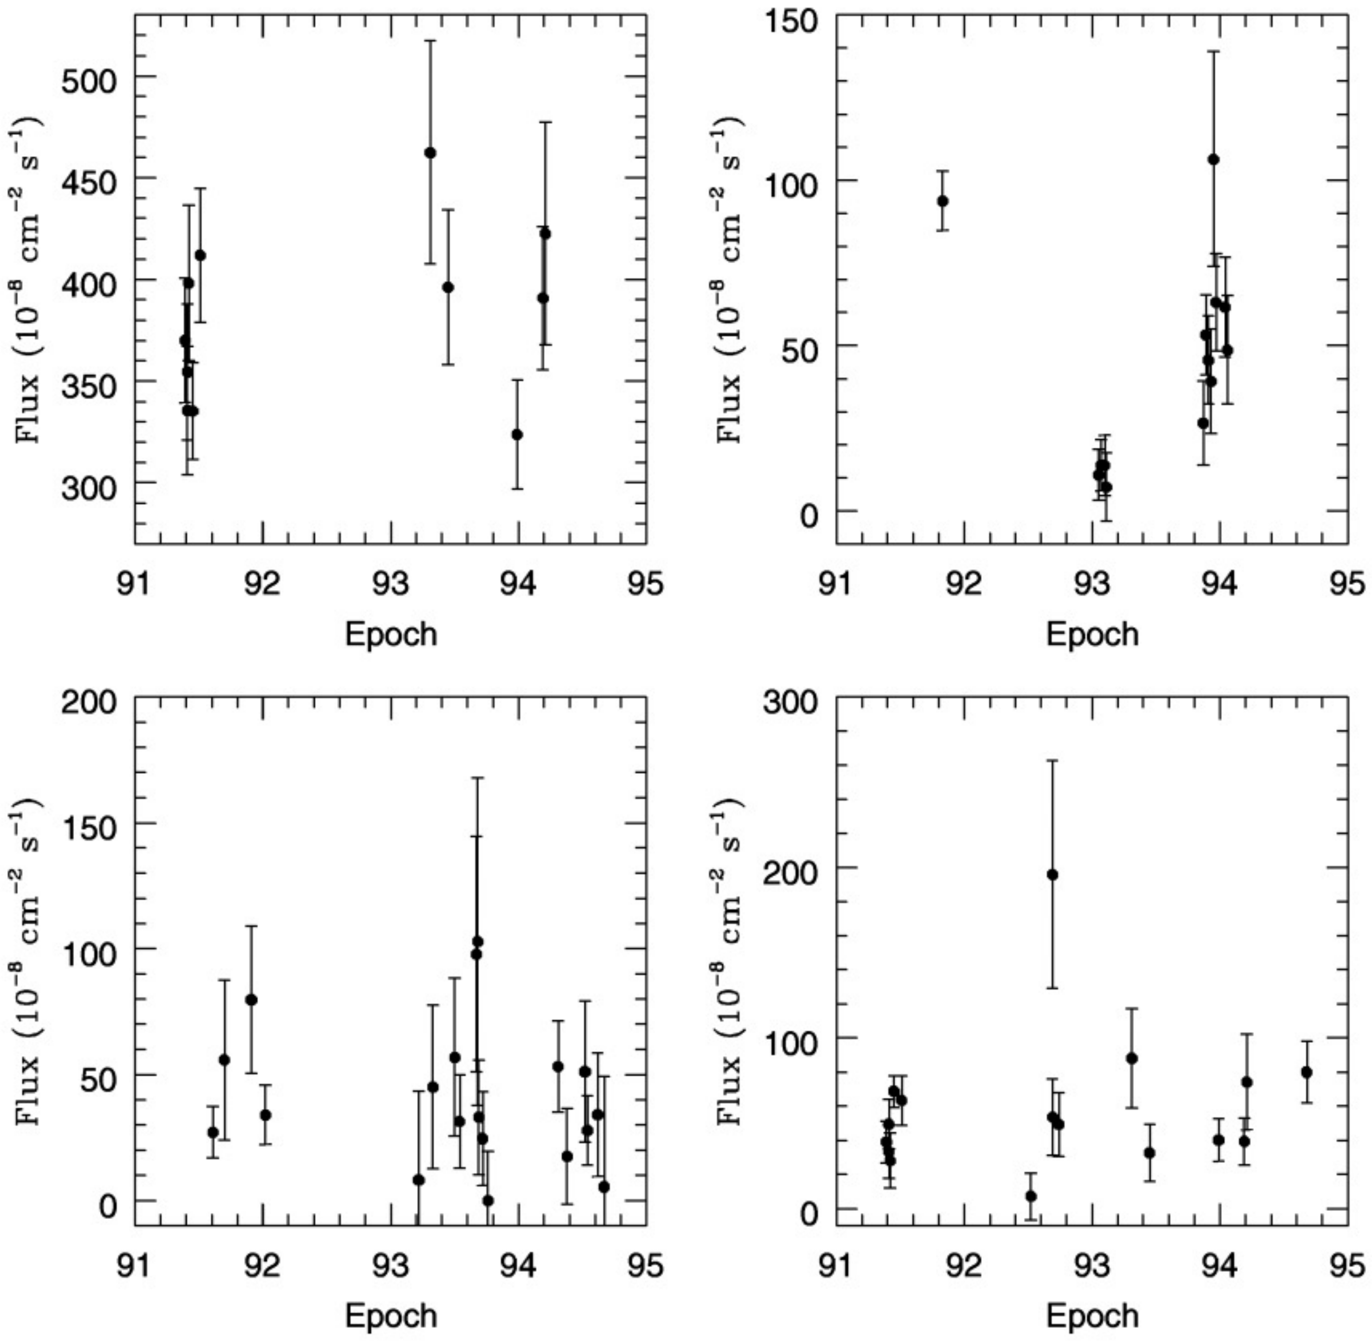

McLaughlin + 1996

Binning schemes

- Regular binning – fixed intervals
- Adaptive binning – adjust intervals to have comparable signal strength
- Bayesian blocks – avoid (or minimize) binning and find intervals of constant rate

Common model adjustments for light curves

- Fix spectral shape parameters for background sources to average values
- Fix all parameters for faint background sources to average values
  - For very short intervals, it may help to use a very simple model that ignores faint sources. Recall that 4FGL uses 8 years of data.
  - Very short intervals may also be good candidates for an unbinned likelihood analysis, i.e. each photon is a bin.
- Leave diffuse component normalizations free. The fit optimizer needs enough freedom in the model parameters to match the data.
  - In short intervals, interstellar emission (Galactic diffuse) and isotropic component may be ambiguous. It may help to fix one to average values
- Target photon index – free or fixed? Depends on what you want to measure.
- Consider removing time intervals if a nearby, bright flare of a source contaminates the measurement or if the Sun or Moon pass close to the target source.
- Check the analysis caveats for systematics that can show up in short intervals, e.g. short timescale effects for photons entering the instrument at different azimuthal angles.

A few temporal effects in Fermi LAT data

- ~96 minute orbital period and ~192 minutes for 2 orbits
- ~1 day modulations
- 53.4 orbital precession period 53.4
- 27.3 day lunar modulations (time periods when the Sun is close to a source can be removed from an analysis)
- ¼ year 91 days
- 1 year Solar modulations (time periods when the Sun is close to a source can be removed from an analysis)

https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_caveats_temporal.html

Some Important Light Curve Diagnostics

- Did the fit converge in all bins? The optimizer will return an answer, but did it find a good answer?
  - Check fit quality flags
  - Is a parameter value at a limit? This is a common cause of convergence issues.
  - There could be bins without data or with very limited data that provide less information to the fit.
- Reality checks
  - Test statistic values; Npred values; Photon index values (if free)
  - Diffuse normalization
  - Flux errors: should be proportional to flux. If the fit does not converge well they may be too small
  - Sensitivity: does the measurement make sense for expected LAT sensitivity?

Fit Quality

The default optimizer used for the likelihood fit in fermipy is Minuit. (You can choose other optimizers using configuration and fit method arguments.)

Minuit returns a fit quality code to indicate convergence of the fit and the accuracy of the error matrix.

0: Error matrix not calculated at all

1: Diagonal approximation only, not accurate

2: Full matrix, but forced positive-definite (i.e. not accurate)

3: Full accurate covariance matrix

Checking Flux Errors

Flux error/flux can't do better than the underlying count statistics. Comparing flux error/flux to Sqrt(npred)/npred exposes outliers.

This happens when the error matrix is not accurate. At some point the optimizer stops trying parameters and returns current values whatever they may be.

Could be caused by things like

- Initial model is too far from the best parameters
- Overconstrained - Parameter or parameters hit a constraint
- Too much freedom – optimizer is struggling to find a minimum because there is too much degeneracy and not enough information
- Too little freedom – not enough freedom in the model to adjust it to the data
- Something in the model is bad – spectral shape or spatial assumption wrong or a parameter is fixed to a bad value.

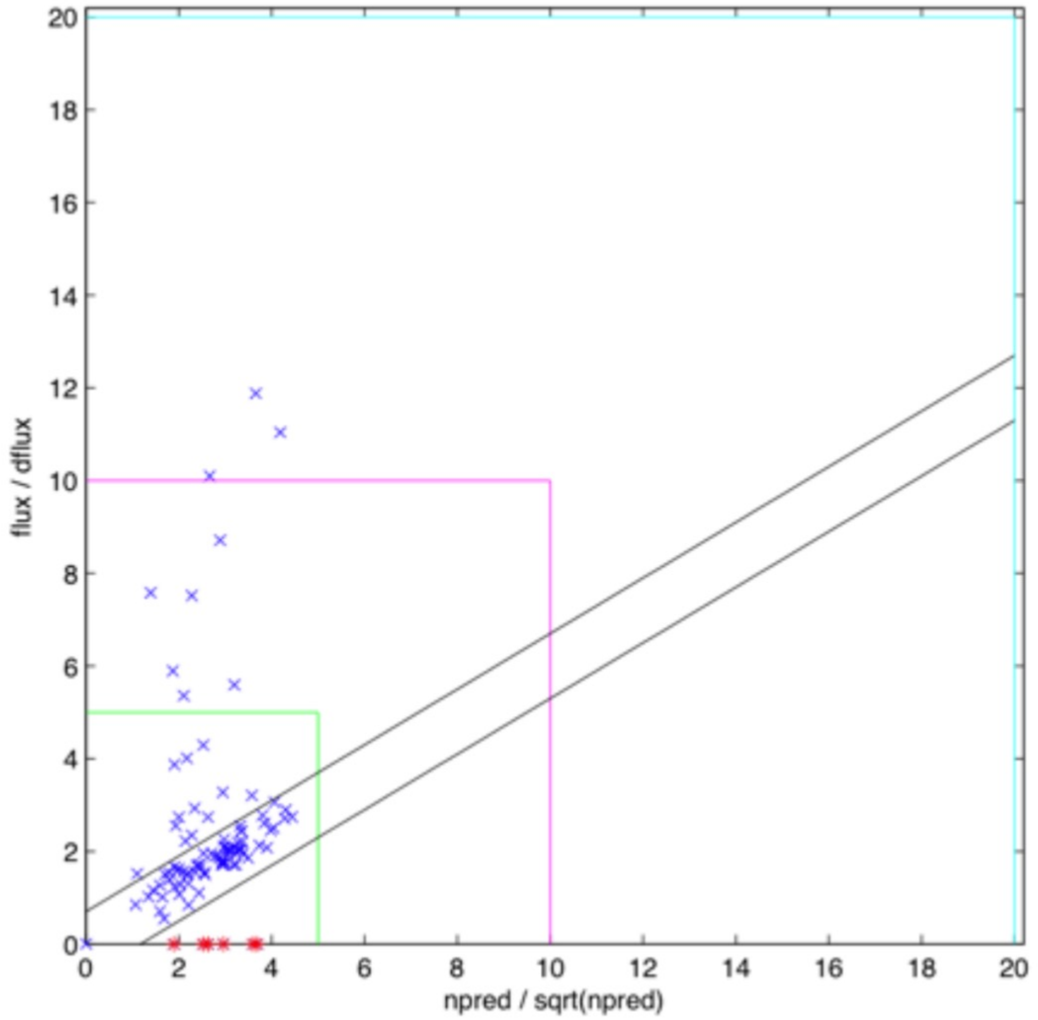

Exposure: Was the source visible and was LAT taking data?

Use Fermitools gtbin + gtexposure to give a quick estimate of daily exposure at the location of NHer 2021.

LAT exposure for a source at the 1-day cadence varies a little depending on location, survey mode, and orbital precession.

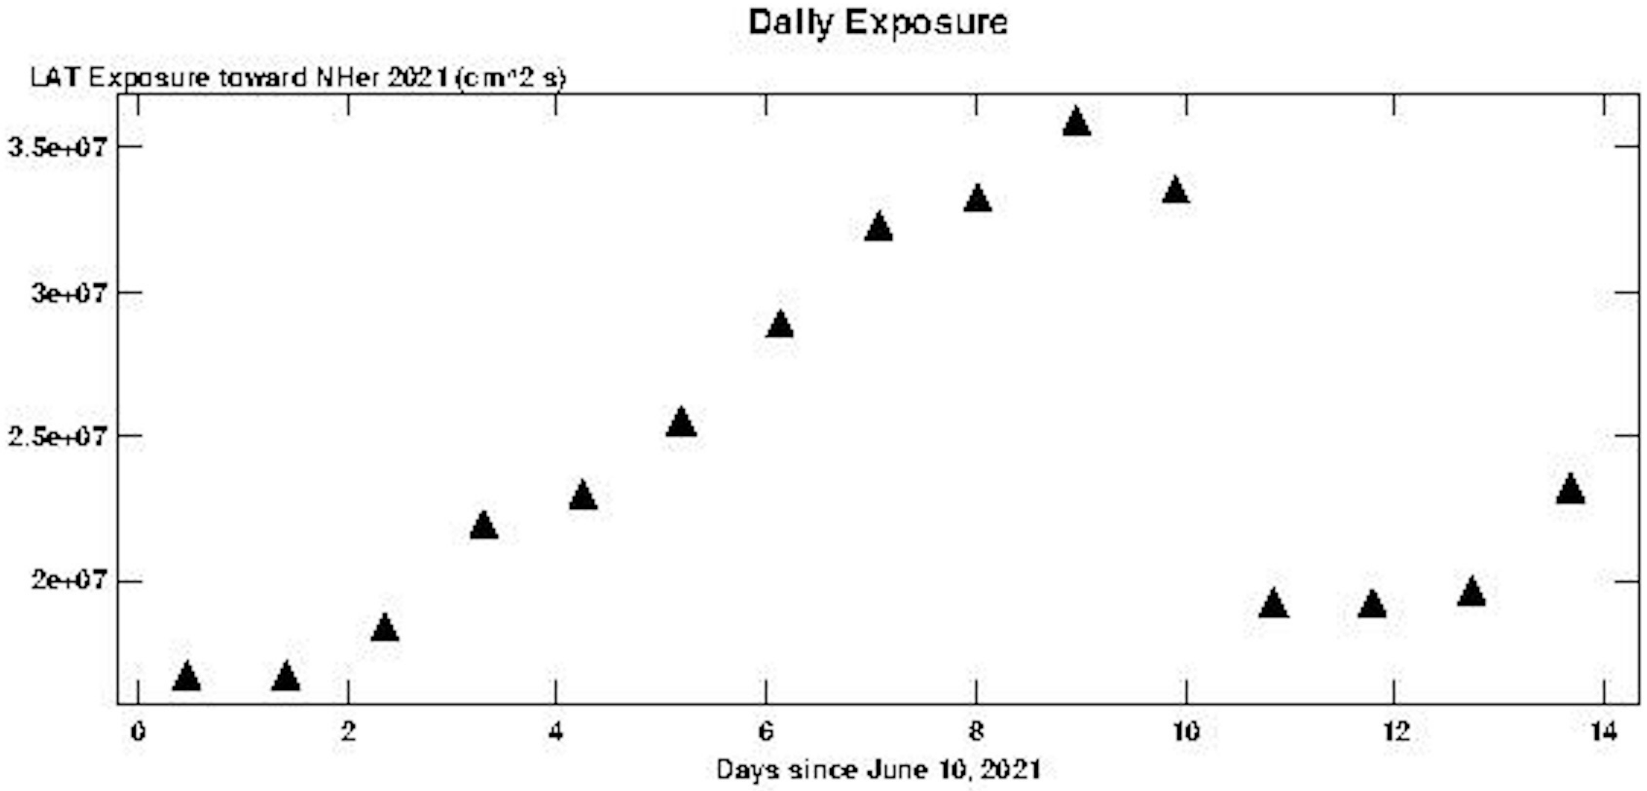

Fermi timeline posting at FSSC has history of observations and a tab for non-science intervals: https://fermi.gsfc.nasa.gov/ssc/observations/timeline/posting/

LAT Sensitivity

LAT Performance page

Point source sensitivity

The first plot shows the P8R3_SOURCE_V2 ten-year integral flux sensitivity for an isolated point source. The integral sensitivity is evaluated as the minimum flux above 100 MeV (in ph cm^{-2} s^{-1}) to obtain the 5-sigma detection in 10 years of observation in survey mode, assuming a power law spectrum with index 2. Quantitatively, the 5-sigma limit is taken as the likelihood test statistic of 25 and at least 10 photons. The sensitivity calculation is based on an exposure map from the first ten years of the mission (August 2008-August 2018) produced with a zenith angle cut of 100 degrees and the standard gtmktime filter selection (DATA_QUAL>0 && LAT_CONFIG==1). This calculation takes into account the loss of observing time due to SAA passages and the instrumental deadtime. Only diffuse backgrounds are considered in the calculation (the effect of the point sources is ignored). The calculation uses gll_iem_v07.fits and iso_P8R3_SOURCE_V2.txt for the galactic IEM and isotropic templates respectively. The semi-analytical prescription is described in the appendix of the 2FGL paper.

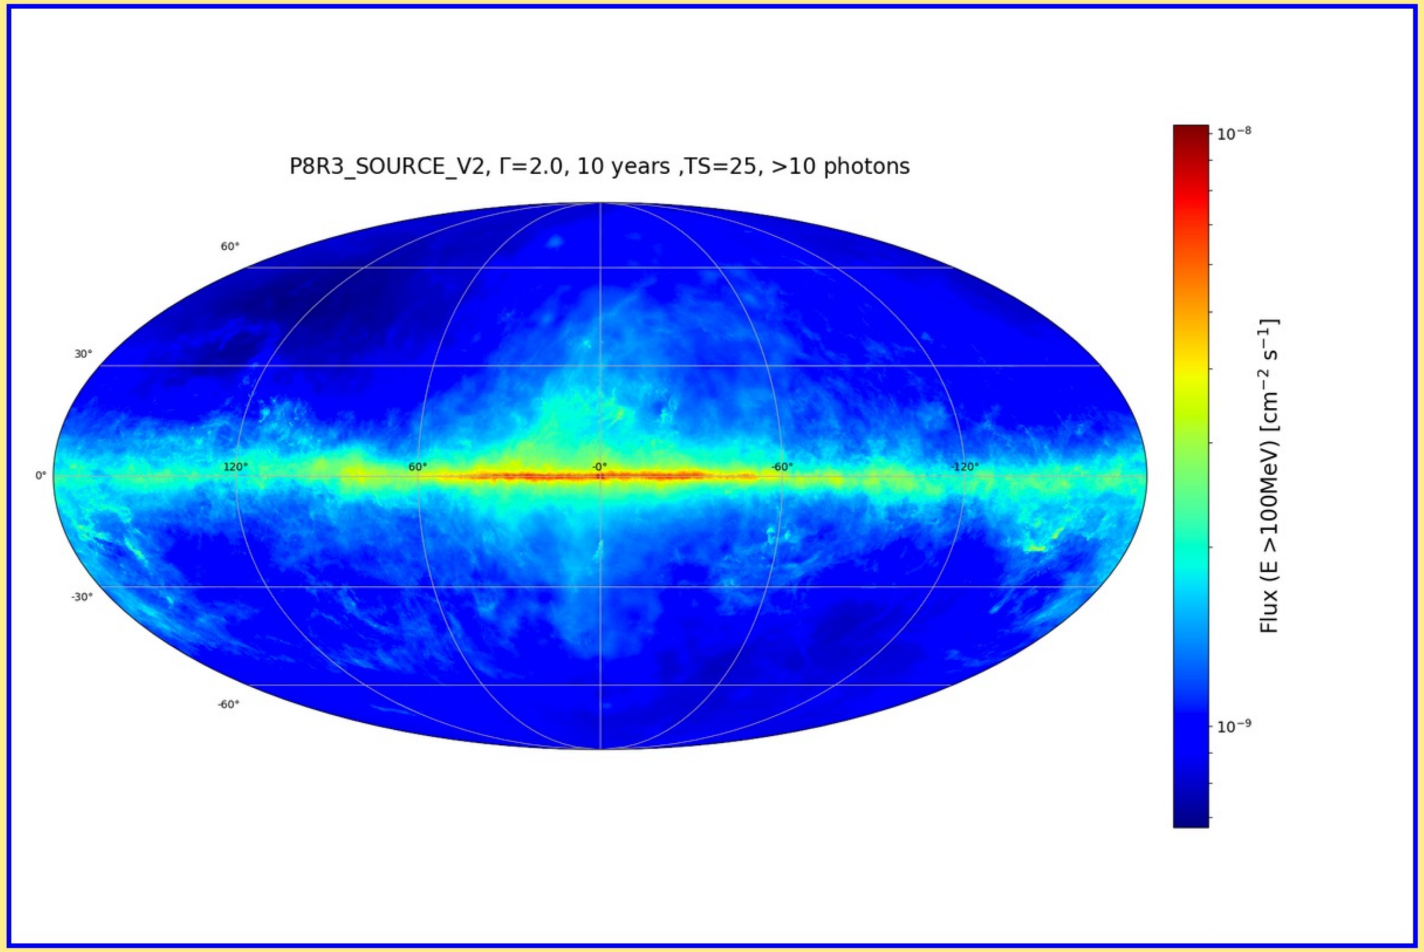

LAT Example – Finding Nova Her 2021

- Prior: the 4FGL catalog represents the gamma-ray sky.
- Hypothesis: a new gamma-ray point source at the location of Nova Her 2021, a Galactic nova that appeared in optical observations on June 11, 2021.
- Does the gamma-ray data from LAT support the hypothesis?
  - Fit the 4FGL catalog to my data set.
  - Fit the 4FGL catalog + a point source at the position of Nova Her 2021.
  - Likelihood ratio test gives a test statistic for the model containing Nova Her as compared to the model without it, TS = 19 (npred=34.7) in my initial data set.
    - Are these nested? Yes! I can vary the flux normalization for Nova Her 2021 to 0 to make the new model match the baseline model

Nova Her 2021 Light Curve – first look

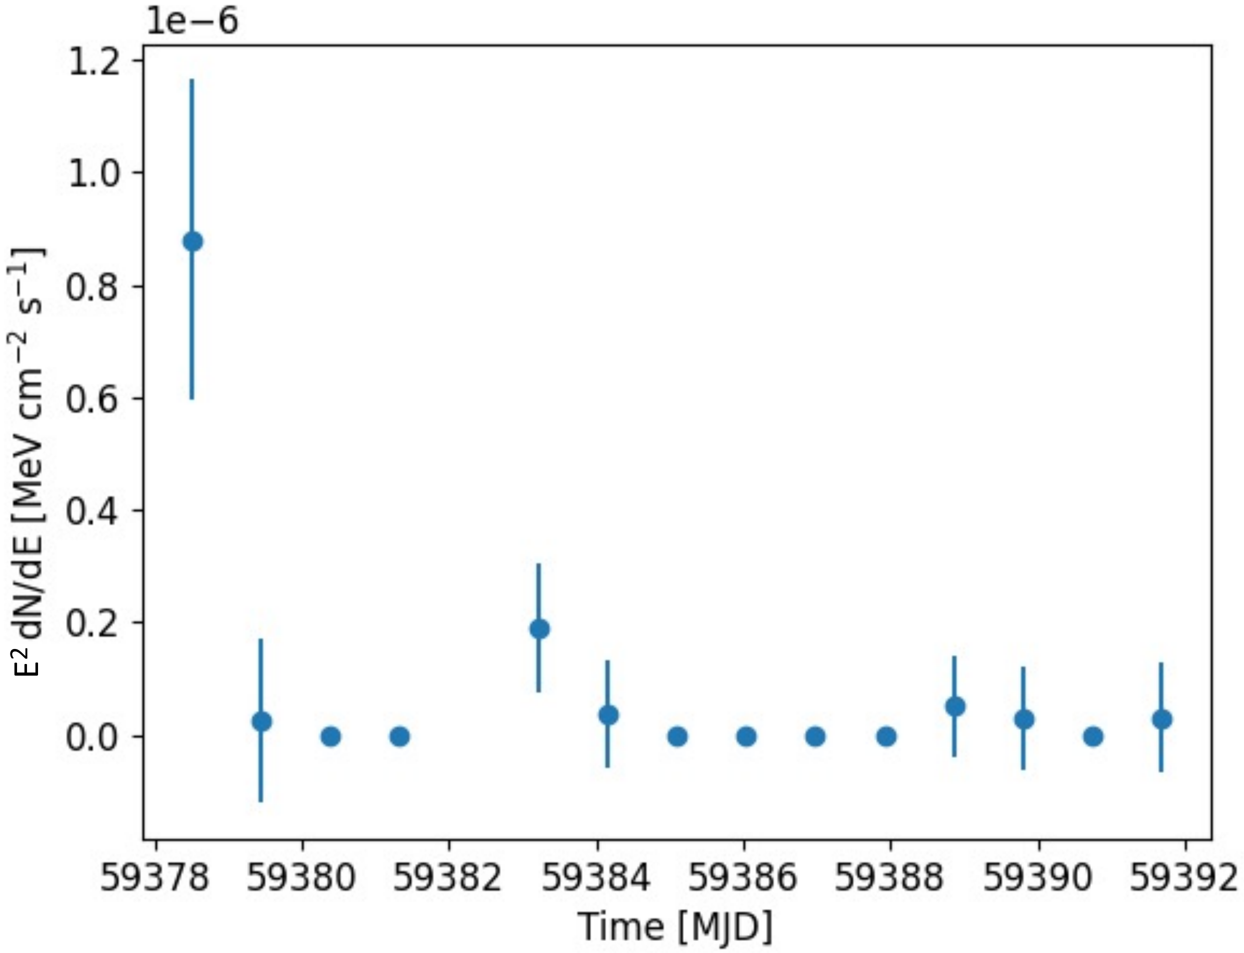

Good news:
June 13 may have a solid detection. Flux is reasonable given LAT's typical 1-day sensitivity.

Bad news:
The nova faded quickly. These other points are not detections and upper limits should be calculated. I need to improve my script.
Also should add earlier data.

Default flux value returned by the lightcurve method is integrated over the energy range for my analysis, 100 MeV to 100 GeV.

Light curve details: bookkeeping

- You can create an analysis script that runs a likelihood analyses for each time interval that you want in your light curve.
  - This requires organizing a large number of files. Each time bin has a unique exposure. This means we have to build likelihood inputs (binned exposure maps and source model maps) for each time bin.
- Fermipy helps by automating the setup and calculations for the likelihood analysis.
  - If you get errors about fit success, this is because there is an issue for handling bins with no data. You can find a fix for this issue at the fermipy github until it there is a fixed in a tagged release.
  - To make light curves, you probably want to use a python script run from the command line and not a notebook.
    - Put the python commands from your notebook in a text file and run it from the command line in FermiBottle, e.g. *python runLC.py &> LCoutput.txt*
    - For command line scripts, skip commands that display something to the screen or change to write plots to a file, e.g. use the matplotlib.pyplot method savefig("figname.png")

Notes on analysis in FermiBottle

- Kernel dying in Jupyter notebook? Your container may be running out of resources. I increased the memory for my docker container from 2 GB to 4 GB. (Use Docker Desktop preference settings instructions for your OS to do that.)
  - Likelihood analysis tends to take a lot of memory because it stores some items in memory to speed up calculations.
- Running analysis as a python script is usually faster than in a Jupyter notebook. You can write out fit objects in the script and then load them into a notebook to work on plots and look at outputs interactively.
- For computationally intensive analysis with many steps or bins, it is easiest to use a cluster of several computer nodes.

Week 4 Group Activity:
Light Curves gone wrong!

- Go to Gather and pick a room. If there are more than 6-8 people, pick a different room to balance numbers.
- Introduce yourselves to each other.
- Pick one person to present your group discussion when we return to the main meeting room.
- Discuss the light curve questions on the slides posted.
- If you can’t get into Gather, come to the main Teams meeting for discussion.