University of CINCINNATI

PISCES

(Parameter Inference with Systematic Covariance and Exact Statistics)

J Hewes

Meeting of the Division of Particles and Fields 2021

13th July 2021

Poisson log likelihood test statistic

$$-2\log\mathscr{L} = 2\sum_{i}^{N}\left[m_i - x_i + x_i \log\left(\frac{x_i}{m_i}\right)\right]$$

- Correct statistical treatment at low statistics.
- Systematic uncertainties reweight simulated spectrum and incur χ^2 penalty.
- Both physics parameters and systematic uncertainties profiled by fitter.
    - Computationally intensive for a large number of degrees of freedom.
- Impact of systematic uncertainties on spectrum can be visualised.

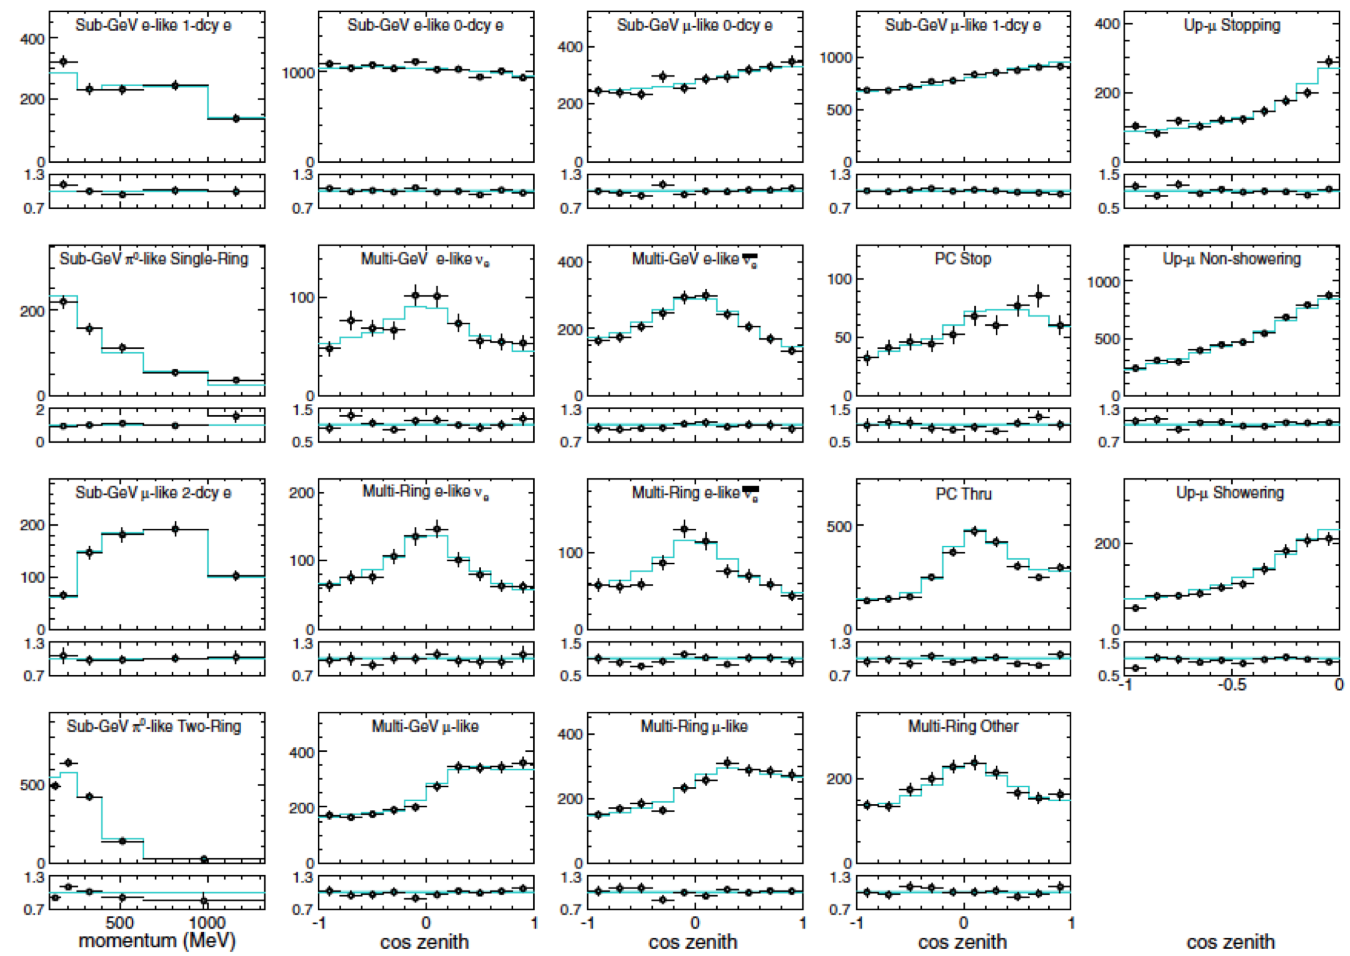

Super-K 3-flavour results (fit contains 520 analysis bins, 155 systematic uncertainties)

Gaussian multivariate χ^2 test statistic

$$\chi^2 = (x_i - \mu_i) V_{ij}^{-1} (x_j - \mu_j)$$

- Incorrect statistical treatment at low statistics.
- Systematic uncertainties encoded into covariance matrix.
- Only physics parameters profiled by fitter.
  - Adding more systematic uncertainties to the matrix does not affect fit time.
- No way to visualise impact of systematic uncertainties.

MINOS & MINOS+ sterile neutrino analysis (arxiv:1710.06488)

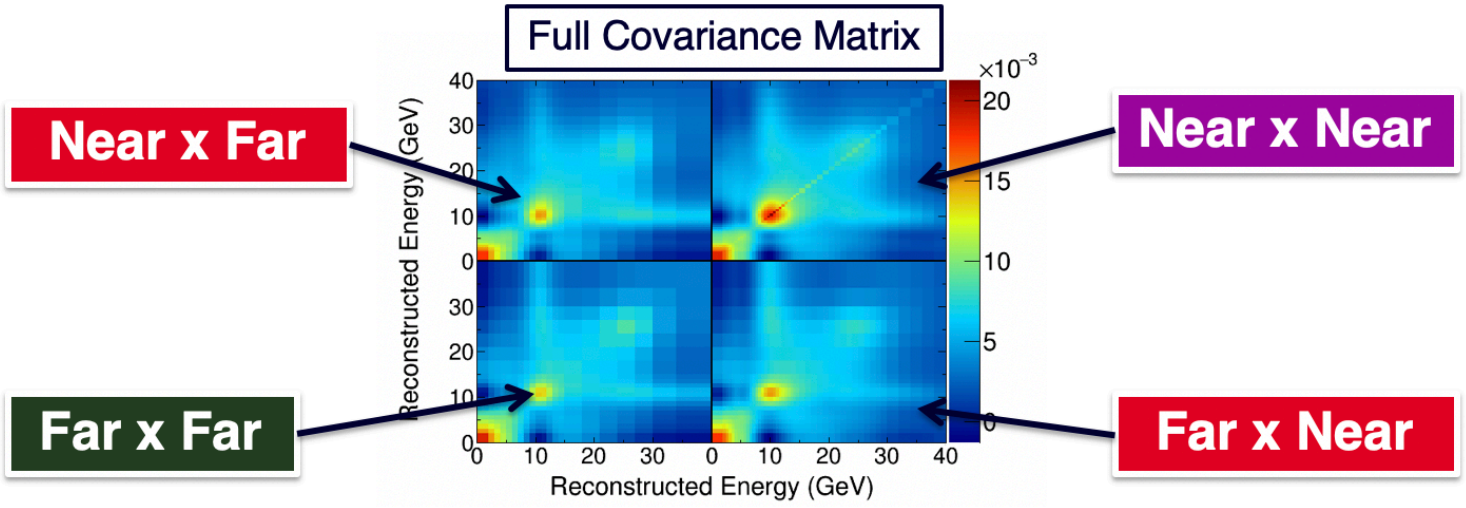

Why PISCES?

- Both the Gaussian multivariate and Poisson log-likelihood test statistics have benefits and drawbacks.

- However, it is also possible to construct a **hybrid approach that leverages the advantages of both**.
  - One can form a Gaussian log-likelihood, and then perform transformations to decouple the statistical and systematic components.
  - One can then replace the statistical component of the Gaussian likelihood with the Poisson formalism.
  - Finally, one utilises the covariance matrix to perform a fast minimisation over the ensemble of systematic uncertainties, finding an optimal systematic pull in each analysis bin.

- The resulting hybrid method is **fast and memory efficient**, **treats statistical uncertainties correctly** at low statistics, **provides systematic pulls** in each analysis bin, and **does not become exponentially slower** as the number of systematic degrees of freedom increases.

PISCES

- Statistical uncertainty is provided by a comparison of the data to the systematically shifted prediction via a Poisson likelihood:

$$\chi^2_{\text{stat}} = 2\sum_i^N \left[\left(\sum_\alpha^M \mu_{\alpha i} s_{\alpha i} \right) - x_i + x_i \log \left(\frac{x_i}{\sum_\alpha^M \mu_{\alpha i} s_{\alpha i}} \right) \right]$$

- An additional penalty term is applied to penalise the systematic uncertainties for pulling away from nominal:

$$\chi^2_{\text{syst}} = \sum_{ij}^N \sum_{\alpha\beta}^M (s_{\alpha i} - 1) V^{-1}_{\alpha i \beta j} (s_{\beta j} - 1)$$

- The final test statistic is the combination of the **Poisson likelihood statistical term** and the **Gaussian multivariate systematic term**:

$$\chi^2 = \chi^2_{\text{syst}} + \chi^2_{\text{stat}}$$

i = analysis bin
α = beam component
μ = nominal prediction
s = systematic shift
x = data
V = covariance matrix

PISCES

- Statistical uncertainty is provided by a comparison of the data to the systematically shifted prediction via a Poisson likelihood:

$$\chi^2_{\text{stat}} = 2\sum_i^N \left[\left(\sum_\alpha^M \mu_{\alpha i} s_{\alpha i} \right) - x_i + x_i \log \left(\frac{x_i}{\sum_\alpha^M \mu_{\alpha i} s_{\alpha i}} \right) \right]$$

Exact Statistics

- An additional penalty term is applied to penalise the systematic uncertainties for pulling away from nominal:

$$\chi^2_{\text{syst}} = \sum_{ij}^N \sum_{\alpha\beta}^M (s_{\alpha i} - 1) V^{-1}_{\alpha i \beta j} (s_{\beta j} - 1)$$

Systematic Covariance

- The final test statistic is the combination of the **Poisson likelihood statistical term** and the **Gaussian multivariate systematic term**:

$$\chi^2 = \chi^2_{\text{syst}} + \chi^2_{\text{stat}}$$

Parameter Inference

i = analysis bin
α = beam component
μ = nominal prediction
s = systematic shift
x = data
V = covariance matrix

Toy example

- Construct samples that represent a pair of selections, each applied in a high-statistics context and a low-statistics context.
  - In the context of neutrino physics, this can be interpreted as something similar to a ν_μ and ν_e selection applied in a near and far detector.

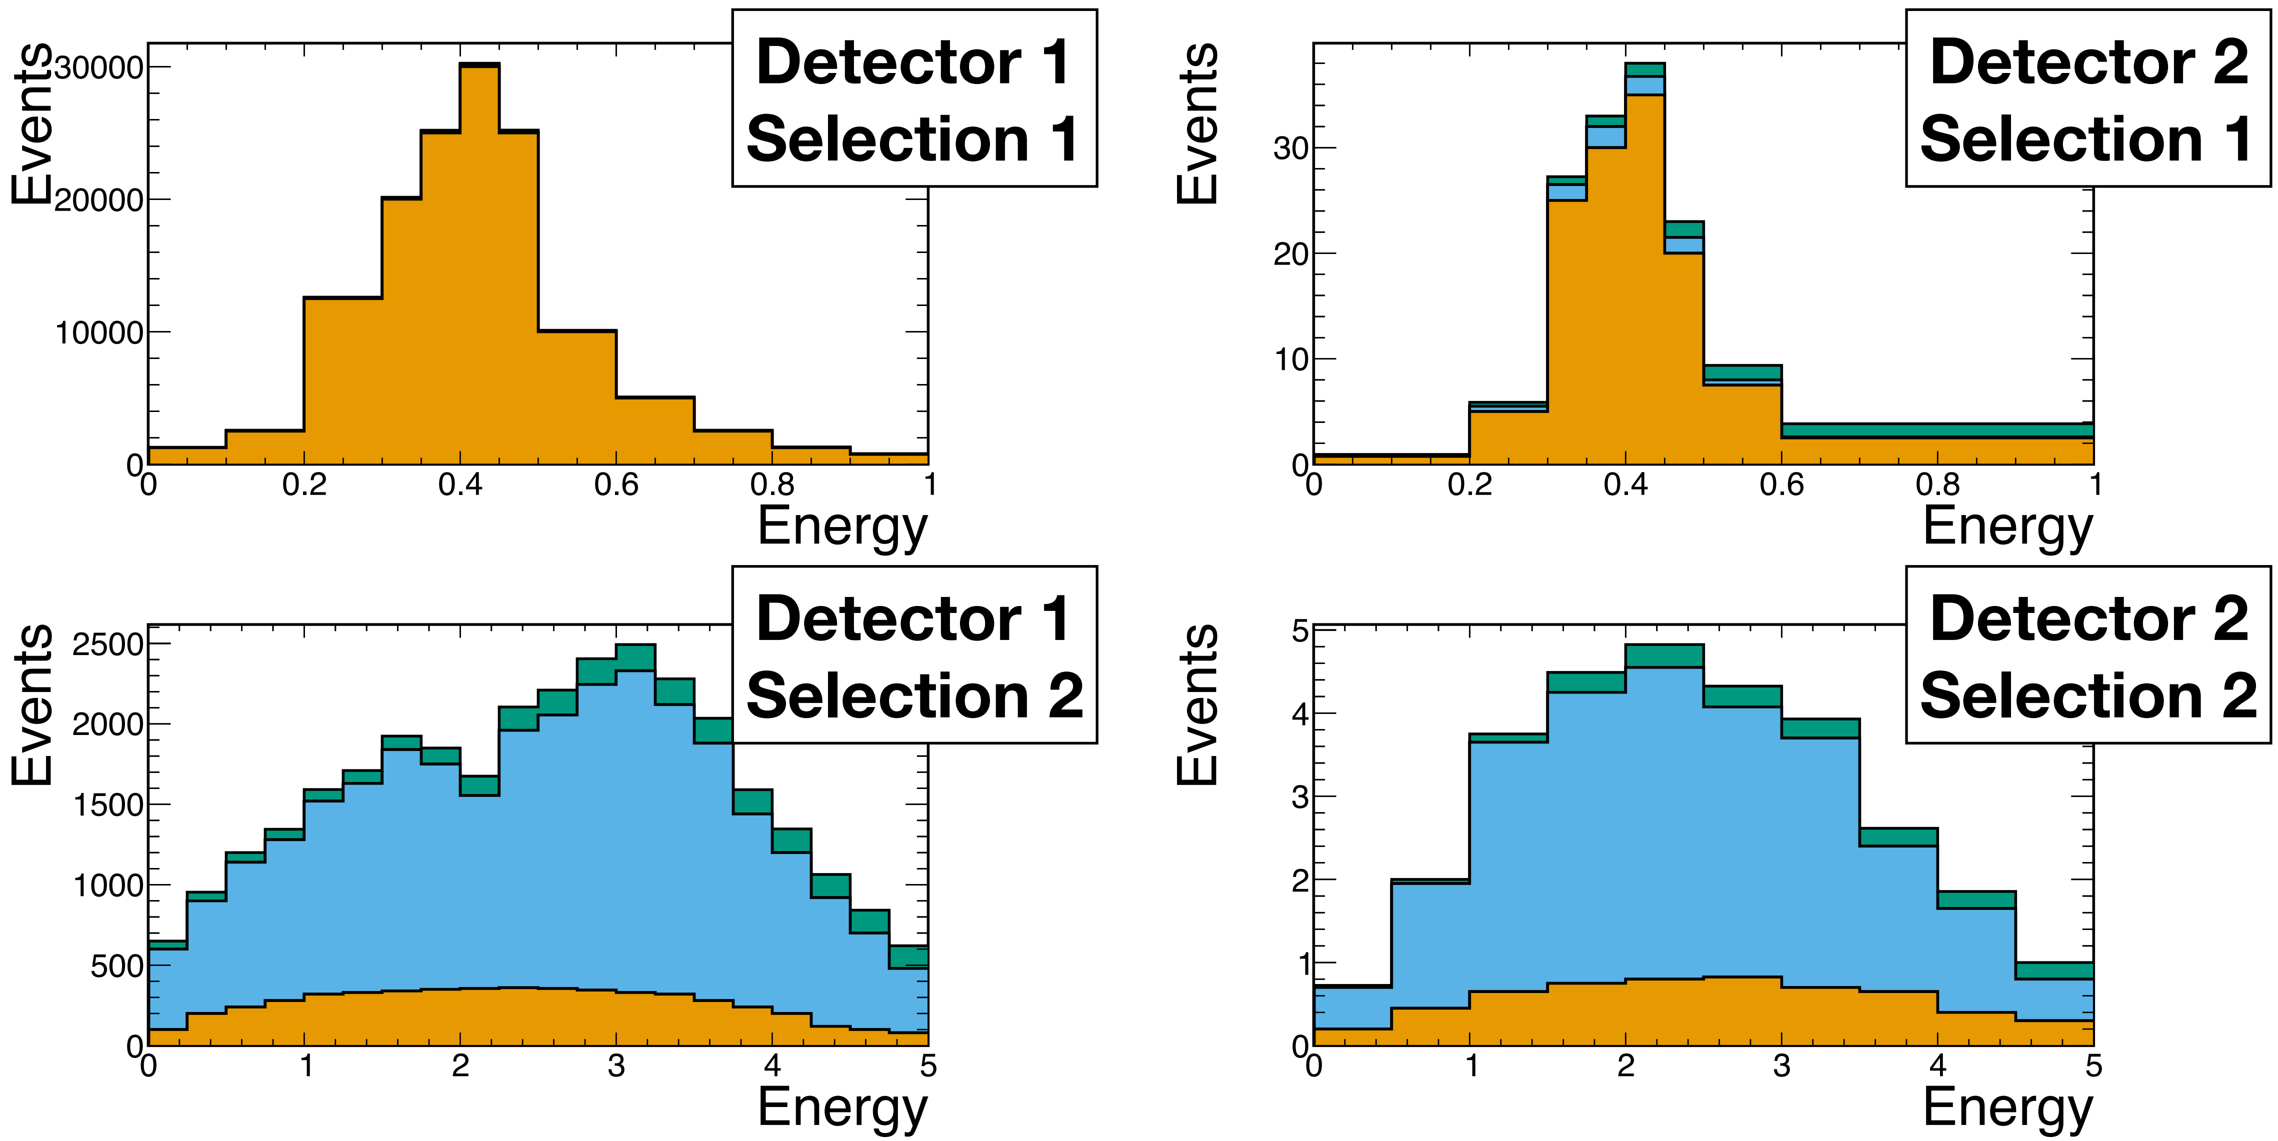

Systematic uncertainties

- Define toy systematic uncertainties to fit to.
  - 5% normalisation uncertainty, uncorrelated between samples but fully correlated between channels within each sample.
  - Define cross-section-like uncertainties drawn from Landau PDFs with different means, widths and amplitudes.
    - Each uncertainty affects only a single channel, but is fully correlated between samples.

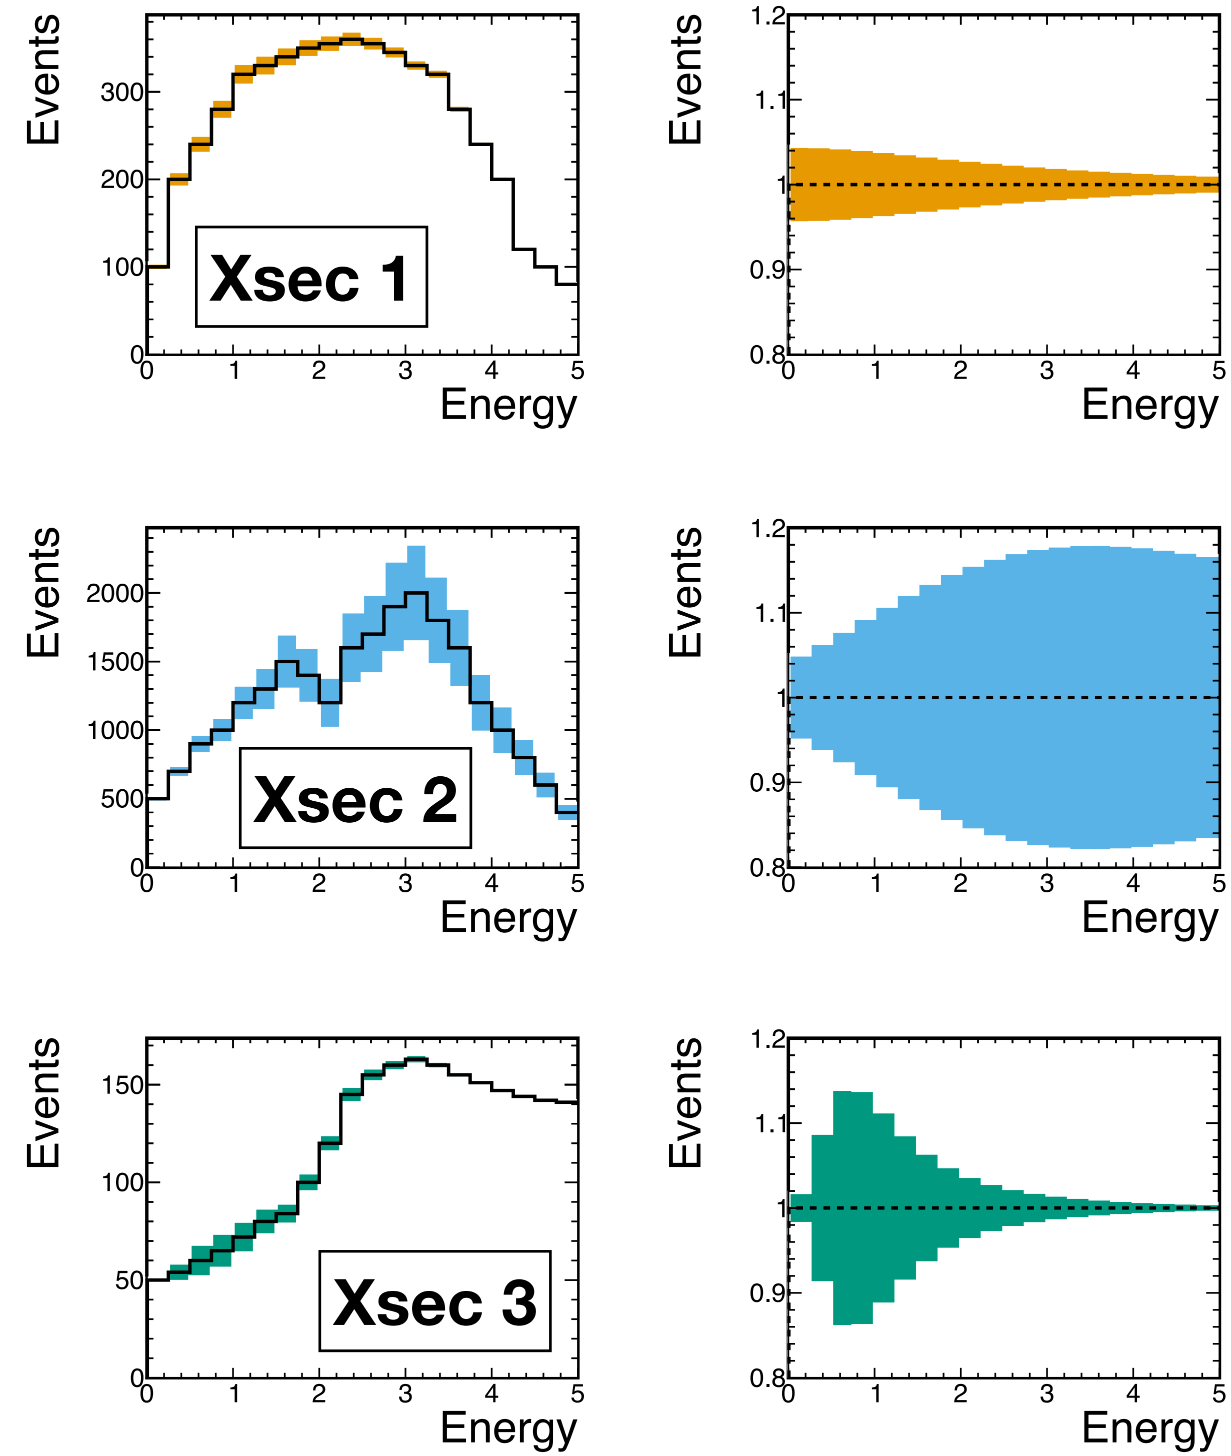

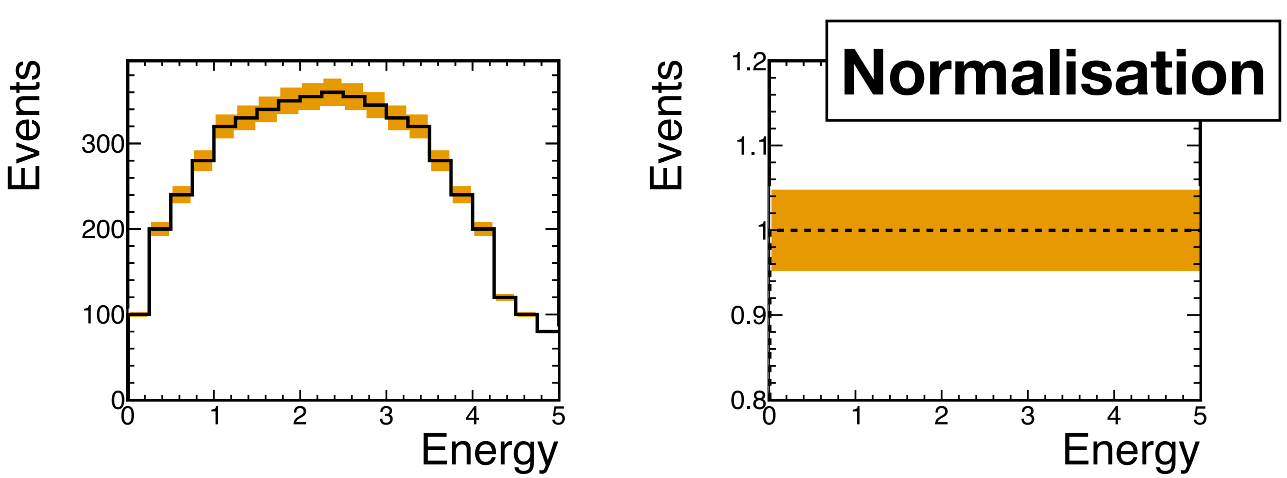

Covariance matrix

- We can now encode all of our systematic uncertainties into a covariance matrix between analysis bins in each sample.
  - For PISCES, we also bin independently in each channel.

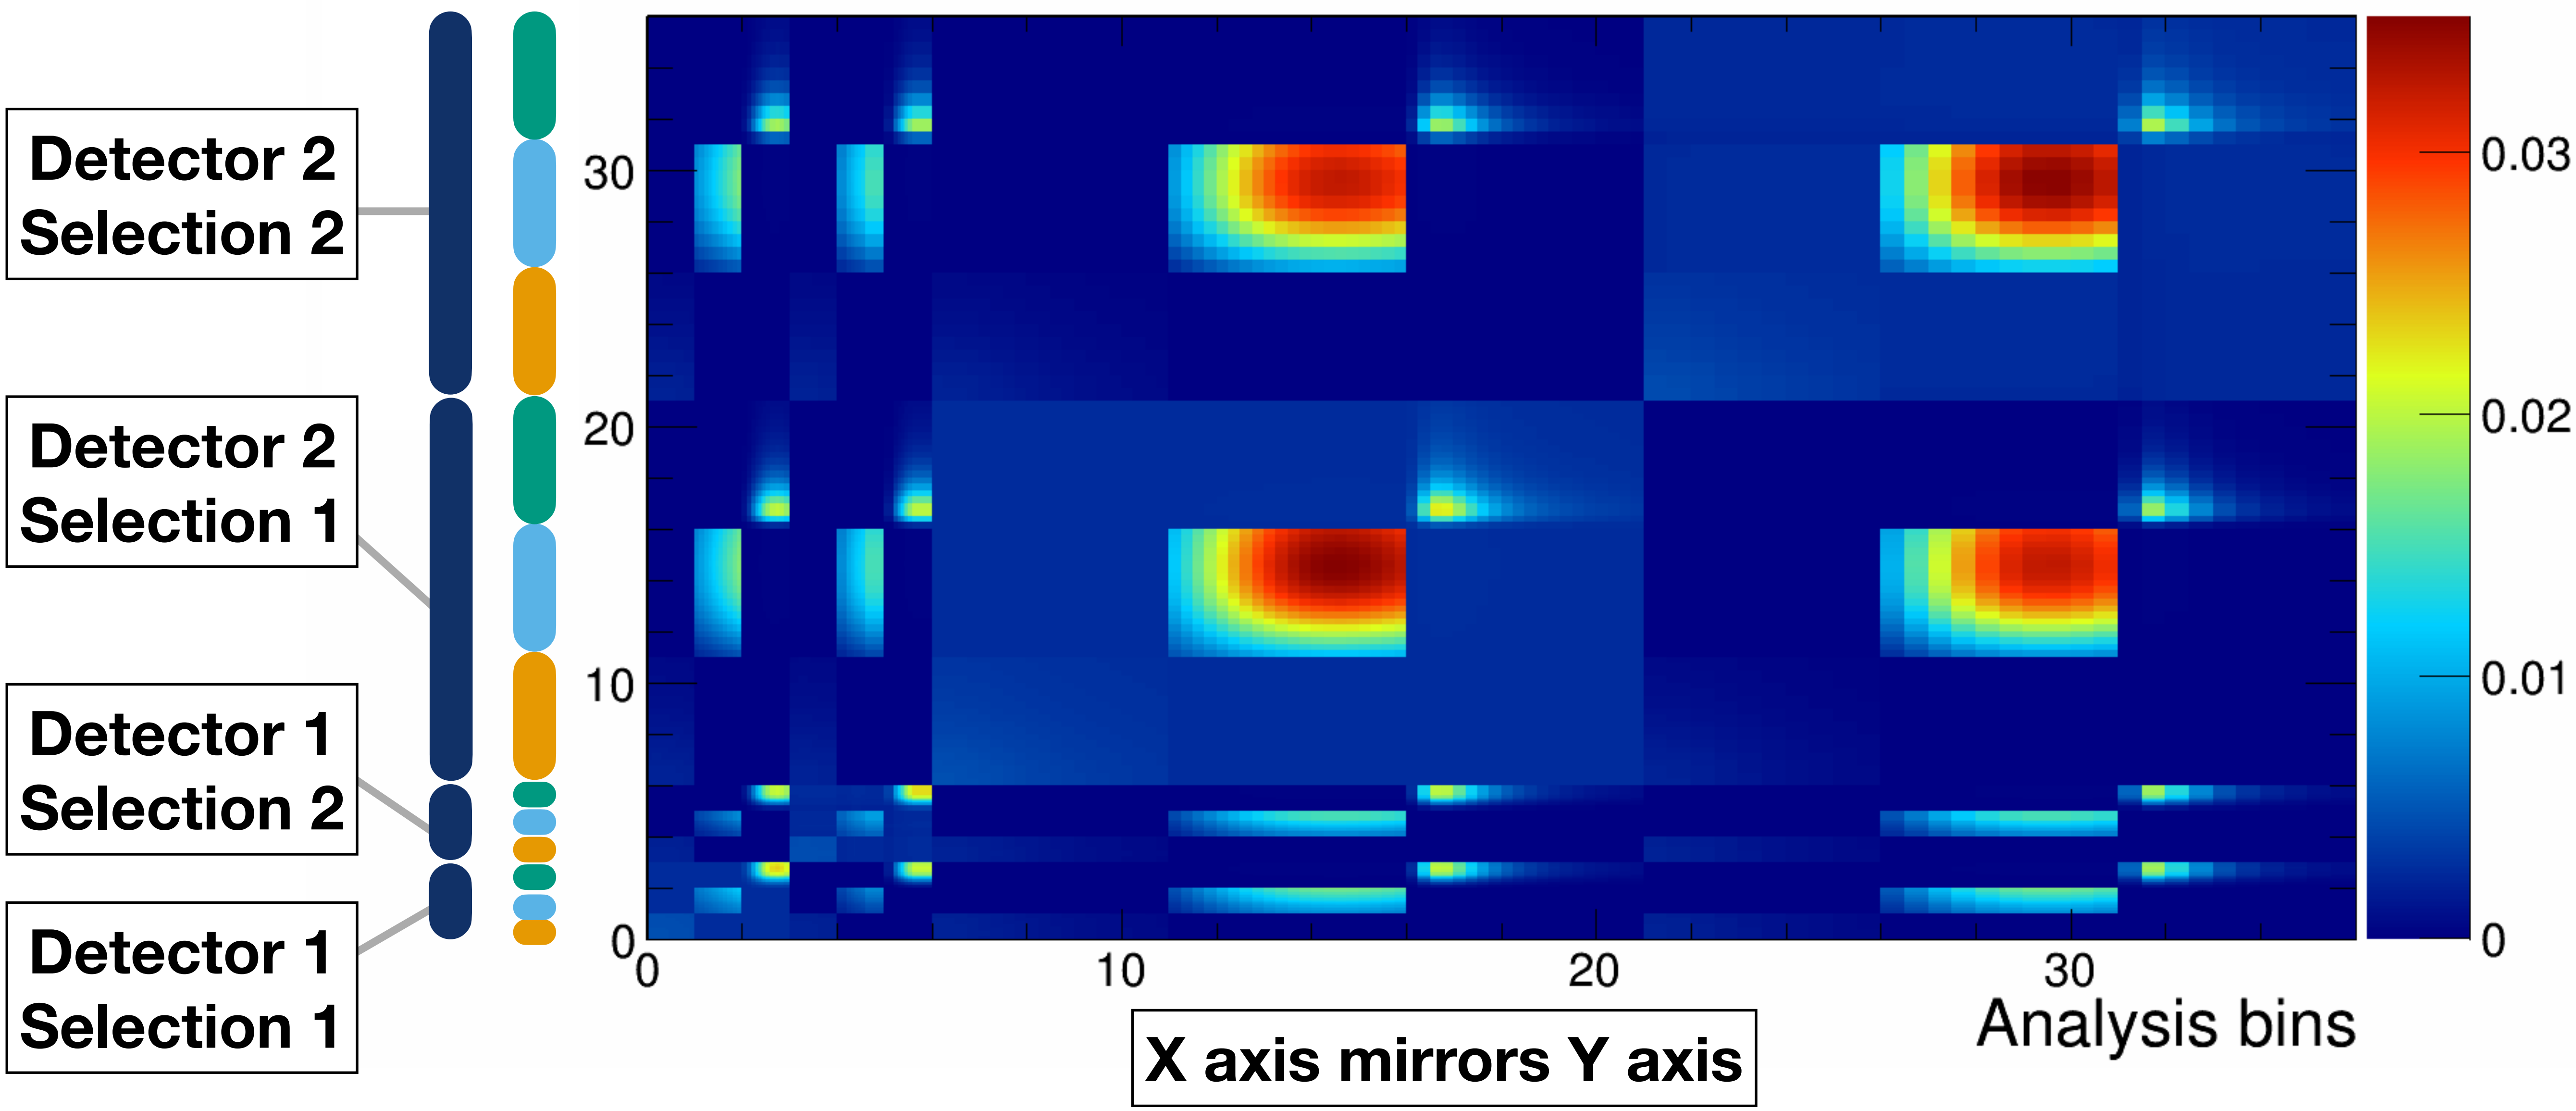

Demonstration

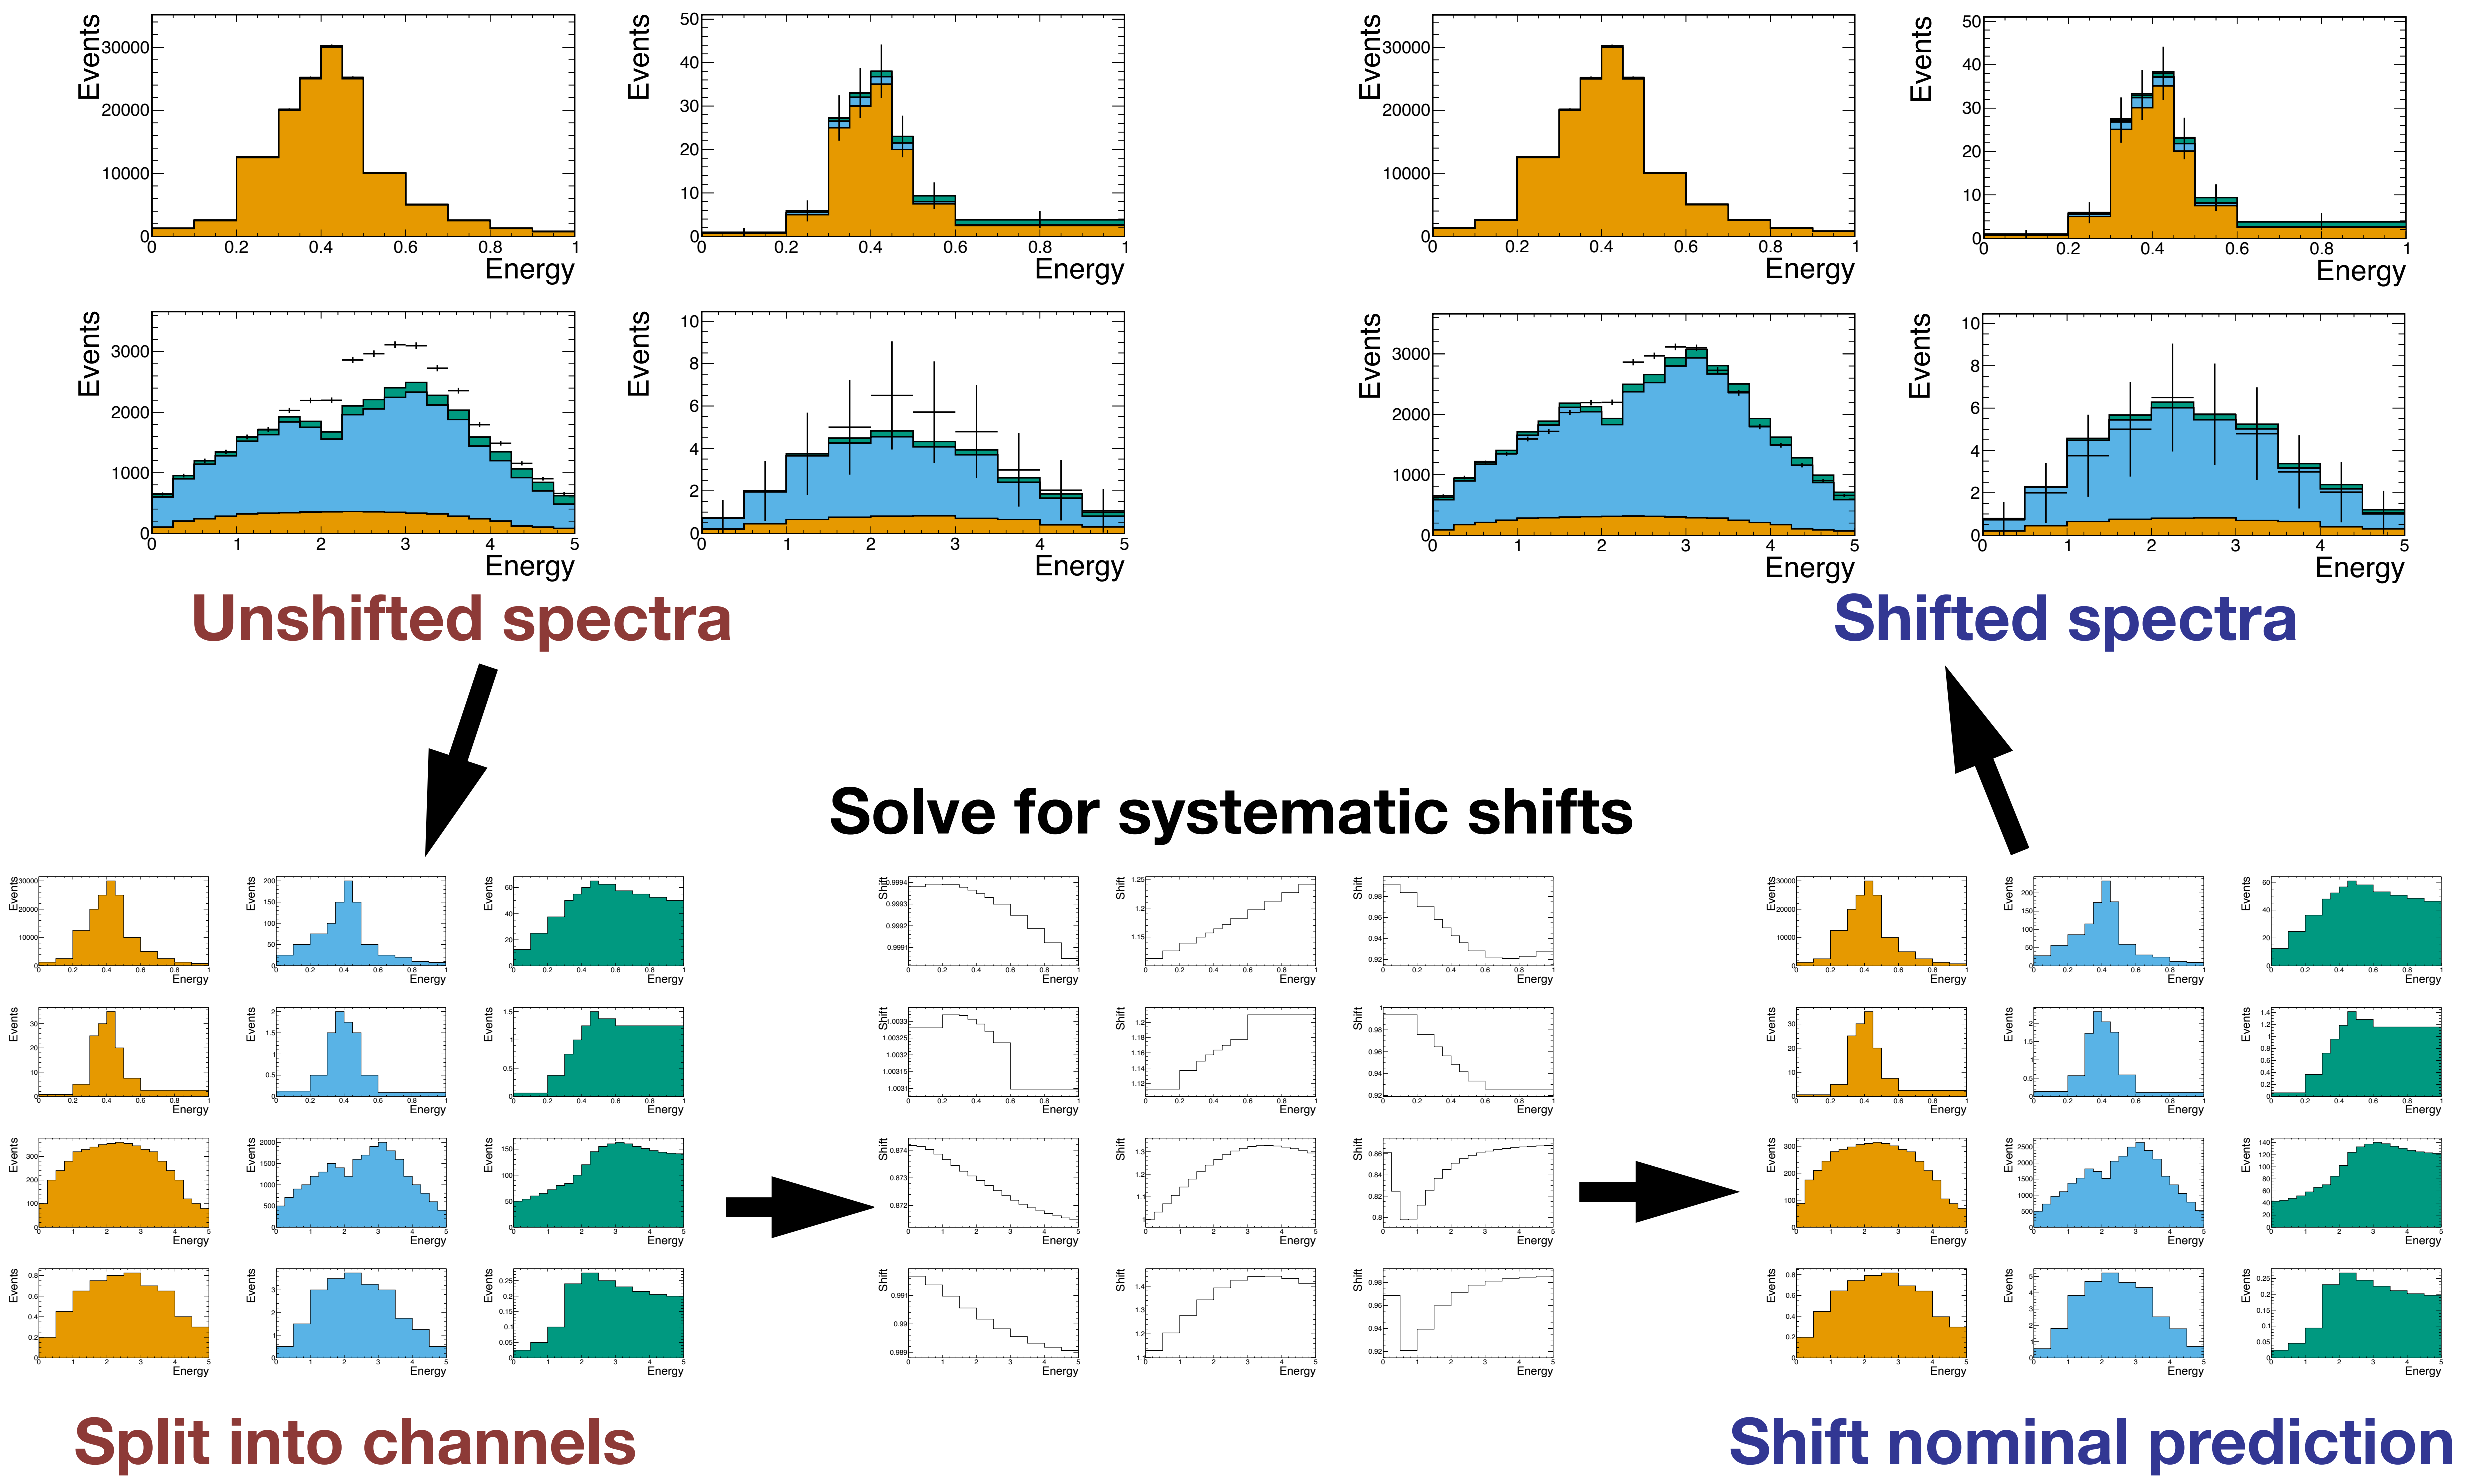

Summary

- PISCES is a novel test statistic for performing fits to low-statistics data.

- It combines the benefits of the likelihood and covariance matrix techniques:
    - Efficient for a large number of systematics.
    - Simple to perform joint fits across many samples.
    - Provides systematic pulls for each analysis bin.
    - Correct treatment of statistical uncertainties.

- PISCES is a fast and memory-efficient technique.
    - Natural candidate for GPU optimisation, which should speed it up even further.

- This method is currently being utilised in the context of neutrino oscillation analyses in NOvA and DUNE, but can be applied to a broad range of analyses in neutrino physics and more broadly in HEP.

Backup

Covariance matrix

- Covariance matrix with all channels summed for each sample.

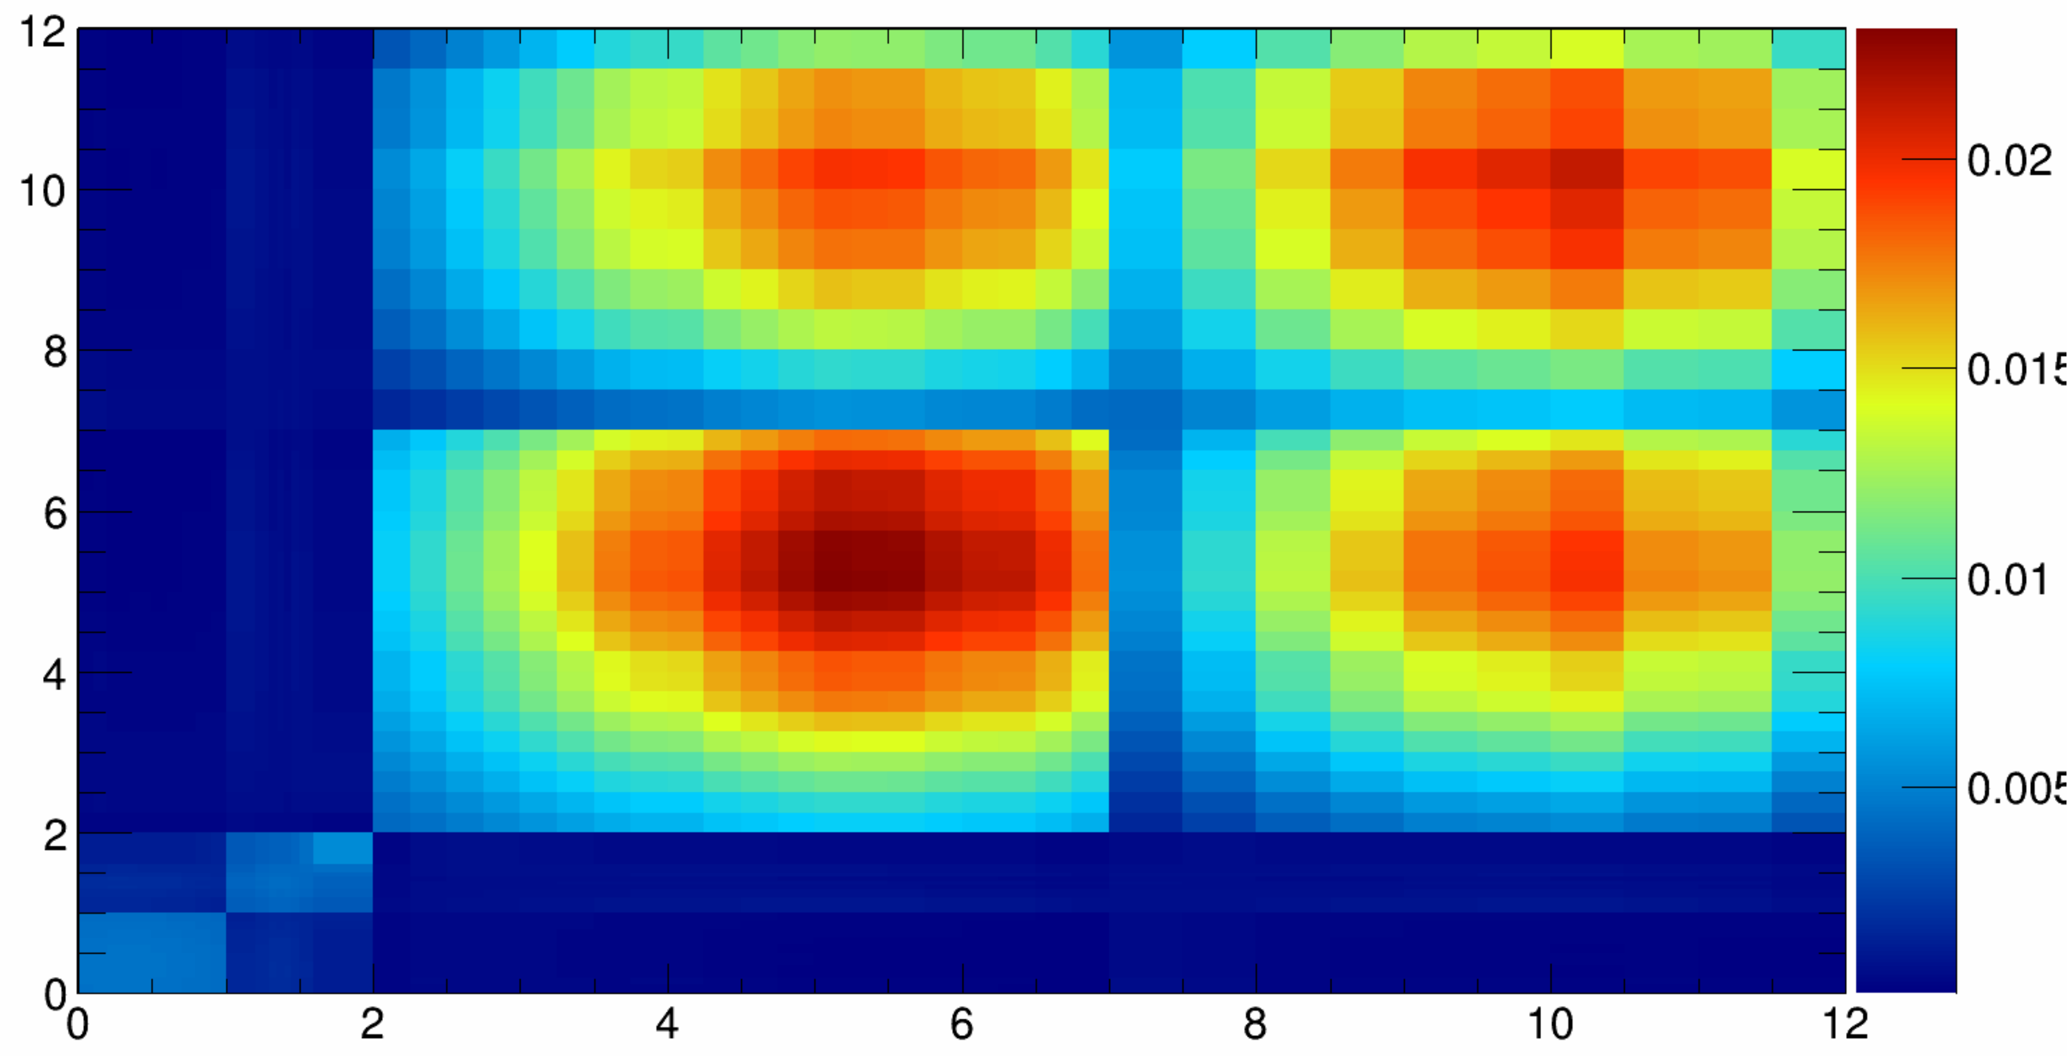

Gaussian multivariate likelihood

- The Gaussian log-likelihood expression is

$$\mathscr{L} = \frac{1}{\sqrt{(2\pi)^N |\Sigma|}} \exp\left(-\frac{1}{2}\sum_{ij}(x_i - \mu_i)\Sigma_{ij}^{-1}(x_j - \mu_j)\right)$$

$$-2\ln\mathscr{L} = \sum_{ij}(x_i - \mu_i)\Sigma_{ij}^{-1}(x_j - \mu_j) + \log|\Sigma| + N\log(2\pi)$$

where Σ is a covariance matrix containing both systematic and statistical uncertainties, μ is the predicted spectrum and x is the data spectrum.

- One can then decouple the statistical and systematic components

$$\chi^2_{\text{Gaussian}} = \sum_{ij}\left[(x_i - m_i)U_{ij}^{-1}(x_j - m_j) + (\mu_i - m_i)V_{ij}^{-1}(\mu_j - m_j)\right] + \log|\Sigma|$$

where m is the predicted spectrum with systematic shifts applied, and U and V are covariance matrices encoding the statistical and systematic uncertainties, respectively.

Gaussian multivariate likelihood

- Now that the statistical and systematic components of the Gaussian likelihood are decoupled, we can then replace the statistical part with its Poisson equivalent

$$\chi^2_{\text{Hybrid}} = 2\sum_i \left[m_i - x_i + x_i \log\left(\frac{x_i}{m_i}\right)\right] + \sum_{ij} (\mu_i - m_i)V_{ij}^{-1}(\mu_j - m_j) + \log|V|$$

- We then transform the absolute V_{ij} matrix into a channelwise relative matrix $V_{\alpha i \beta j}$ in order to eliminate the $\log|V|$ term, as described on the next slide

$$\chi^2 = 2\sum_i^N \left[\left(\sum_\alpha^M \mu_{\alpha i} s_{\alpha i}\right) - x_i + x_i \log\left(\frac{x_i}{\sum_\alpha^M \mu_{\alpha i} s_{\alpha i}}\right)\right] + \sum_{ij}^N \sum_{\alpha\beta}^M (s_{\alpha i} - 1)V_{\alpha i \beta j}^{-1}(s_{\beta j} - 1)$$

where the product of the nominal prediction μ and the fractional systematic shift s has been substituted for the systematically shifted prediction m.

Oscillation channel matrix

$$-2\log\mathscr{L} = \sum_{ij}(x_i-\mu_i)\Sigma^{-1}_{ij}(x_j-\mu_j)\boxed{+\log|\Sigma|+}N\log(2\pi)$$

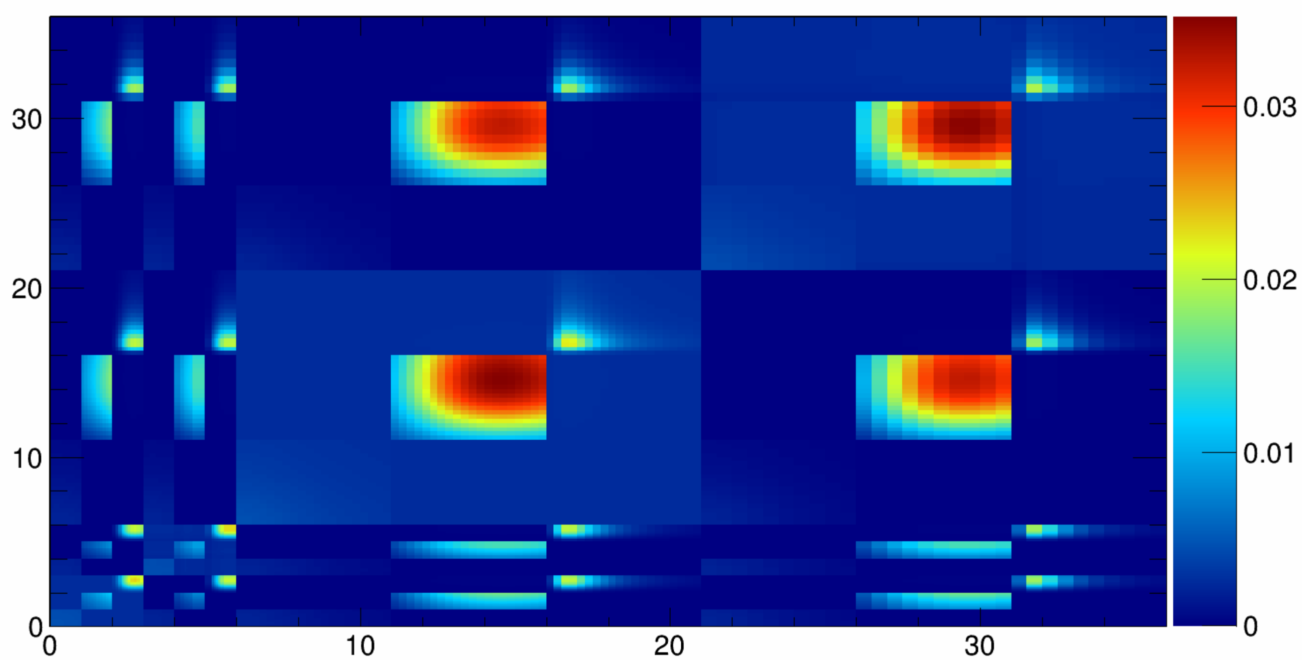

- The χ^2 in this technique contains a term which is dependent on the determinant of the covariance matrix.
- In order for this technique to be valid, this matrix must remain constant.
- In this new method, the matrix is **fractional** and systematic pulls are calculated **independently** for each beam component.
- This means the matrix is fixed, and the matrix determinant is a constant term that can be neglected.

Gaussian multivariate likelihood

- For each MINUIT χ^2 evaluation, optimal systematic pulls are found iteratively using Newton's method.

- The gradient of one systematic pull, and the Hessian matrix of a pair of pulls, with respect to the χ^2 are calculated:

$$\frac{\partial \chi^2}{\partial s_{\gamma k}} = 2\left(\mu_{\gamma k} - \frac{\mu_{\gamma k} x_k}{\sum_{\alpha}^{M} \mu_{\alpha k} s_{\alpha k}} + \sum_{\alpha i}^{MN} (s_{\alpha i} - 1) V^{-1}_{\alpha i \gamma k}\right)$$

$$\frac{\partial^2 \chi^2}{\partial s_{\gamma k} \partial s_{\delta l}} = 2(1+\lambda)\left(\frac{\mu_{\gamma k} \mu_{\delta l} x_k}{\left(\sum_{\alpha}^{M} \mu_{\alpha k} s_{\alpha k}\right)} \delta_{kl} + V^{-1}_{\gamma k \delta l}\right)$$

- This system is then solved by Cholesky decomposing and solving the linear system using Eigen routines.

Regularisation

- Under real conditions, solving this set of coupled equations can occasionally become unstable.

- Several techniques have been developed for mitigating these instabilities:
  - For difficult minimisations that do not converge quickly, we utilise the Levenberg-Marquardt algorithm to take smaller steps when solving for systematic pulls.
  - For difficult minimisations, Newton's method will occasionally attempt to pull analysis bins negative whiile solving.
    - Any bins that want to pull negative are masked off and removed from minimisation.
    - Once the remaining system has settled, these bins are reintroduced and the minimisation is repeated.