AGATA Week 2024

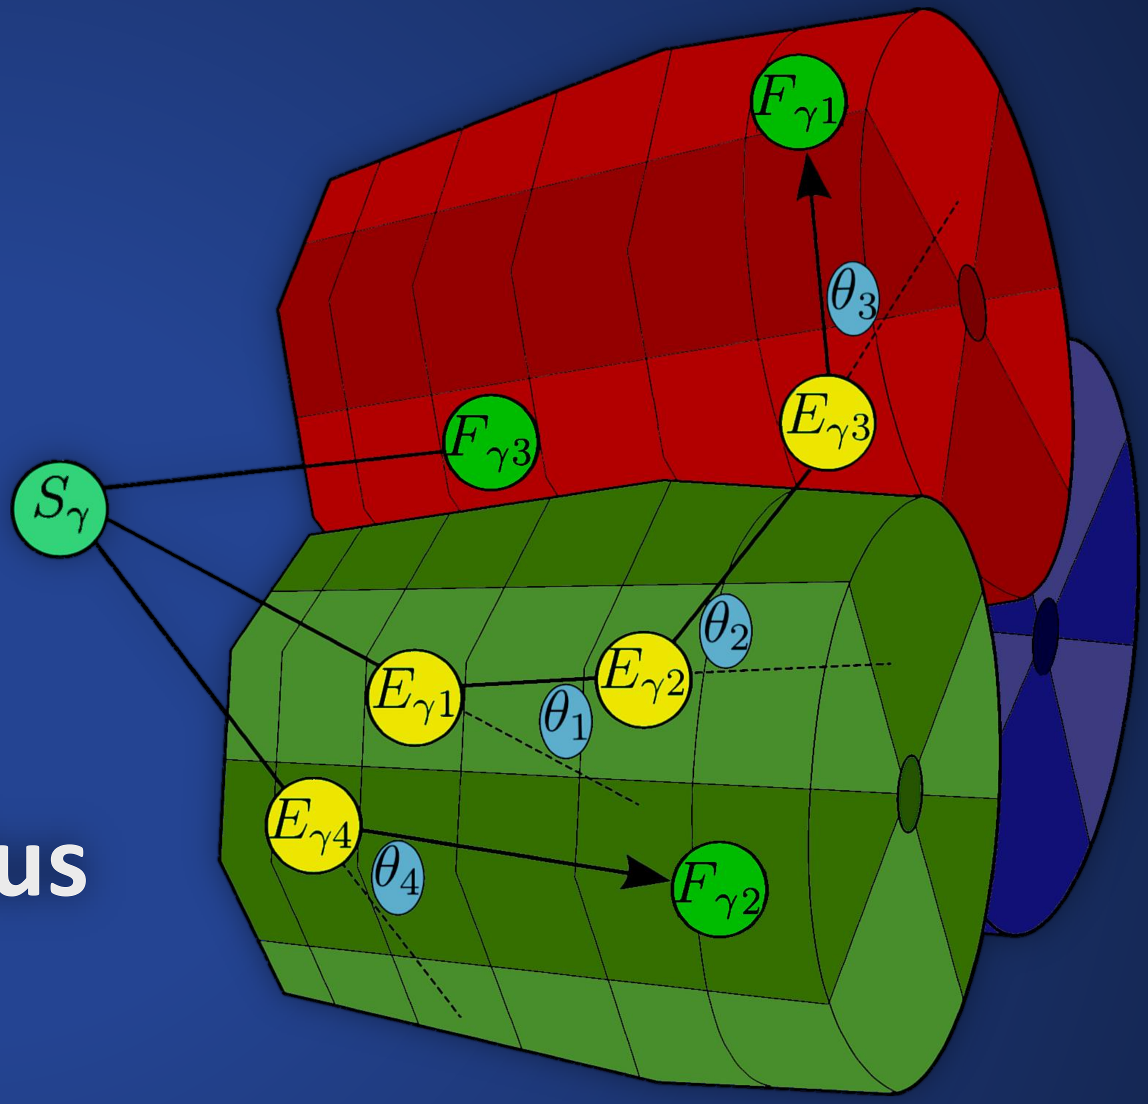

Liverpool R&D: (ML/AI) Status

Dr. Fraser Holloway

Fraserh@liverpool.ac.uk

UNIVERSITY OF LIVERPOOL

Science & Technology Facilities Council

AGATA ADVANCED GAMMA TRACKING ARRAY

Presentation Overview

- AGATA signal response
- Parametric vs Optimisation based PSA
- Recap of SIMPLEX PSA algorithm
- Initial results & issues
- Masking within AGAPRO
- Dual-Segment PSA
- Crystal-wide Optimisation
- Intelligent Preprocessing

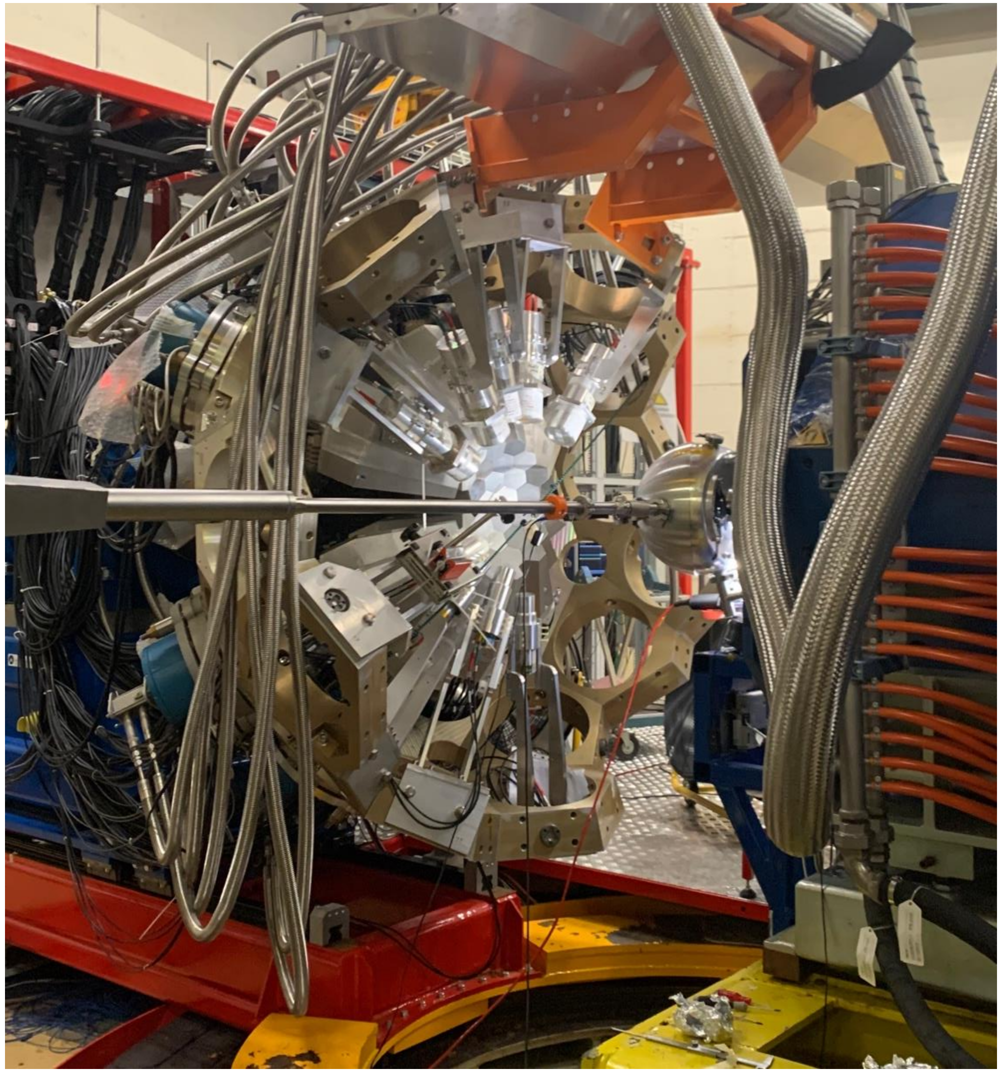

AGATA at LNL (22/10/23)

High-Fold Signals in AGATA

- High-Fold signals are separated and processed individually when possible.
- Conventionally this is achieved using segment windows, renormalisation & recursive subtraction.
- This only works on well-separable signals, typically a few segments apart.

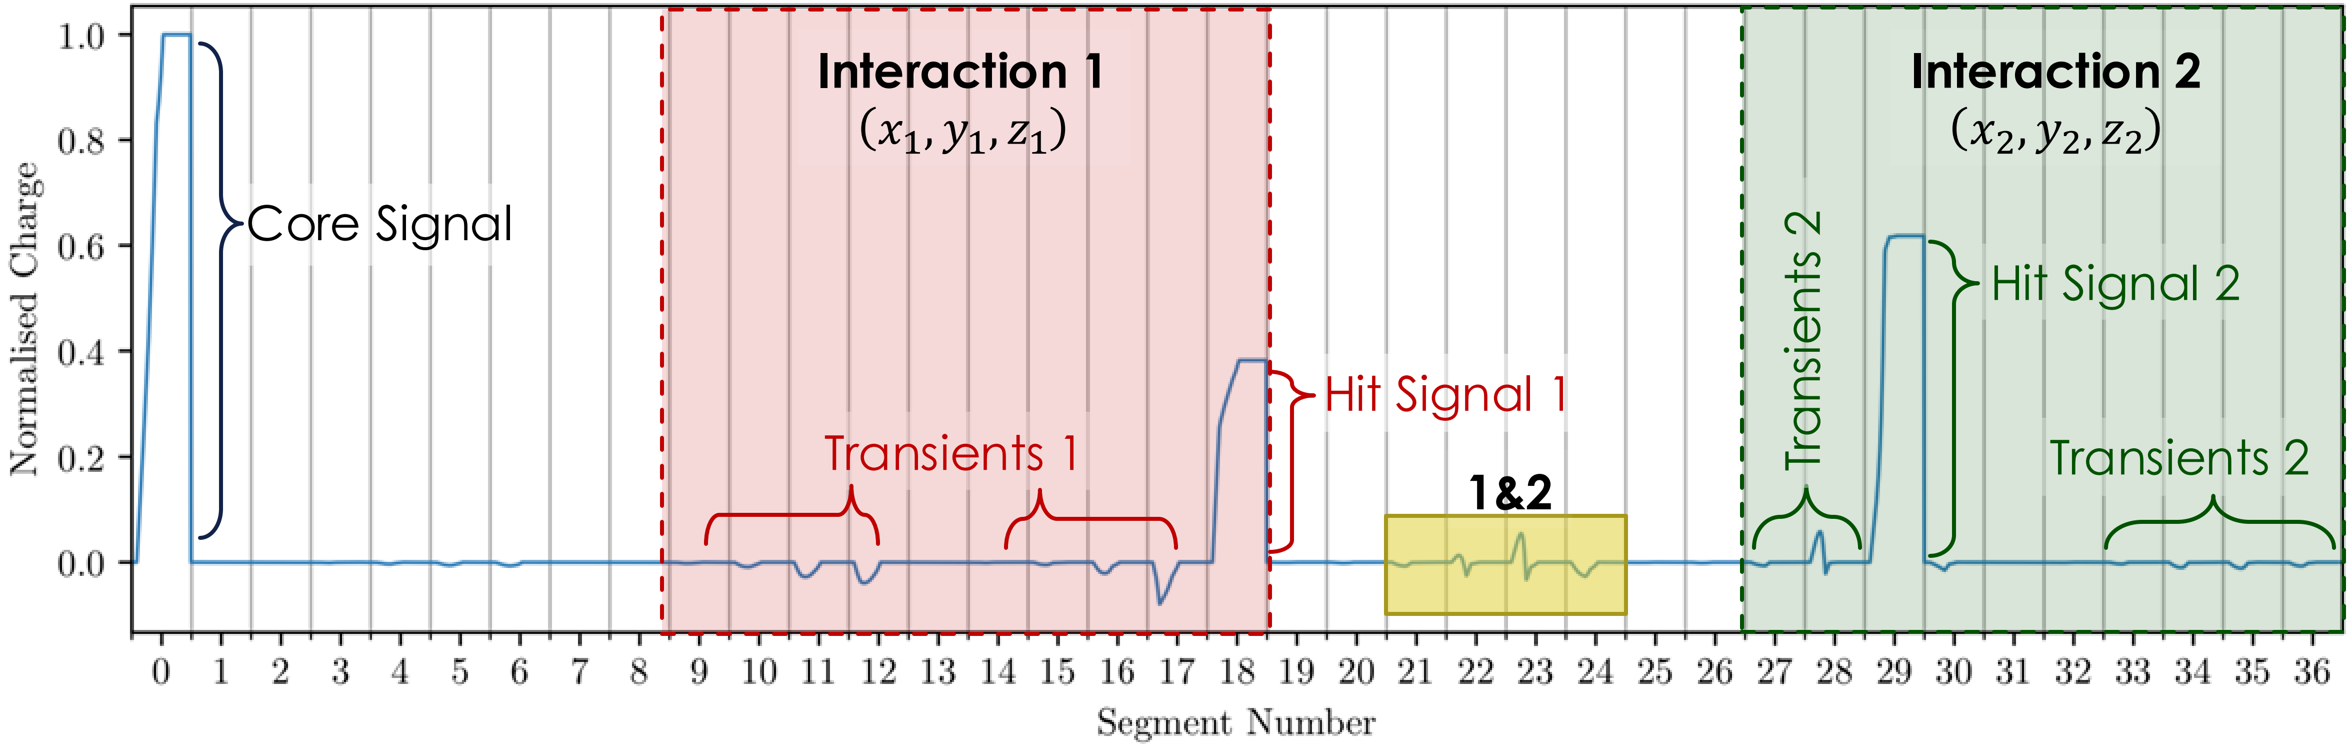

Multiple Interactions Within the Same Segment

- These can look near-identical to Fold-1 signals and will regularly fool parametric, PSA & ML methods.
- The underlying signals are heavily entangled and cannot be solved individually.
- They have massive implications for tracking (e.g. a [60°,30°] scatter pair looking like a 90° scatter).

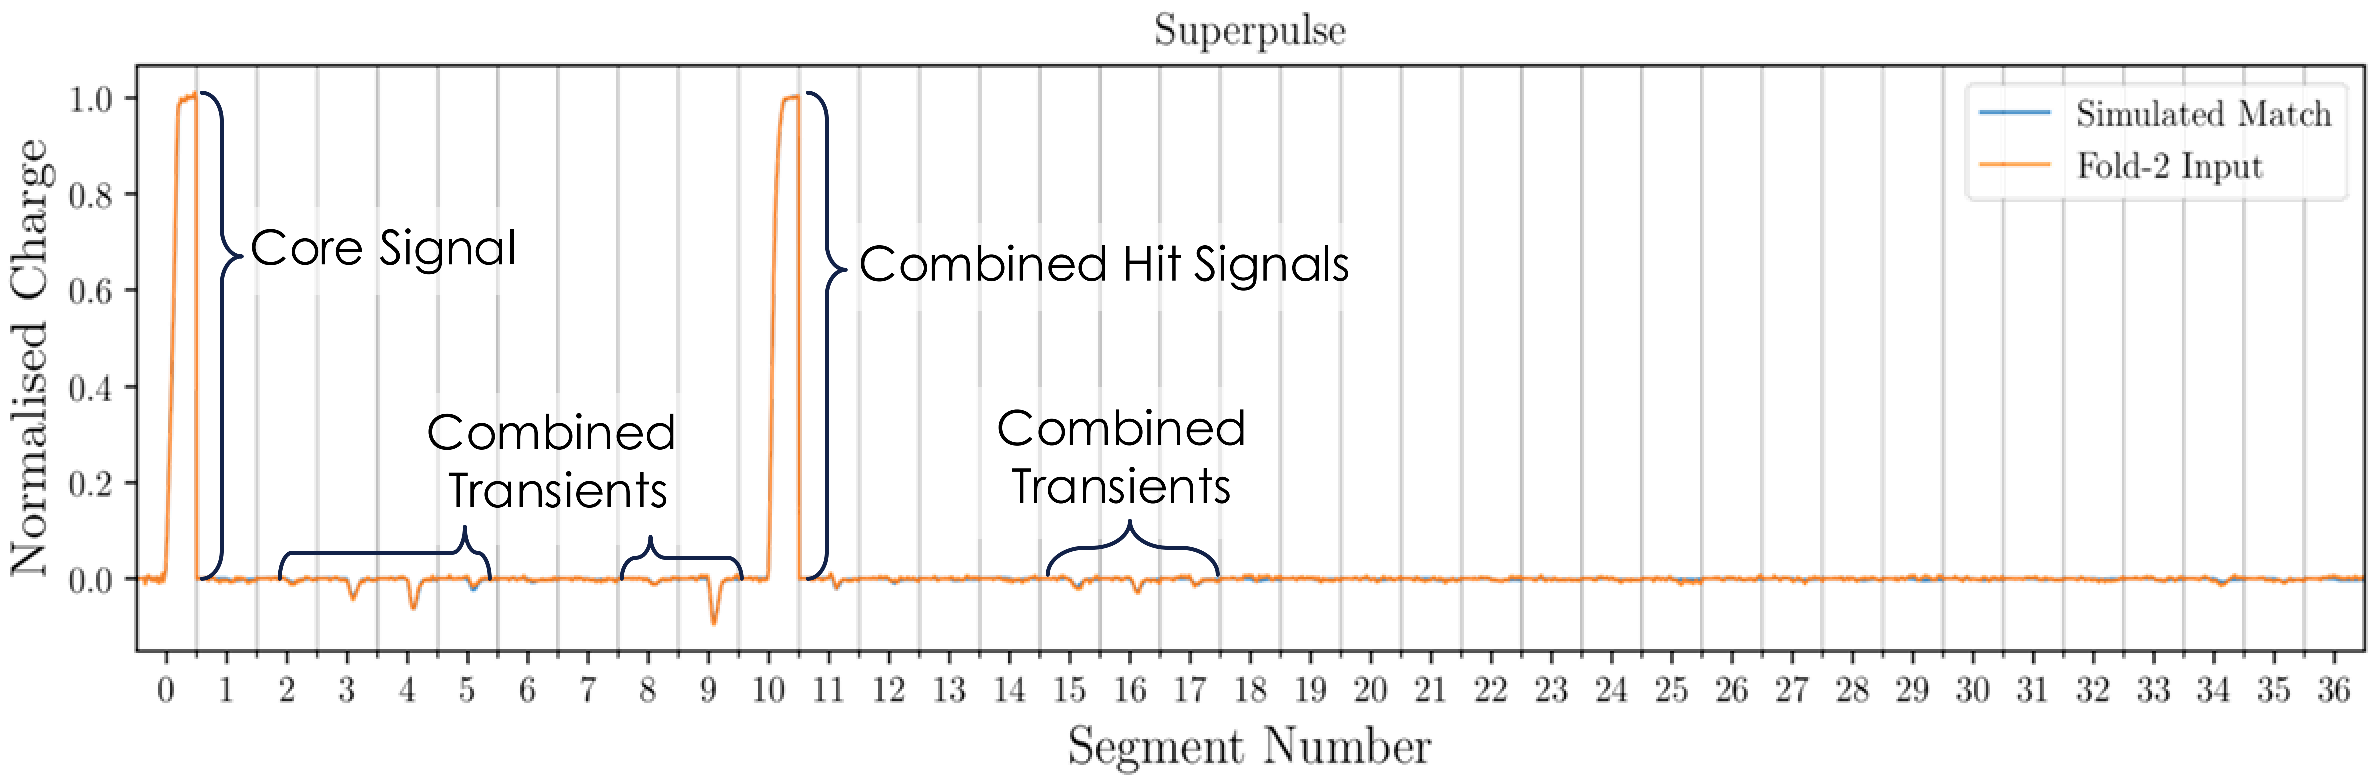

Pulse Shape Analysis for AGATA

Within AGATA Pulse Shape Analysis is used to quickly compare experimental signals (M) against simulated signals (S) using a Figure of Merit (FoM).

$$FoM = \sum_{i=1}^{37}\left(\sum_{j=1}^{121}(|M(j)_i - S(j)_i|)^2\right)$$

The position of the closest match in the basis is then used for the predicted interaction position and given to the GRT.

Current PSA utilises a 2-part grid search (AGS), initial optimisation of a coarse grid followed by a fine-grid search around the local optima.

This work focusses on solving 2-interactions within the same segment for AGATA.

Fraser.Holloway@liverpool.ac.uk

Parametric Vs Non-Parametric PSA

Parametric PSA concerns features that can be easily derived ($T_{10\text{-}90}$, ICA)

- These methods are easy to program, simple to apply to ASICs & FPGAs
- They are typically surjective (many positions have the same value).
  - Insufficient for use in γ–ray localisation & tracking.
- Apply poorly to multi-interaction PSA.

Non-parametric PSA concerns features that are more complex:

- Components of the signal (e.g. from PCA/SVD or AI)
- Can be used to form regression predictors (e.g. in CZT)
  - Inaccurate but fast.
- Often used in optimisation (kNN) for localisation:
  - Accurate but slow.

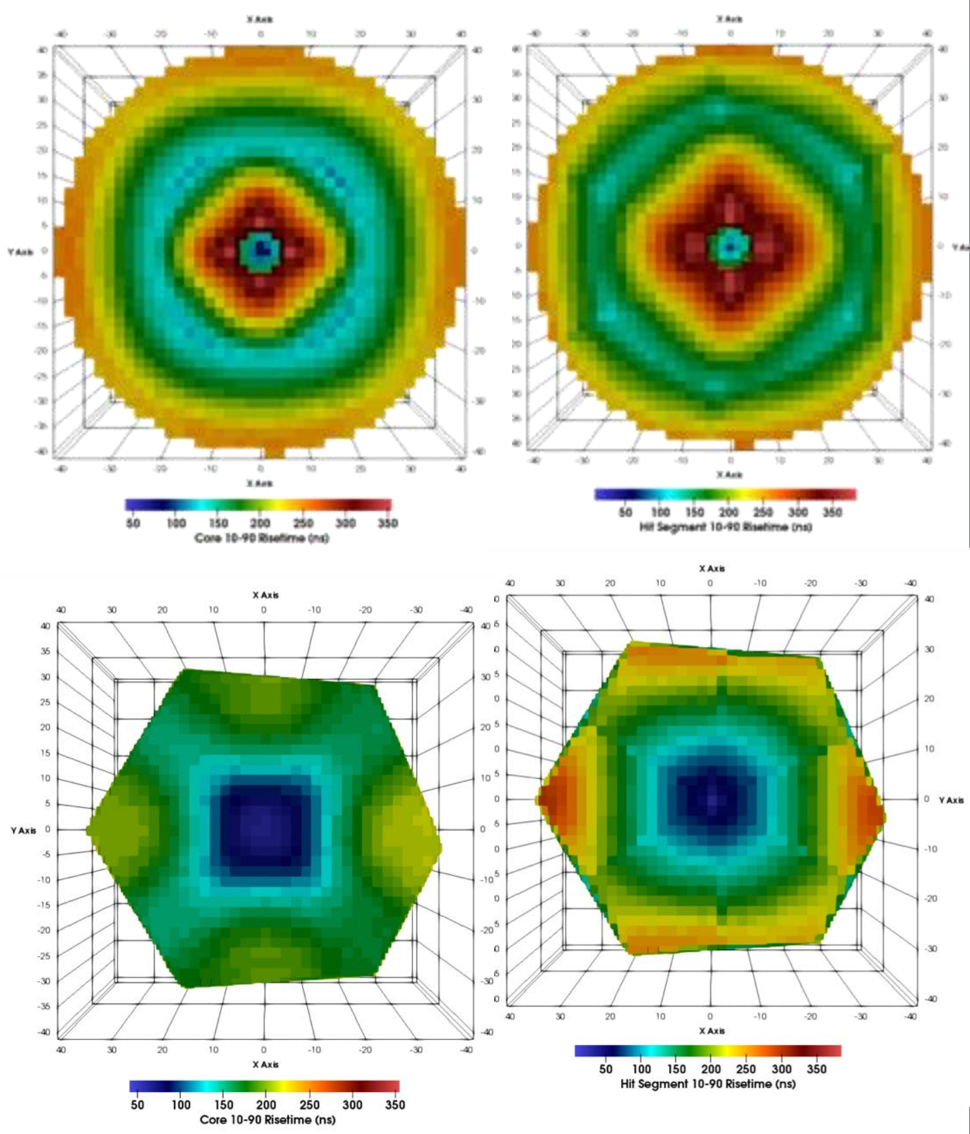

fraser.holloway@liverpool.ac.uk

Non-parametric Responses – Coax Detectors

- Monolithic Coaxial detectors show some signal variation with position, most signals form a continuous response manifold in $\mathcal{R}_n$.
- This can be observed with covariance matrices.
- Alternatively, we can use manifold learning.

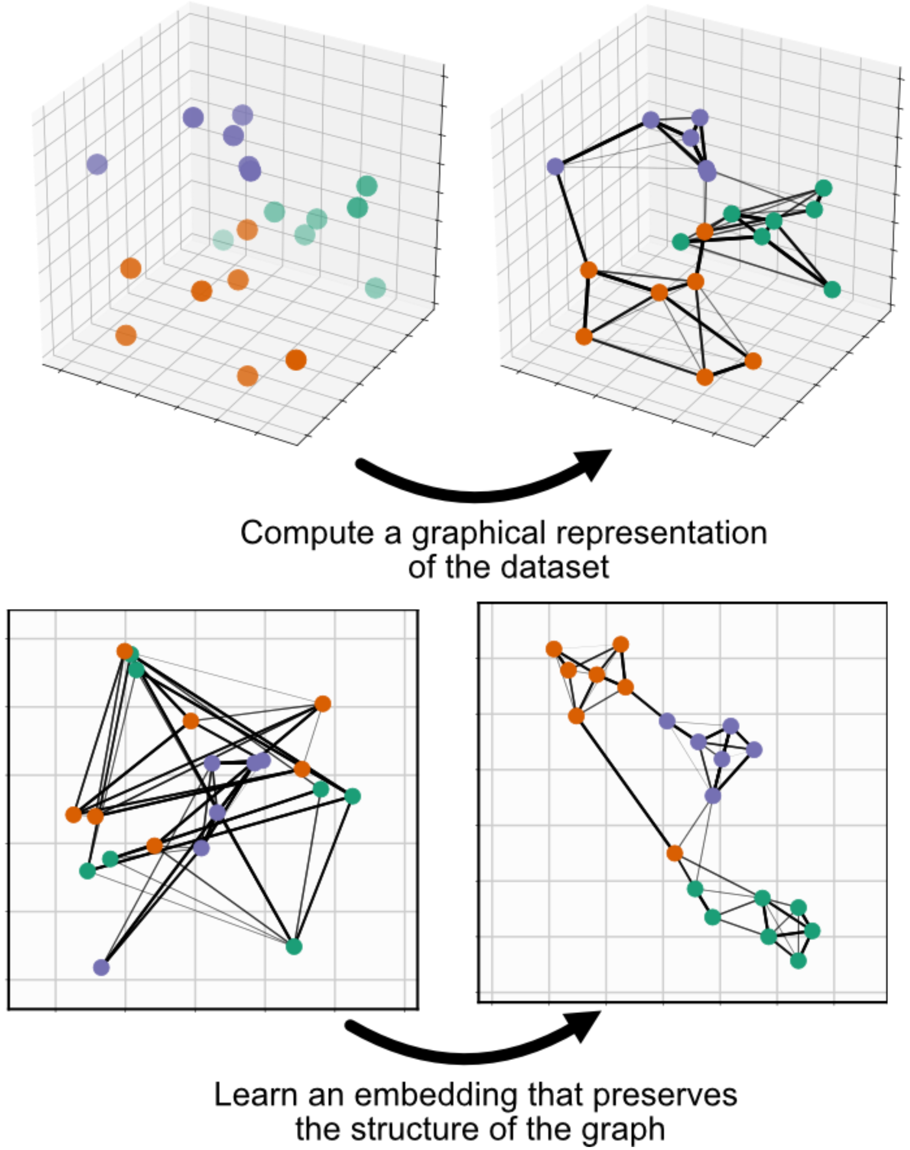

AGATA Response

- Segmentation within AGATA provides a wealth of information within the signal response.
- Detector response goes from temporal to spatiotemporal.
- Segment responses concentrate in submanifolds.
- Strong radial and azimuthal dependence observed.
- Transient signals provide additional fidelity (x, y, z).

- These effects allow for ~ 3 interaction PSA within a segment.
  - Upper limit of ~ 9 interactions across a crystal.

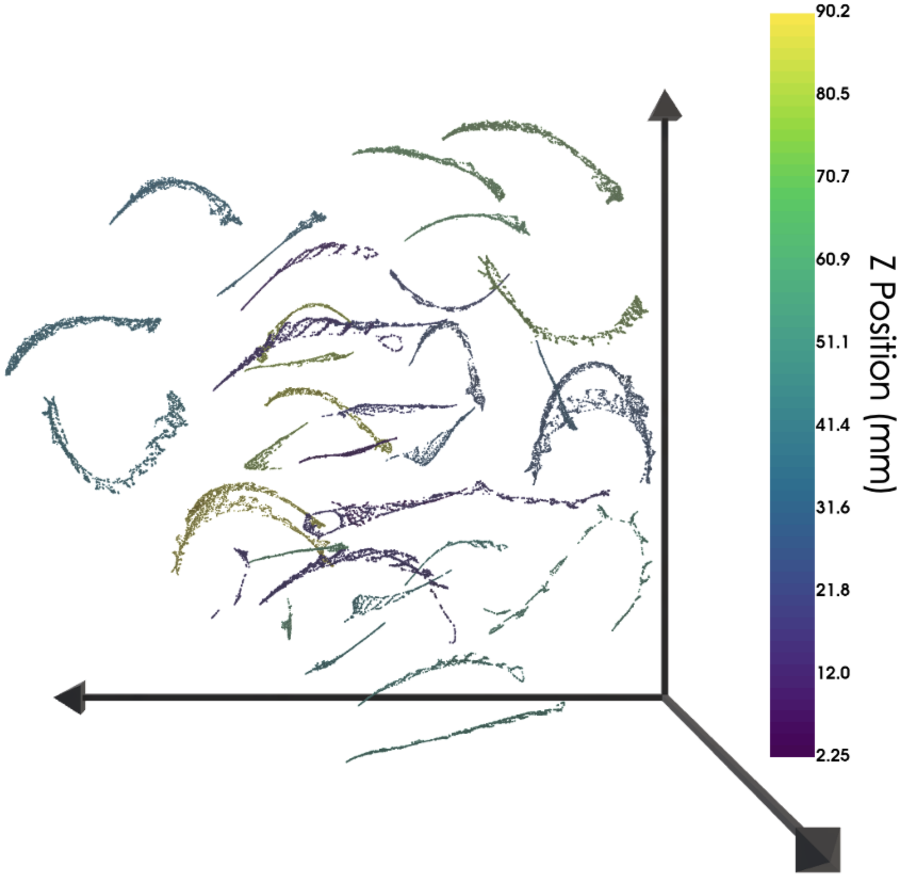

Graph Structures & Accelerated PSA

As part of my work I investigated the use of graph structures to accelerate PSA, they have several marked benefits:

- Data-driven ∴ no parametric methods (e.g. T_{90}) needed.
- Self-organised ∴ no rigid grid structure required.
- Hierarchical ∴ adaptive resolution.

In my work I profiled several methods for accelerated PSA:

- Space Partitioning (exact kNN): kD-Tree, M-Trees, etc.
- Coarse Indexing (approximate kNN): HNSWLib, FAISS.

Recommendations from this work are being implemented.

- My current work is focussed on multiple-interactions.

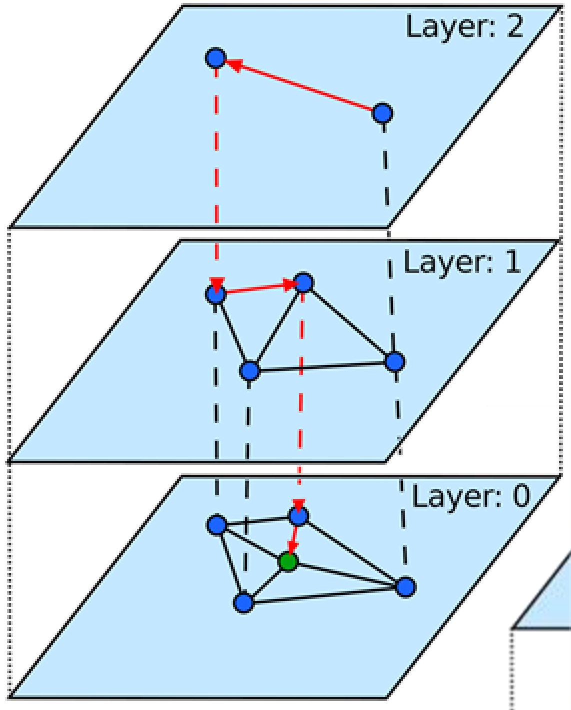

$HNSW$ Structure

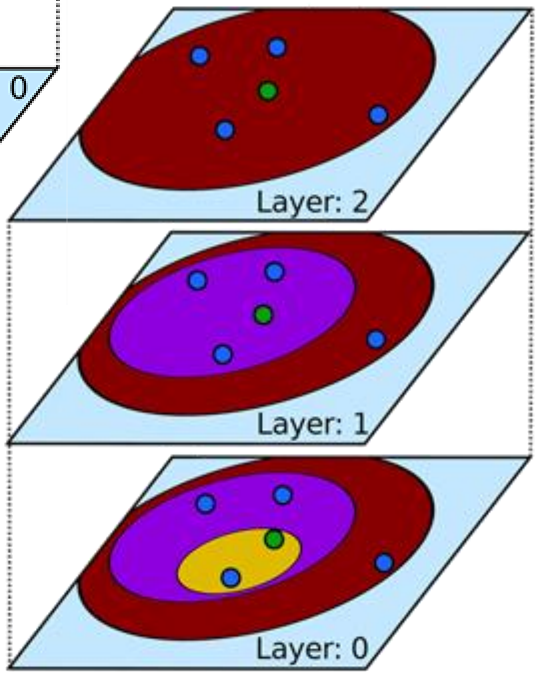

M-Tree Structure

Fraser.Holloway@liverpool.ac.uk

Multiple-Interaction PSA

Multiple interaction events make up a significant fraction of the data collected by AGATA.

Around 30% of events in AGATA are comprised of events where multiple interactions occur in the same segment.

- This will vary as a function of γ-ray energy

In these cases, the PSA tends to fit a weighted average of the positions.

- Preference to predict barycentre at segment centres.

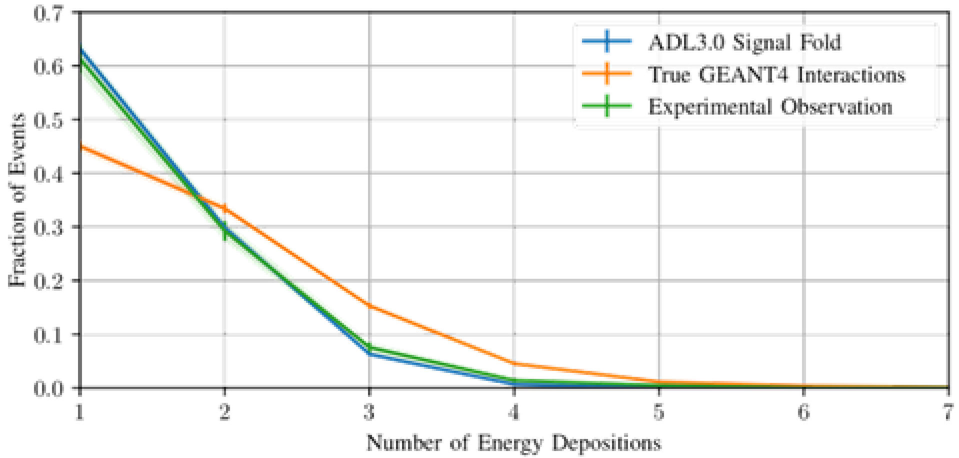

Evaluation of multi-interaction events using conventional PSA is not feasible.

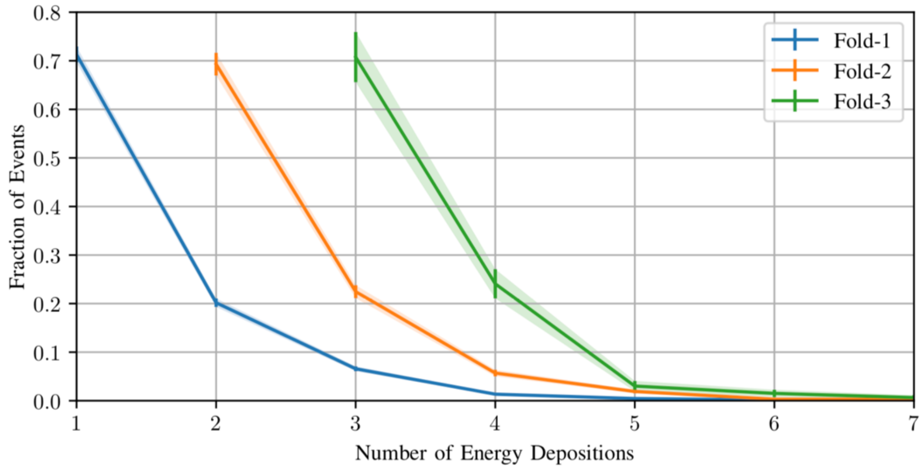

Initial implementation of GRETA algorithm was performed but a flaw was found in their approach that is currently irreconcilable.

- Currently working with the GRETA team to see if it's fixable.

Fraser.Holloway@liverpool.ac.uk

Simplex De-mixing in High Dimensions

- A new convention was developed to evaluate multi-interaction signals using distance geometry and simplex de-mixing.
  - $\overrightarrow{Op_1} \cdot \overrightarrow{qO} = \overrightarrow{Op_2} \cdot \overrightarrow{qO} = \overrightarrow{p_1p_2} \cdot \overrightarrow{qO} = 0$
  - $\angle Op_1p_2 \ \& \ \angle Op_2p_1 < 90°$
- This convention allows for direct fitting of multi-interaction signals.
  - Also works recursively, allowing for 9-interaction fits.
- Exhaustive combinations are required to find the best solution:
  - ~1,100,000 combinations for 2 interactions.
  - ~550,000,000 combinations for 3 interactions.
  - ~10^{30} combinations for 9 interactions.
- Not feasible to do on a CPU exhaustively real-time.
  - Somewhat possible on a GPU.
  - How about using graph acceleration?

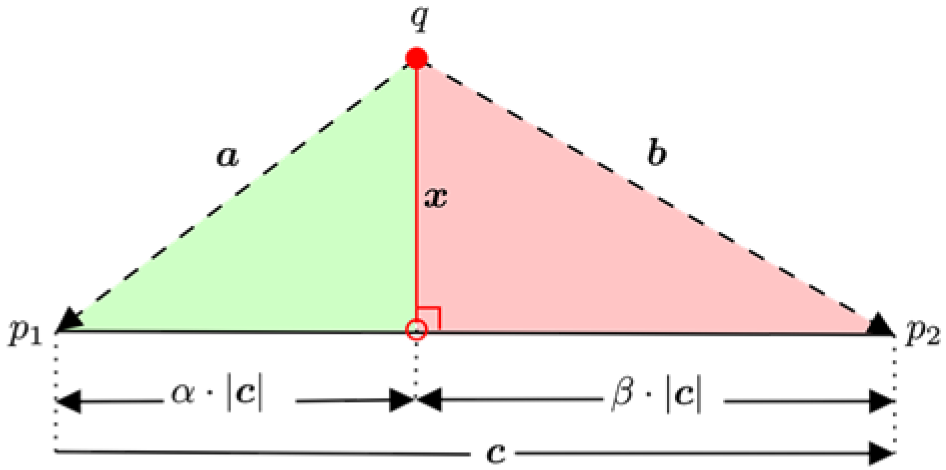

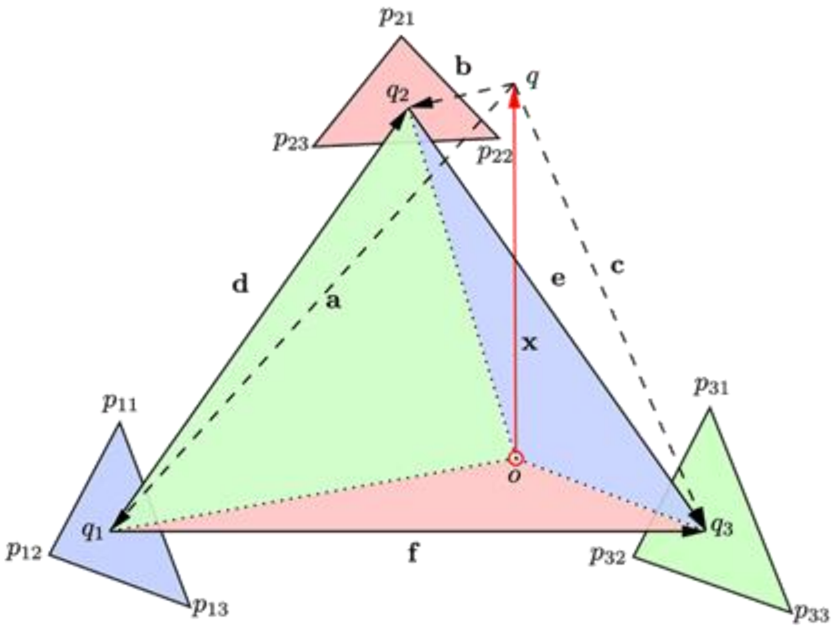

Hypergraph-accelerated PSA

- 2-interaction solutions exist as edges that connect between nodes of two orthogonal trees.
- Lower levels have more nodes ∴ more edges.
- Just like how nodes can have child nodes edges can have child edges.
  - The validity of an edge applies recursively.
    - If an edge is invalid, all child edges are also invalid.
- This allows for us to cull large sections of the graph.

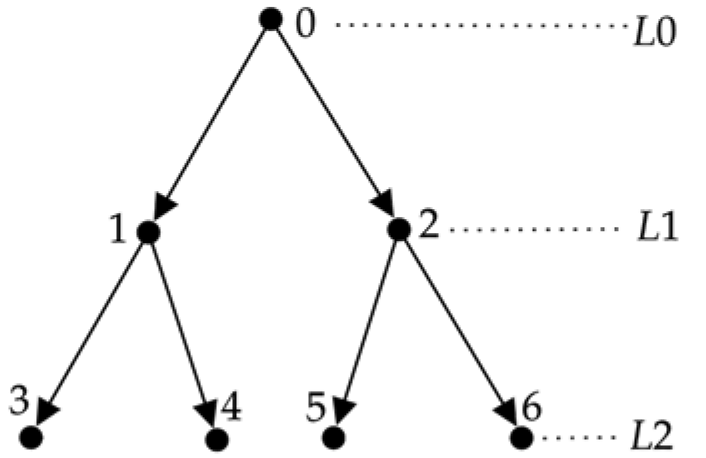

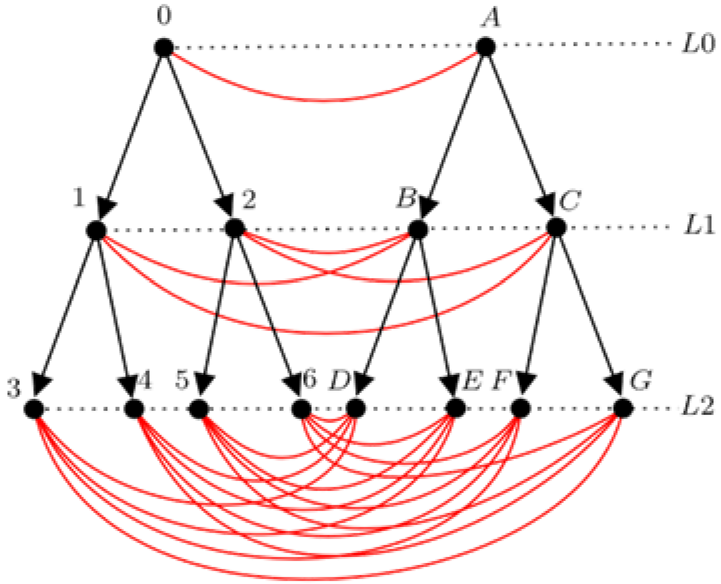

Fraser.Holloway@liverpool.ac.uk

Hypergraph-accelerated PSA

Two hypergraph-accelerated methods have been developed:

- Hyper-M-Tree
    - Recursive covering (Exact kNS search).
- Hyper-Hierarchical New Small World Graphs.
    - Greedy simplex traversal (Approximate kNS search).

These overcome conventional limitations of the existing method:

- Native support for PSA uncertainty
- Self-learned hierarchy.
- Support for dynamic resolution (FEM) bases.

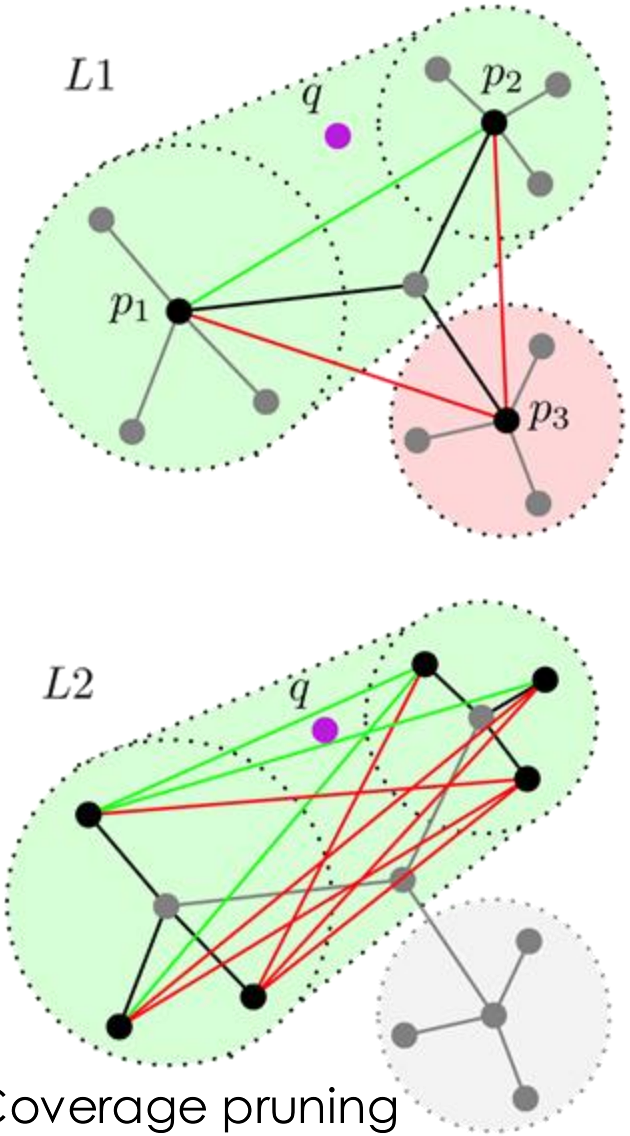

Coverage pruning

Fraser.Holloway@liverpool.ac.uk

General PSA Update

SIMPLEX and GRETINA PSA codes have been implemented within AGAPRO;

- Exhaustive & AGS approaches available.
- Rudimentary support for uncertainty propagation.
- Accelerated using SIMD instructions.

A flaw was found in the precomputation step W.R.T masking of distances

- Investigating the remedy uncovered other issues unrelated to the new algorithm.
- Currently working on a more comprehensive PSA pipeline.
  - Better handling of single vs multi-interaction events & Recursive Subtraction & Time-shifting.
  - Integrates a more intelligent preprocessing using ML.
  - Will have a crystal-wide optimiser for the full superpulse.

fraser.holloway@liverpool.ac.uk

Assessing PSA performance

- A dataset was extracted from E680 at GANIL, fission ^{98}Zr nucleus with a strong γ-ray transition at 1.2 MeV.
- Fragments were emitted with velocities of $\beta \approx 10$, making the resolution of such transition strongly dependant on the quality of the Doppler correction.
  - Doppler correction only applies to first interaction.
- PSA also will heavily influence tracking accuracy.
  - Should influence peak prominence.

Important note:

- The spectral performance of AGATA is the combined effect of PSA, GRT & simulated bases accuracy.
  - GRT has yet to be tuned for the new PSA.
  - Bases are being recalculated at higher fidelity.
  - **This is all with the current flaws, I am working on fixes.**
  - **Performance will likely change in the future.**

fraser.holloway@liverpool.ac.uk

First Hit Predictions – Front Faces

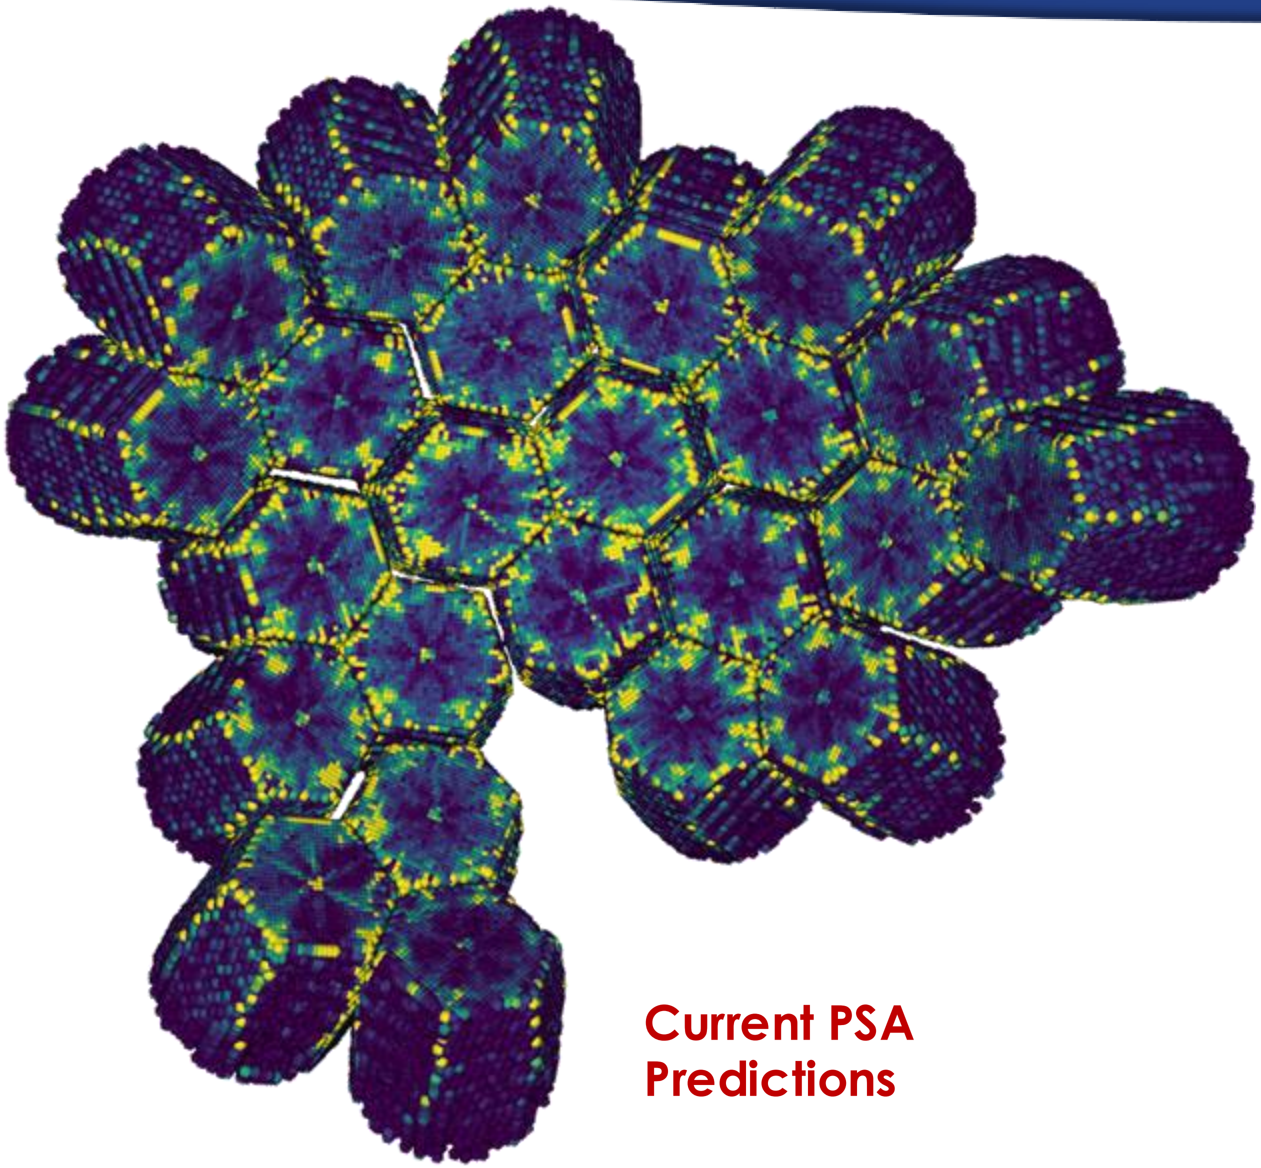

Current PSA Predictions

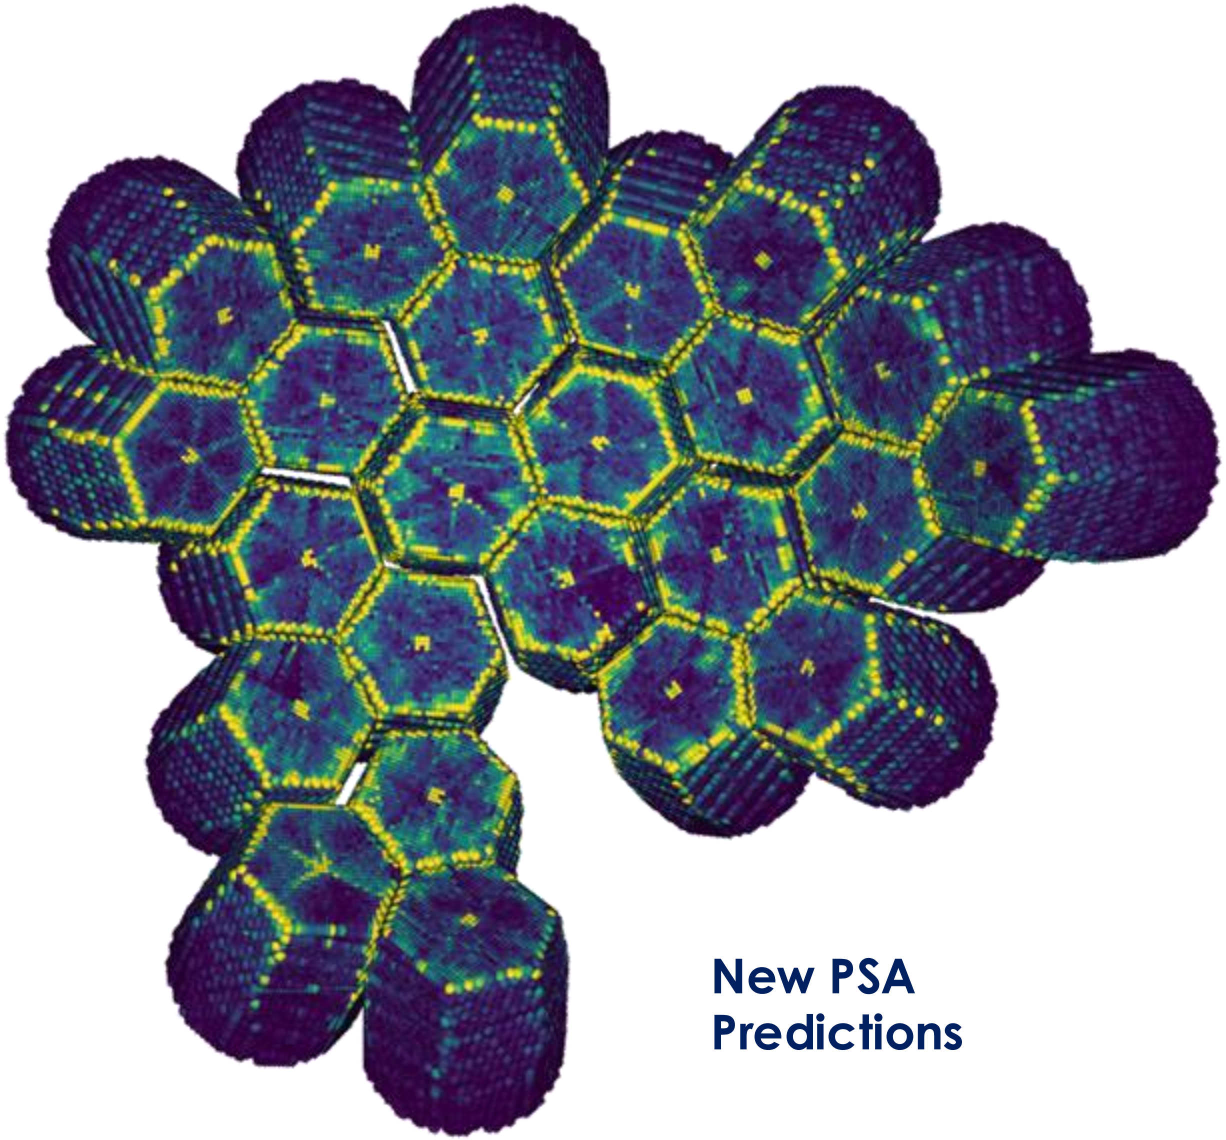

New PSA Predictions

First Hit Predictions – Back Faces

Existing PSA Predictions

New PSA Predictions

Interaction Positions

Observed number of interactions changes drastically with the new algorithm

Prevalence of single-interaction signals is reduced significantly, average track length increased.

Previously 44% of tracks were P.E single-site

Now only 17% are P.E single-site

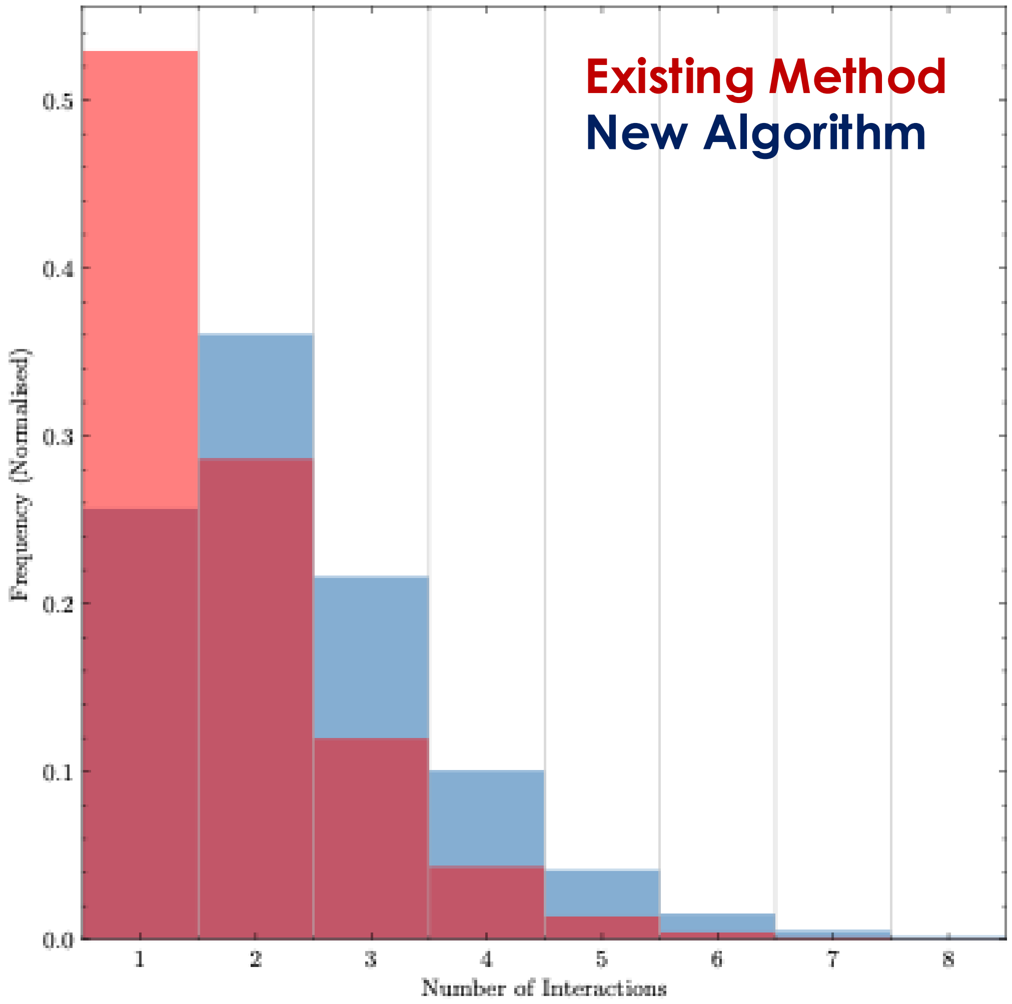

Fraser.Holloway@liverpool.ac.uk

Dynamic Masking in AGAPRO

Whenever a signal is processed in AGAPRO a window is created to allow PSA to process events independently.

- This occurs roughly 2/3 of the time.
- Comes in two forms `hmask` & `lmask (& nmask)`:
  - `hmask` is pre-event, depends on the hit segment.
    - `neighbours = -2`
    - 36 static masks that are well-known.
  - `lmask` is live, depends on the other events in the crystal.
    - Removes other net charge signals.
    - Could remove neighbour transients, currently disabled.
    - Around 3500 unique masks
- Because `lmask` changes for each event the precomputed distances are directly affected.
  - This is problematic, there are 35 million of them.

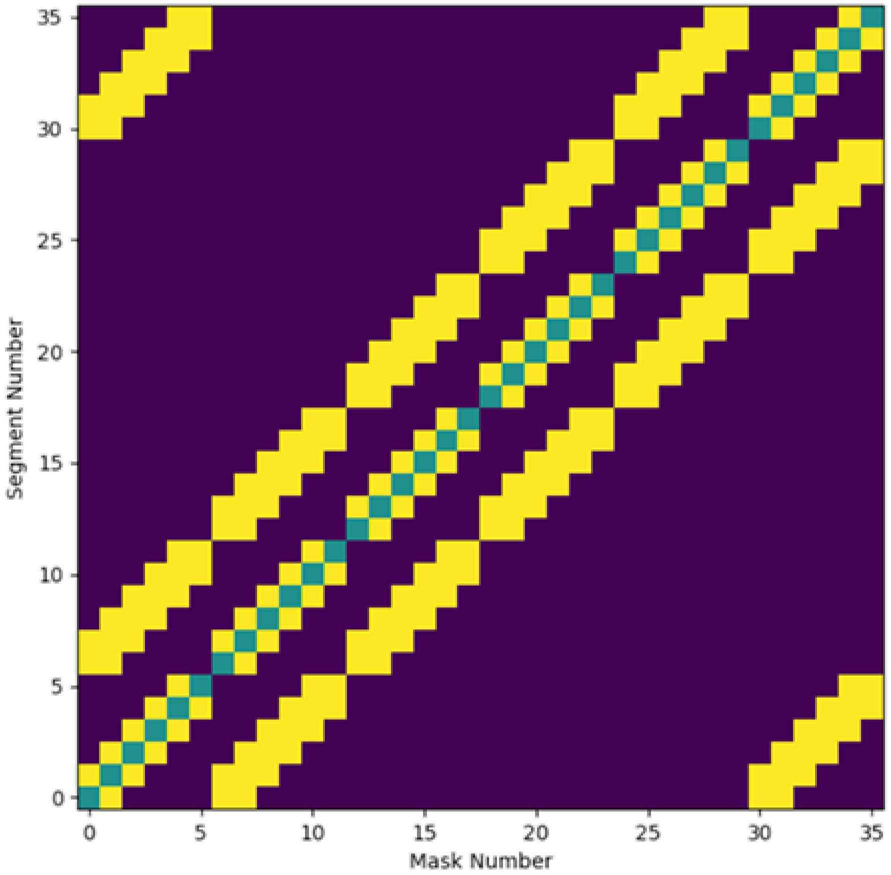

Observed hmask

fraser.holloway@liverpool.ac.uk

Dynamic Masking in AGAPRO

Whenever a signal is processed in AGAPRO a window is created to allow PSA to process events independently.

- This occurs roughly 2/3 of the time.
- Comes in two forms `hmask` & `lmask` (`& nmask`):
  - `hmask` is pre-event, depends on the hit segment.
    - `neighbours = -2`
    - 36 static masks that are well-known.
  - `lmask` is live, depends on the other events in the crystal.
    - Removes other net charge signals.
    - Could remove neighbour transients, currently disabled.
    - Around 3500 unique masks
- Because `lmask` changes for each event the precomputed distances are directly affected.
  - This is problematic, there are 35 million of them.

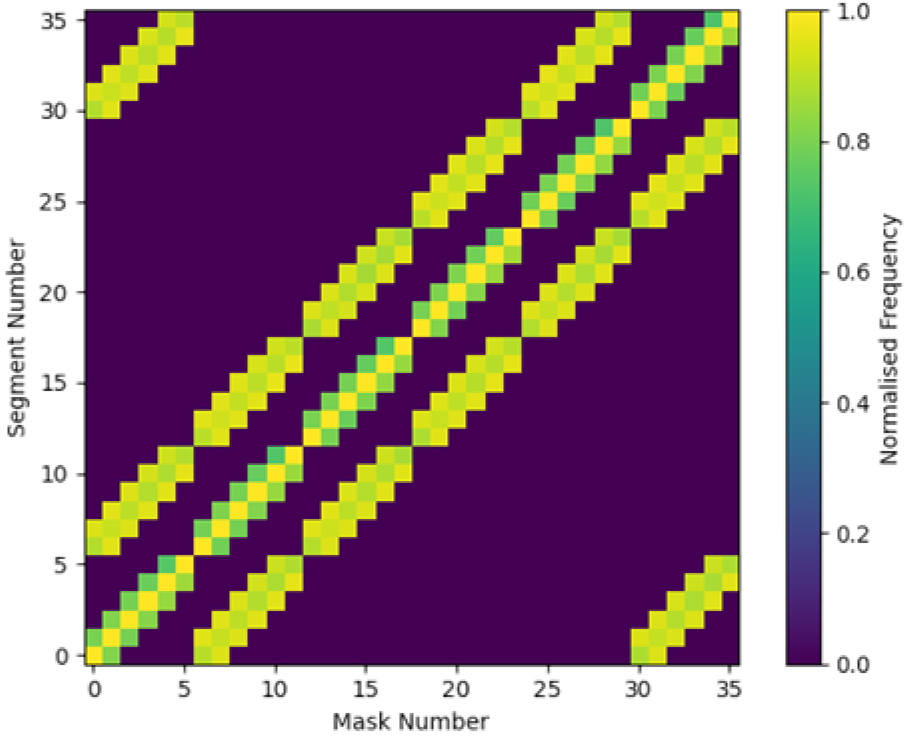

lmask Frequency (normalised)

fraser.holloway@liverpool.ac.uk

Memory Constraints

Precomputed distances are a key part of the multi-interaction algorithms, vastly outnumber novel distances.

- Segment 6 has ~2,000 novel distances, ~2,000,000 precomputed.
- Distance computation is expensive, 1-2,000 float operations each time.
- Having to re-compute d_c for each event is not feasible to do real-time.

Thankfully $L2^2$ is commutative, can be split into sub-distances:

- d_c can be precomputed for each segment, summed once the `lmask` is determined
  - Requires significant memory increase of ~36x
- Could be optimised using Boolean Covering to find prime implicants using Quine-McCluskey algorithm.
- d_c can be precomputed for `hmask` and its components
  - Calculate `hmask` distance and then subtract `lmask` components
    - Slightly better memory efficiency, only 8x previous version

- Current RAM utilisation is 6 GB/crystal, could be reduced to ~2.5 GB

fraser.holloway@liverpool.ac.uk

Dual-Segment PSA

Occurrence of hits that influence `lmask` is around 63%, vast majority occur in neighbouring segments.

- Even without removing neighbours of coincident events the lmask is significantly reduced.
  - Some PSA is run on a single segment, not great.
  - Presence of transients is not accounted for, will affect PSA.
    - Removing the neighbour segments altogether is worse.

- 2-interaction PSA is not explicitly limited to a single segment.
  - Can be modified to run on the full crystal.
    - Significant memory requirements again, ~32x larger
  - Can be modified to run on overlapping windows.
    - Slightly better, 8-12x larger.

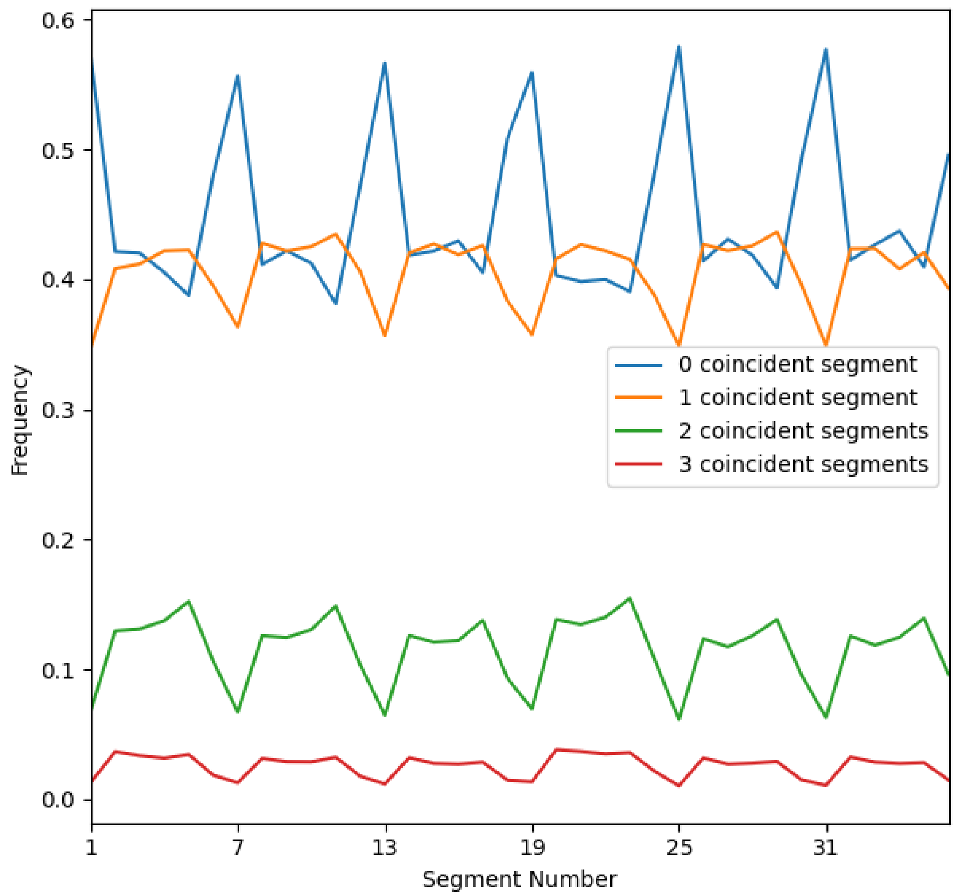

fraser.holloway@liverpool.ac.uk

Crystal-wide PSA Post-processing

- Within AGAPRO PSA the superpulse is separated into windows, and processed at the segment level.
  - Masking minimises contributions from other interactions but is not perfect.
    - No way to account for transients on net charge signal.
    - Fitting errors affect subsequent fits.
  - A global fit across multiple clusters isn't feasible.
    - Mathematically possible, but time and memory consuming.
- Formalism for a generalised solution for n-interactions was developed.
- This allows for a crystal-wide solution to be found for the full superpulse, not just each window.
- The best 1 & 2 interaction solutions for each segment can be passed to a global fitter.
  - A global solution is then found from the reduced set.
  - Allows for crystal-wide re-optimisation of energies, even within the same segment.

fraser.holloway@liverpool.ac.uk

Intelligent PSA Preprocessing

- 2-interaction PSA is a time-consuming process.
- Assessment of 1 vs 2 interactions is performed after PSA search via fallback vetoes:
  - Interaction separation.
  - Energy fraction & absolute energy.
  - χ^2 reduction.

- It'd be useful to skip the 2-interaction PSA if it's unlikely to give a good result.

- Can we infer what the expected χ^2 , energy fractions & separation are before PSA?
  - If so, then can we only run 2-interaction on events that will show an improvement?
  - Can this information provide a better starting guess for PSA?

fraser.holloway@liverpool.ac.uk

Self-Organising Maps

- Self-Organising Maps are a graph-ANN that use competitive learning to fit maps to an underlying manifold of data.
- The map is initialised and slowly iterates to form a topologically constrained approximation of the data.
- Nodes in the graph then learn from their immediate surroundings via inference.
- Often used for clustering and feature detection.
- I've adapted them for polylinear regression.

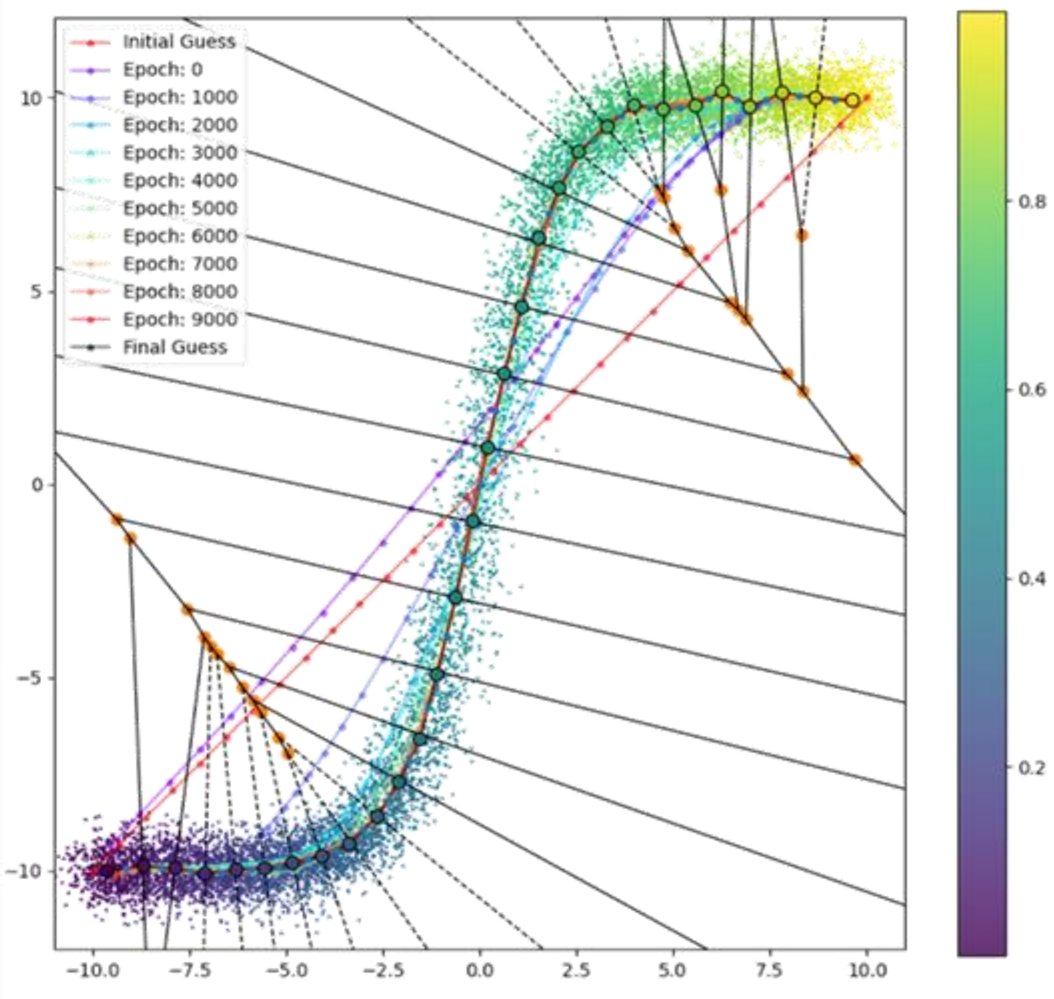

fraser.holloway@liverpool.ac.uk

SOMs for Pseudo-Parametric PSA

- SOMs can be augmented with existing analytical PSA techniques to perform regression.
- Map is converted into a simplicial complex.
- This allows for a poly-linear segmentation of the response space.
- Deviations from the map spline can be considered noise.
  - Significant deviation could be considered multi-interaction.

- Analytical code is easy to accelerate, map is easily tuneable.
  - Multithreaded capable of running at 3 MHz.
  - Can be implemented on FPGAs, Digitisers.

- Should work for multi-input regression:
  - Estimate both regression parameters and energy fractions.
  - Maths is very long, requires 4-interaction PSA.

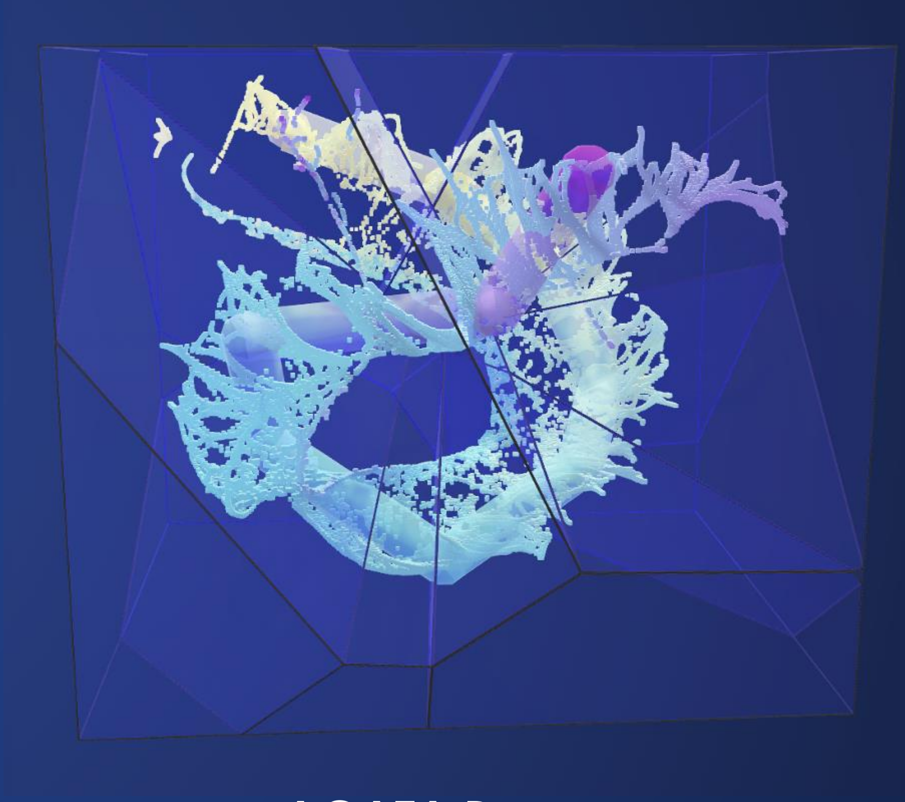

AGATA Response (UMAP Embedding)

Performance on Coaxials

- SOM was initially trained on simulated AGATA basis (A005).
- Further refinement using experimental data.
- Inference then applied to learn radial position.
  - Should be able to learn t_0 offset & neutron damage.

- Initial experimental results look promising:
  - Worse performance than kNN, better than parametric.
  - Limitations in interpolation, will be fixed with GSOMs.

- Implementation of FPGA-native code ongoing.

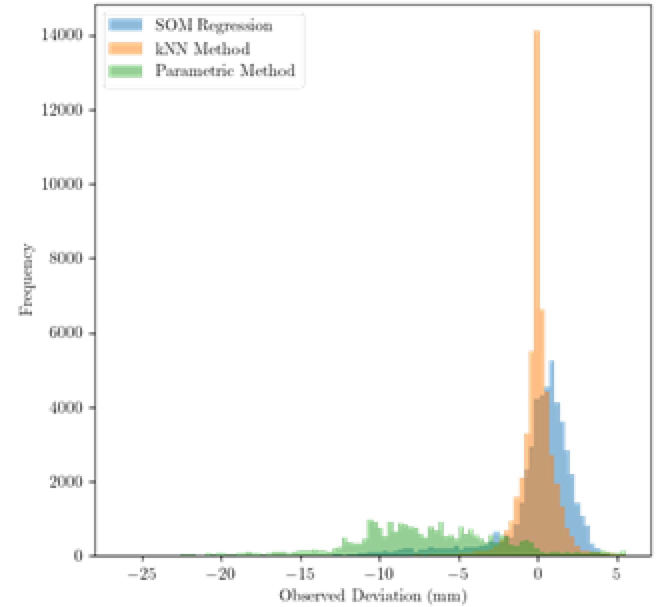

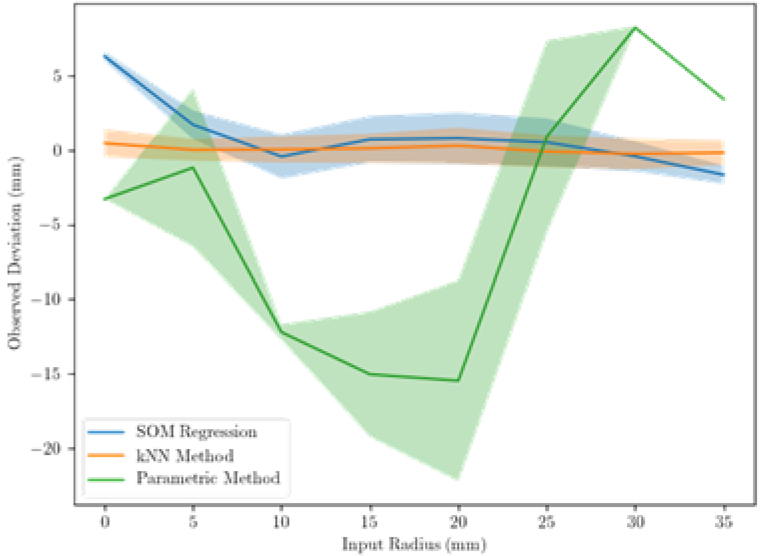

fraser.holloway@liverpool.ac.uk

ML-Assisted Interaction Veto

- The number of interactions can be inferred by estimating the χ^2 for approximations of the underlying manifold.
- 1-interaction SOM approximation requires determining best simplex interpolation in the map.
  - Essentially conventional 2-interaction PSA.
- 2-interaction SOM approximation requires determining the best interpolation of 2 interpolations.
  - Essentially conventional 4-interaction PSA.
    - Solving of a 5-simplex, (hyper-hyper tetrahedron).
    - Moderately difficult.

- If the χ^2 for the 2-interaction approximation is significantly better we can assume the full search will do the same.
- If the χ^2 for the 2-interaction approximation fails the vetoes we can assume that the full search will fail to find a good solution either.

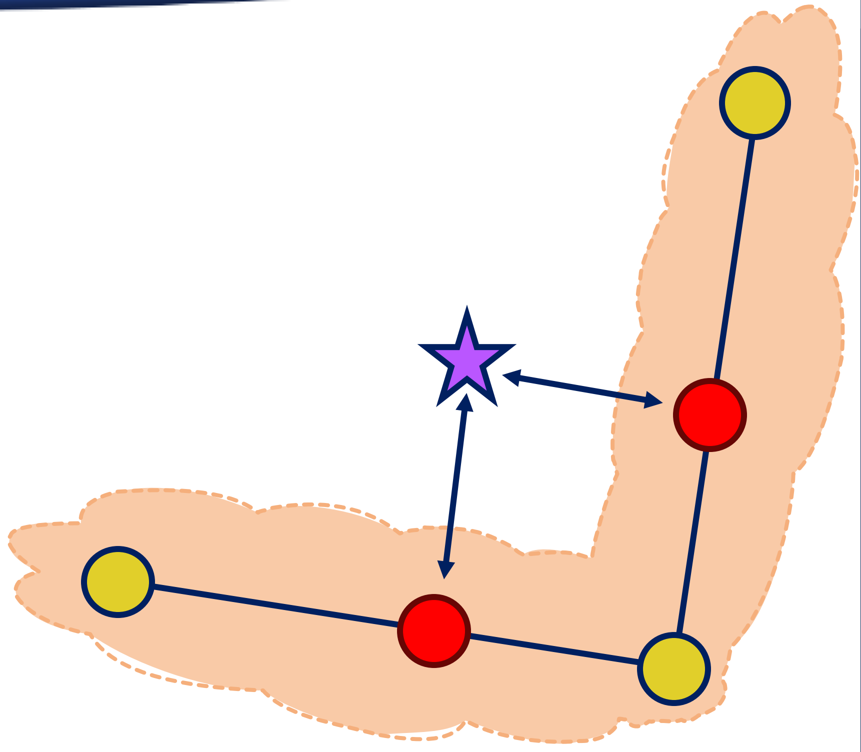

Signal Manifold
SOM Nodes
Experimental Query
1-int Approximation
2-int Approximation

fraser.holloway@liverpool.ac.uk

ML-Assisted Interaction Veto

- The number of interactions can be inferred by estimating the χ^2 for approximations of the underlying manifold.
- 1-interaction SOM approximation requires determining best simplex interpolation in the map.
  - Essentially conventional 2-interaction PSA.
- 2-interaction SOM approximation requires determining the best interpolation of 2 interpolations.
  - Essentially conventional 4-interaction PSA.
    - Solving of a 5-simplex, (hyper-hyper tetrahedron).
    - Moderately difficult.

- If the χ^2 for the 2-interaction approximation is significantly better we can assume the full search will do the same.
- If the χ^2 for the 2-interaction approximation fails the vetoes we can assume that the full search will fail to find a good solution either.

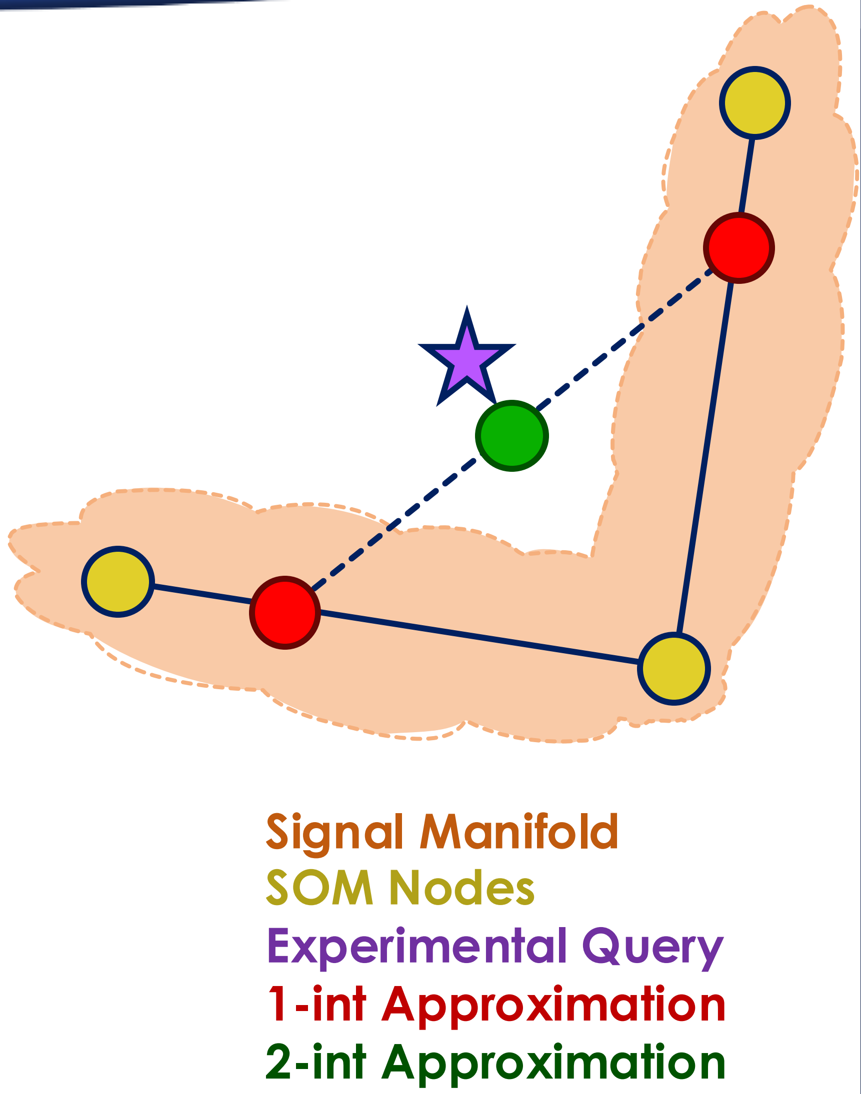

fraser.holloway@liverpool.ac.uk

Intelligent Search Structures

- Initial testing of the veto looks promising but needs proper implementation in AGAPRO.

- Initial predictions of the energy fractions & indices are ignored, could be factored into further PSA.
- Estimation of t_0 offset would reduce iterative fitting.
- Ideally should inform predictions of further PSA.
  - Could be used to build better search structure.

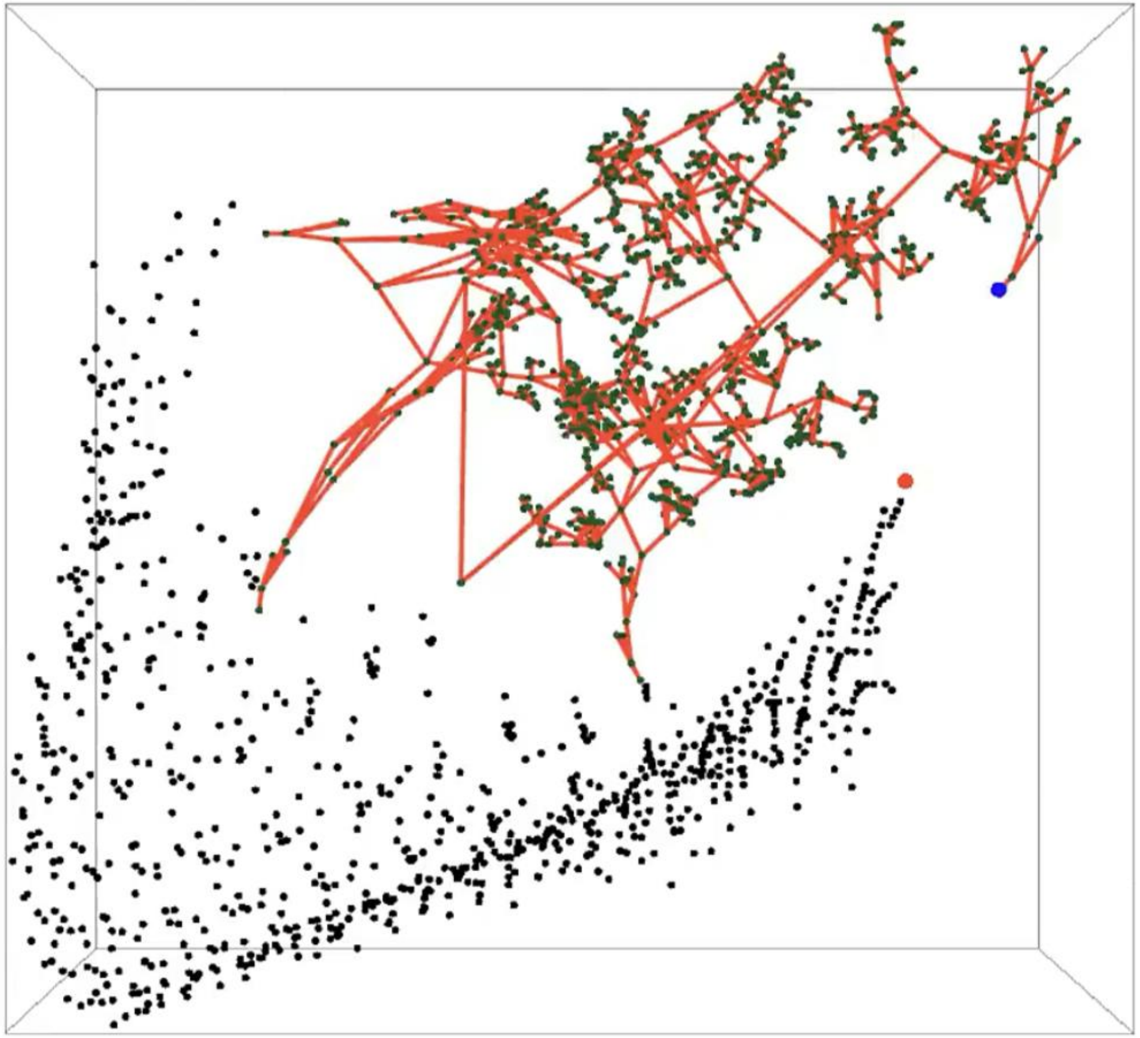

fraser.holloway@liverpool.ac.uk

Thanks for Listening

Any Questions?

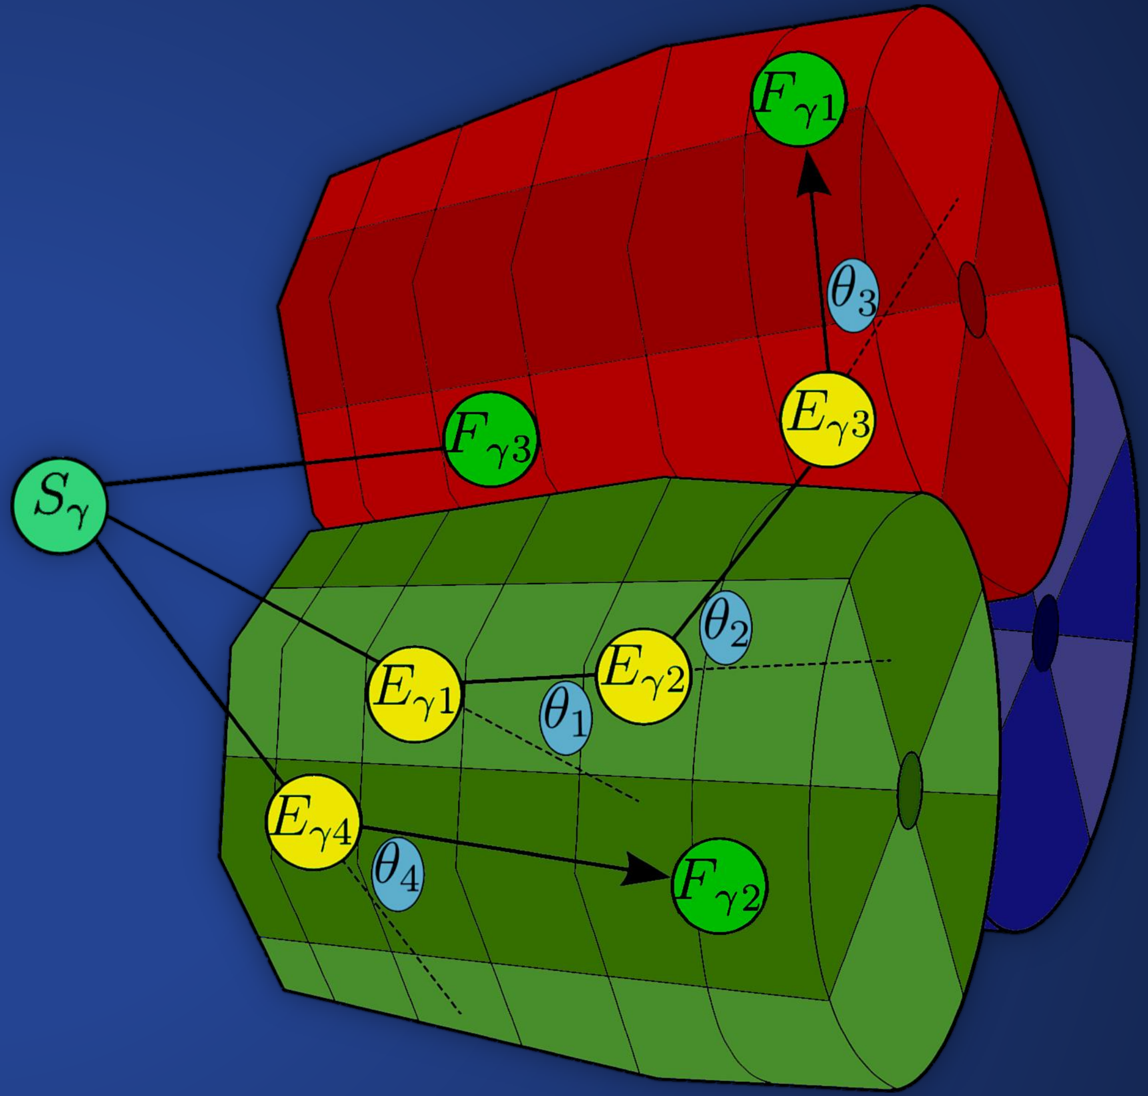

Dr. Fraser Holloway

Fraser.holloway@liverpool.ac.uk

UNIVERSITY OF LIVERPOOL

Science & Technology Facilities Council

AGATA ADVANCED GAMMA TRACKING ARRAY