Anthropomorphic simulation of a low-cost PET scanner

Ryan Atkin, Jemma Bagg, Simon Carthew, James Keaveney

South African Institute of Physics Conference
University of the Western Cape
9 July 2026

UCT-CERN

OPPENHEIMER MEMORIAL TRUST

AIMS African Institute for Mathematical Sciences SOUTH AFRICA

UNIVERSITY OF CAPE TOWN
IYUNIVESITHI YASEKAPA • UNIVERSITEIT VAN KAAPSTAD

Positron Emission Tomography (PET)

- Study the metabolic activity in the body
- Used for cancer treatments, Alzheimer's and heart studies

F^{18}-Fluorodeoxyglucose (FDG)
A glucose analogue

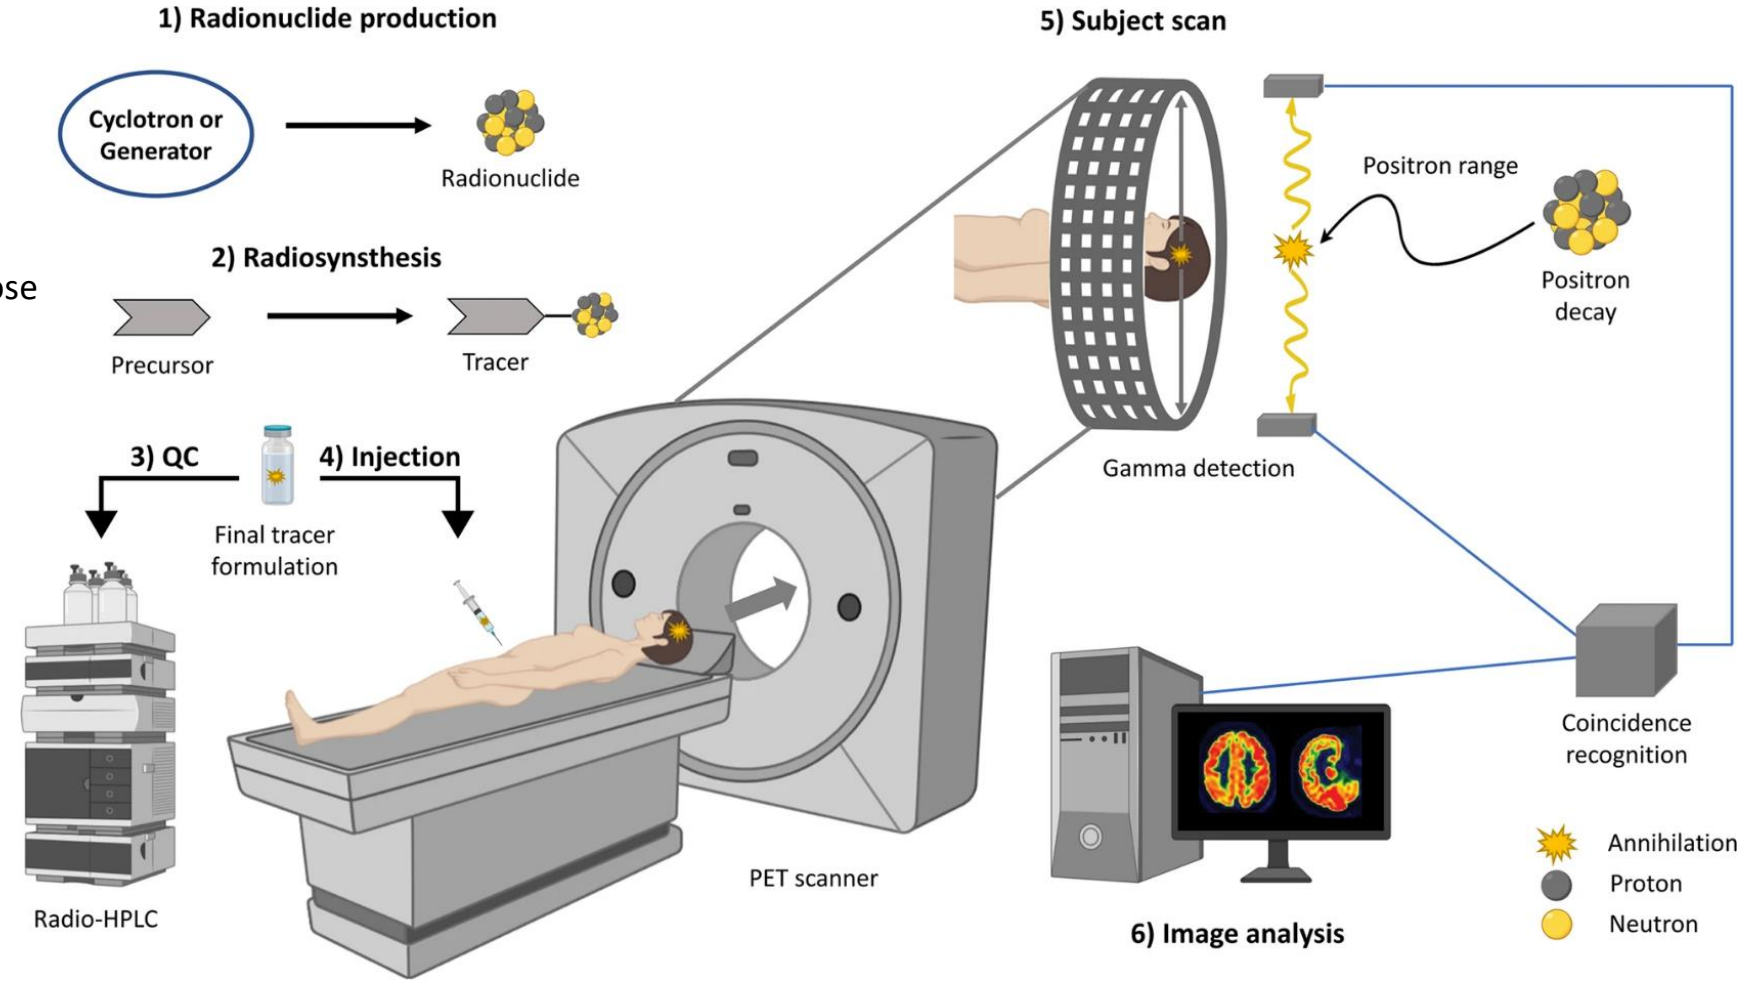

Why low cost

- WHO recommends 2 PET scanners per 1 million people
- SA would require ~120 scanners, however we only have 21 with 5 in the public sector
- With aging populations and increasing prevalence of cancer, improving access to PET scans will be vital

Province	State	Private	Research
Eastern Cape	-	1	-
Free State	-	1	-
Gauteng	3	7	1
KwaZulu-Natal	1	3	-
Limpopo	-	-	-
Northern Cape	-	-	-
Western Cape	1	1	2

Why low cost

- WHO recommends 2 PET scanners per 1 million people
- SA would require ~120 scanners, however we only have 21 with 5 in the public sector
- With aging populations and increasing prevalence of cancer, improving access to PET scans will be vital

Province	State	Private	Research
Eastern Cape	-	1	-
Free State	-	1	-
Gauteng	3	7	1
KwaZulu-Natal	1	3	-
Limpopo	-	-	-
Northern Cape	-	-	-
Western Cape	1	1	2

- Important step in designing a scanner is simulating it
- To this end, we want to create a simulation framework from gamma source, through readout, to basic image reconstruction
  - Will allow easy scanner design testing in future
- Geant4-based, OpenGate10 simulation framework

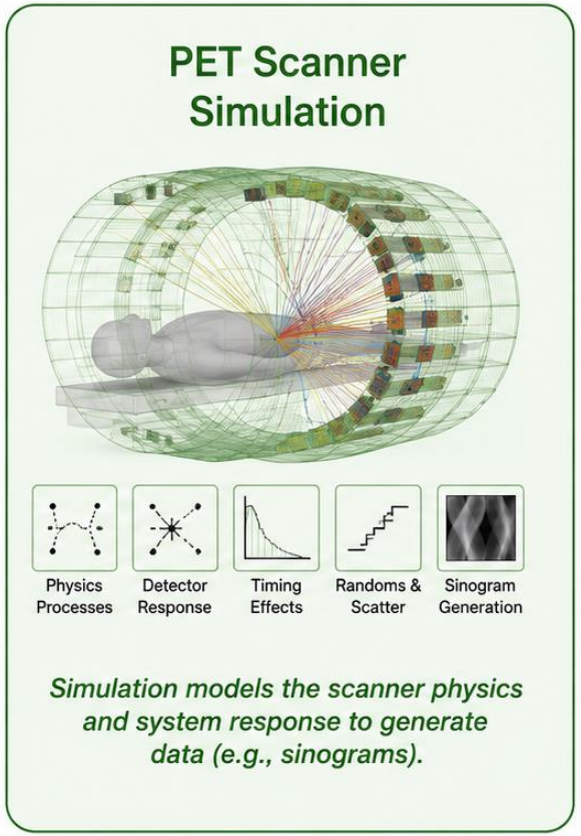

Anthropomorphic brain

https://zenodo.org/records/8045458

- CT images of the head, integrated through 3rd dimension
- Defines the material of our brain volume in the simulation

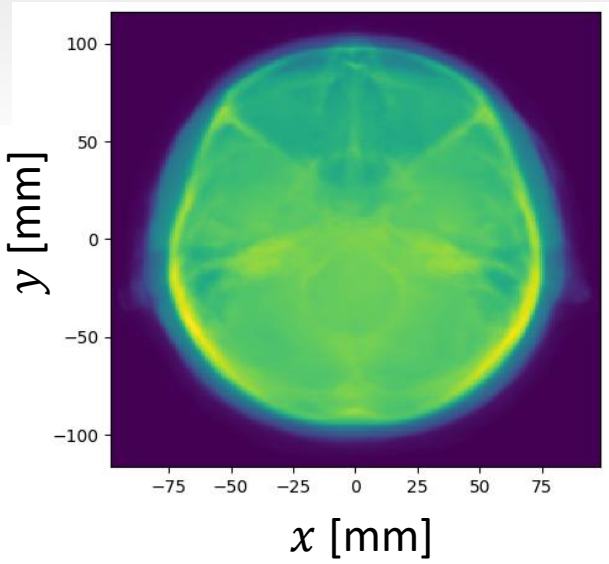

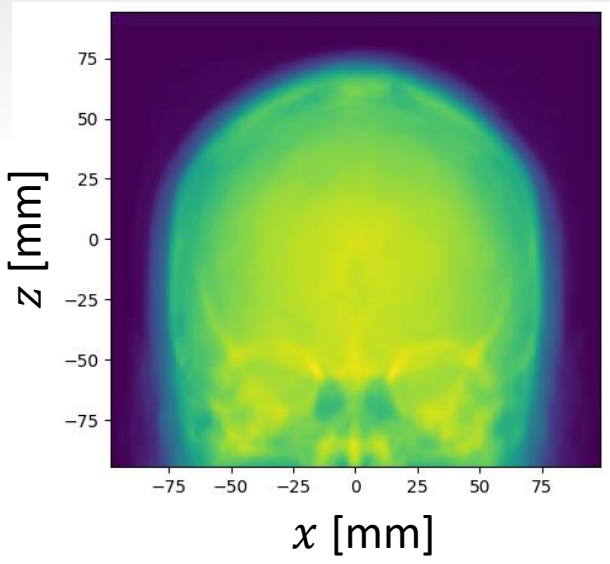

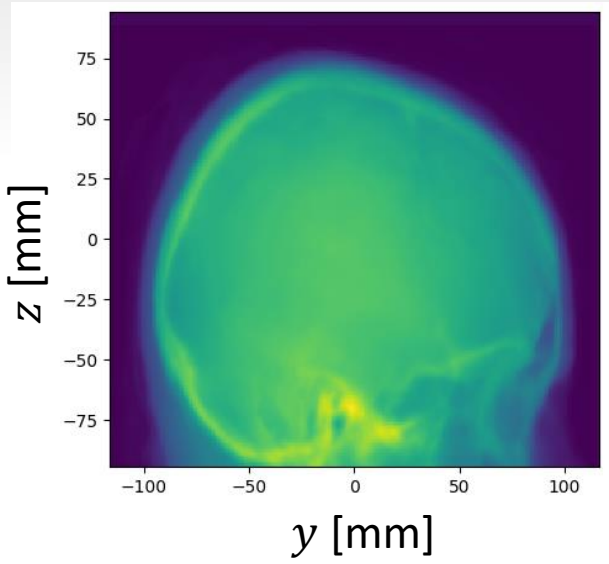

- FDG uptake images of the brain, sliced in 3rd dimension
- Defines the location and rate of the FDG source in the simulation

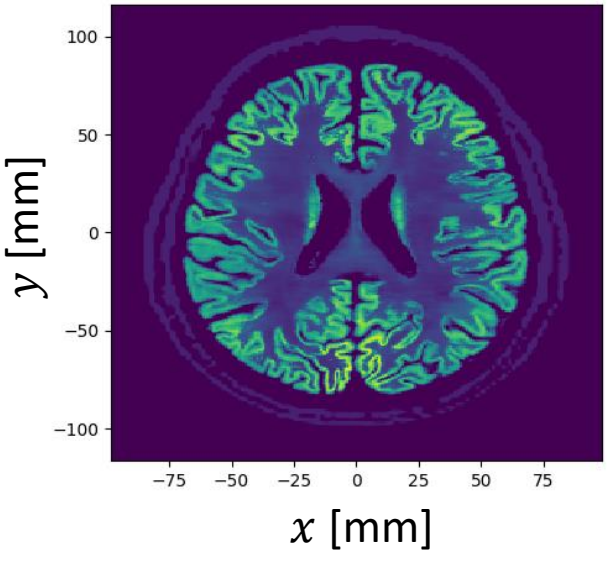

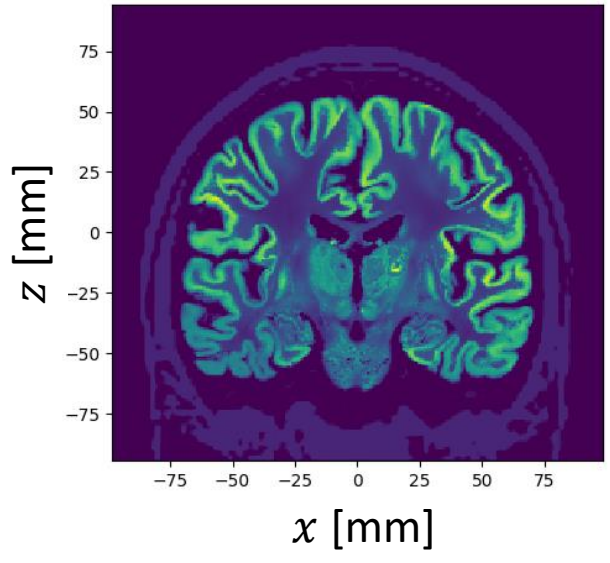

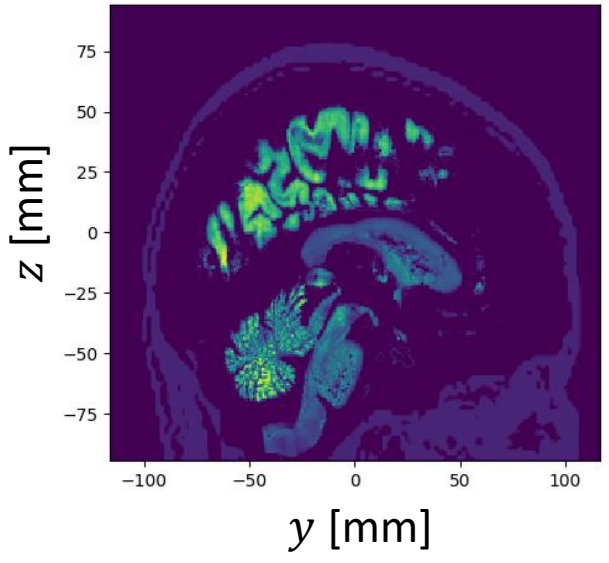

PicoPET current design

- Two rings of modules in axial direction
  - 10 mm between each ring
  - inner radius of 160 mm
- 32 modules per ring
  - 25.68 x 25.68 x 15.2 mm in size
- 8x8 LYSO scintillating crystals per module
  - each separated by 0.08 mm of reflective film
- Crystals are 3.12 x 3.12 x 15 mm in size

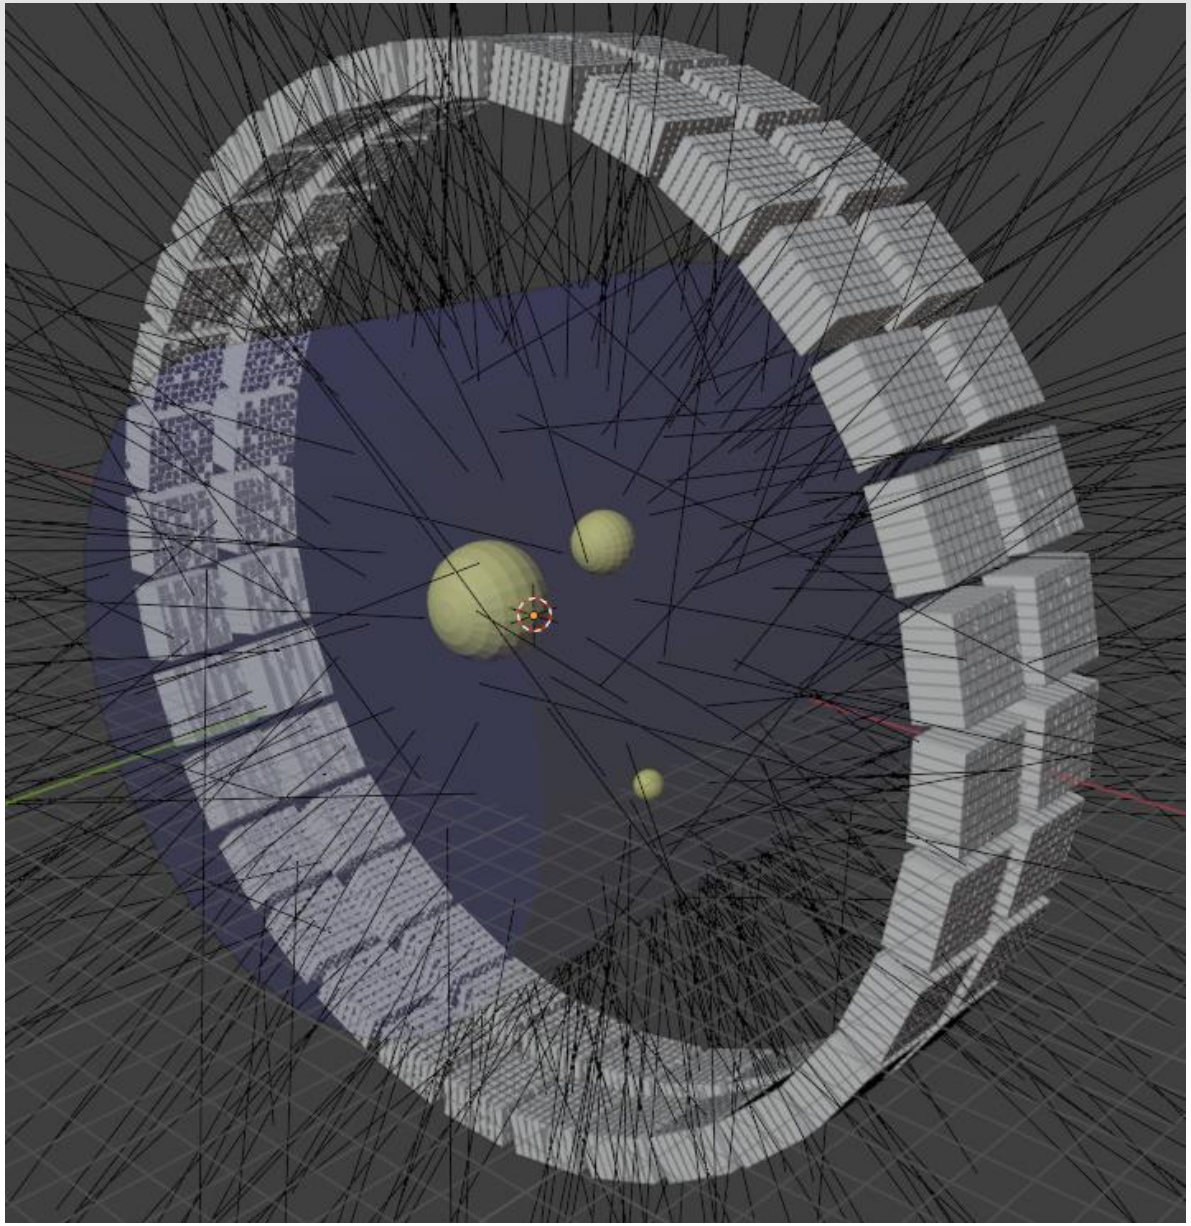

Digitization of energy deposits

Digitizer approximates the output of the electronics in the real scanner

Done in following order:

- Per event, get all the energy deposits in the crystal volumes
- Find the weighted sum per module of the energy deposits
  - Assign a position uniformly at random within that crystal

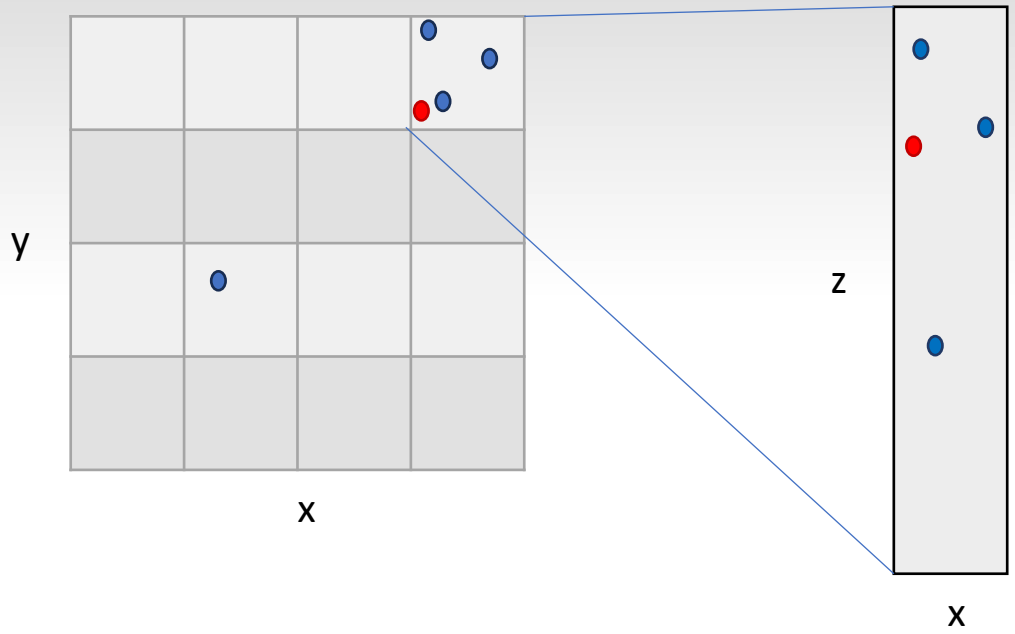

Digitization of energy deposits

Digitizer approximates the output of the electronics in the real scanner

Done in following order:

- Per event, get all the energy deposits in the crystal volumes
- Find the weighted sum per module of the energy deposits
  - Assign a position uniformly at random within that crystal
- Smear the energy via a gaussian with:
  FWHM = $0.112 * \sqrt{511 * E}$
- Smear the timing of the first hit via gaussian with:
  FWHM = 220 ps
- Apply an energy window: 450 < E < 650 keV

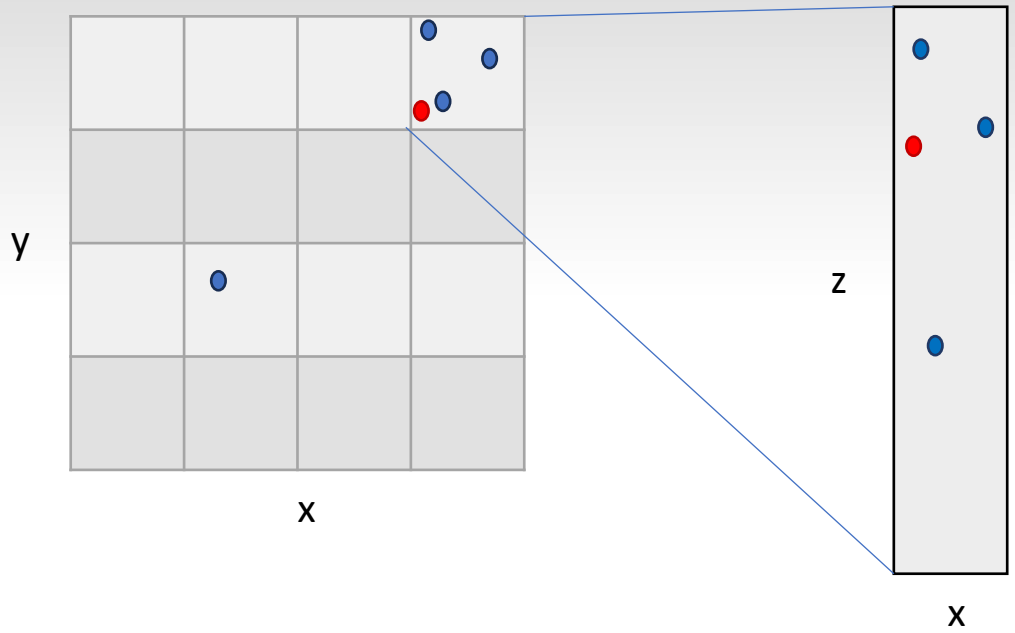

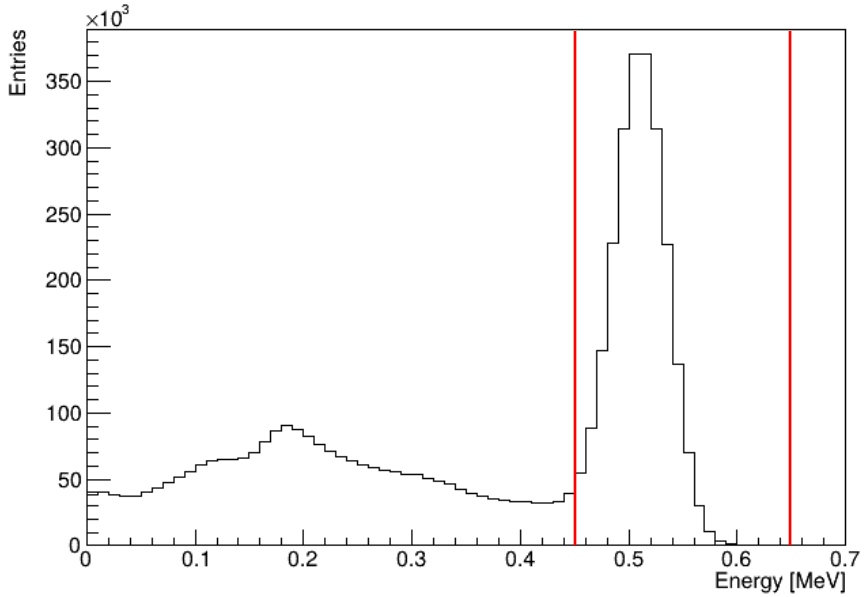

Lines of responses (LoRs)

- LoRs formed by linking two gammas found to be coincident in time window of 1.5 ns
- Defined by the angle and signed length of the distance of closest approach
  - $\theta \in [0,180^{\circ}]$
  - $s = x_0 \cos\theta + y_0 \sin\theta$

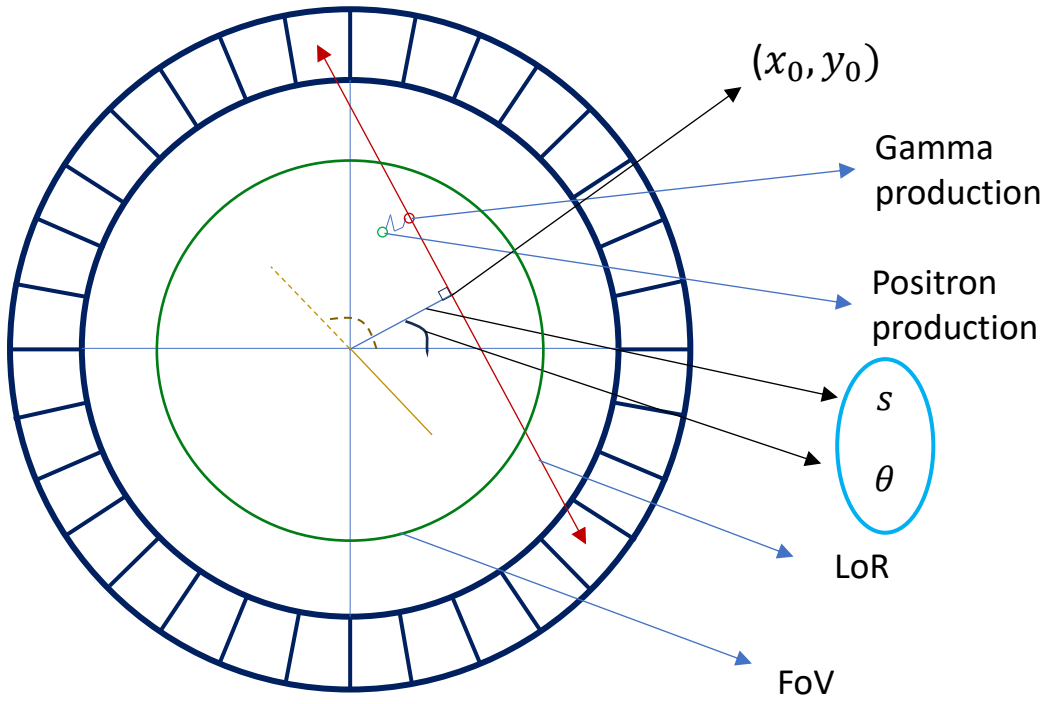

Lines of responses (LoRs)

- LoRs formed by linking two gammas found to be coincident in time window of 1.5 ns
- Defined by the angle and signed length of the distance of closest approach
  - $\theta \in [0,180^{\circ}]$
  - $s = x_0 \cos\theta + y_0 \sin\theta$

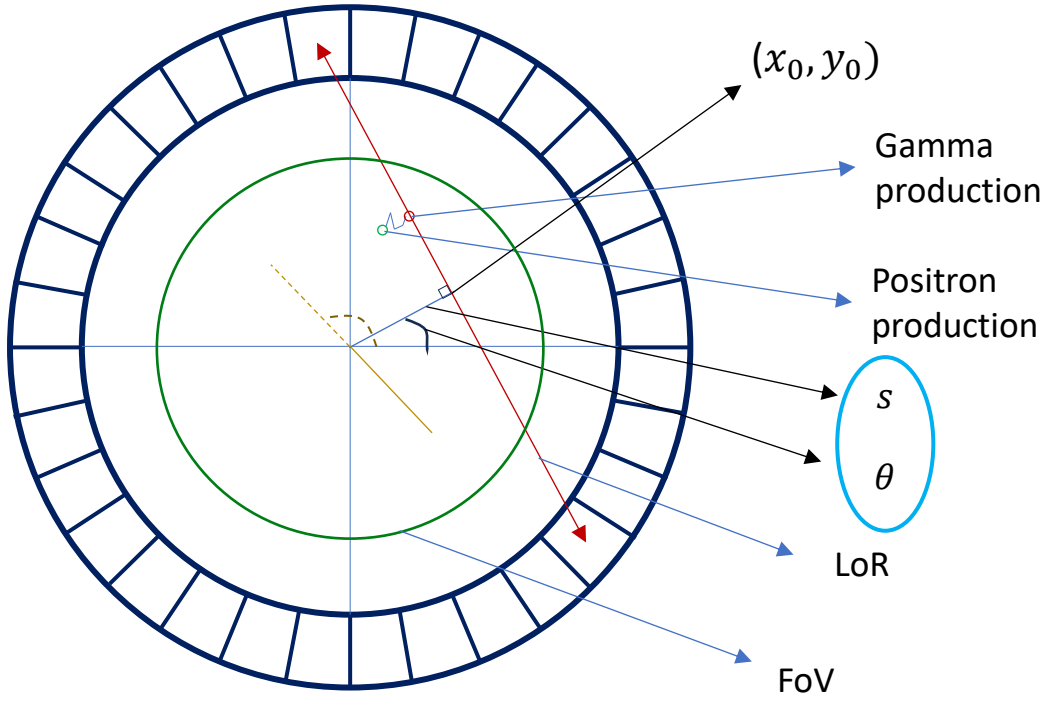

- Use the s and θ to form a sinogram
  - So named since a point source will trace out a sinusoid in the sinogram space
- Sinogram is used to perform image reconstruction

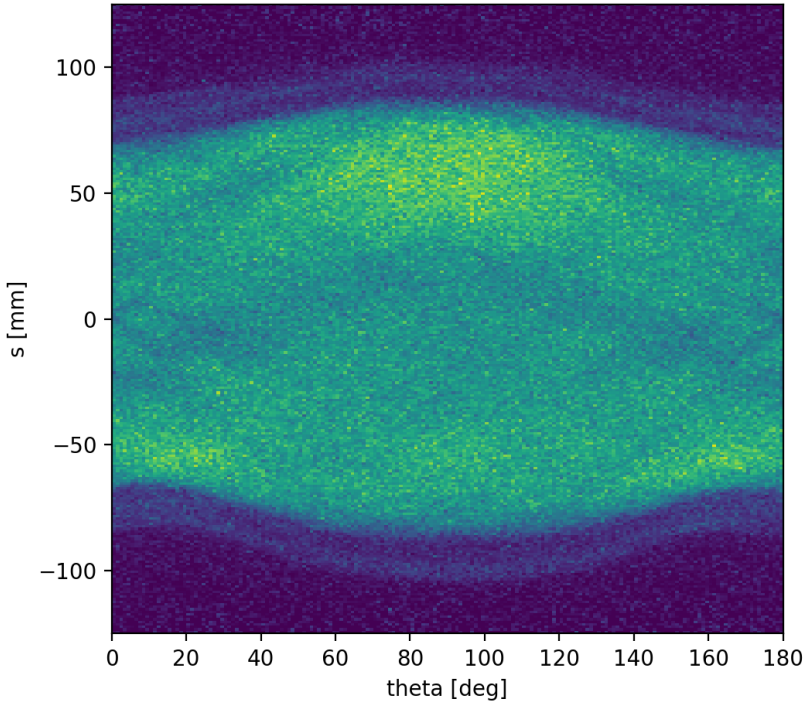

Attenuation

https://zenodo.org/records/8045458

- Shows how much the gammas are likely to be attenuated in different regions as they pass through the head
- Mostly corrects for the gammas that were absorbed in the phantom
- Usually obtained from the CT scan

Z-slice of the μ image

$u = \exp(-\int \mu dl)$

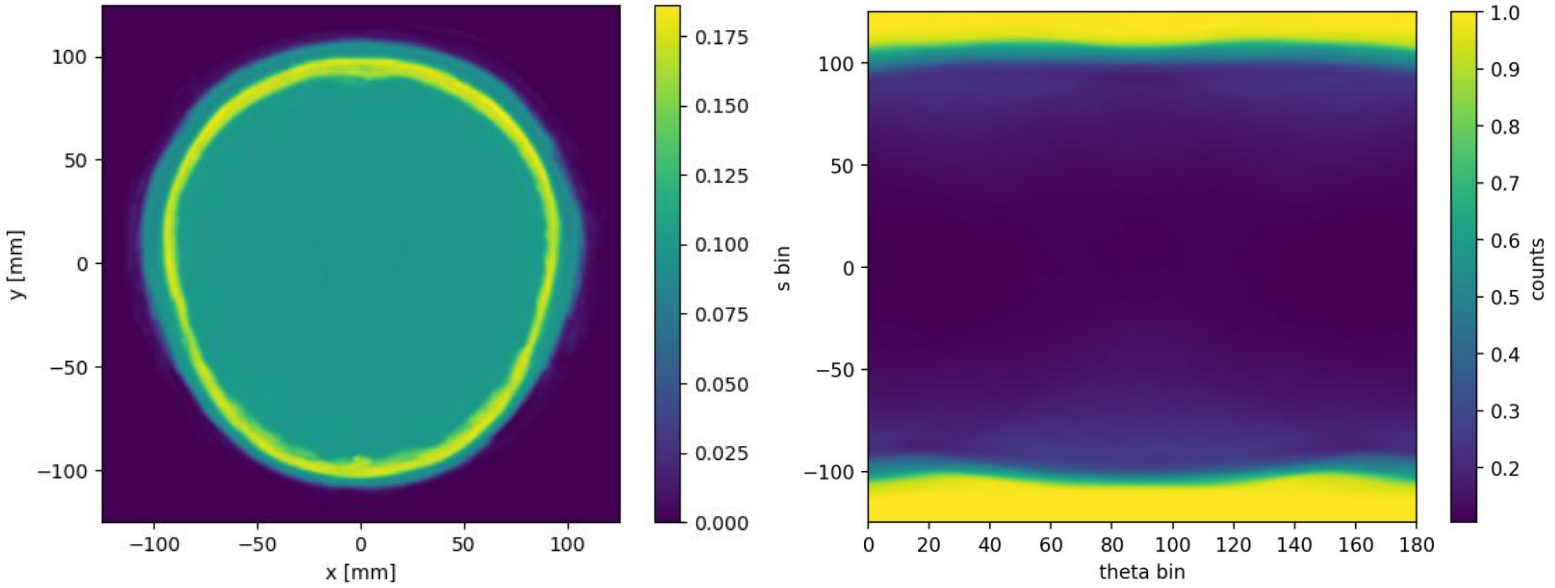

Attenuation and Normalization

https://zenodo.org/records/8045458

- Shows how much the gammas are likely to be attenuated in different regions as they pass through the head
- Mostly corrects for the gammas that were absorbed in the phantom
- Usually obtained from the CT scan

- Corrects for regions in the scanner that have low sensitivity
- Obtained by creating a cylindrical source in the middle of the scanner with a radius of 125 mm
- Covers the full FOV, and so creates all possible LoRs

Z-slice of the μ image

$u = \exp(-\int \mu dl)$

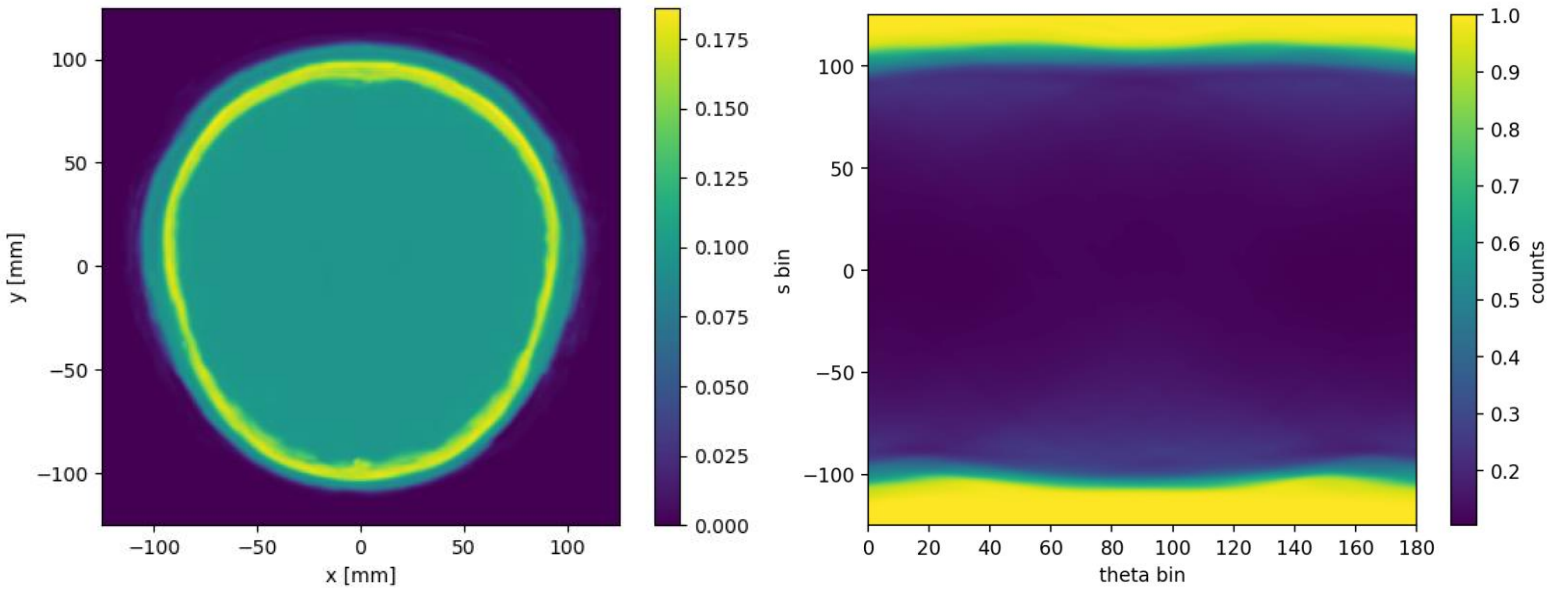

Normalization Sinogram

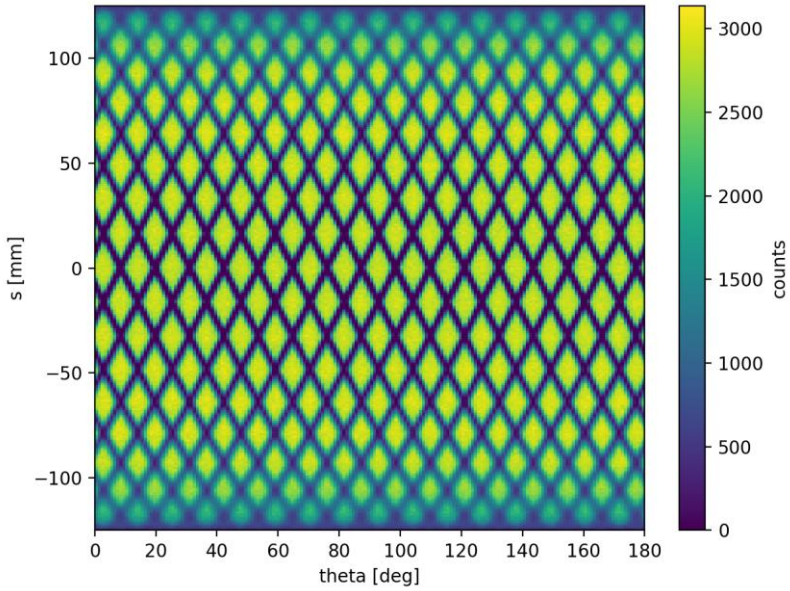

Randoms and Scatters

- Measured coincidences (Prompts) come from three sources:
  - True coincidences (T), Random Coincidences (R), Scattered coincidences (S)
  - $P = T + R + S$

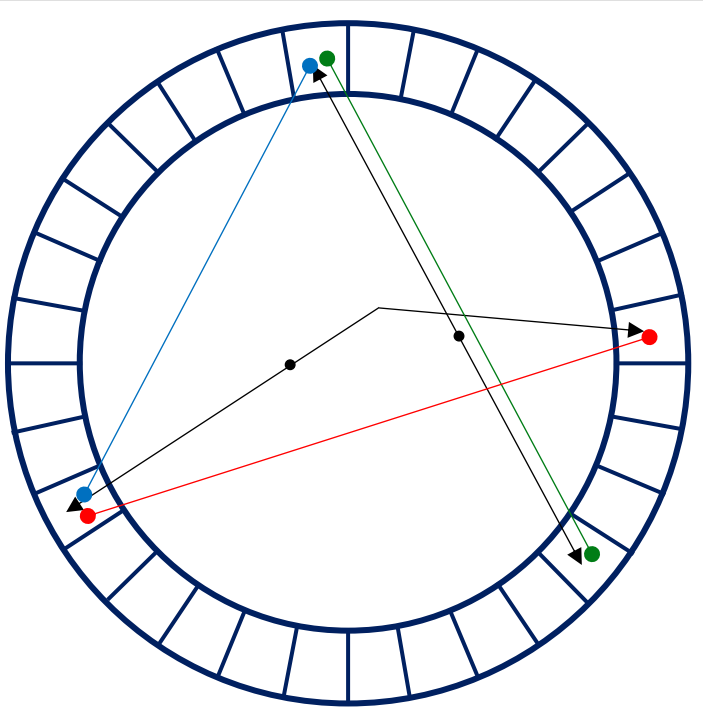

Randoms and Scatters

- Measured coincidences (Prompts) come from three sources:
  - True coincidences (T), Random Coincidences (R), Scattered coincidences (S)
  - $P = T + R + S$
- Randoms can be estimated using a delayed window method
  - Uniform in time

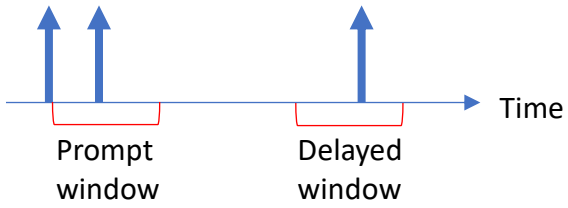

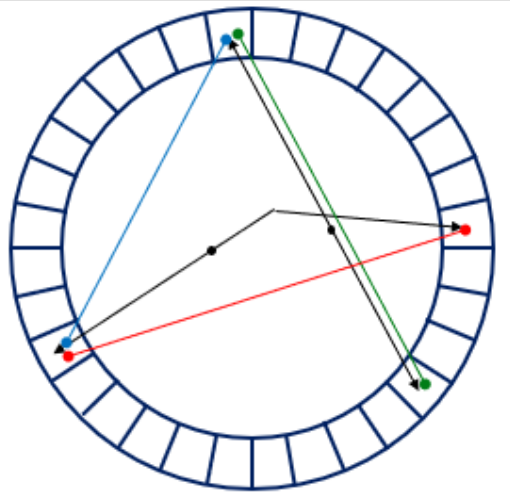

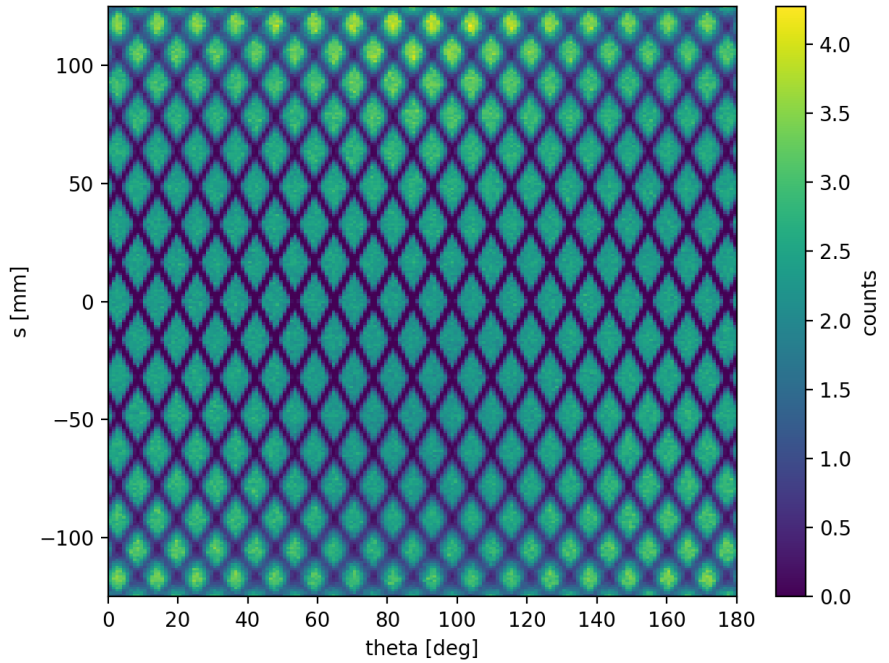

Randoms and Scatters

- Measured coincidences (Prompts) come from three sources:
  - True coincidences (T), Random Coincidences (R), Scattered coincidences (S)
  - $P = T + R + S$
- Randoms can be estimated using a delayed window method
  - Uniform in time
- Still debating exactly how to estimate scatters, so for the sake of this talk, the scatters are defined using the truth information
  - Coincidences formed from gammas whose $E_k \neq 511\ \mathrm{keV}$ as it enters the module are considered scatters

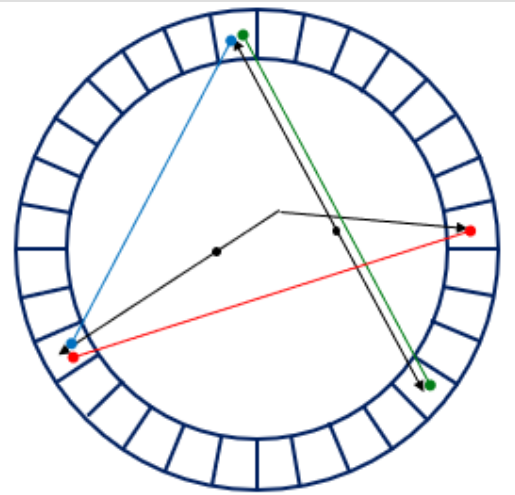

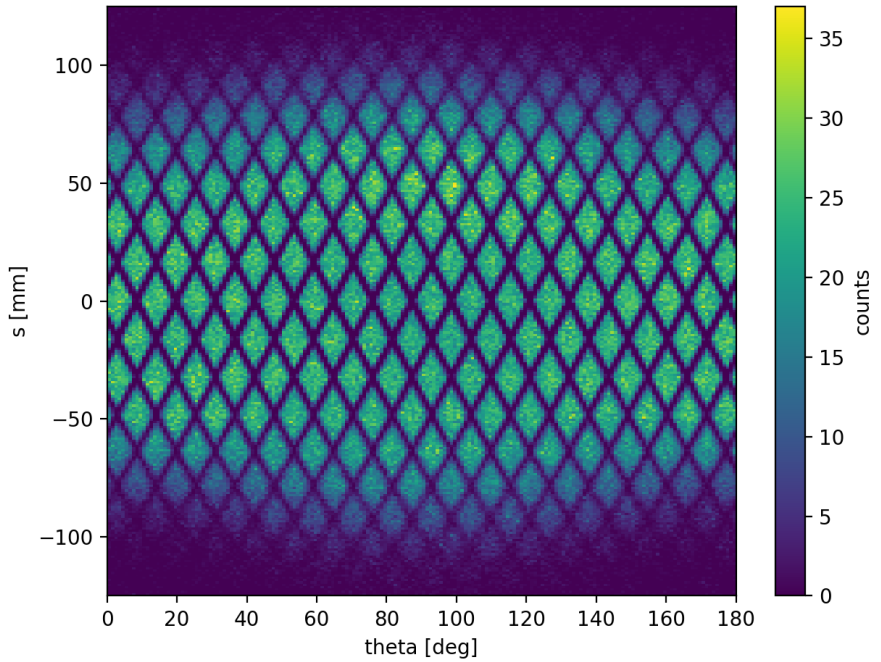

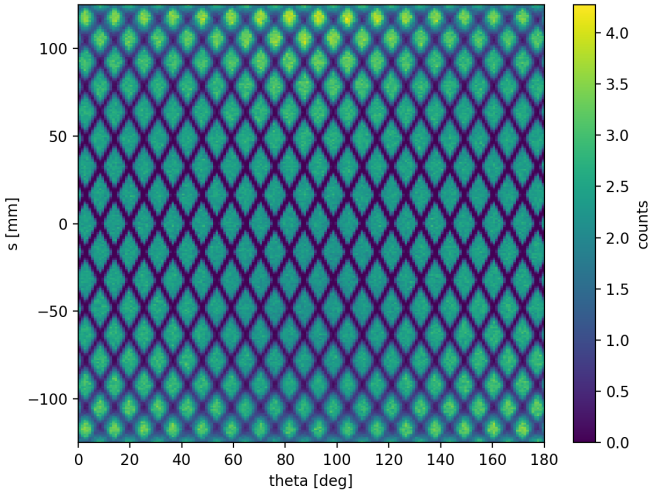

Z-slicing

- Ideally, we would do 3D Image reconstruction
  - However, computationally expensive and is more complicated than 2D reconstruction
- Instead, we do 2D reconstruction in axial slices using the full 3D data

Z-slicing

- Ideally, we would do 3D Image reconstruction
  - However, computationally expensive and is more complicated than 2D reconstruction
- Instead, we do 2D reconstruction in axial slices using the full 3D data

- Assign every LoR to every slice with a corresponding weight
  - $w = \frac{dz}{B-A}$
  - $dz \equiv$ axial distance LoR covers that slice within the FoV
  - A, B are the z points where the LoR crosses the FoV

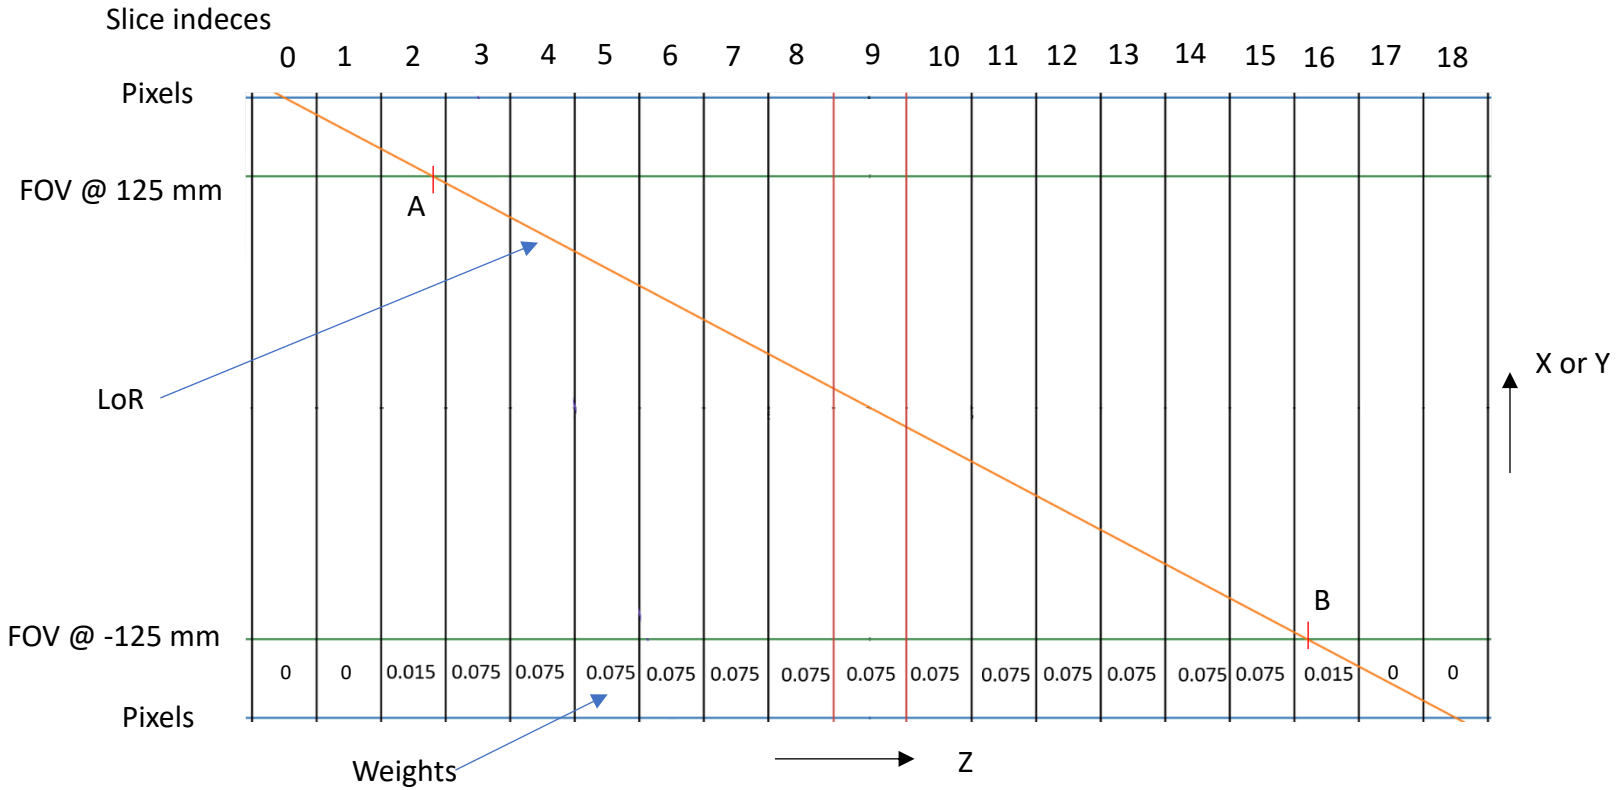

MLEM

Maximum Likelihood Expectation Maximization is an iterative approach

- y = input sinogram
- $x^k = k^{th}$ iteration of our image guess
  - Starts as uniform image, i.e., $x^0 = \mathbf{1}_{m,n}$
- A = forward projection matrix
  - Going from image to sinogram
- A^T = back projection matrix
  - Going from sinogram to image
- $C = N * u$
  - $u \equiv$ exp(- attenuation sinogram)
  - $N \equiv$ normalisation sinogram
- $b = r + s$ is the sinogram contribution due to "background" gammas
  - $r \equiv$ sinogram due to randoms
  - $s \equiv$ sinogram due to scatters

$$x^{k+1} = x^k \frac{A^T\left(C * \frac{y}{C * A(x^k) + b}\right)}{A^T(C)}$$

MLEM

Maximum Likelihood Expectation Maximization is an iterative approach

- y = input sinogram
- $x^k = k^{th}$ iteration of our image guess
  - Starts as uniform image, i.e., $x^0 = \mathbf{1}_{m,n}$
- A = forward projection matrix
  - Going from image to sinogram
- A^T = back projection matrix
  - Going from sinogram to image
- $C = N * u$
  - $u \equiv$ exp(- attenuation sinogram)
  - $N \equiv$ normalisation sinogram
- $b = r + s$ is the sinogram contribution due to "background" gammas
  - $r \equiv$ sinogram due to randoms
  - $s \equiv$ sinogram due to scatters

$$x^{k+1} = x^k \frac{A^T\left(C * \frac{y}{C * A(x^k) + b}\right)}{A^T(C)}$$

$A^T(N * u)$

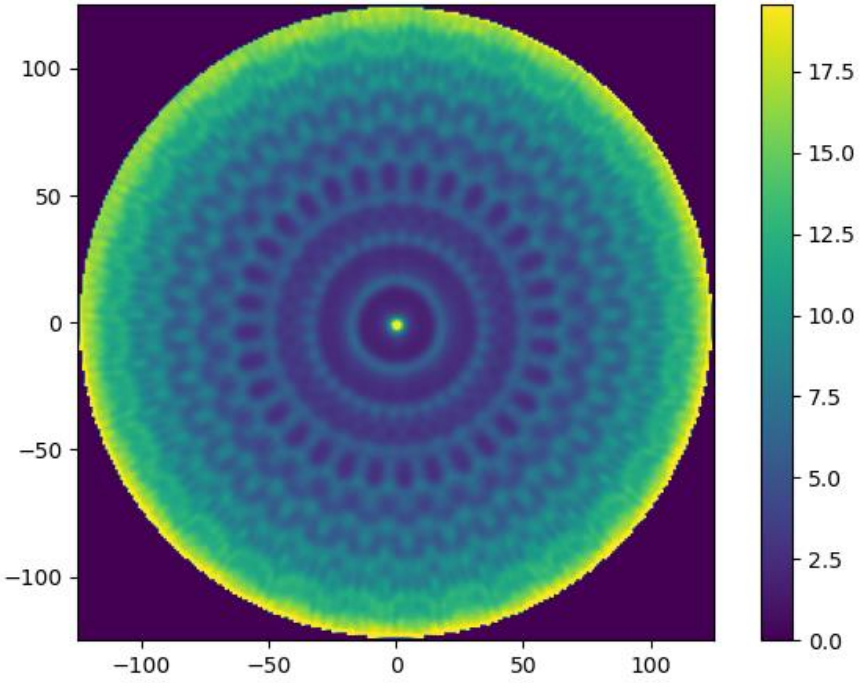

MLEM reconstruction

- 20 minute scan with 20 MBq of activity
- Sliced for $-17.84\ \mathrm{mm} < z < -14.64\ \mathrm{mm}$
- 10 iterations of the MLEM algorithm

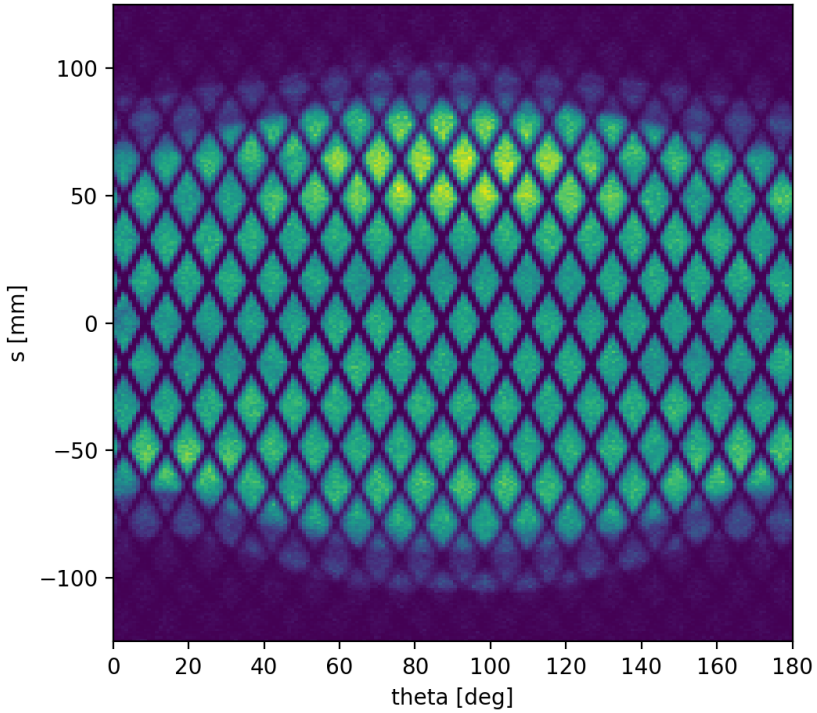

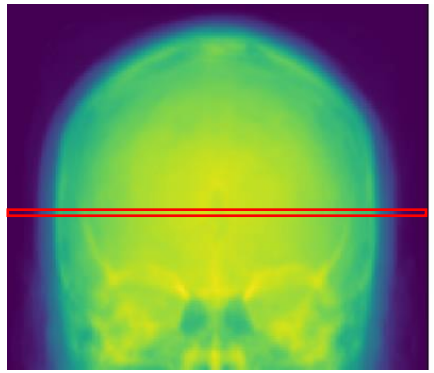

MLEM reconstruction

- 20 minute scan with 20 MBq of activity
- Sliced for $-17.84\ \text{mm} < z < -14.64\ \text{mm}$
- 10 iterations of the MLEM algorithm

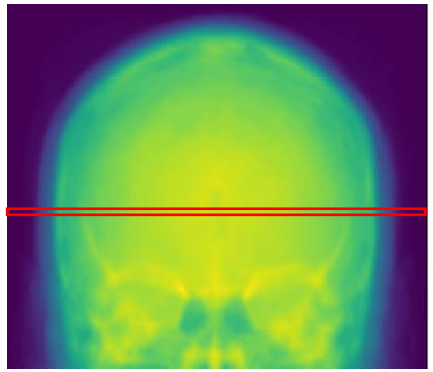

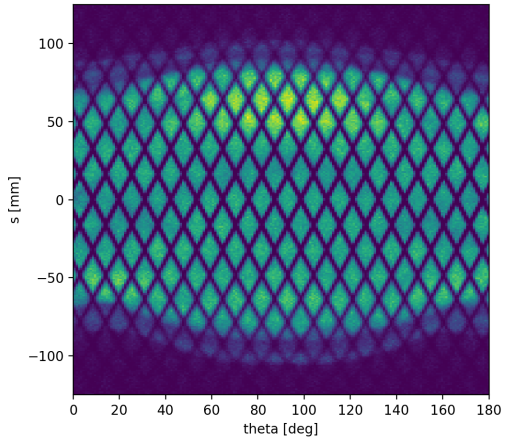

True FDG image

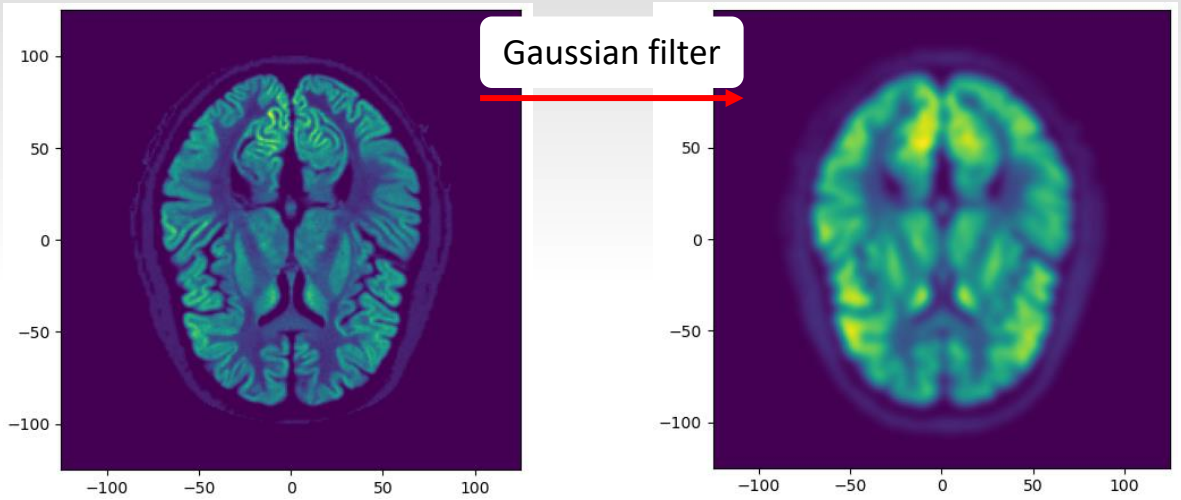

MLEM reconstruction

- 20 minute scan with 20 MBq of activity
- Sliced for $-17.84 < z < -14.64$
- 10 iterations of the MLEM algorithm
- MSE and SSIM shown in title of reconstructed image

True FDG image

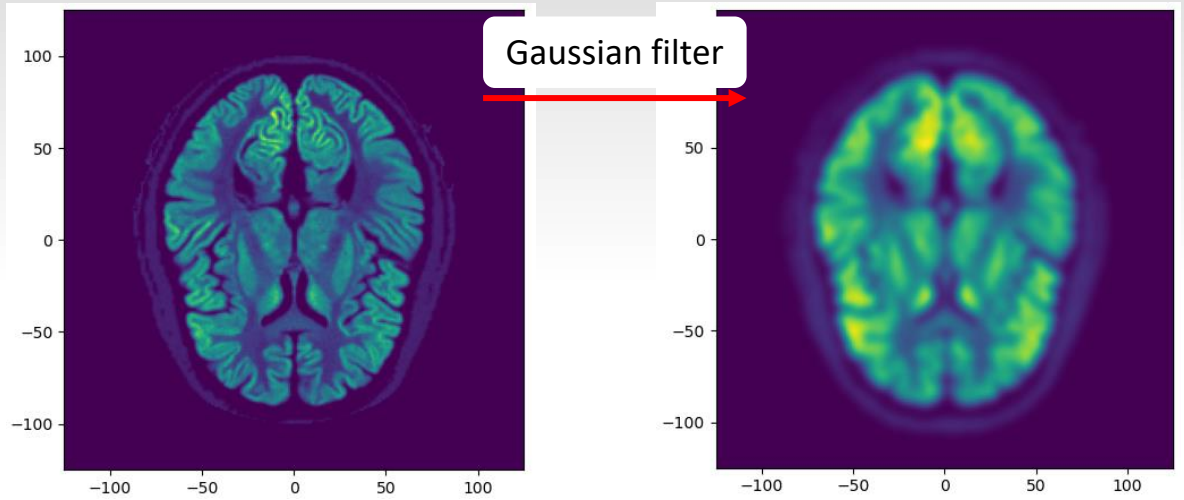

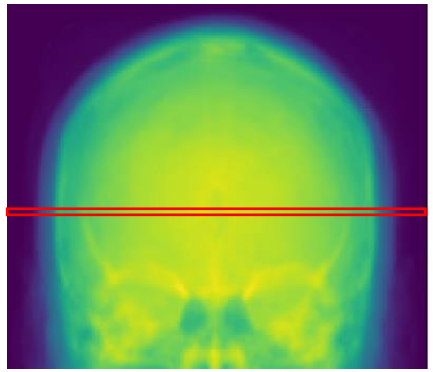

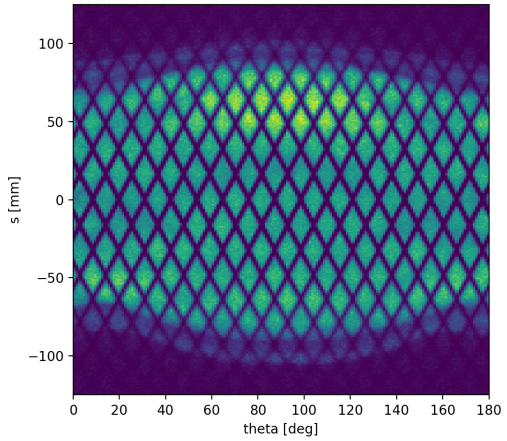

Example of a real PET scan

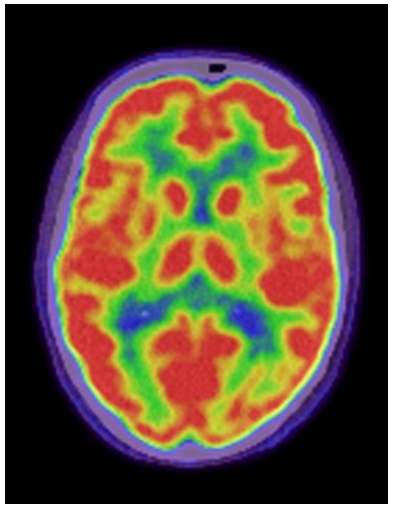

https://radiopaedia.org/cases/normal-brain-pet

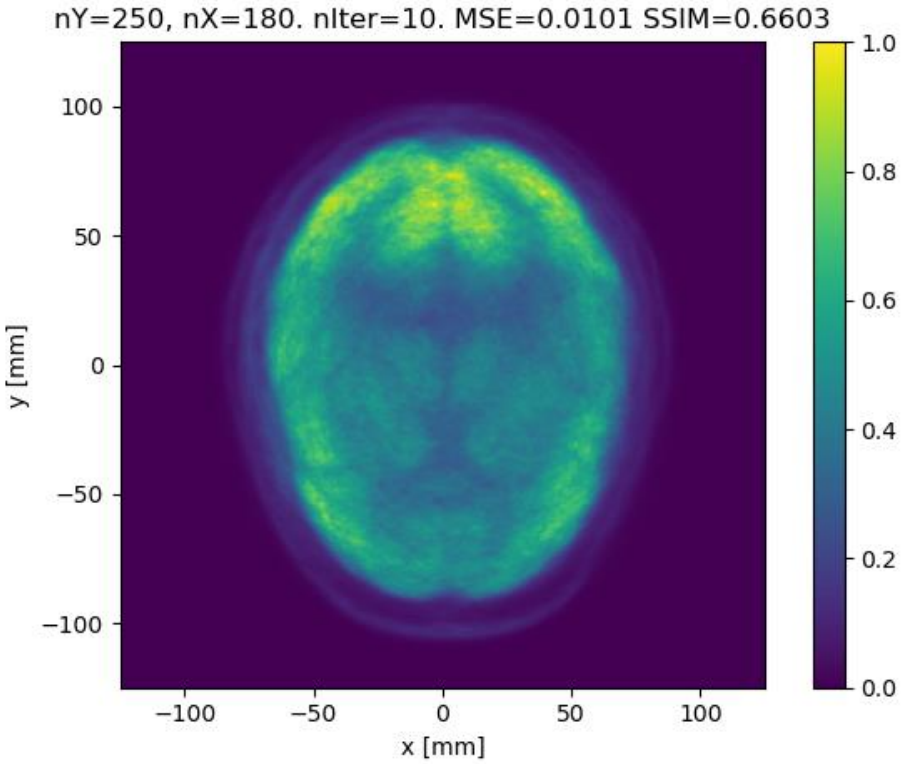

MLEM reconstructions

No corrections

Attenuation, normalisation

Attenuation, normalisation and backgrounds

10 iterations

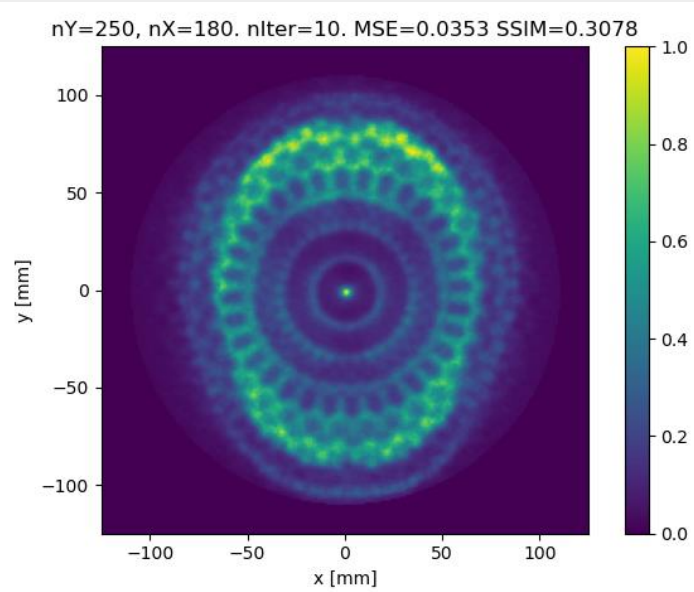

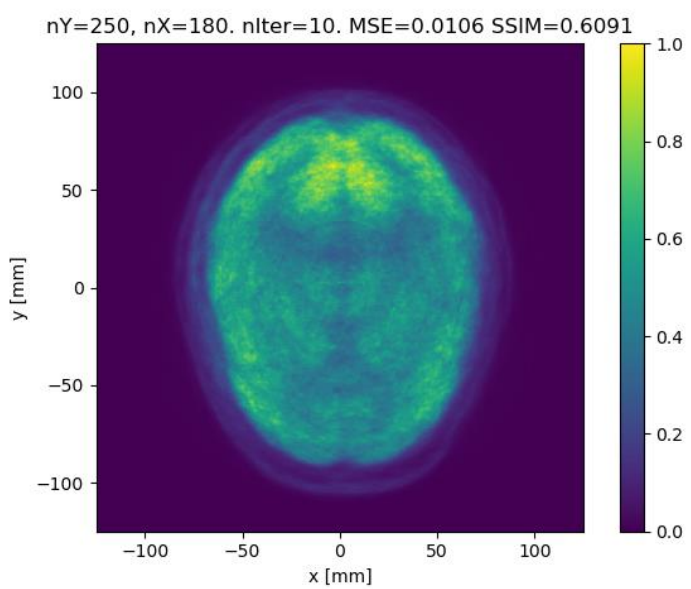

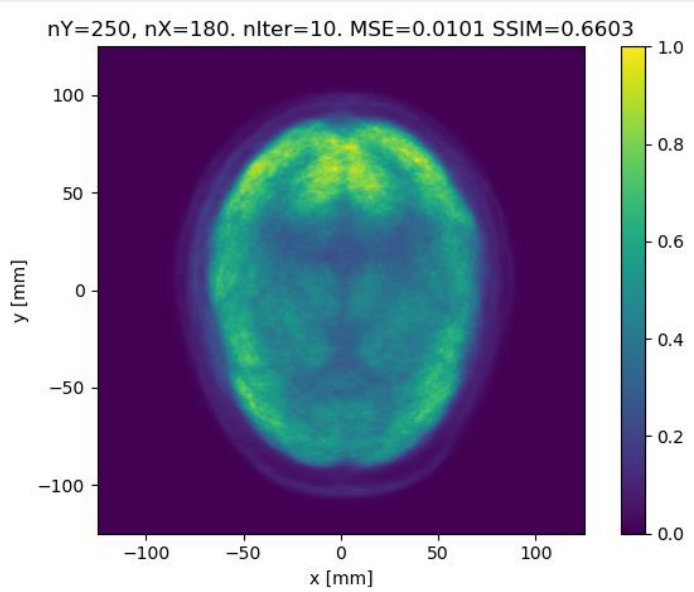

30 iterations

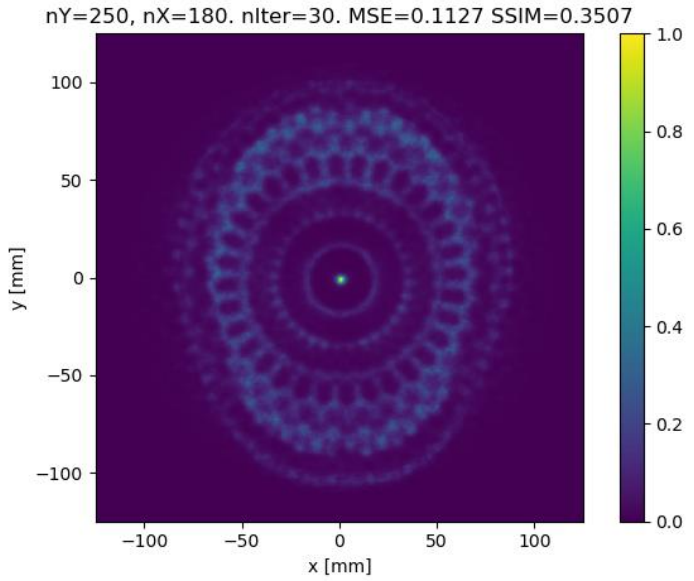

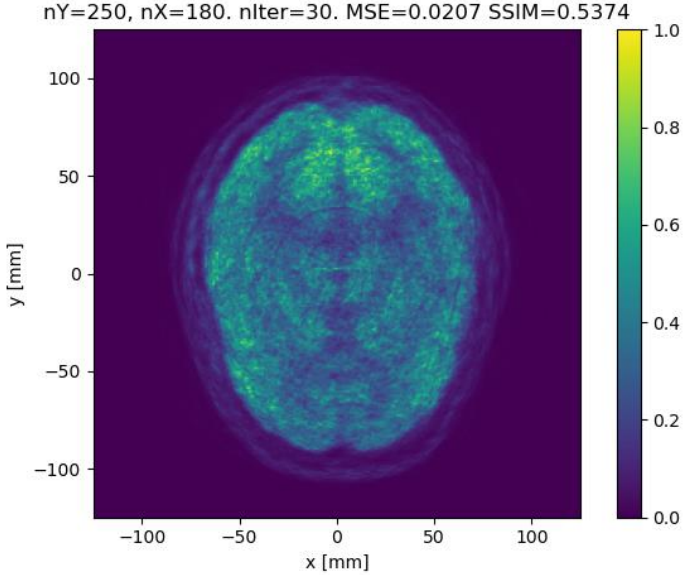

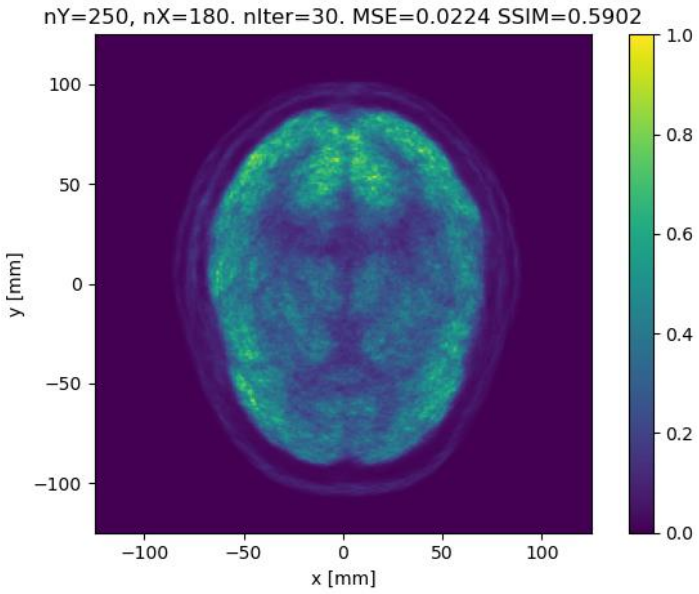

Conclusion

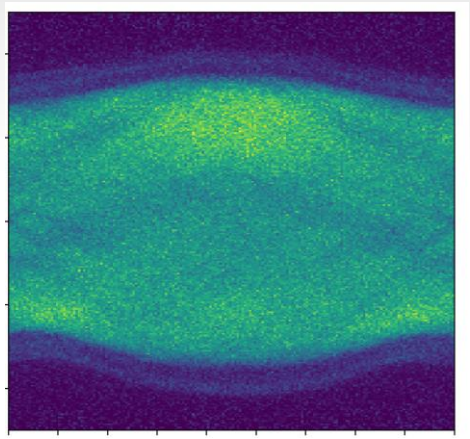

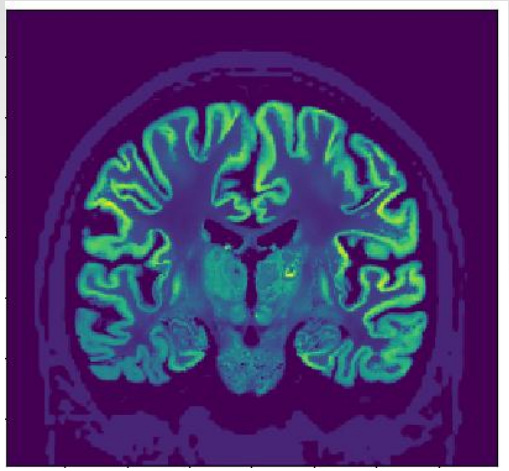

- Designing a low-cost PET scanner
- Creating a basic, yet full simulation framework for a PET scan, using an anthropomorphic brain
- After taking into account several corrections, we have promising results
- Still need to sort out scatters estimation
- Test different designs

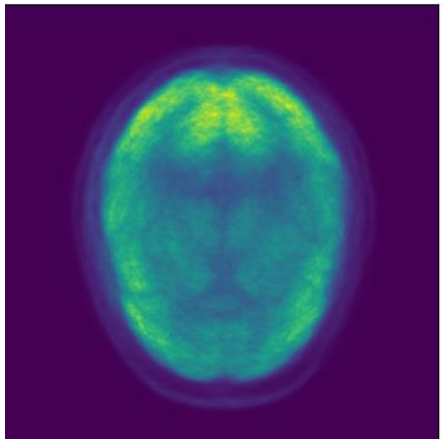

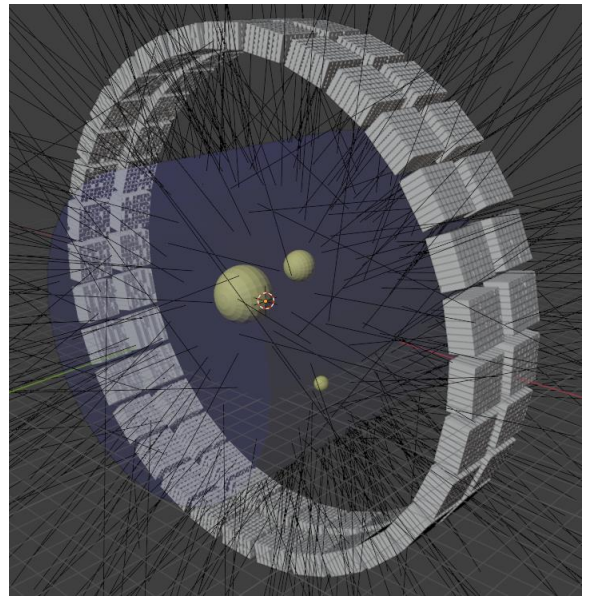

Extra Slides

Development

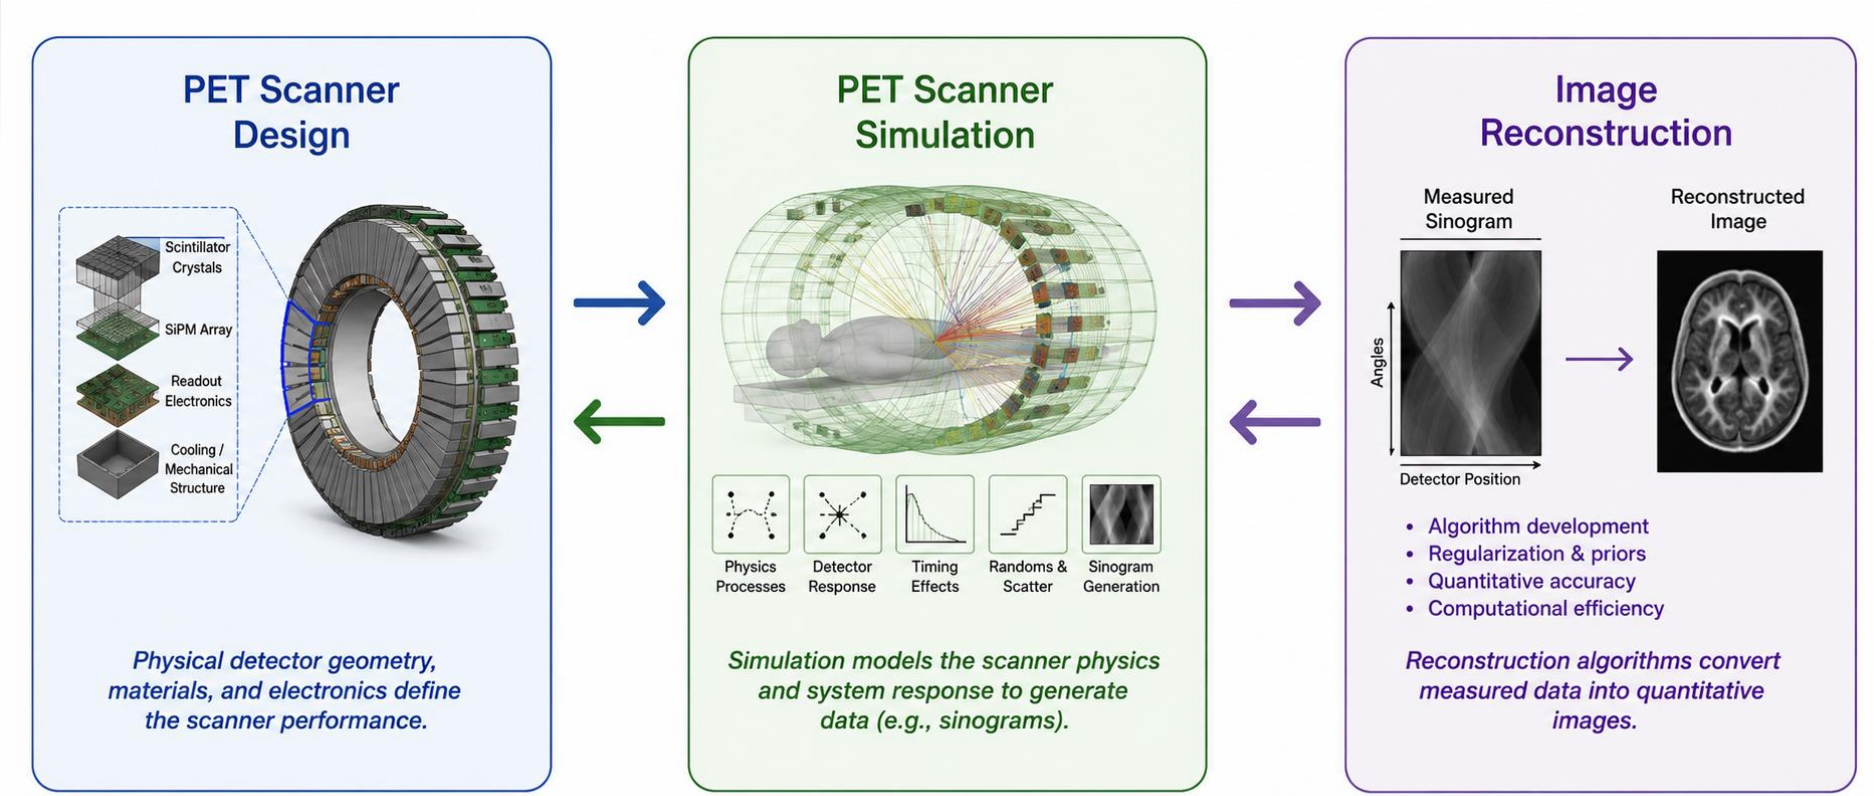

- Lies in-between the detector design and the image reconstruction research
- Informed by both the design and image reconstruction, while also influencing the progress of the two
- Provides an understanding of the scanner long before it exists

Sinograms

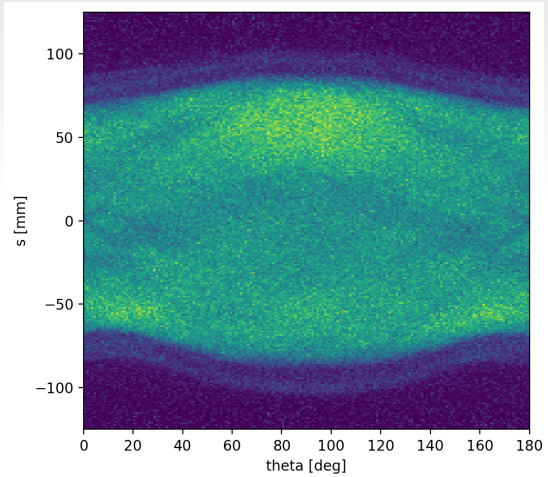

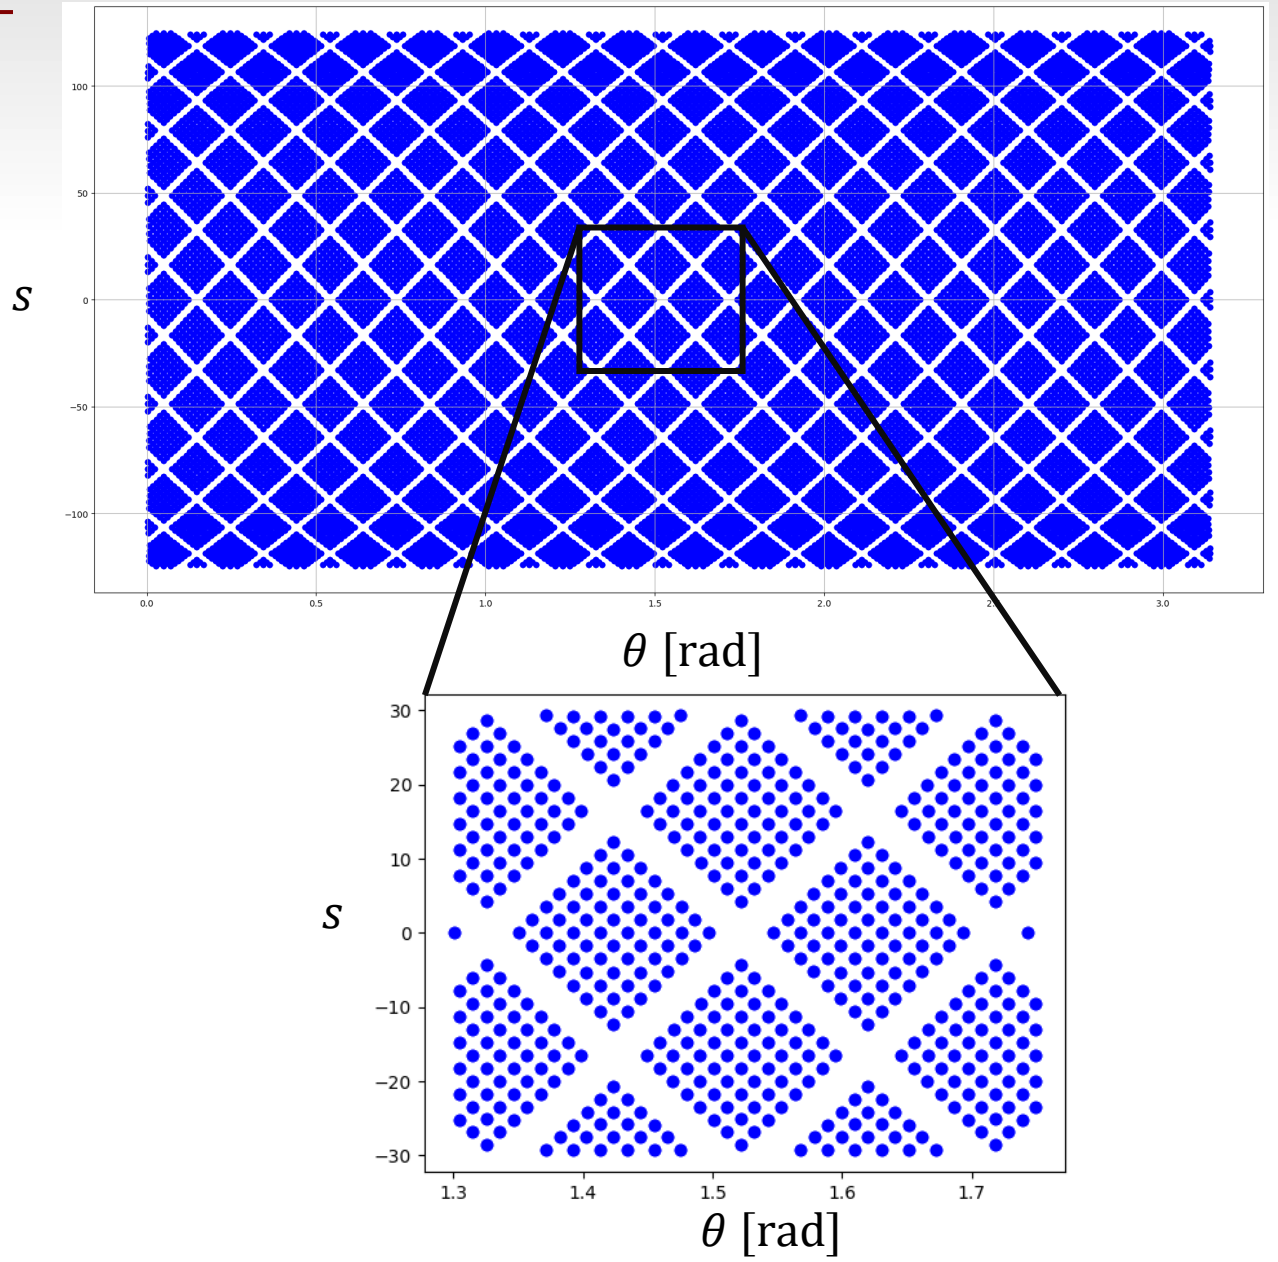

- Create a 2D histogram from (θ, s)
  - Currently bins are 250 in s, and 180 in θ
  - So each bin is $1^{\circ} \times 1\text{mm}$
- With our current PicoPET_v2 scanner:
  - We have 32 inter-module gaps per ring
  - Basic digitisation also discretises the (θ, s) to a point
- Three things to deal with:
  - The empty bins due to the inter-module gaps
  - Dealing with the discretisation of the hits
  - How to slice the data in the axial direction (z)

Dealing with gaps

- Crystal smearing applied before forming the LoR
- Take the centre of the pixel, uniformly smear it within extent of the pixel
  - The hit could have been anywhere in the pixel, so smear it as such
- Removes any gaps within the module

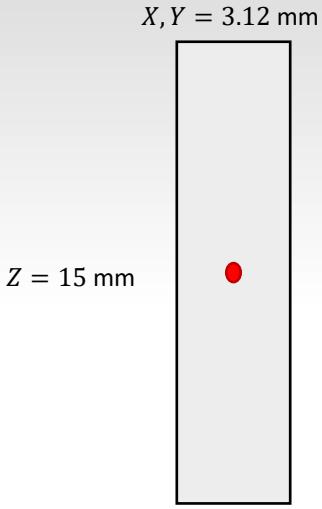

- After crystal smearing and sinogram formation, find all the bins in the sinogram with zero entries
- Perform simple interpolation across θ to fill in the gaps

Bin values > 0 set to 1

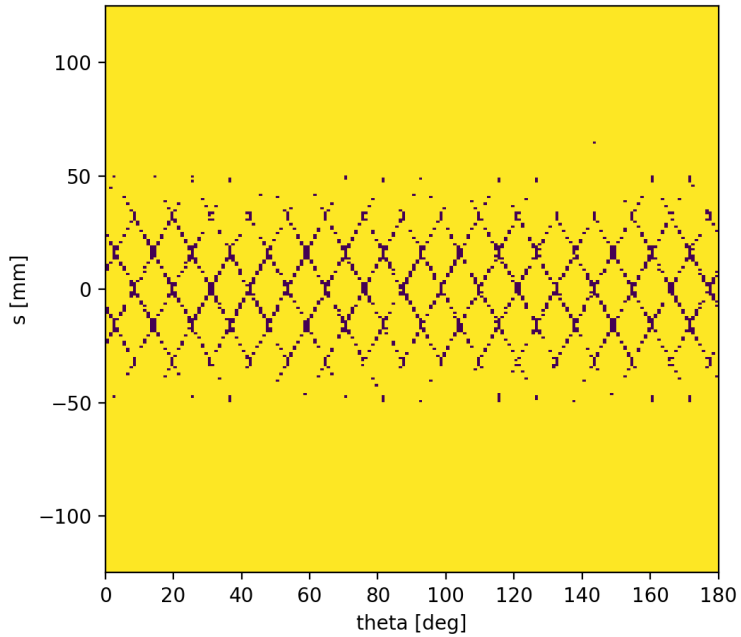

Original results

- Very first results with no corrections applied
- Did 10 iterations of MLEM
- Inter-module gaps and the discretisation show up as rings in the reconstructed image

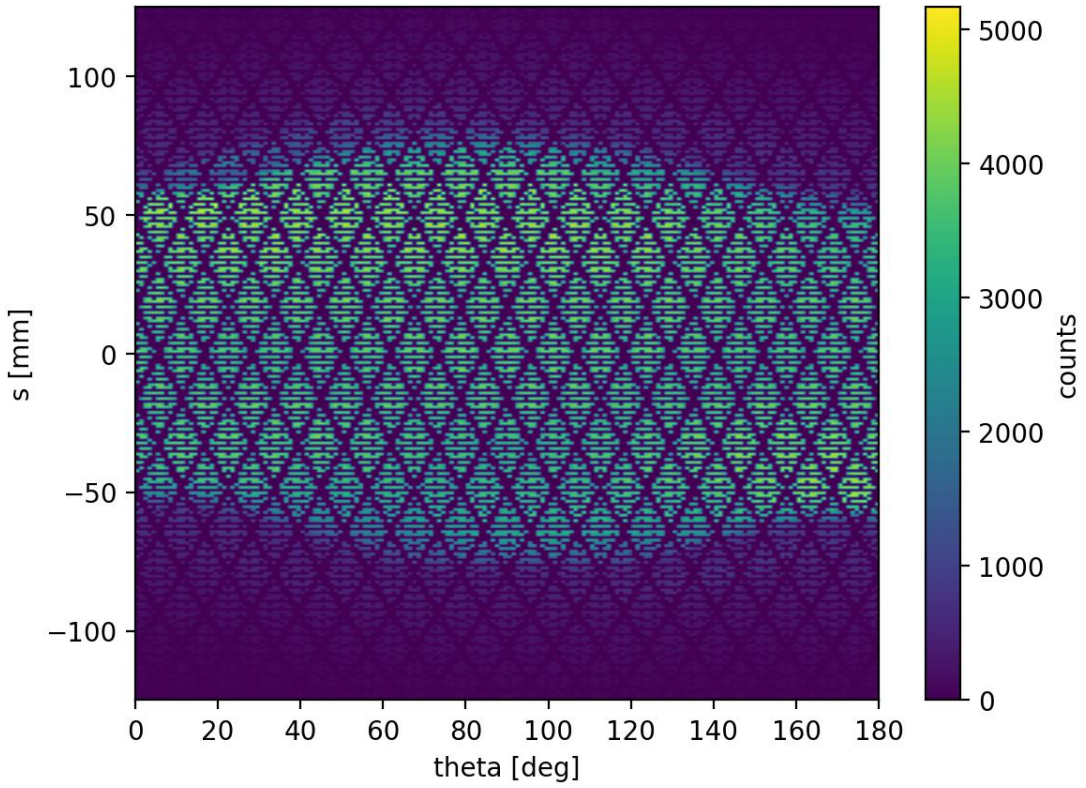

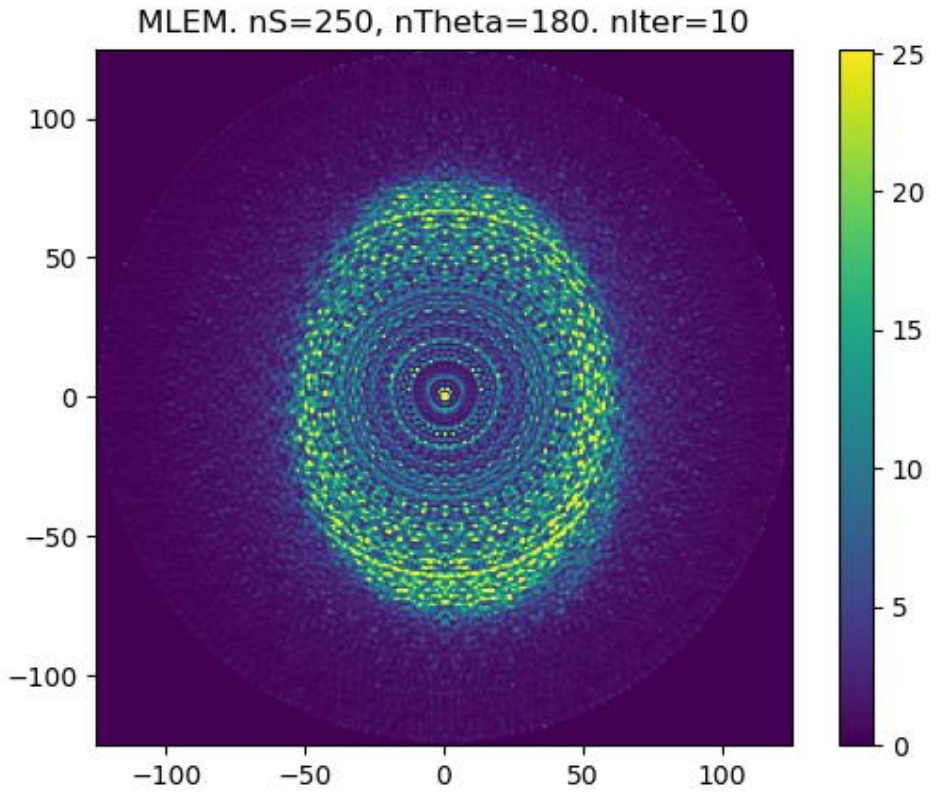

CT volume

$$HU = 1000 * \frac{\mu - \mu_{water}}{\mu_{water}}$$

$\mu \equiv$ linear attenuation co-efficient of the voxel

This CT image ranges has $-1\,073.6201 < HU < 1\,851.5774$

```
ct.voxel_materials = [
    [-2000, -950, "G4_AIR"],
    [-950, -300, "G4_LUNG_ICRP"],
    [-300, -100, "G4_ADIPOSE_TISSUE_ICRP"],
    [-100, 100, "G4_TISSUE_SOFT_ICRP"],
    [100, 300, "G4_MUSCLE_SKELETAL_ICRP"],
    [300, 800, "G4_B-100_BONE"],
    [800, 3000, "G4_BONE_CORTICAL_ICRP"],
]
```

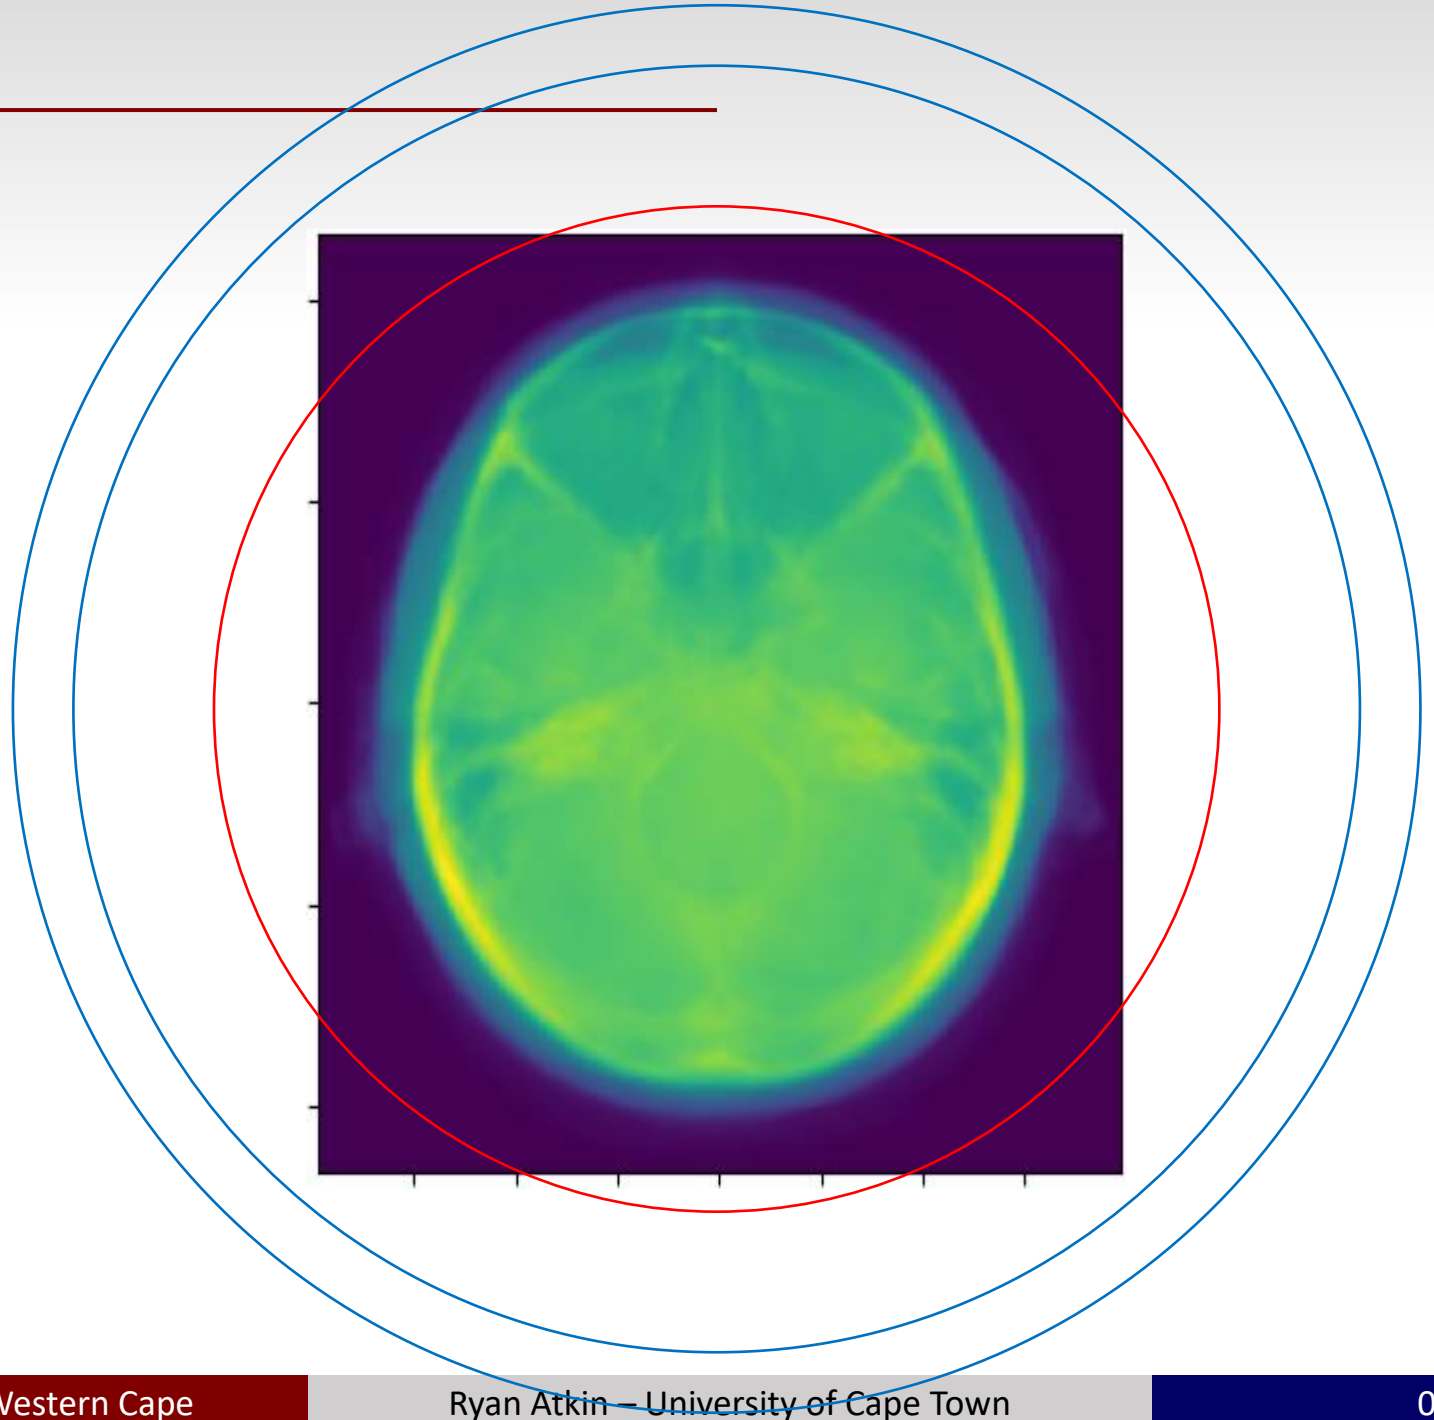

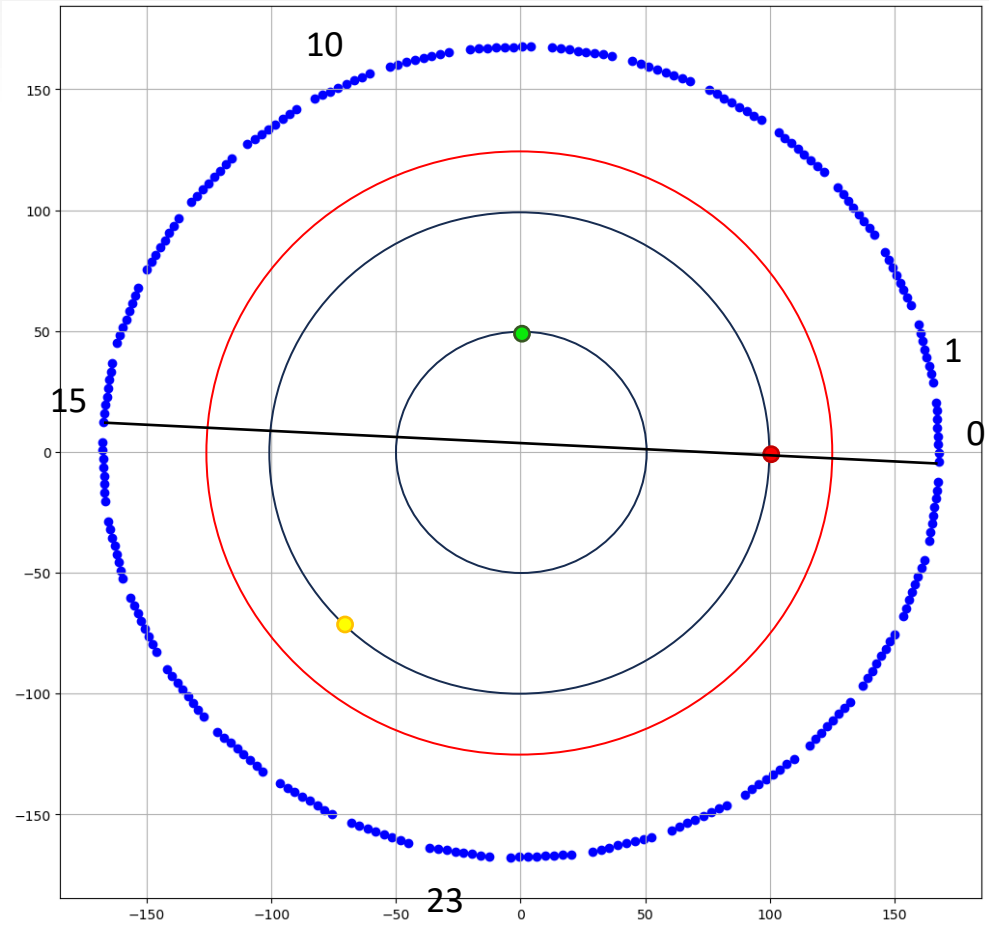
10
1
0
15
23

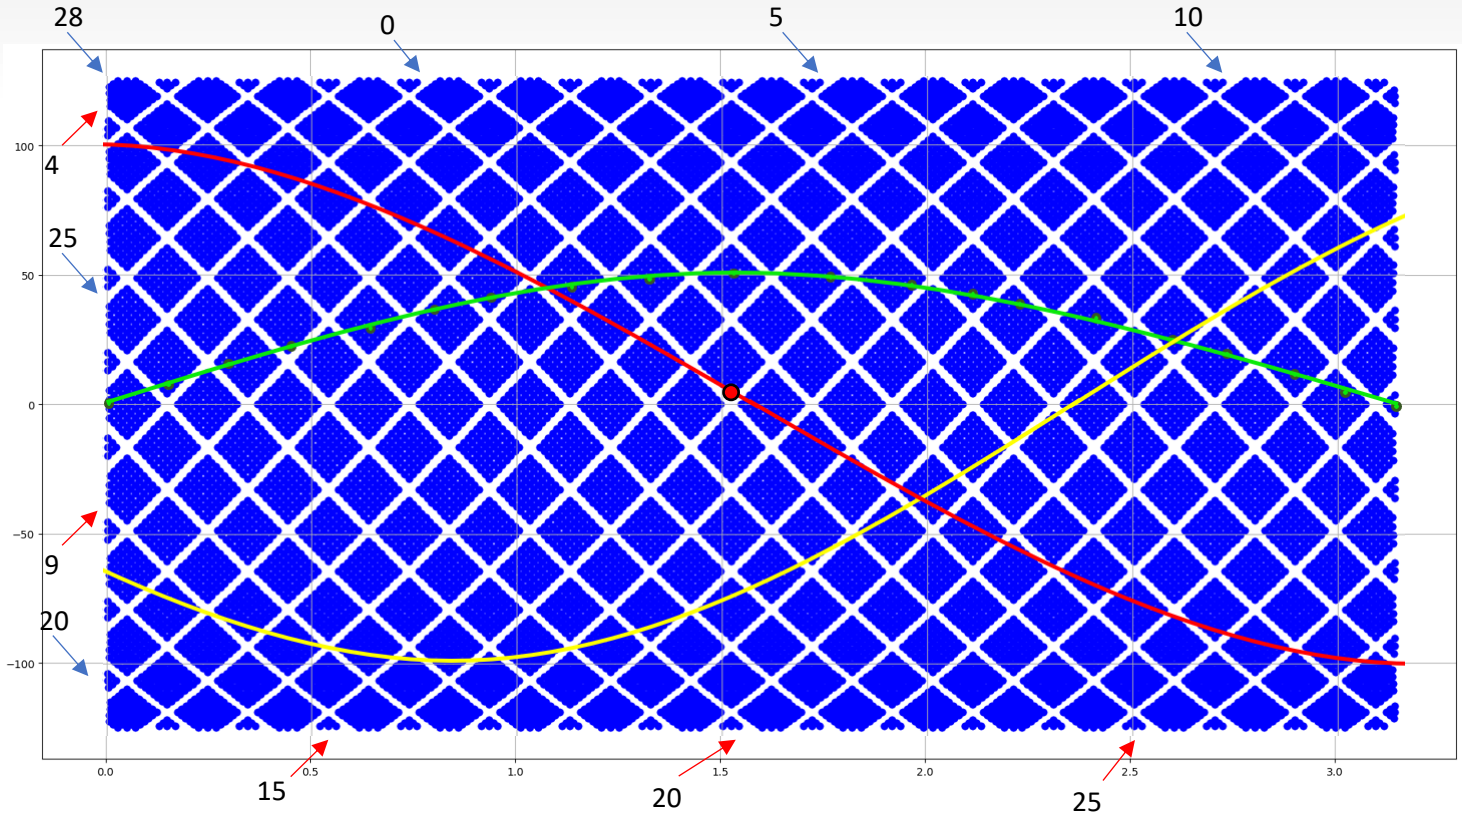
28
0
5
10
4
25
9
20
15
20
25

Crystal Smearing

- Applied before forming the LoR
- Discretising the hit to the centre of the crystal adds uncertainty anyway
- Smearing the readout position thus doesn't add any new uncertainty
- Two smearing options:
    1. Take the max hit, add gaussian smear in Z with FWHM=3 mm, and gaussian smear in X and Y with FWHM=2 mm
        - Trying to replicate something James is working on
        - Native digitizers in OpenGate can't select the first hit in the module, so using maximum energy hit
    2. Take the centre of the pixel, uniformly smear it within extent of the pixel
        - The hit could have been anywhere in the pixel, so smear it as such
        - X and Y are uniform, however, Z could be improved by smearing with a function that favours the "front" of the pixel
- Crystal smearing removes any gaps within the module

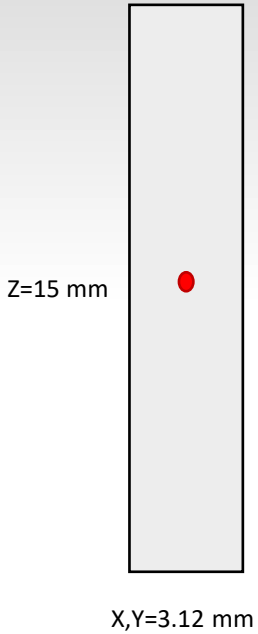

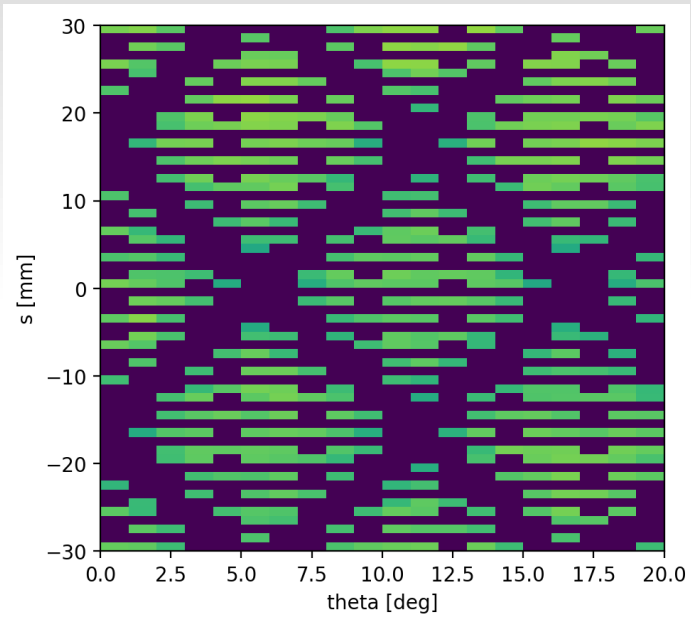

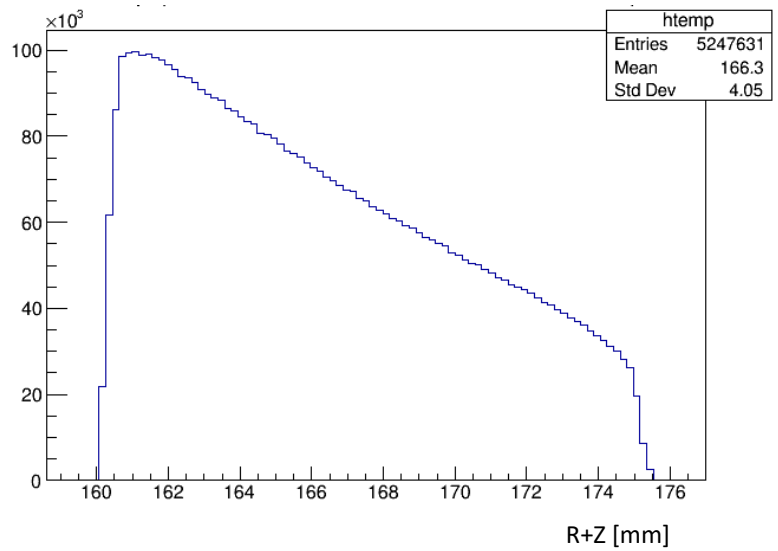

Phillips Scanner (OpenGate)

- 23 040 crystals compared to 4064, and slightly larger crystals
- Just over double the radius
- Greater coverage in Z

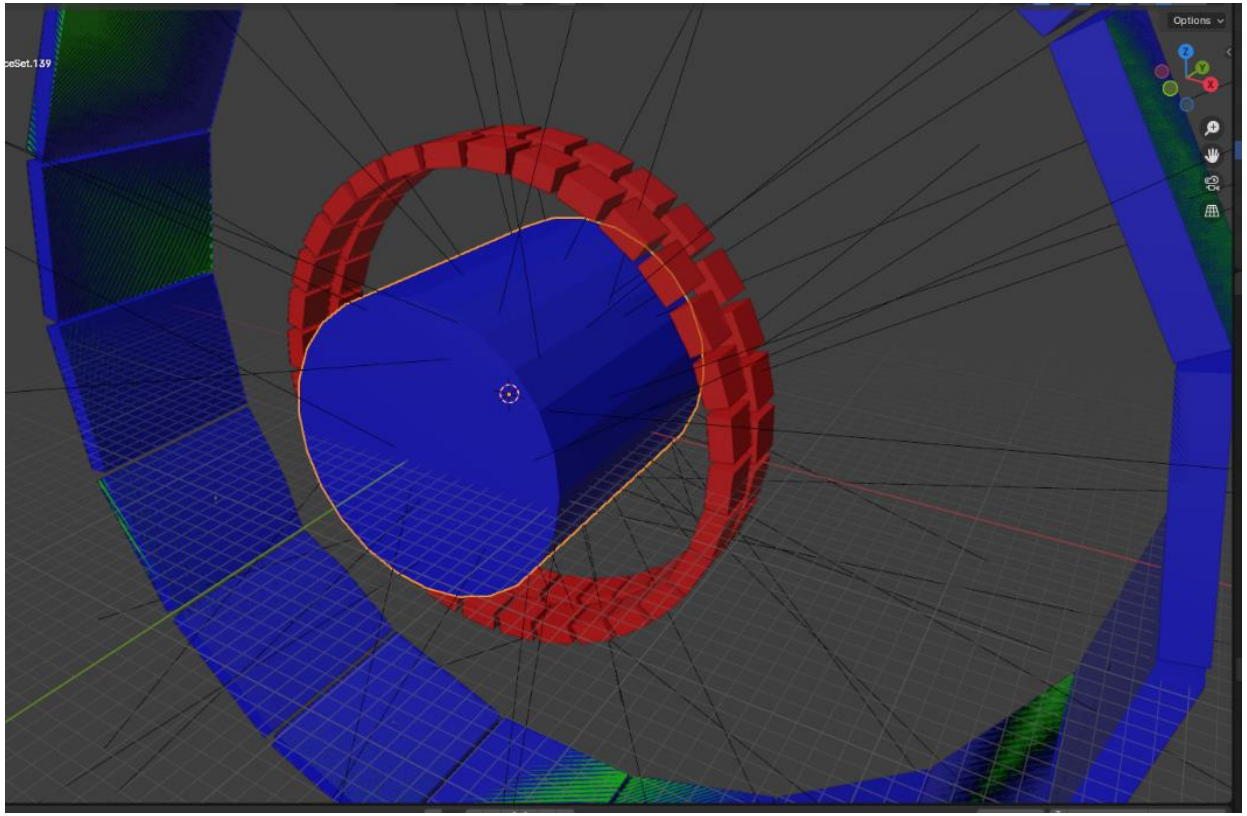

PicoPET_v2

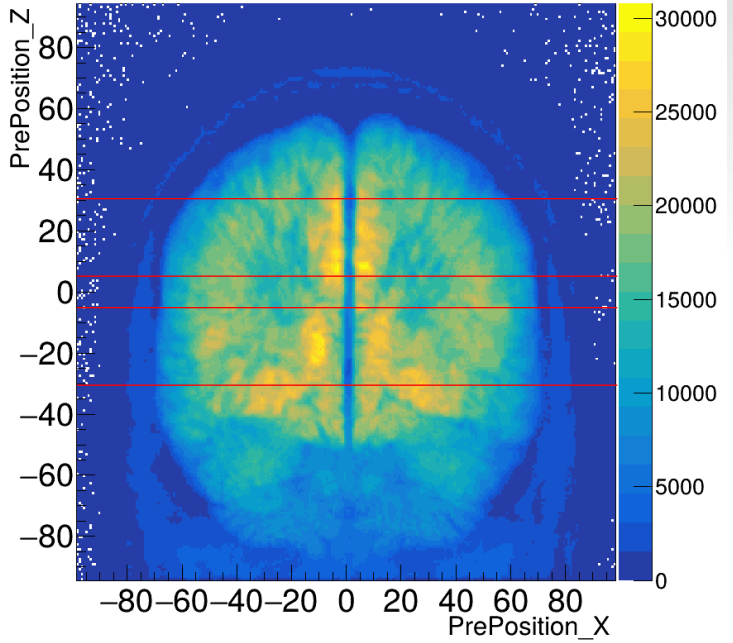

Phillips

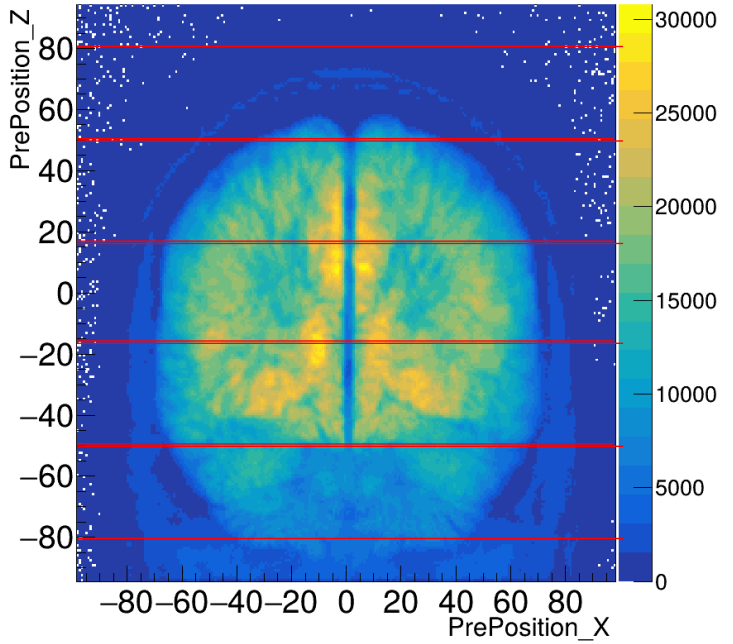

Phillips

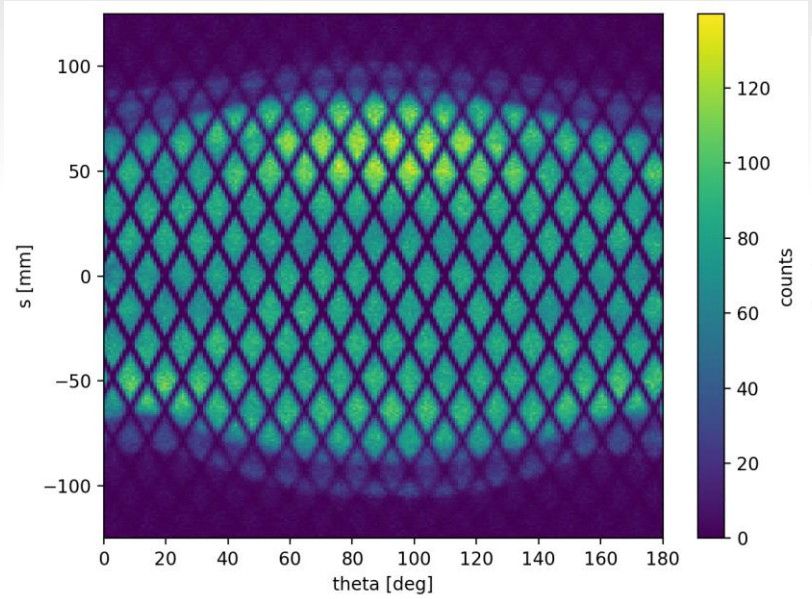

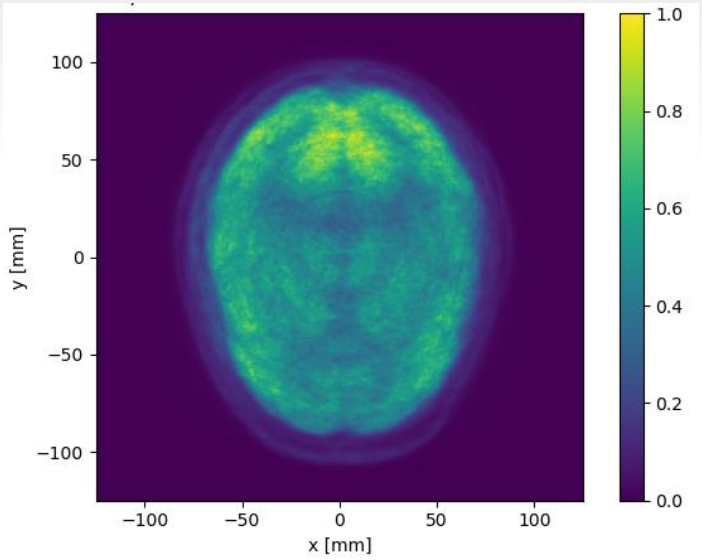

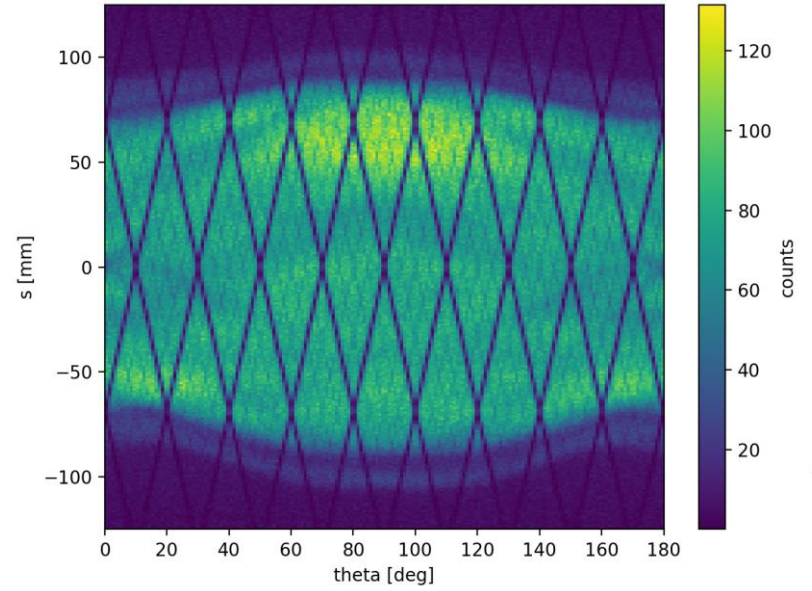

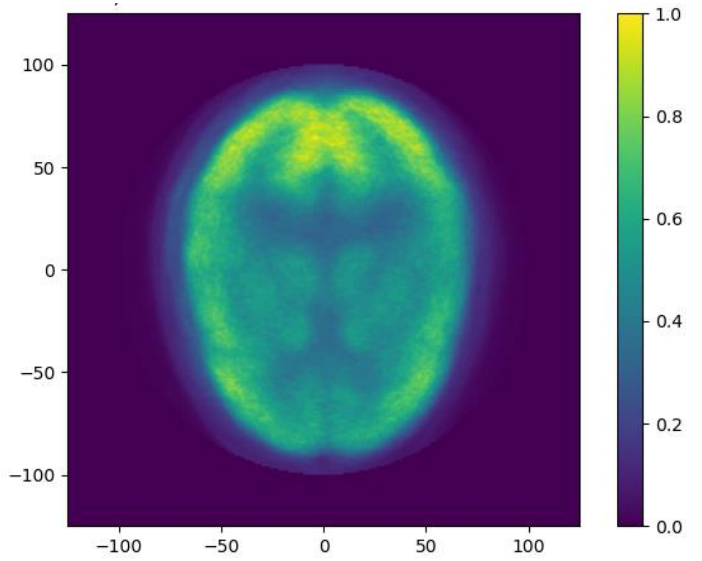

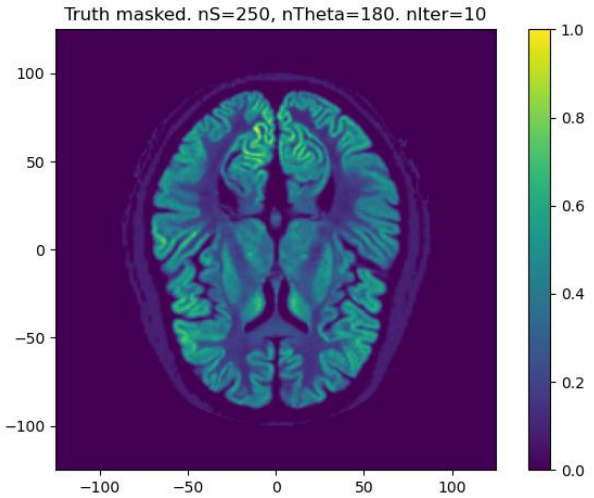

PicoPET_v2 with smear
20 min at 20 MBq
21 913 886 events in total
7 814 875 (35.53%) filled this slice
integral of 1 552 545.23

Both had:

- Same digitization
- Same reconstruction method

Phillips with smear
10 min at 10 MBq
40 279 660 events in total
11 622 757 (28.86%) filled this slice
integral of 2 085 853.25

Z-slicing

- Ideally, we would do 3D Image reconstruction
  - However, this can be computationally expensive and is more complicated than the FBP used in the 2D MLEM
- Instead, we do 2D reconstruction using the full 3D data
- Split the data into axial slices, the same thickness as a pixel
- Assign every LoR to every slice with a corresponding weight
  - $w = \frac{dz}{B-A}$
  - $dz \equiv$ axial distance LoR covers that slice within the FoV
  - A, B are the z points where the LoR crosses the FoV
- Perform 2D reconstruction per slice
  - Normalisation sinograms built the same way
  - Attenuation sinogram and truth image defined to cover corresponding z range
- Stack slices to get 3D image
  - Can apply axial smoothing after reconstruction
  - Or axial regularization during MLEM iterations to improve combining in z

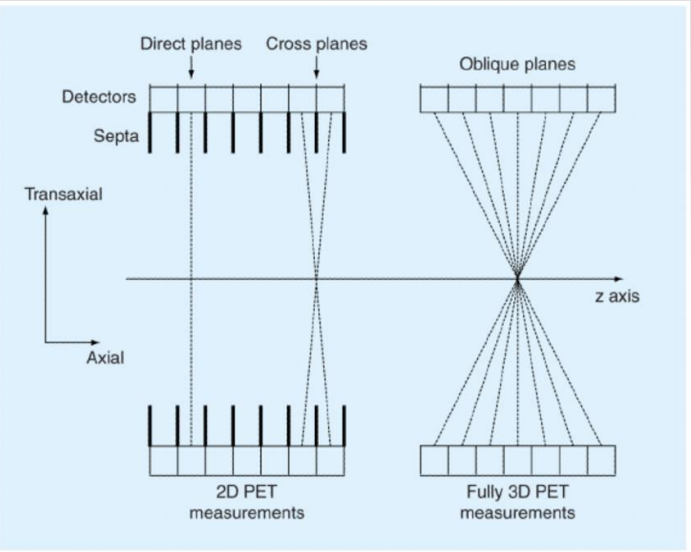

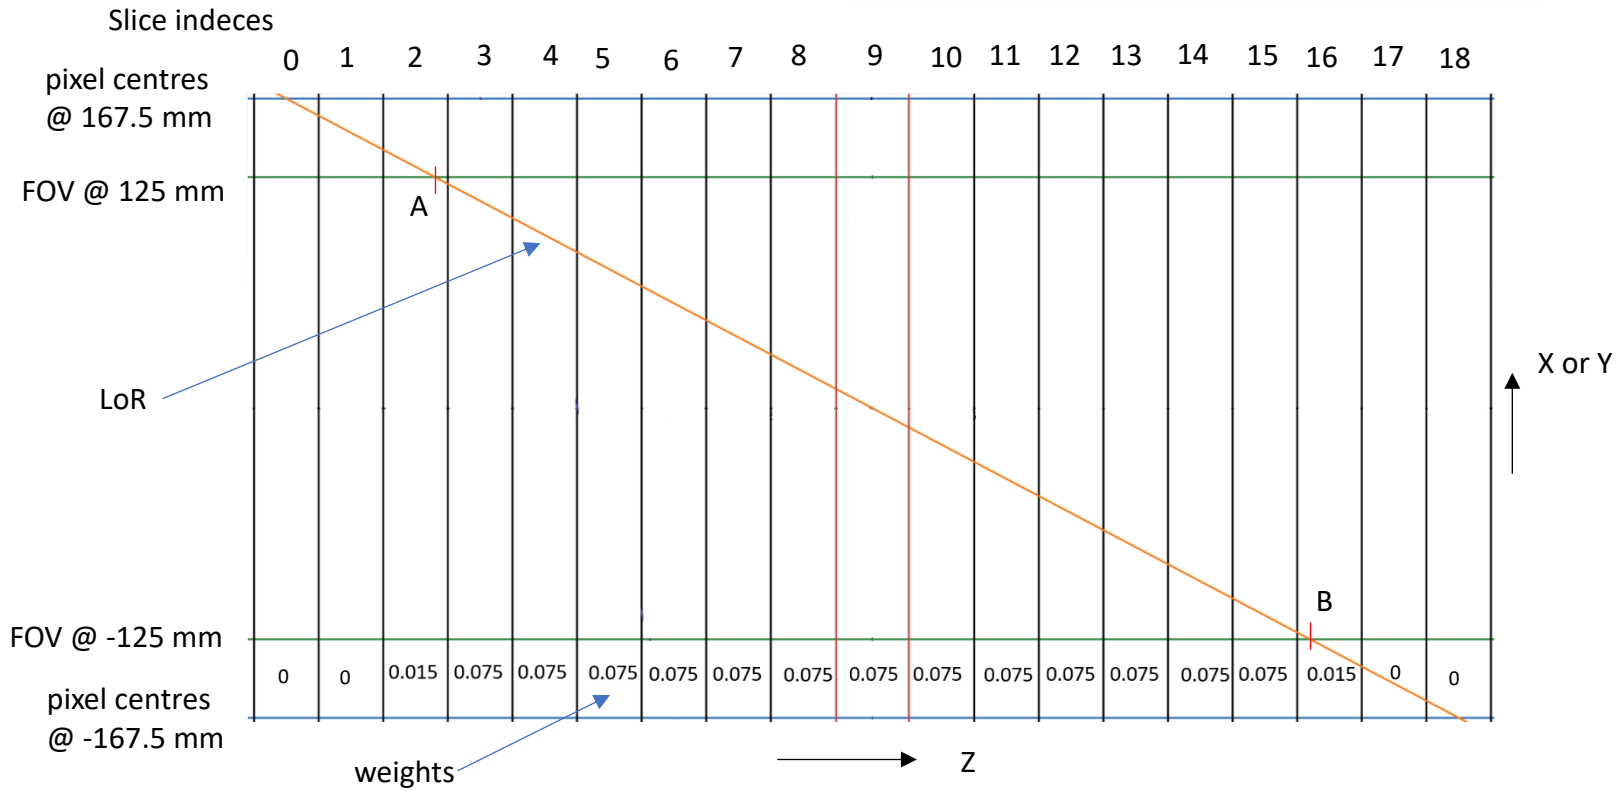

- Looking at a theta slice of $0^{\circ} < \theta < 3^{\circ}$
- No z-slicing yet
- $P - R = T + S$

Prompts

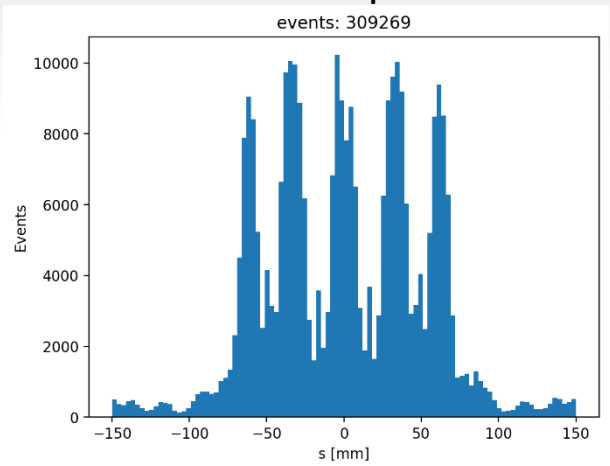

Randoms

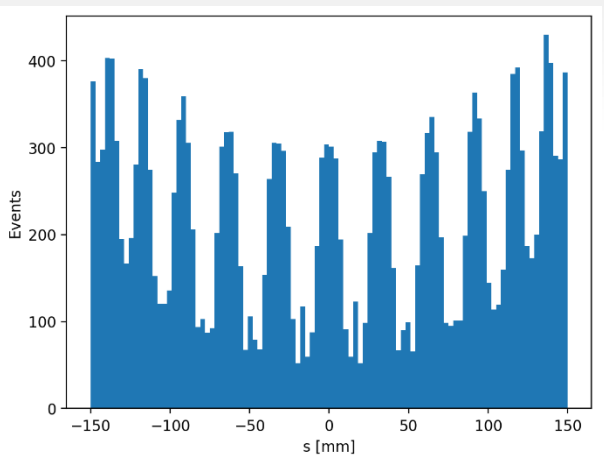

Normalisation

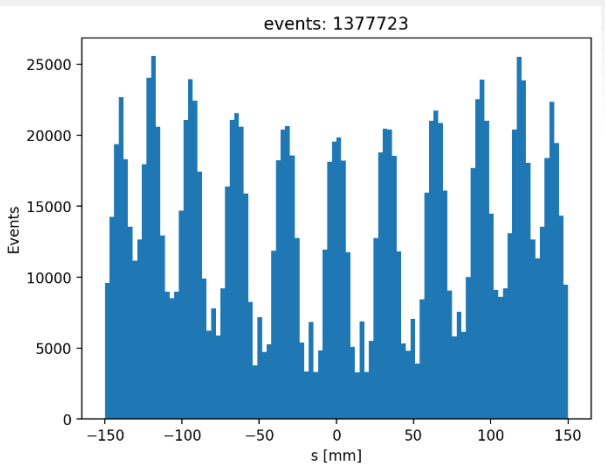

True scatters

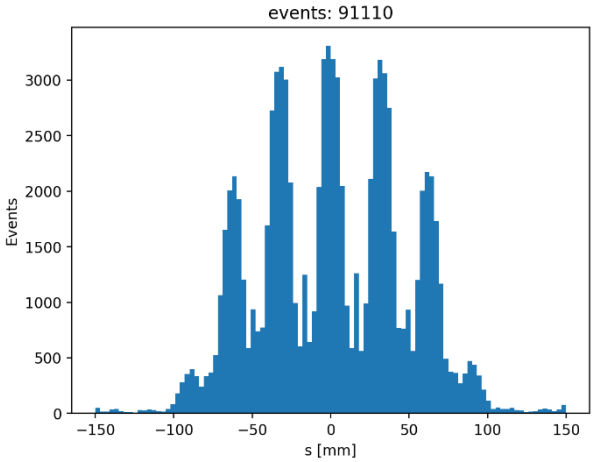

True scatters divided by normalisation

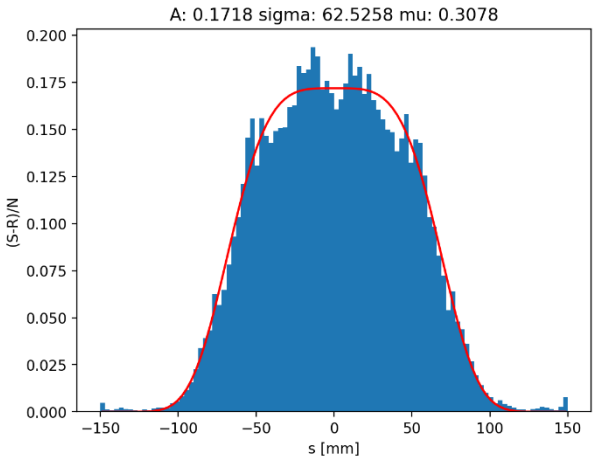

$\frac{P - R}{N}$

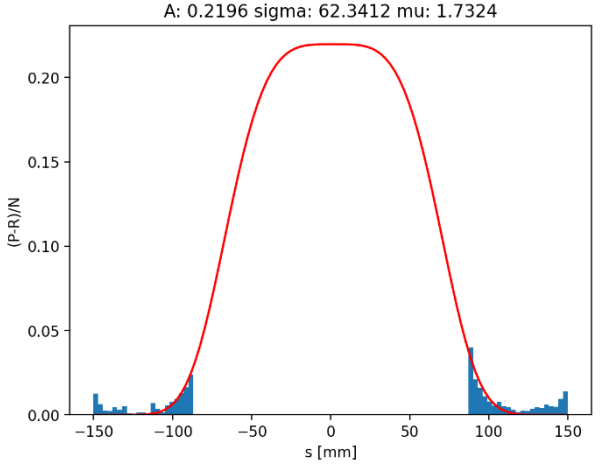

$N * \frac{P - R}{N}$