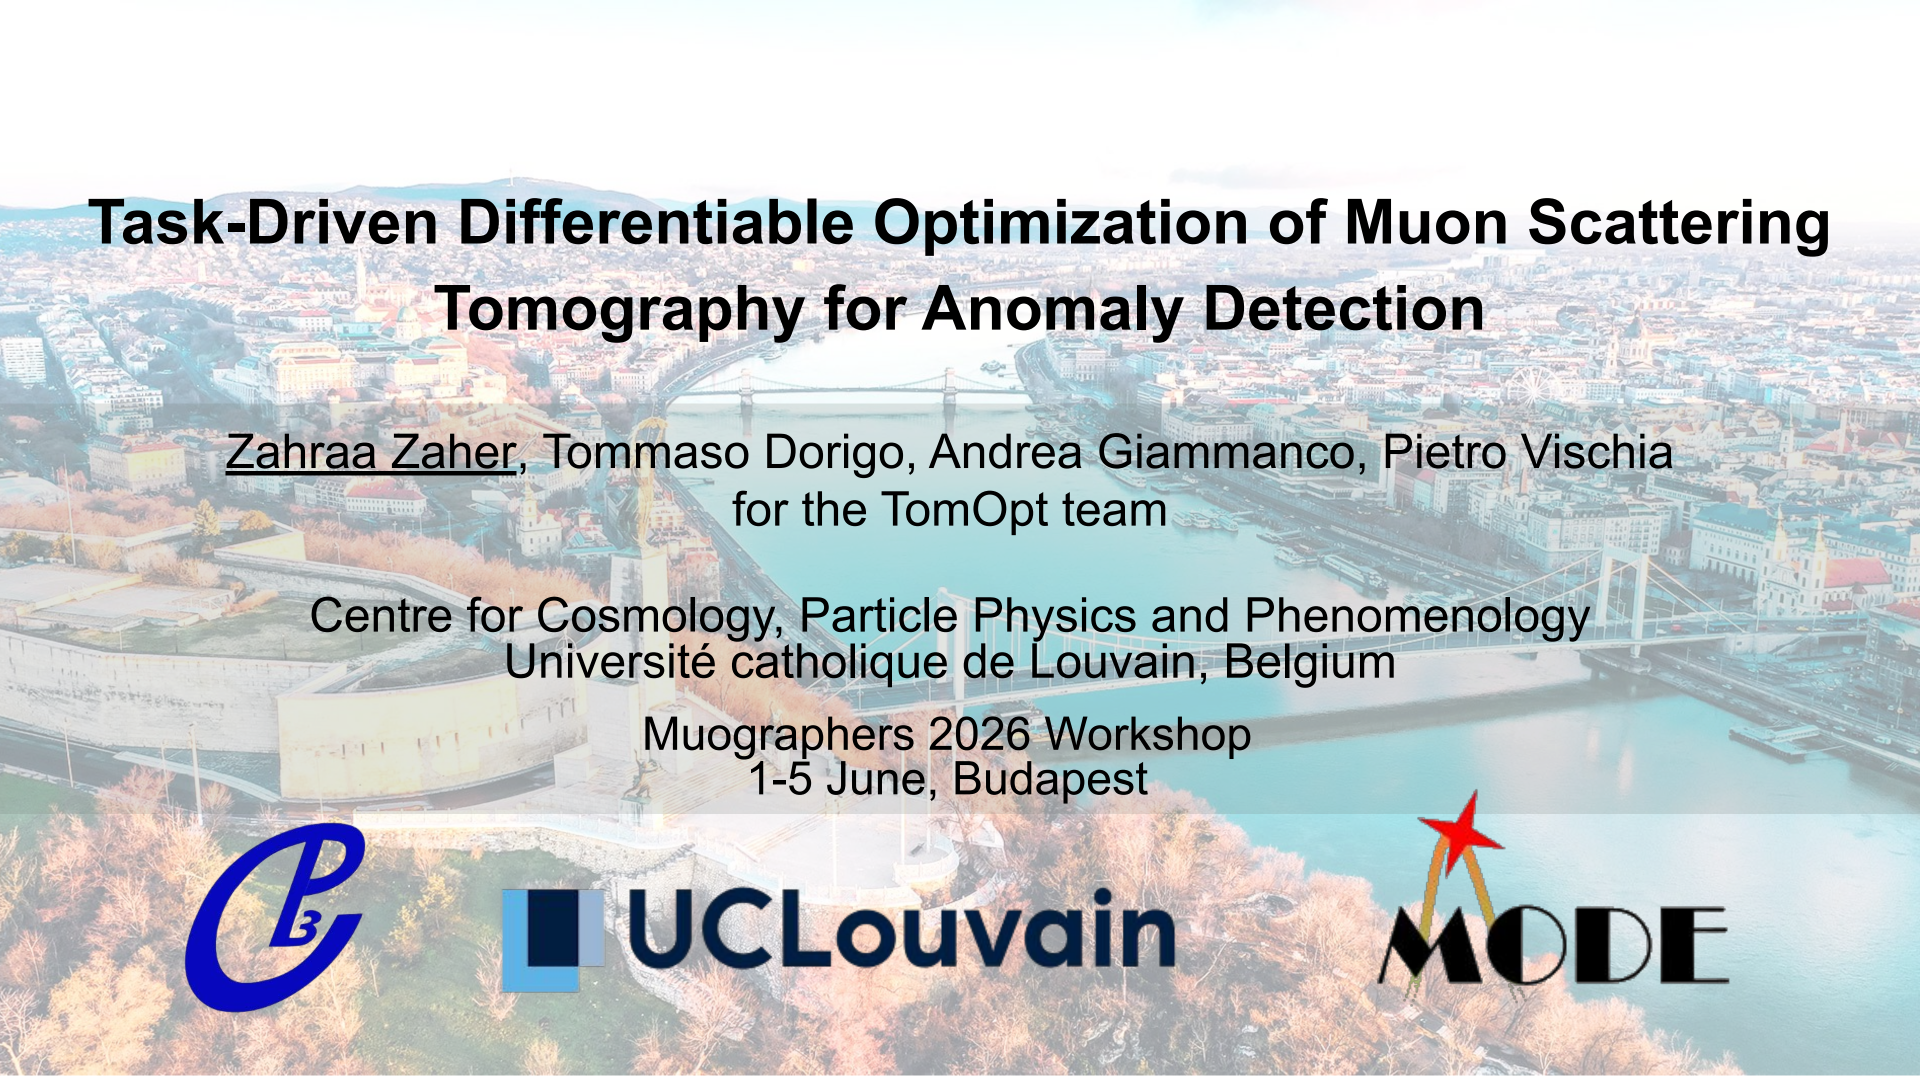

Task-Driven Differentiable Optimization of Muon Scattering Tomography for Anomaly Detection

Zahraa Zaher, Tommaso Dorigo, Andrea Giammanco, Pietro Vischia
for the TomOpt team

Centre for Cosmology, Particle Physics and Phenomenology
Université catholique de Louvain, Belgium

Muographers 2026 Workshop
1-5 June, Budapest

Problem statement |

zahraa.zaher@uclouvain.be

Problem statement |

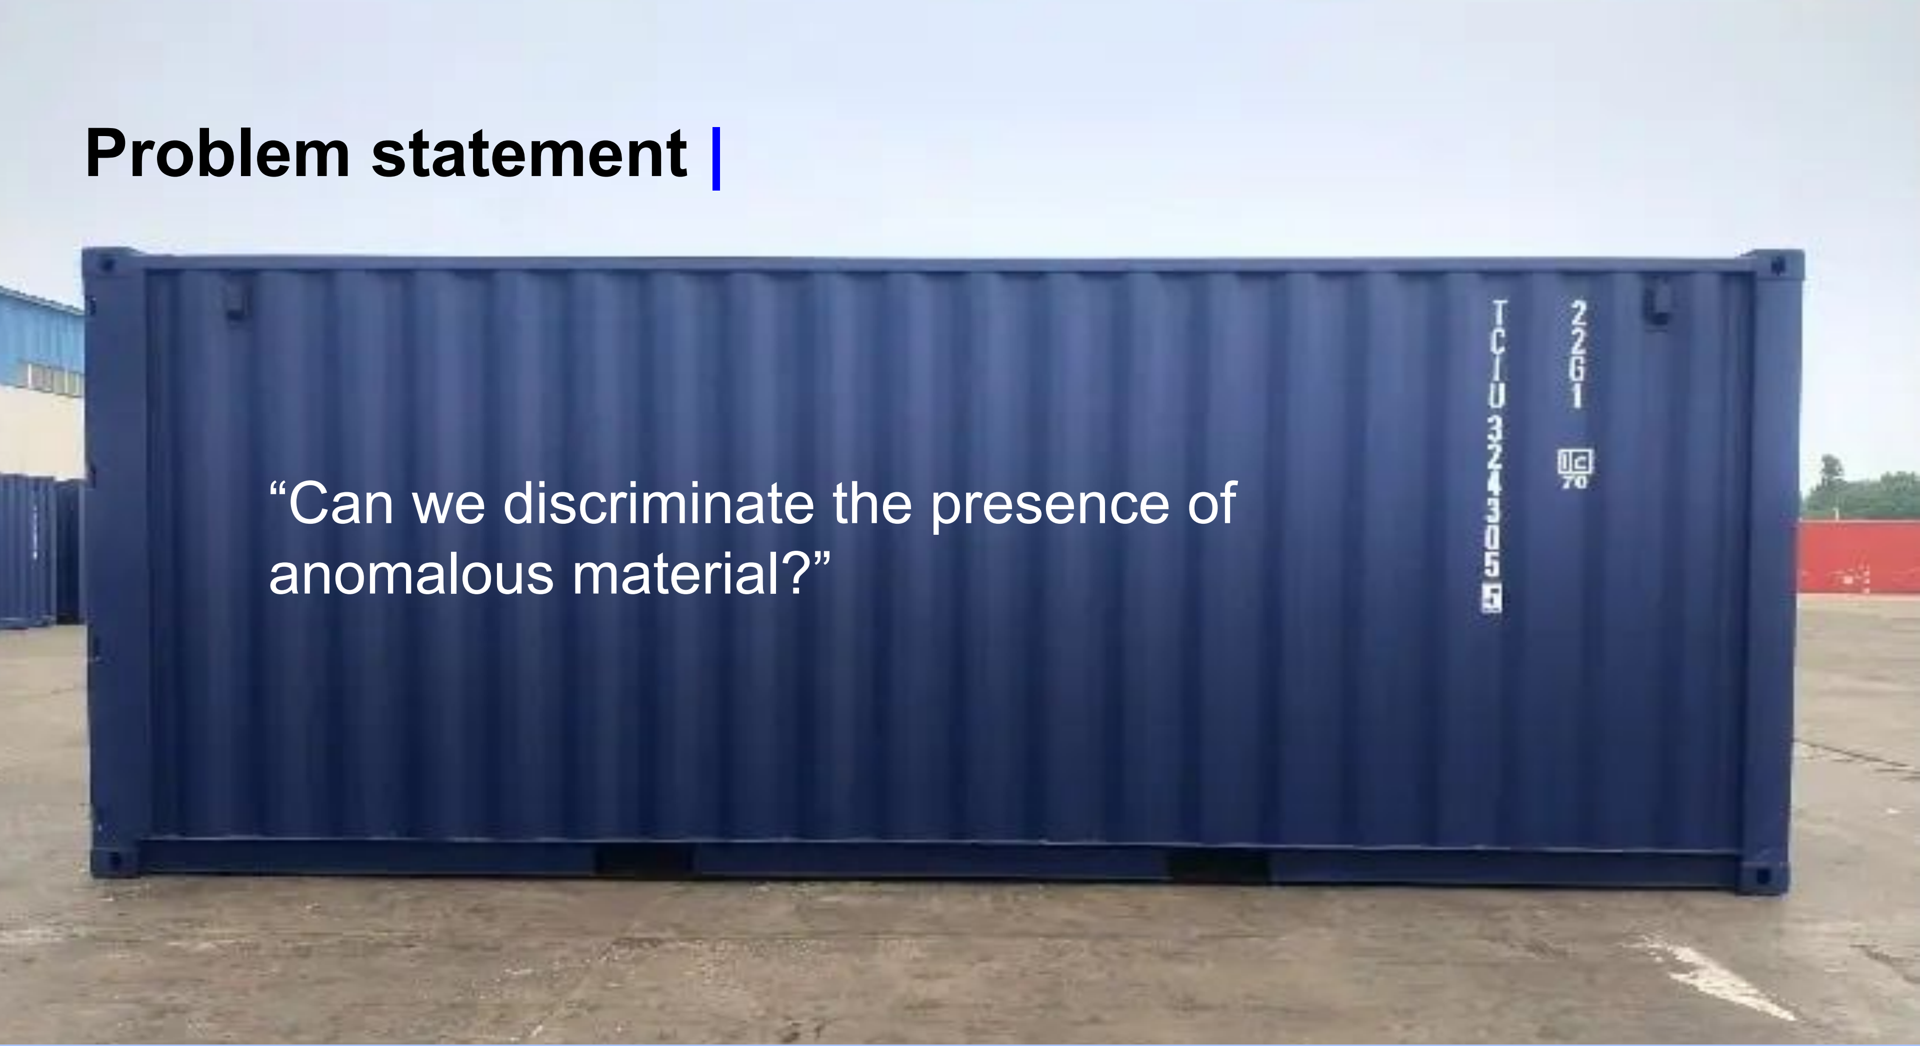

“Can we discriminate the presence of anomalous material?”

zahraa.zaher@uclouvain.be

State-of-the-Art

Solution 1

zahraa.zaher@uclouvain.be

State-of-the-Art

Solution 1

X-ray radiography

Solution 2

State-of-the-Art

Solution 1

X-ray radiography

Solution 2

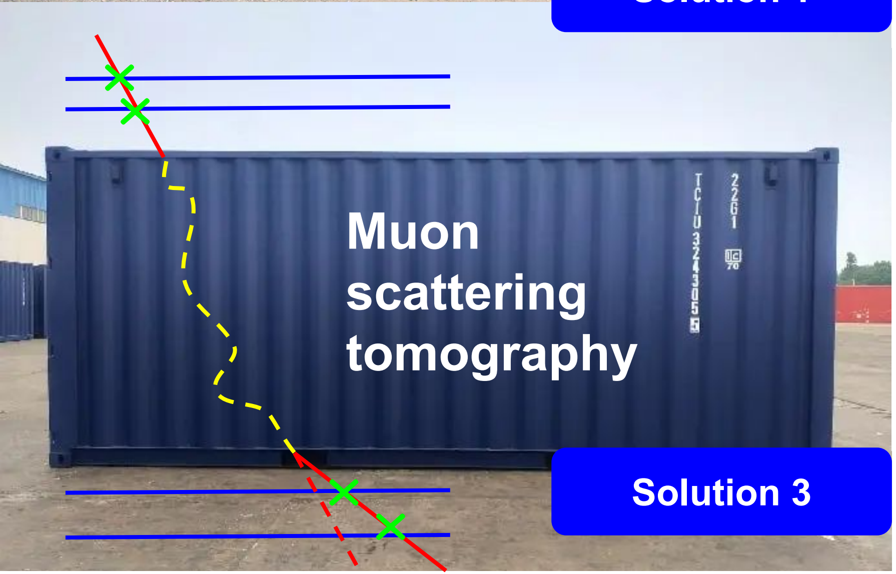

Muon scattering tomography

Solution 3

Operational constraints:

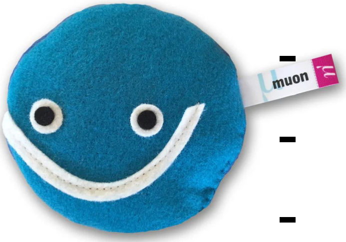

- limited exposure time
- sparse muon statistics
- imperfect detector resolution

State-of-the-Art

Solution 1

X-ray radiography

Solution 2

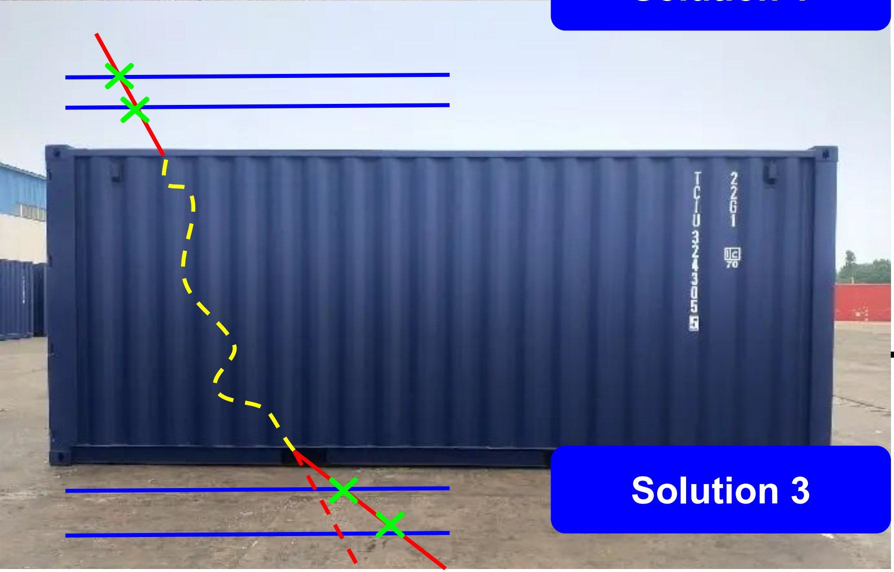

Solution 3

- Image reconstruction
- Anomaly detection

State-of-the-Art |

Image reconstruction

Radiation length inference algorithms

- E.g. Point-of-Closest-Approach (PoCA), statistical, clustering, ..

$$\theta_{RMS} = \frac{13.6\,MeV}{\beta c p}\sqrt{\frac{x}{\boxed{X_0}}}\left[1 + 0.038\,\ln\left(\frac{x}{X_0\,\beta^2}\right)\right]$$

zahraa.zaher@uclouvain.be

State-of-the-Art

Image reconstruction

Radiation length inference algorithms

- E.g. Point-of-Closest-Approach (PoCA), statistical, clustering, ..

$$\theta_{RMS} = \frac{13.6\,MeV}{\beta c p}\sqrt{\frac{x}{\boxed{X_0}}}\left[1 + 0.038\,\ln\left(\frac{x}{X_0\,\beta^2}\right)\right]$$

- Image quality/inference precision relies on **exposure**

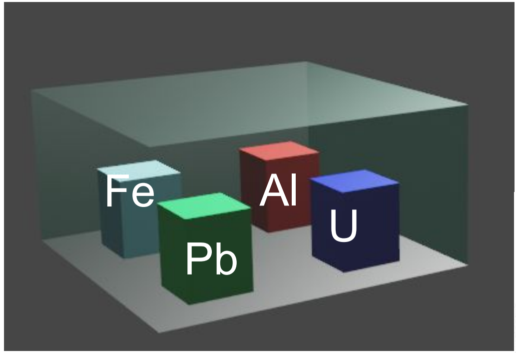

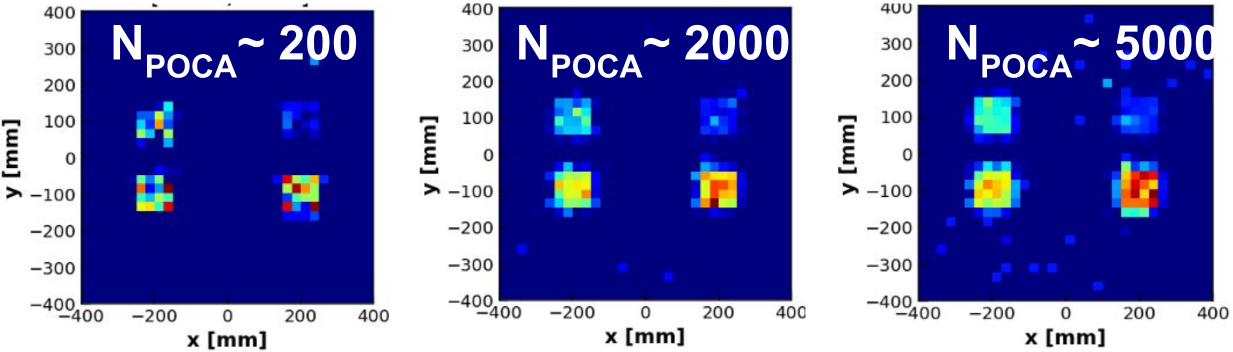

State-of-the-Art |

Image reconstruction

- **Radiation length inference algorithms**
- **Deep learning-based image reconstruction**

E.g. U-Net

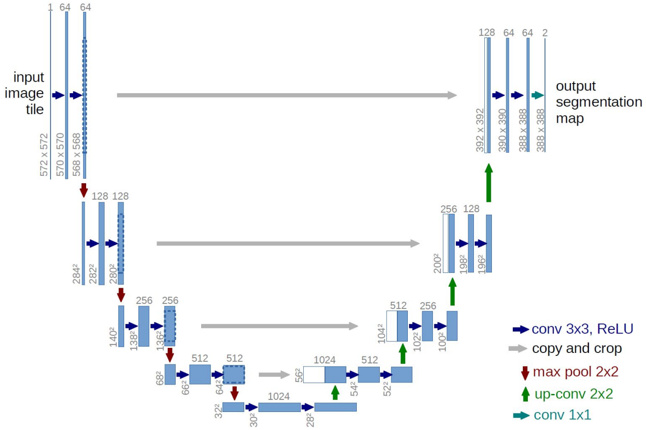

Requires:

- Large datasets
- Dense input grids
- Sufficient muon statistics

zahraa.zaher@uclouvain.be

State-of-the-Art

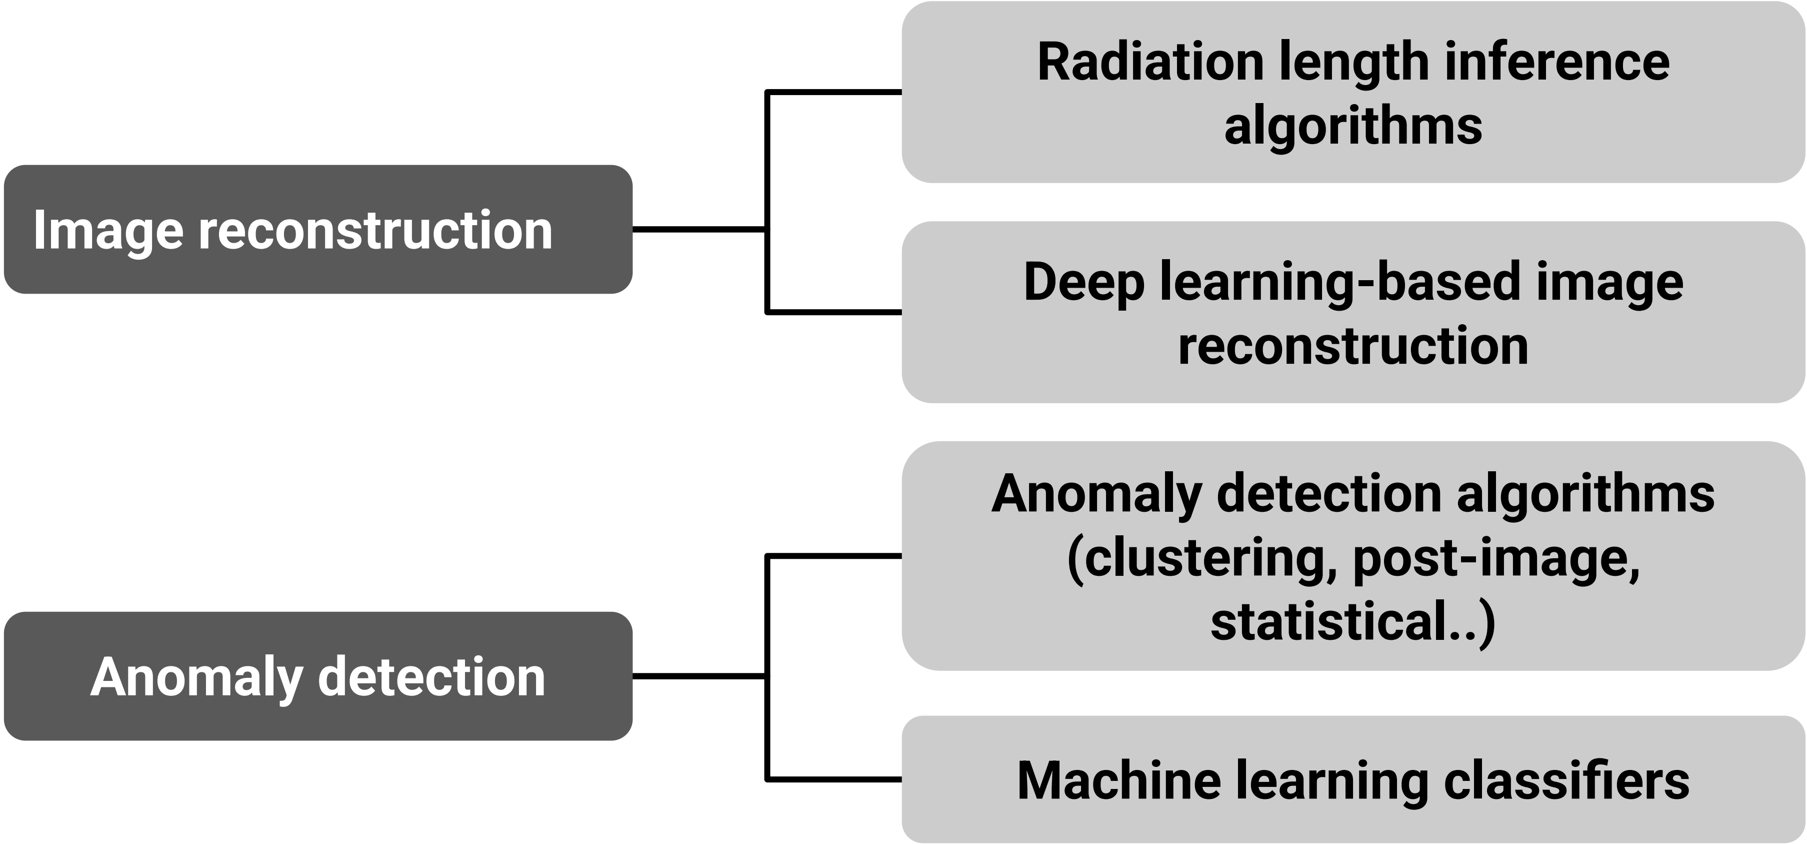

Image reconstruction

- Radiation length inference algorithms
- Deep learning-based image reconstruction

Anomaly detection

- Anomaly detection algorithms (clustering, post-image, statistical..)
- Machine learning classifiers

zahraa.zaher@uclouvain.be

Content |

Image reconstruction

- **Radiation length inference algorithms**
- **Deep learning-based image reconstruction: POCA-NET**

From **sparse PoCA** information, can we:

1. Infer **material X_0** ?
2. Detect anomalies **and optimize** the detector setup to maximize detection?

Anomaly detection

- **Anomaly detection algorithms (clustering, post-image, statistical..)**
- **Machine learning ~~classifiers~~ detector optimization with TomOpt**

Part I |

- **Image reconstruction**
  - **Radiation length inference algorithms**
  - **Deep learning-based image reconstruction: POCA-NET**
- **Anomaly detection**
  - **Anomaly detection algorithms (clustering, post-image, statistical..)**
  - **Machine learning ~~classifiers~~ detector optimization with TomOpt**

zahraa.zaher@uclouvain.be

POCA-NET |

X_0 Inference from sparse PoCA Point Cloud

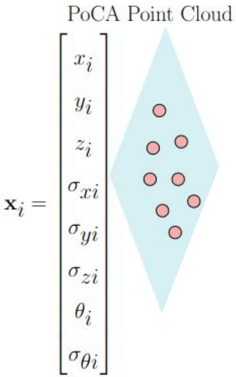

$$\mathbf{x}_i = \begin{bmatrix} x_i \\ y_i \\ z_i \\ \sigma_{xi} \\ \sigma_{yi} \\ \sigma_{zi} \\ \theta_i \\ \sigma_{\theta i} \end{bmatrix}$$

Sparse, unordered point
cloud of PoCA muon events

zahraa.zaher@uclouvain.be

POCA-NET |

X_0 Inference from sparse PoCA Point Cloud

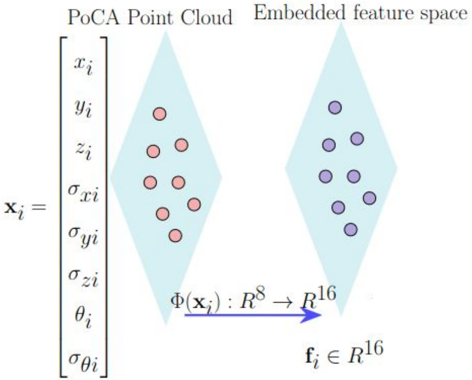

Shared MLP maps each PoCA point to an embedded feature space representation

POCA-NET |

X_0 Inference from sparse PoCA Point Cloud

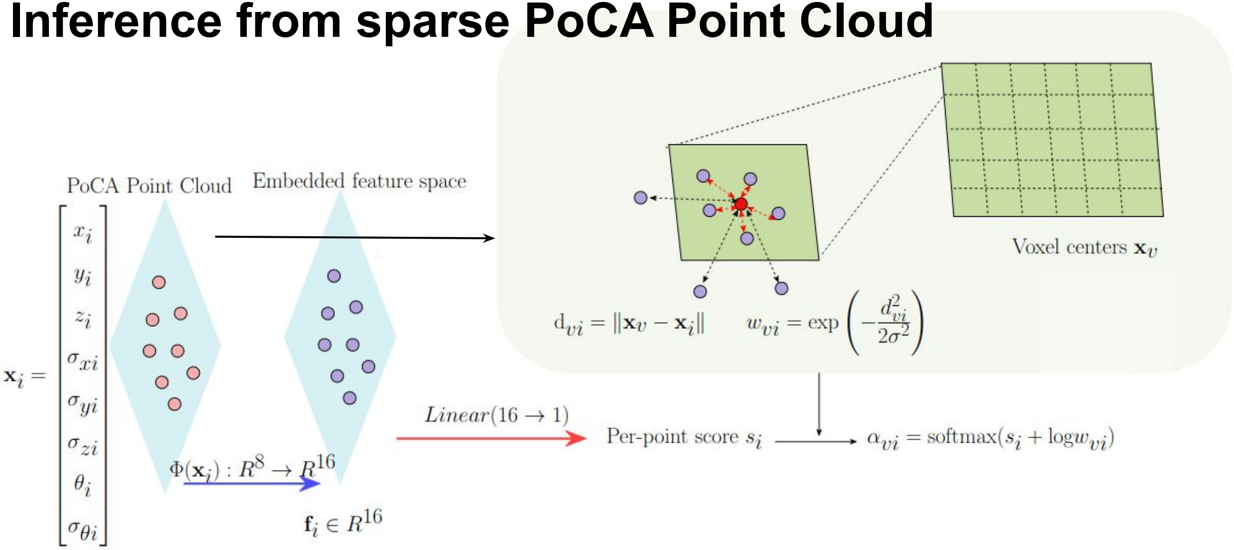

Point-voxel distance weight
Per-point learned attention score

→ Point-voxel weight

POCA-NET |

X_0 Inference from sparse PoCA Point Cloud

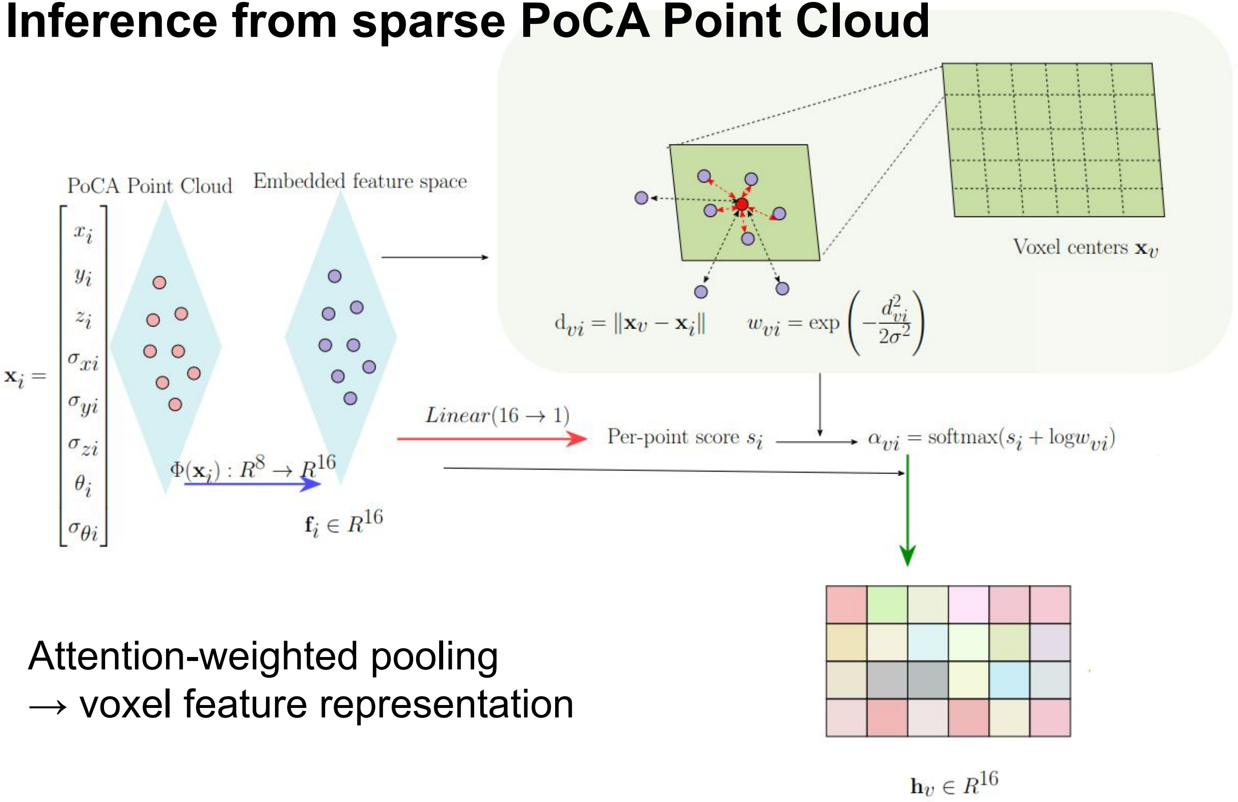

Attention-weighted pooling
→ voxel feature representation

POCA-NET |

X_0 Inference from sparse PoCA Point Cloud

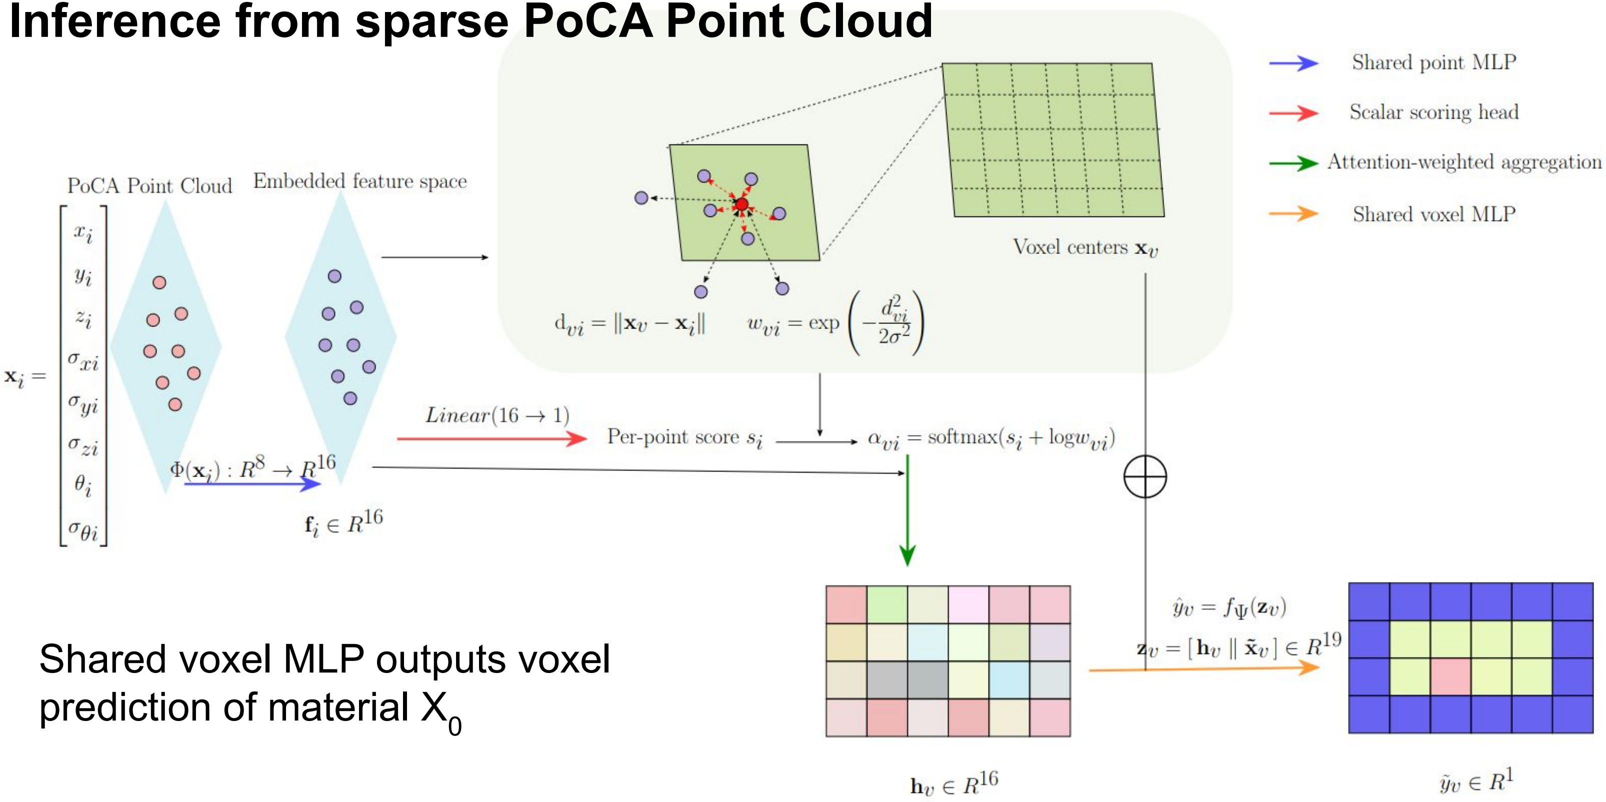

Shared voxel MLP outputs voxel prediction of material X_0

POCA-NET |

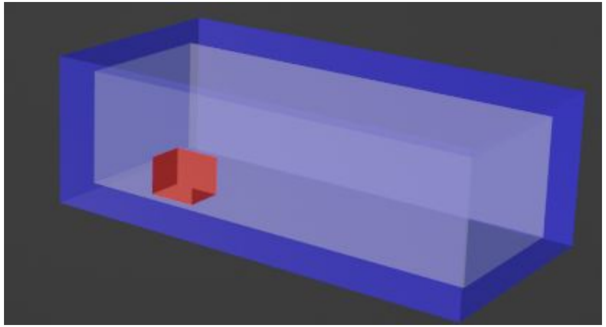

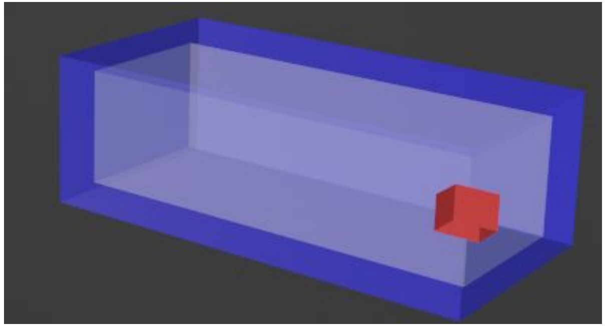

- Dataset created with **TomOpt**
- **Simulated** cargo “container” with steel borders 1m x 1m x 0.4m
- **5 target materials:**
  - Beryllium ($X_0 \approx 35.28$ cm)
  - Aluminum ($X_0 \approx 8.90$ cm)
  - Steel ($X_0 \approx 1.78$ cm)
  - Iron ($X_0 \approx 1.76$ cm)
  - Uranium ($X_0 \approx 0.31$ cm)
- **49 target positions**
- Splits done for **unique material/position**
- train/val/test : **4100/400/400**

zahraa.zaher@uclouvain.be

POCA-NET vs 3D-UNet

- Benchmark Model: A standard 3D-UNet
    - Advantage: better at capturing **local structures** and spatial relationships

- Data Selection: Top 10,000 muons selected based on scattering angles (~4.5 minutes of exposure)

- Optimization: Both models trained to minimize Huber loss:

$$\mathcal{L} = \frac{1}{V}\sum_{v=1}^{V} L_\delta\left(\hat{y}_v - \log X_0^{(v)}\right), \quad L_\delta(r) = \begin{cases} \frac{1}{2}r^2 & |r| \le \delta \\ \delta(|r| - \frac{1}{2}\delta) & |r| > \delta \end{cases}$$

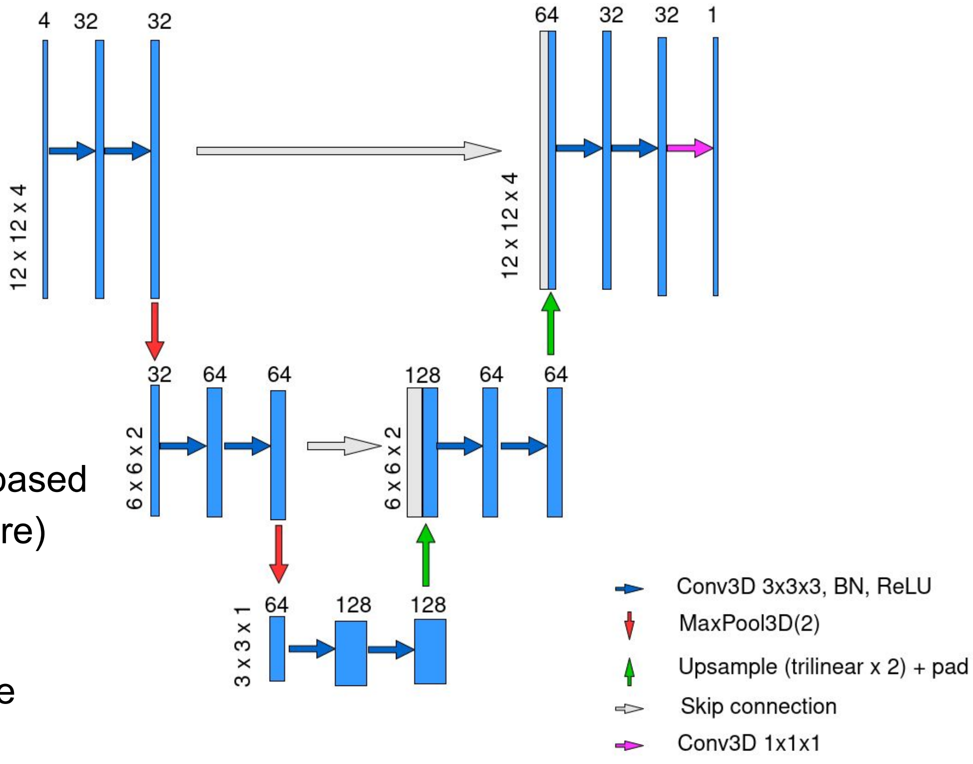

Structural-Similarity-Index-Measure (SSIM) |

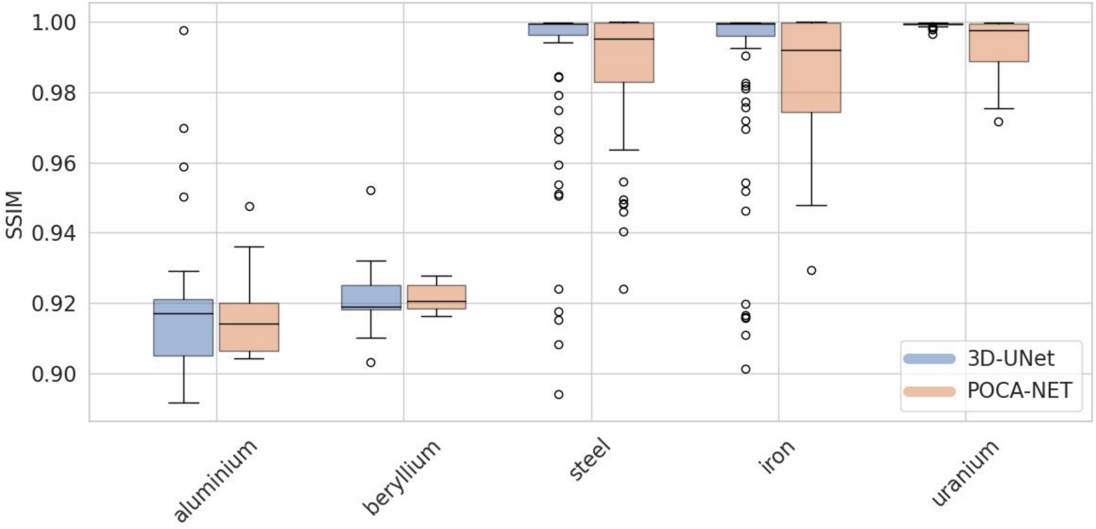

Atomic number Z

Higher is better

Normalized Mean-Squared-Error (NMSE) |

NMSE constrained in target region

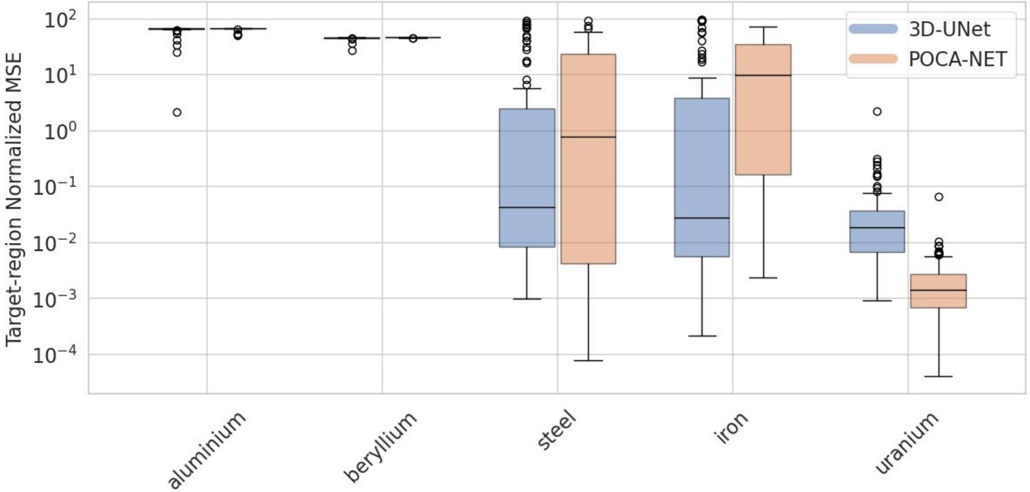

Atomic number Z

Normalized Mean-Squared-Error (NMSE) |

NMSE constrained in target region

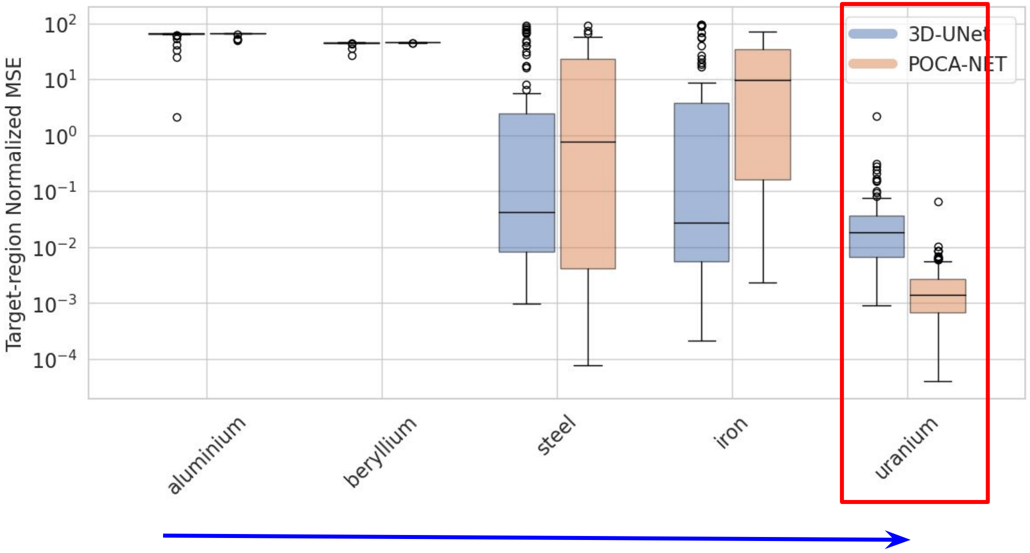

Atomic number Z

POCA-NET

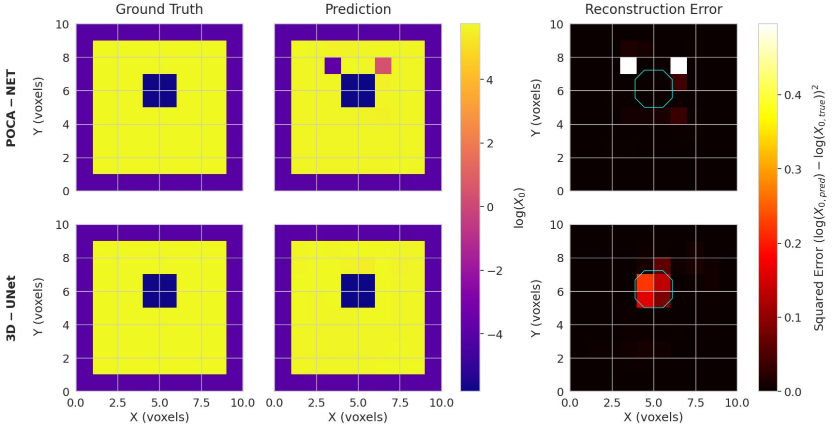

3D-UNET

Normalized Mean-Squared-Error (NMSE) |

NMSE constrained in target region

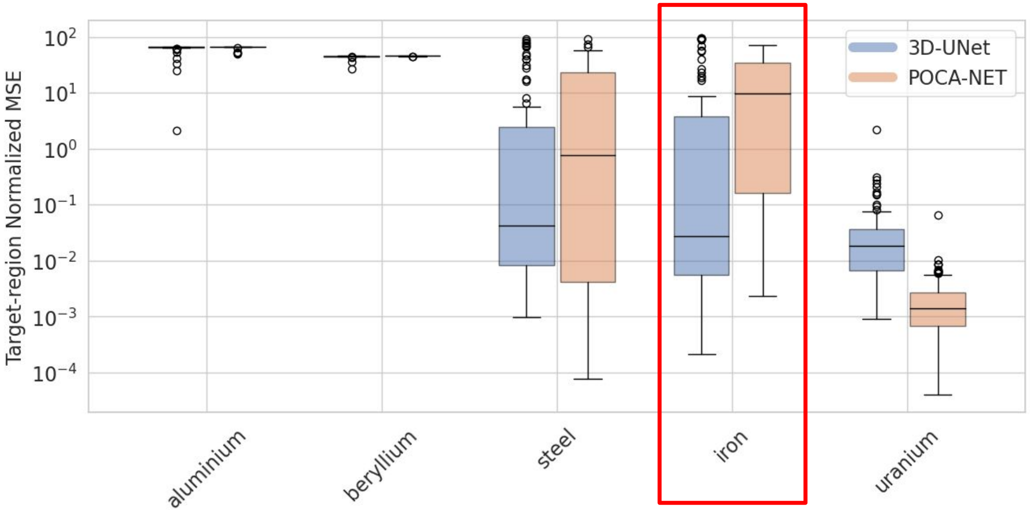

Atomic number Z

POCA-NET

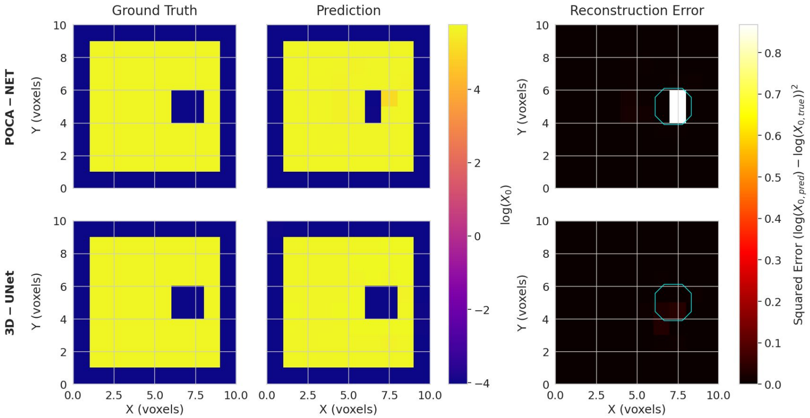

3D-UNET

zahraa.zaher@uclouvain.be

Part II |

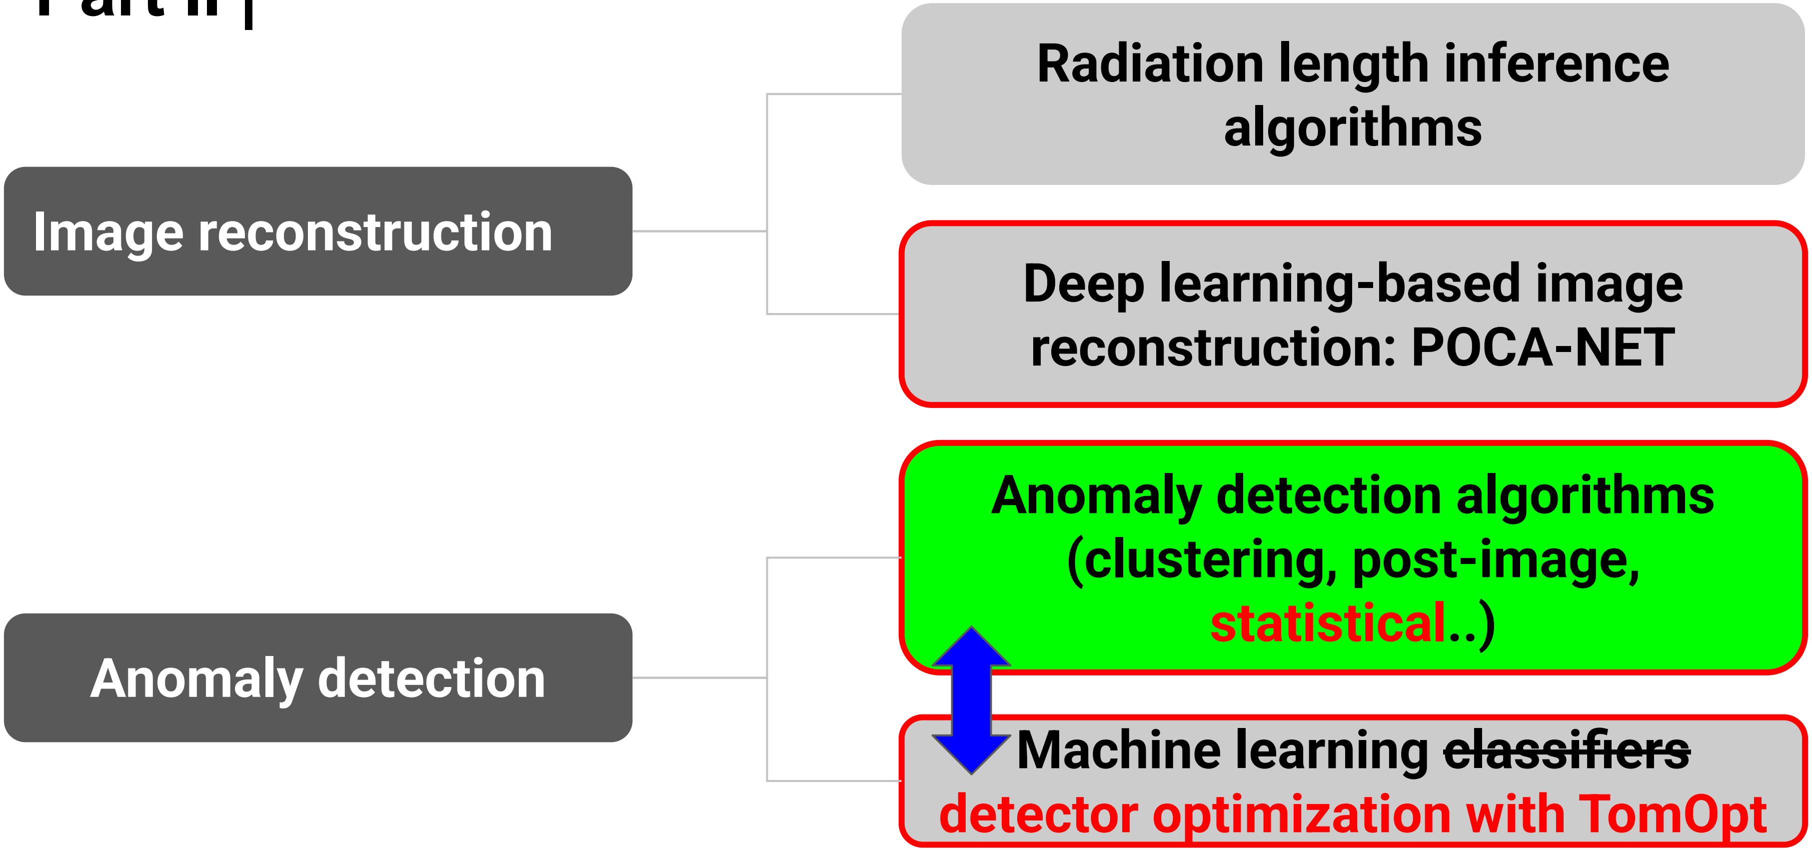

Image reconstruction

- Radiation length inference algorithms
- Deep learning-based image reconstruction: POCA-NET

Anomaly detection

- Anomaly detection algorithms (clustering, post-image, statistical..)
- Machine learning ~~classifiers~~ detector optimization with TomOpt

zahraa.zaher@uclouvain.be

Differentiable Hypothesis test for high-Z anomaly detection |

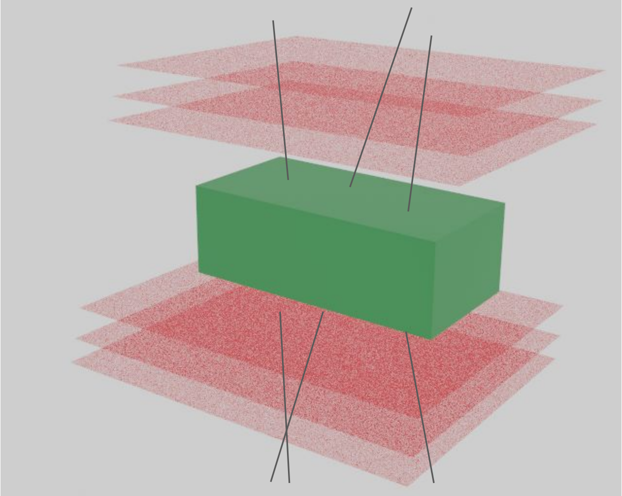

Assumption: Cargo container contains a <u>single material </u>(or materials with <u>similar atomic number Z</u>)

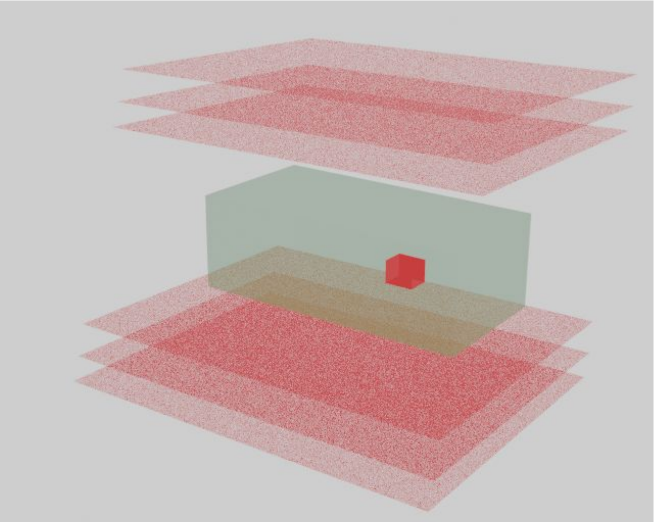

U-target present?

zahraa.zaher@uclouvain.be

Differentiable Hypothesis test for high-Z anomaly detection |

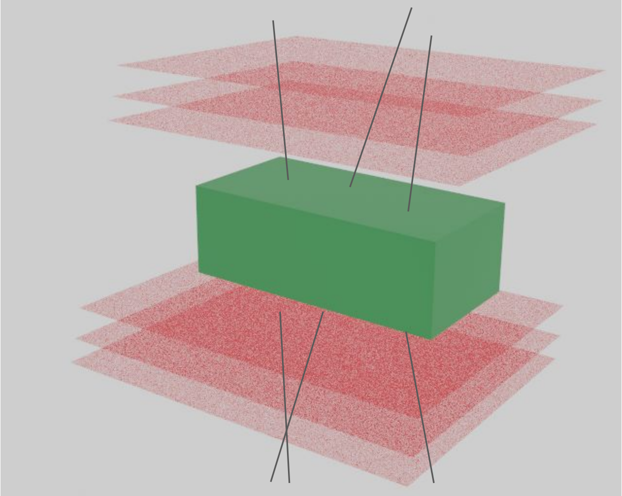

Assumption: Cargo container contains a single material (or materials with similar atomic number Z)

U-target present?

Why differentiable ?

To be compatible with gradient-based parameter optimization

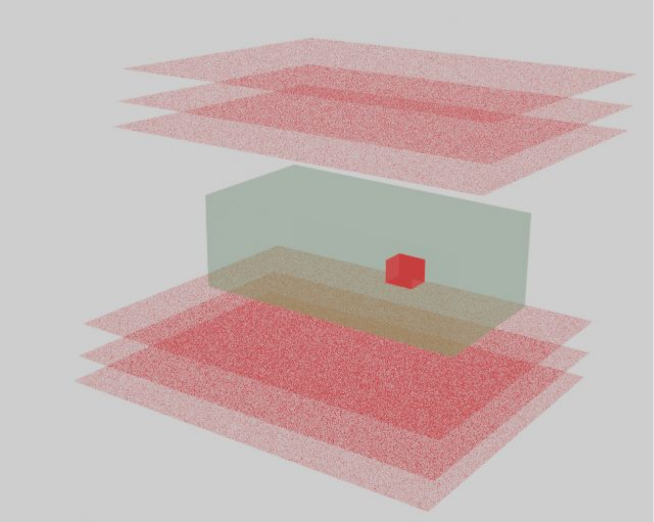

Jensen-Shannon divergence -based test |

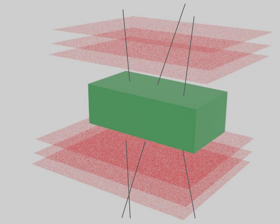

H_0

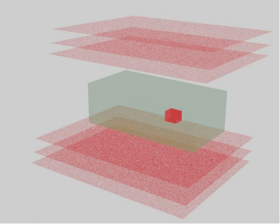

H_1

PoCA reconstruction
→ scattering distribution

Jensen-Shannon divergence -based test |

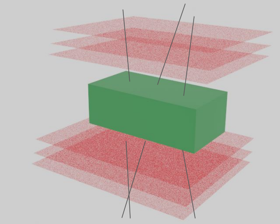

H_0

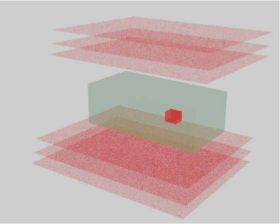

H_1

PoCA reconstruction
→ scattering distribution

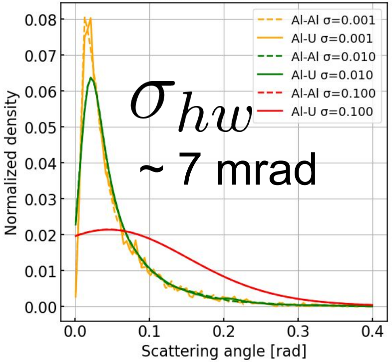

σ_{hw}
~ 7 mrad

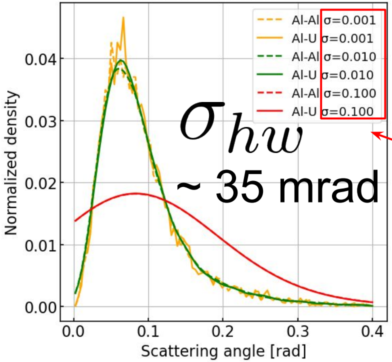

σ_{hw}
~ 35 mrad

σ_{sw}
Bandwidth

Jensen-Shannon divergence -based test |

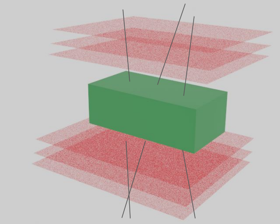

H_0

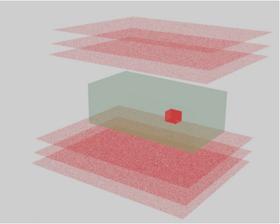

H_1

PoCA reconstruction
→ scattering distribution

Using volumes under H_0

1. Jensen-Shannon divergence wrt mean background
2. T_c at FPR=5%

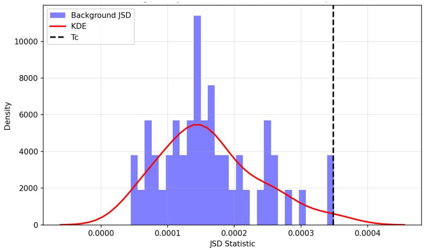

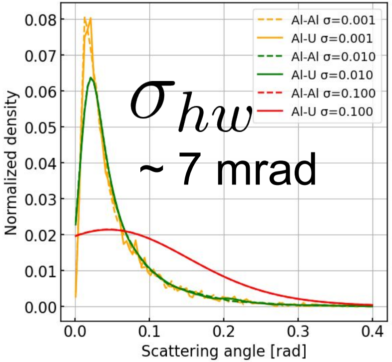

σ_{hw}

~ 7 mrad

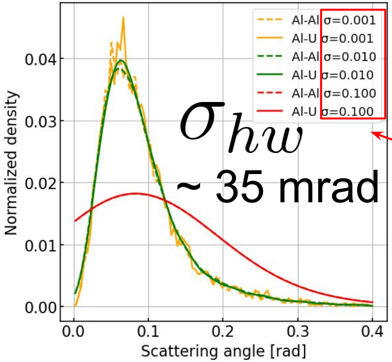

σ_{hw}

~ 35 mrad

σ_{sw}

Bandwidth

Jensen-Shannon divergence -based test |

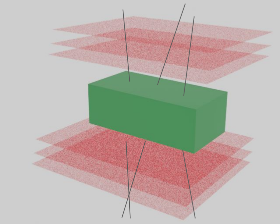

H_0

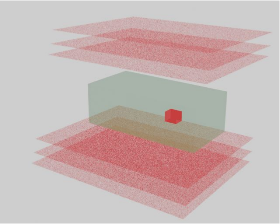

H_1

PoCA reconstruction
→ scattering distribution

Using volumes under H_0

1. Jensen-Shannon divergence wrt mean background
2. T_c at FPR=5%

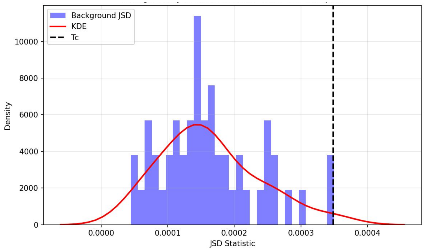

Using volumes under H_1

1. Jensen-Shannon divergence wrt mean background
2. Detections above T_c

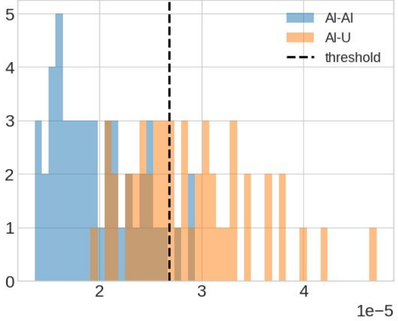

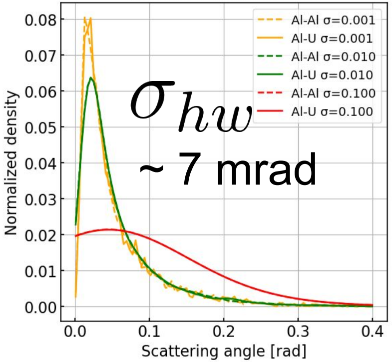

σ_{hw}

~ 7 mrad

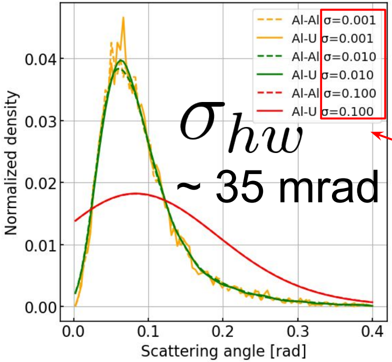

σ_{hw}

~ 35 mrad

σ_{sw}

Bandwidth

Jensen-Shannon divergence -based test |

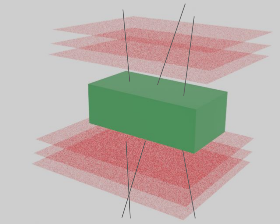

H_0

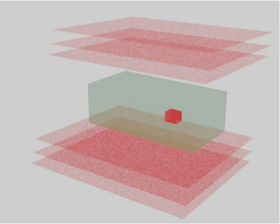

H_1

PoCA reconstruction
→ scattering distribution

Using volumes under H_0

1. Jensen-Shannon divergence wrt mean background
2. T_c at FPR=5%

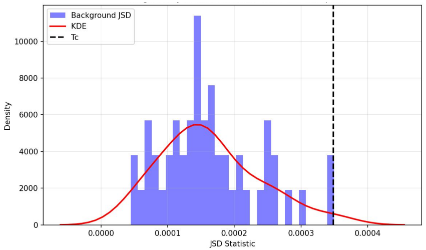

Using volumes under H_1

1. Jensen-Shannon divergence wrt mean background
2. Detections above T_c

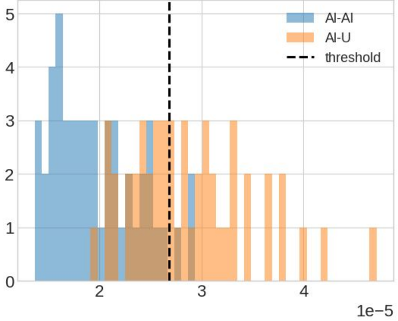

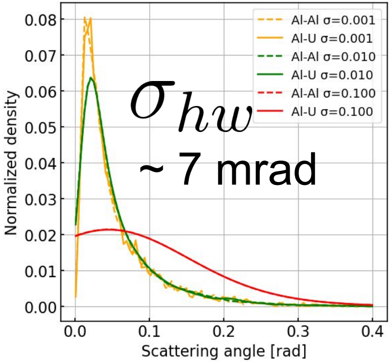

σ_{hw} ~ 7 mrad

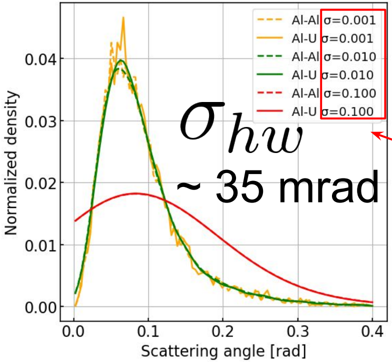

σ_{hw} ~ 35 mrad

σ_{sw}
Bandwidth

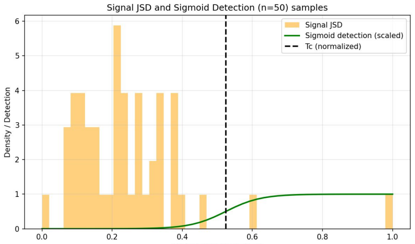

H_0

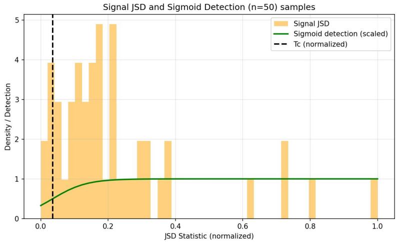

H_1

MC Geant4 simulation |

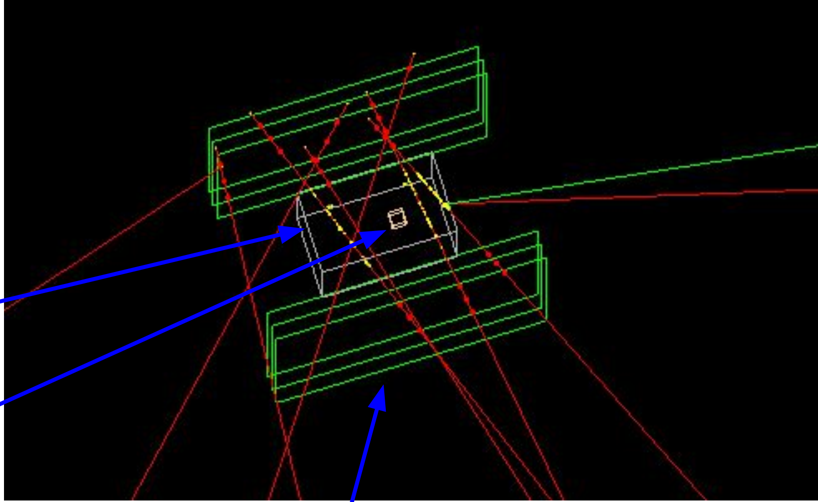

Aluminium block: 1 x 0.8 x 0.6 m^3

Uranium target: 10 x 10 x 10 cm^3

Detector panels: 2x2 m^2 panel

Spatial resolution modeled by Gaussian smearing of hit positions (1 mm , 5 mm)

Separation between panels: 10 cm

- Dataset:
  H_0 : 50 samples
  H_1: 50 samples
- Total ~**2h** of exposure per sample
- Scene integration using B2G4 (more in Felix's talk)

Factors affecting detection power |

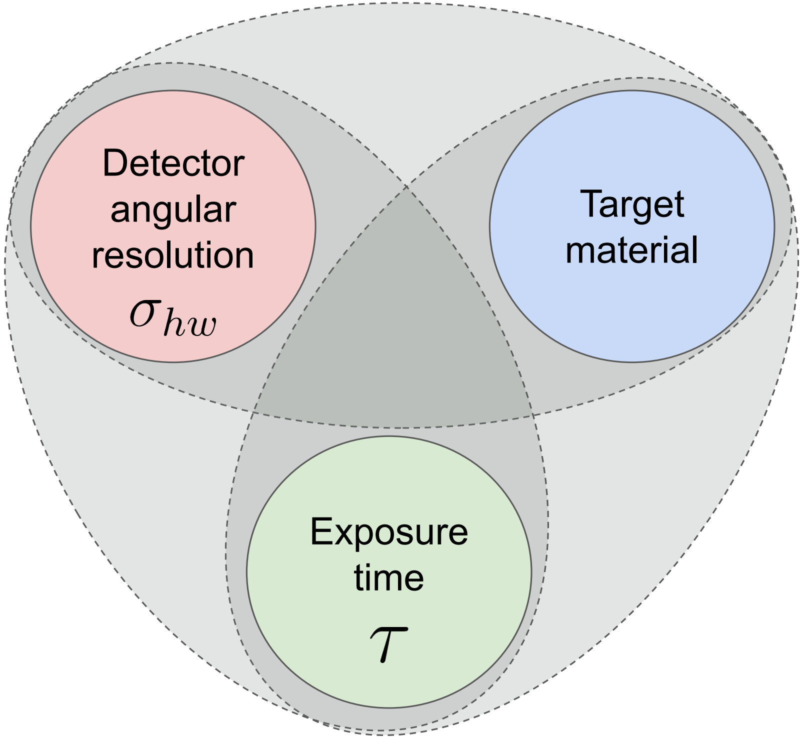

Detector angular resolution

σ_{hw}

Target material

Exposure time

τ

True Positive Rate (TPR) at 5% False Positive Rate (FPR) |

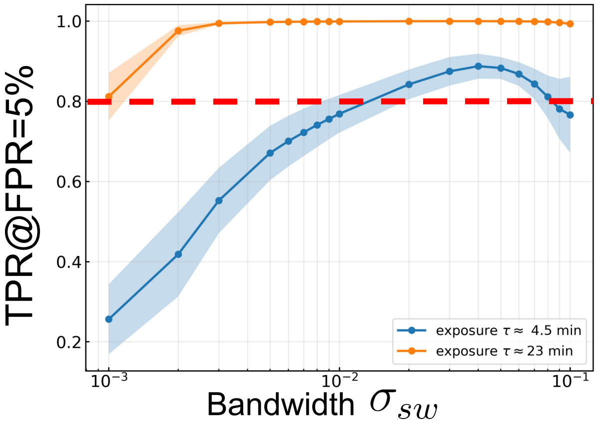

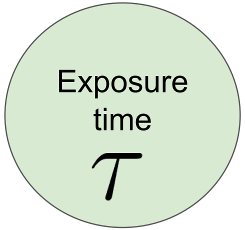

4.5 minutes vs 23 minutes

σ_{hw} ~ 7 mrad

U target

True Positive Rate (TPR) at 5% False Positive Rate (FPR) |

Target material

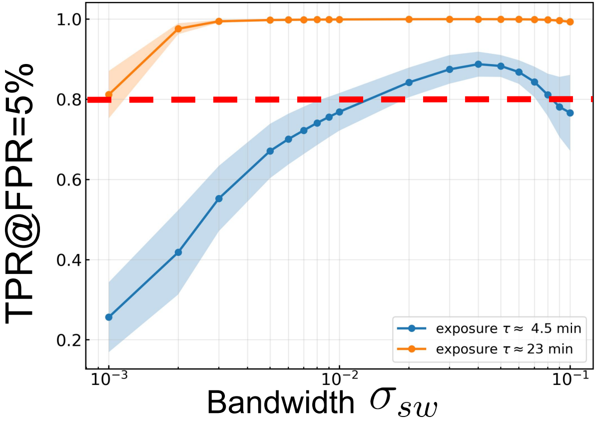

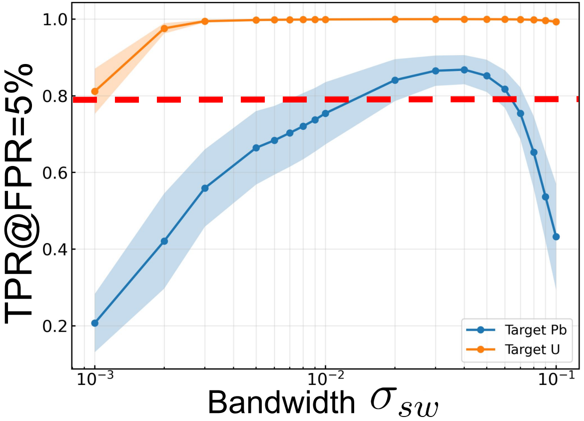

4.5 minutes vs 23 minutes

σ_{hw} ~ 7 mrad

U target

Pb target vs U target

σ_{hw} ~ 7 mrad

τ ~ 23 minutes

When:

- σ_{hw} is reliable
- statistics are abundant

→ sensitivity to ‘bulk’ density estimate (high-Z signature) at large σ_{sw}

True Positive Rate (TPR) at 5% False Positive Rate (FPR) |

Detector angular resolution σ_{hw}

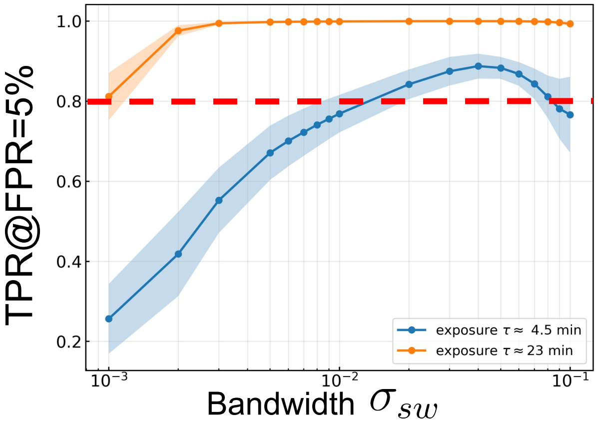

4.5 minutes **vs** **23 minutes**

σ_{hw} ~ 7 mrad

U target

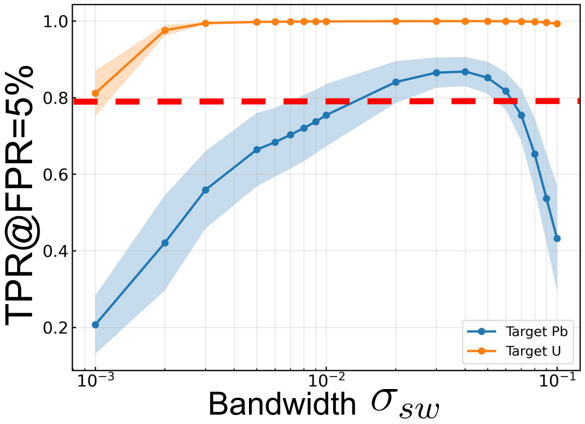

Pb target **vs** **U target**

σ_{hw} ~ 7 mrad

τ ~ 23 minutes

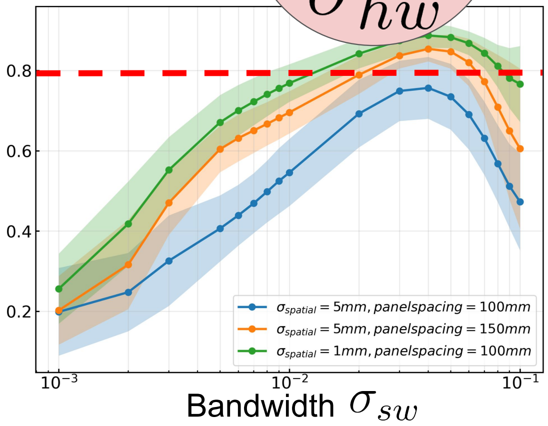

35 mrad **vs** **23 mrad** **vs** **7 mrad**

τ ~ 4.5 minutes

U target

True Positive Rate (TPR) at 5% False Positive Rate (FPR) |

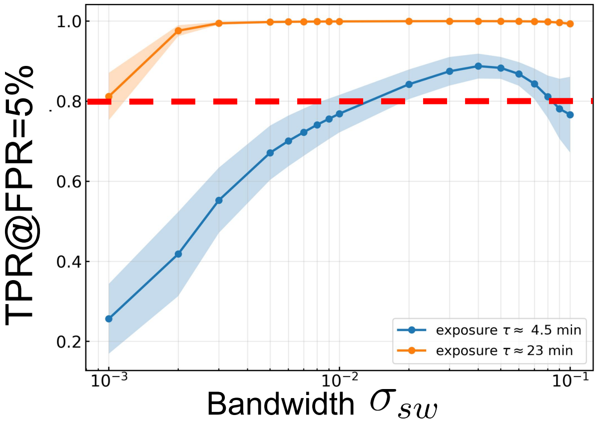

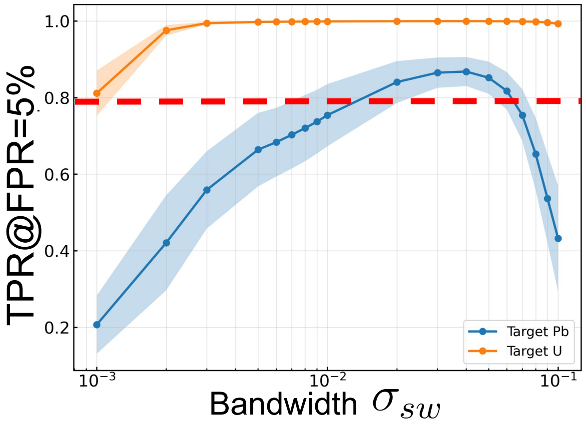

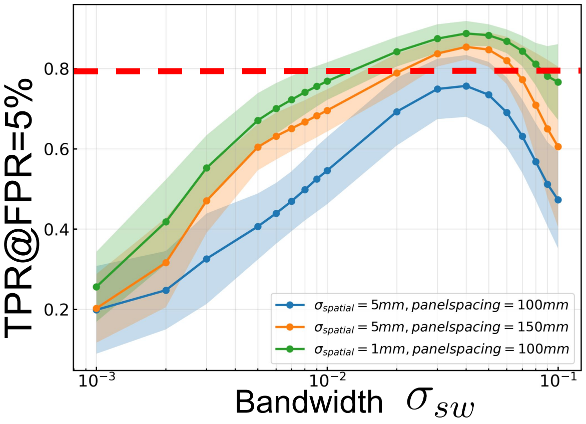

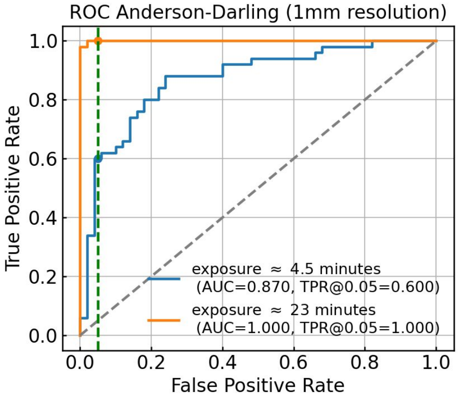

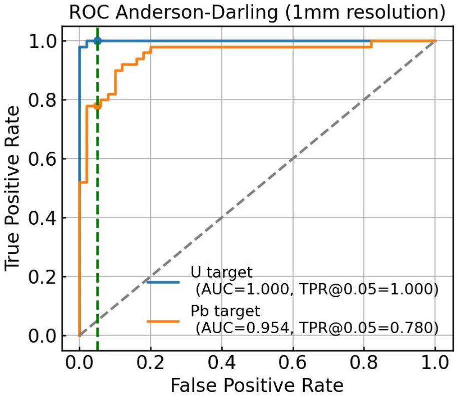

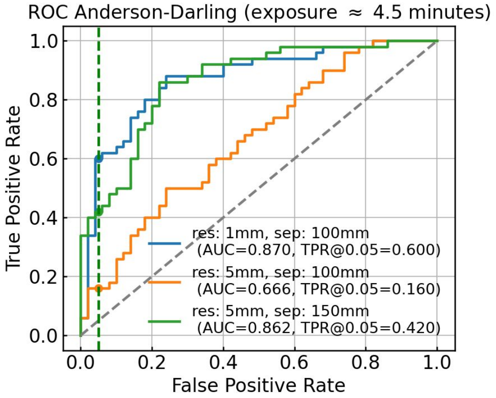

Part III |

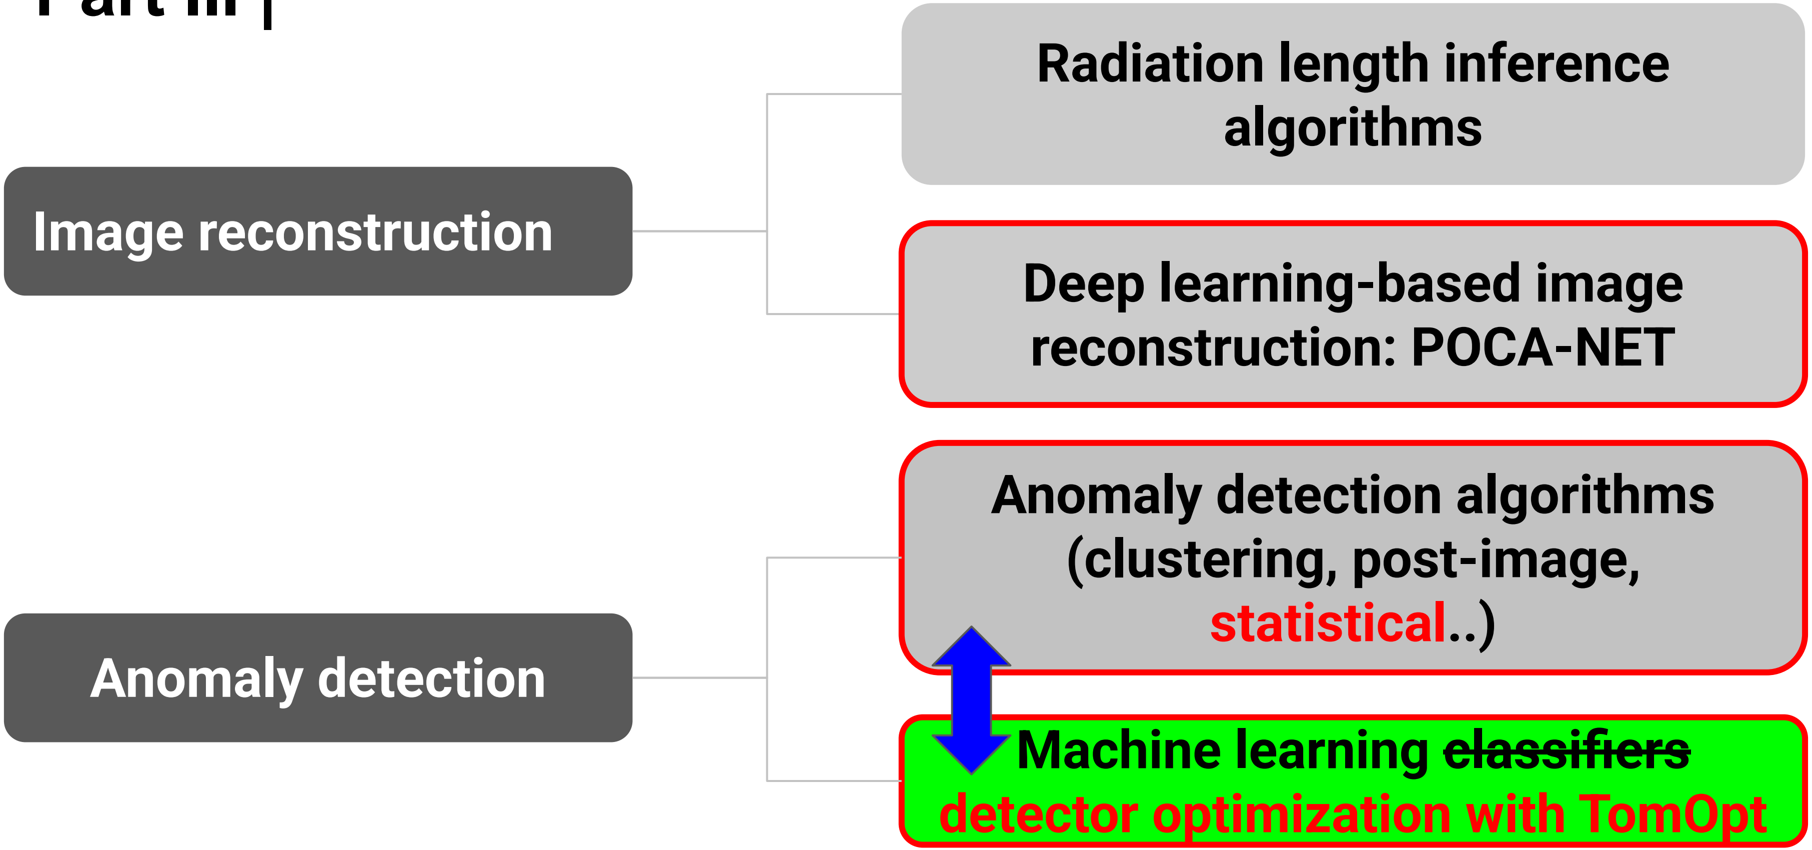

Image reconstruction

- Radiation length inference algorithms
- Deep learning-based image reconstruction: POCA-NET

Anomaly detection

- Anomaly detection algorithms (clustering, post-image, statistical..)
- Machine learning ~~classifiers~~ detector optimization with TomOpt

TomOpt: Gradient-based co-design |

Task : co-optimization of hardware σ_{hw} and software bandwidth σ_{sw}
for anomaly detection of uranium target

TomOpt: Gradient-based co-design |

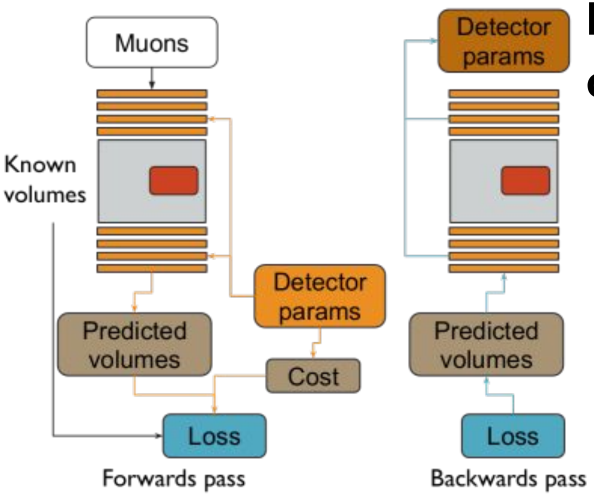

End-to-end MST detector **optimization** pipeline:

- Muon generation
- Mon transport
- Inference
- Inference loss minimization
- Parameter update (e.g. panel size, position, ..)

Task : co-optimization of hardware σ_{hw} and software bandwidth σ_{sw} for anomaly detection of uranium target

TomOpt: Gradient-based co-design | $a_{n+1} = a_n \boxed{- \gamma \nabla \mathcal{L}(a_n)}$

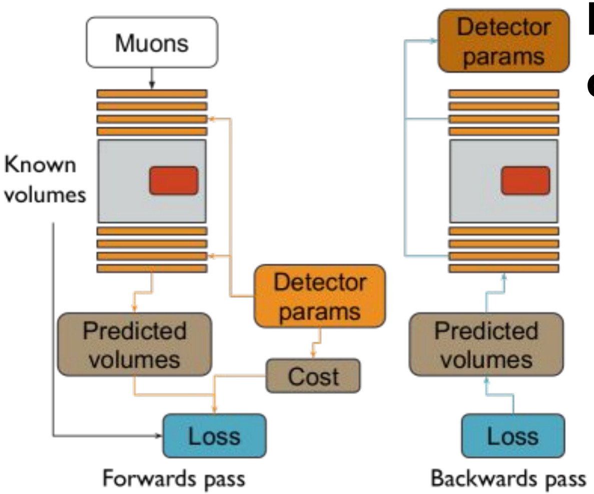

End-to-end MST detector **optimization** pipeline:

- ○ Muon generation
- ○ Mon transport
- ○ Inference
- ○ Inference loss minimization
- ○ Parameter update (e.g. panel size, position, ..)

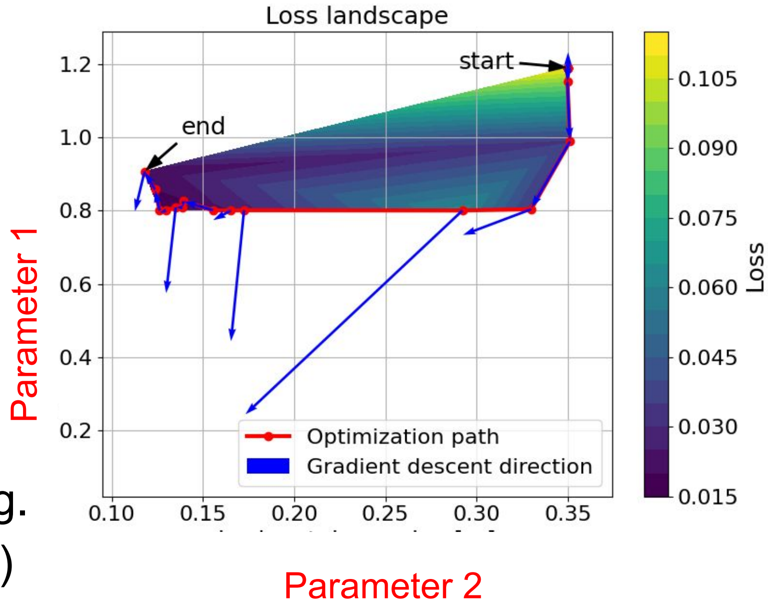

Task : co-optimization of hardware σ_{hw} and software bandwidth σ_{sw}

for anomaly detection of uranium target

zahraa.zaher@uclouvain.be

TomOpt: Gradient-based co-design

Searching for a reference operating point in parameter space

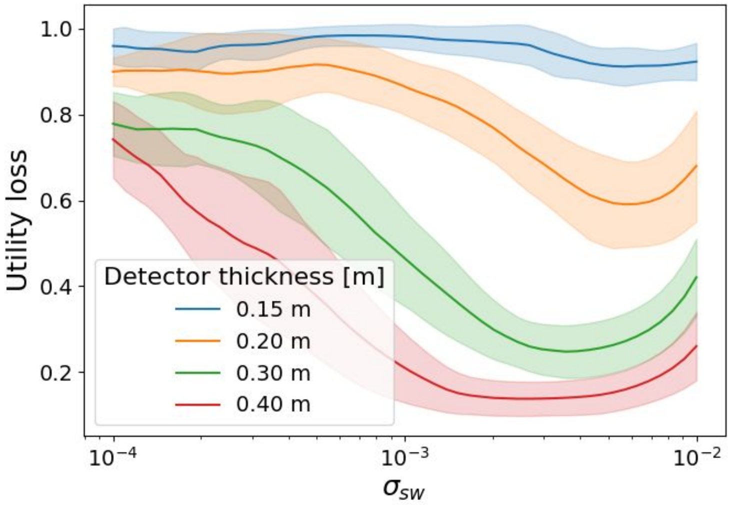

Reference

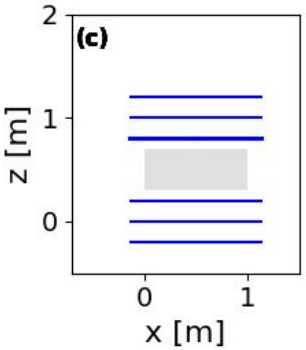

TomOpt: Gradient-based co-design

Initial

~ 20 iterations

Optimized

Reference

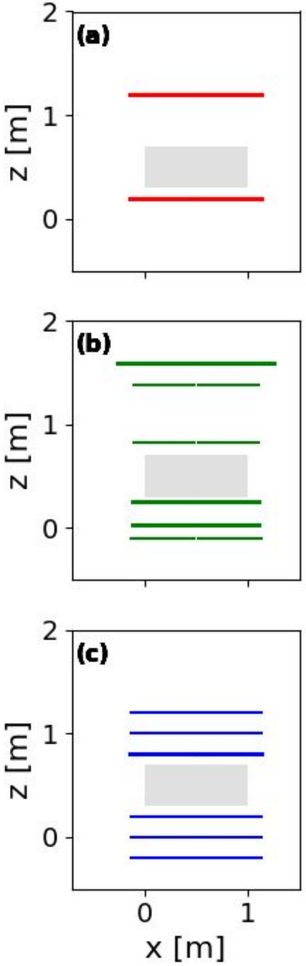

TomOpt: Gradient-based co-design

ROC AUC across different **signal/background** material assignments

Initial

~ 20 iterations

Optimized

Reference

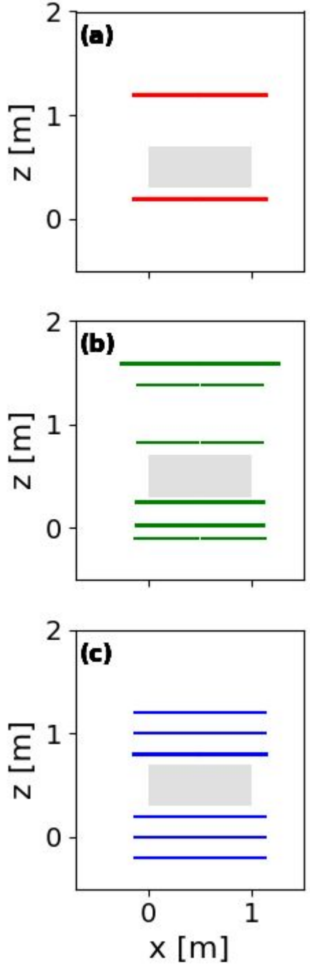

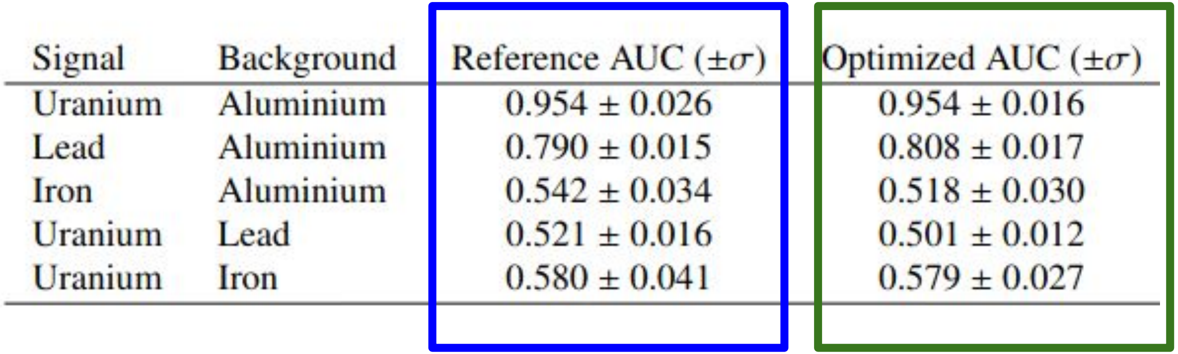

Signal	Background	Reference AUC ($\pm\sigma$)	Optimized AUC ($\pm\sigma$)
Uranium	Aluminium	0.954 ± 0.026	0.954 ± 0.016
Lead	Aluminium	0.790 ± 0.015	0.808 ± 0.017
Iron	Aluminium	0.542 ± 0.034	0.518 ± 0.030
Uranium	Lead	0.521 ± 0.016	0.501 ± 0.012
Uranium	Iron	0.580 ± 0.041	0.579 ± 0.027

https://doi.org/10.48550/arXiv.2603.26613

zahraa.zaher@uclouvain.be

Back to our security problem

Assumption: container contains a single material (or materials with similar Z)

Question:
“Can we discriminate the presence of anomalous material?”

Back to our security problem

Reference scan
(e.g. passing an X-ray scan)

Let's slice our container

Remaining slices are compared against the reference

Back to our security problem

Reference scan
(e.g. passing an X-ray scan)

Let's slice our container

Remaining slices are compared against the reference

If statistic exceeds expected T_c
→ **anomaly**

zahraa.zaher@uclouvain.be

THANK YOU

zahraa.zaher@uclouvain.be

BACKUP

zahraa.zaher@uclouvain.be

Radiation length inference

$$\theta_{RMS} = \frac{13.6\,MeV}{\beta c p}\sqrt{\frac{x}{\boxed{X_0}}}\left[1 + 0.038\,\ln\left(\frac{x}{X_0\,\beta^2}\right)\right]$$

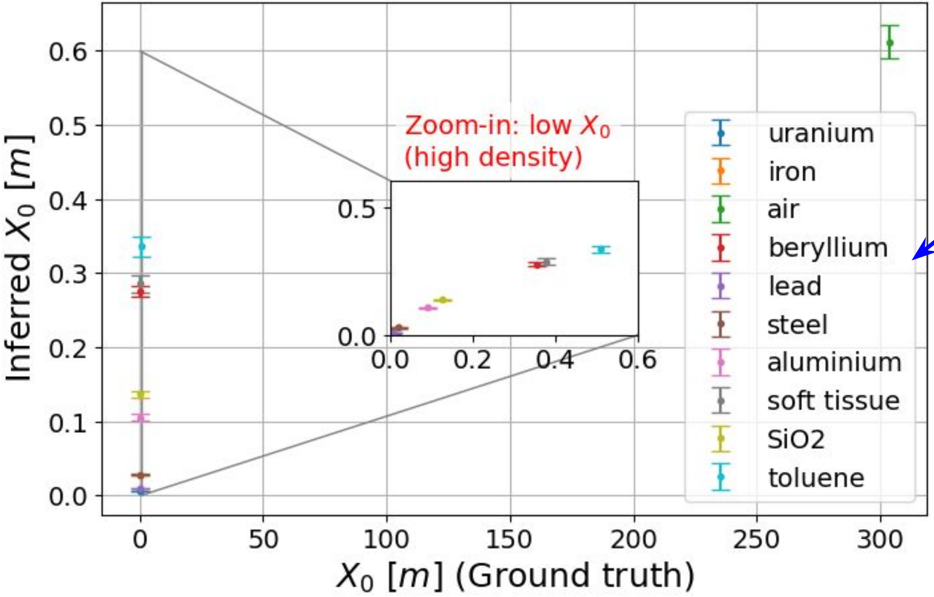

Unshielded blocks

zahraa.zaher@uclouvain.be

POINT-NET architectuee |

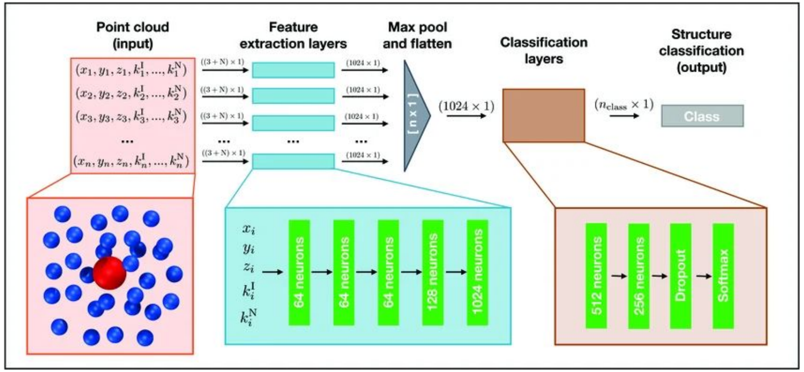

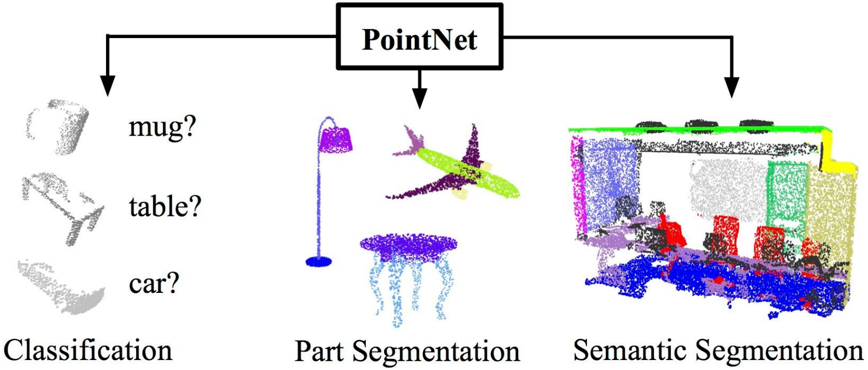

POCA-NET vs 3D-UNet

Global normalized MSE across all voxels (including background)

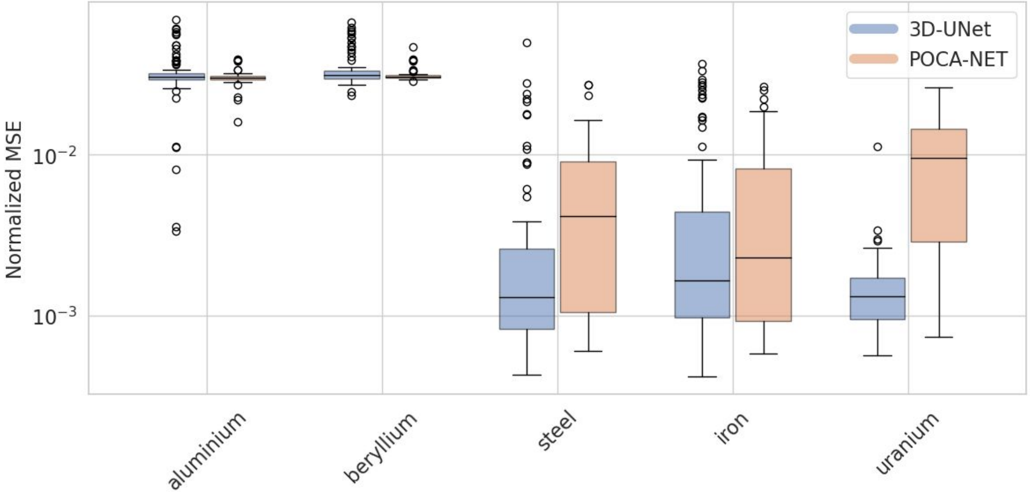

MC Geant4 simulation |

Aluminium block: 1 x 0.8 x 0.6 m^3

Uranium target: 10 x 10 x 10 cm^3

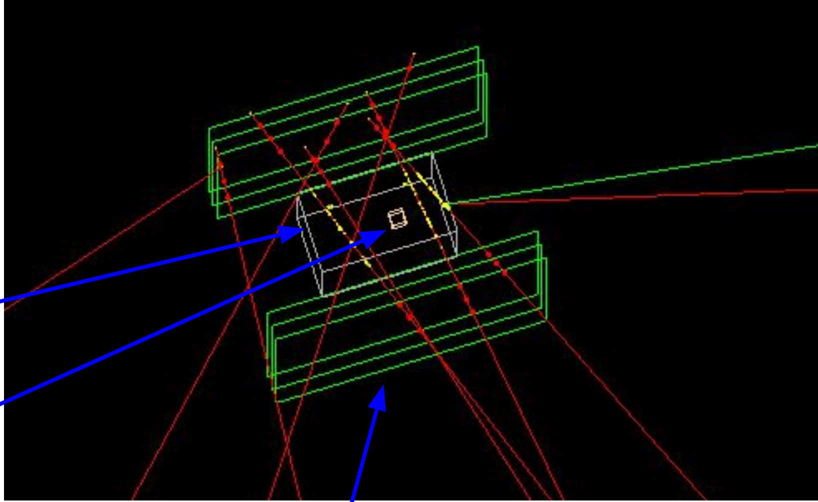

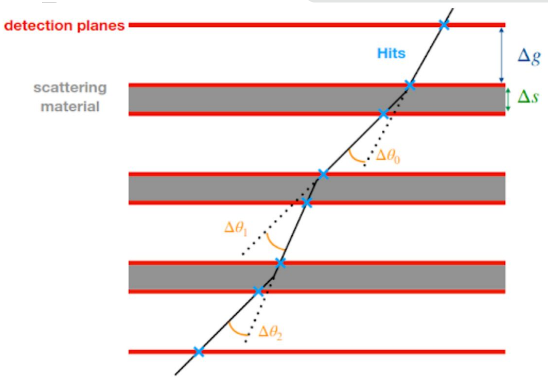

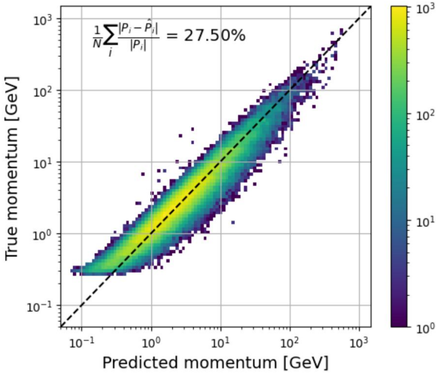

https://doi.org/10.3390/particles8020043

Detector panels: 2x2 m^2 panel

Spatial resolution modeled by Gaussian smearing of hit positions (1 mm , 5 mm)

Separation between panels: 10 cm

Rough **momentum** knowledge in [1, 50] GeV

Background

Signal

difference

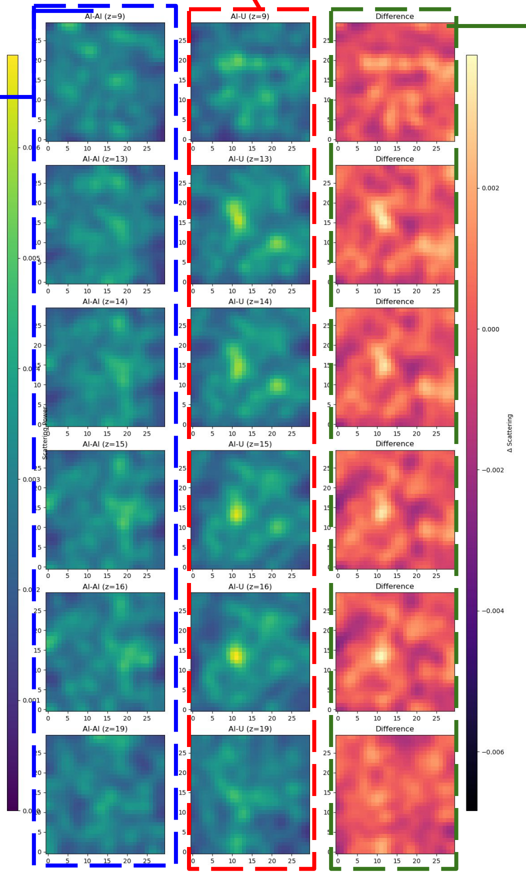

- **Exposure ~4.5 minutes**

- **5 mm resolution**

U-Target block

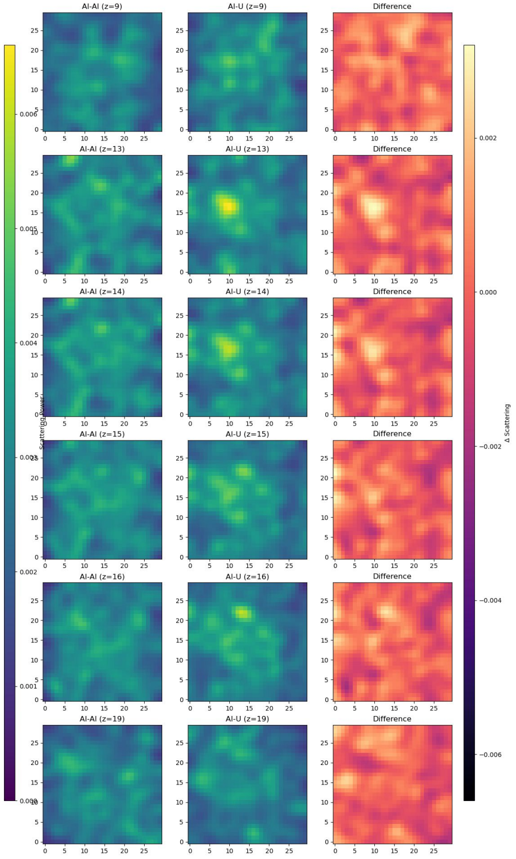

Pb-Target block

Background

Signal

difference

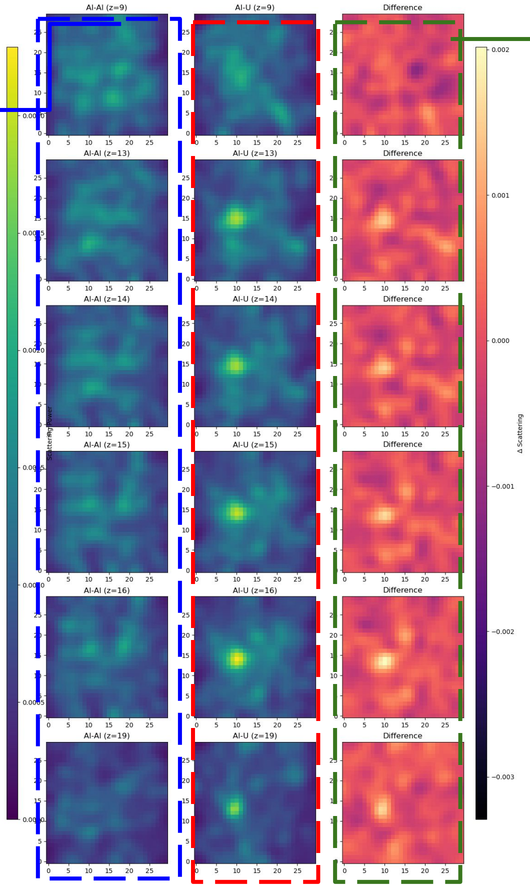

U-Target block

- **Exposure ~4.5 minutes**
- **1 mm resolution**

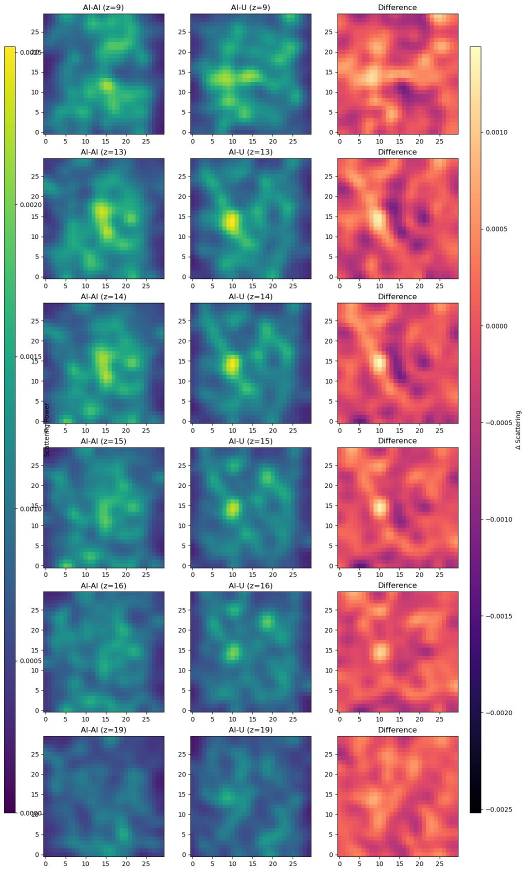

Pb-Target block

zahraa.zaher@uclouvain.be

Shape vs mean sensitivity of JSD test

N_muons = 150k (~ 67 min)

Shape mismatch sensitivity

Collapse to mean-shift (tail sensitivity)

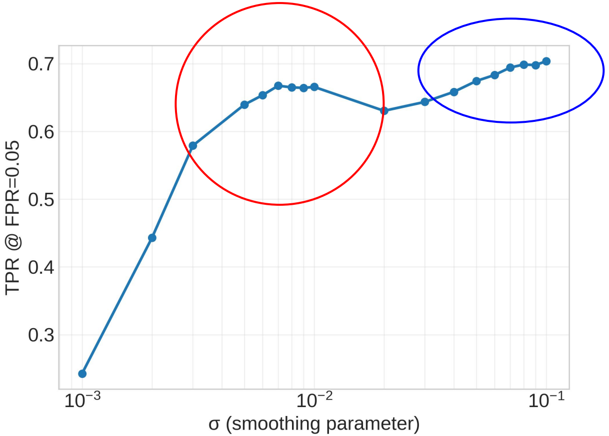

Res = 1 mm
sep = 100mm

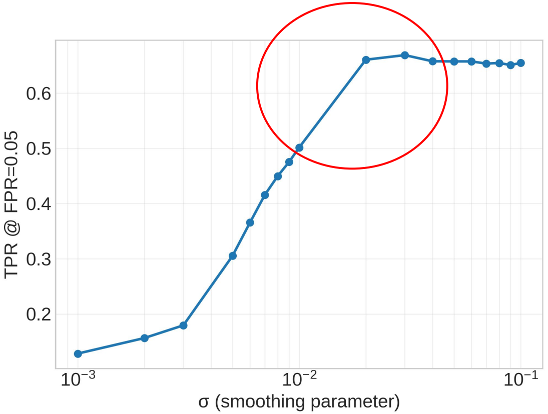

Res = 5 mm
sep = 100 mm

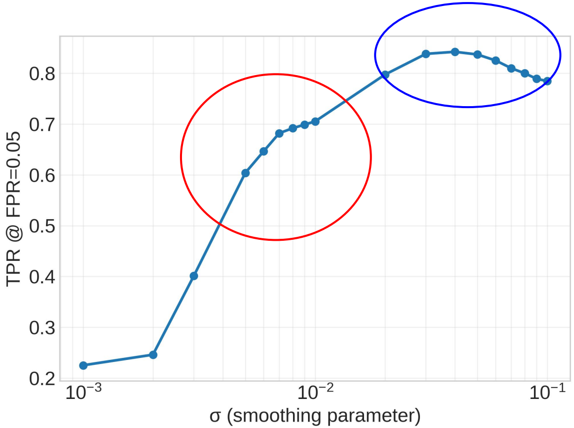

Res = 5 mm
Sep = 150 mm

Anderson-Darling test

Natural cosmic muon energy spectrum

Exposure ~ 67 minutes

1 mm vs 5 mm resolution

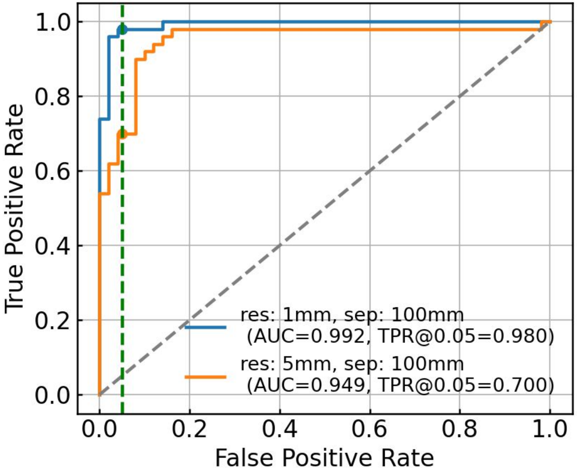

5 mm resolution:
100 mm vs 150 mm panel separation

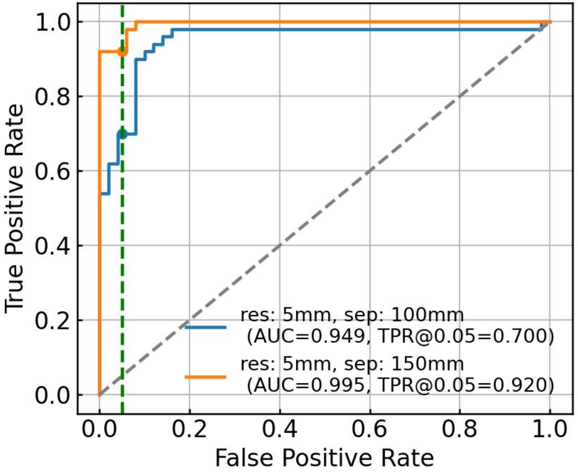

5 mm resolution:
67 minutes vs 90 minutes

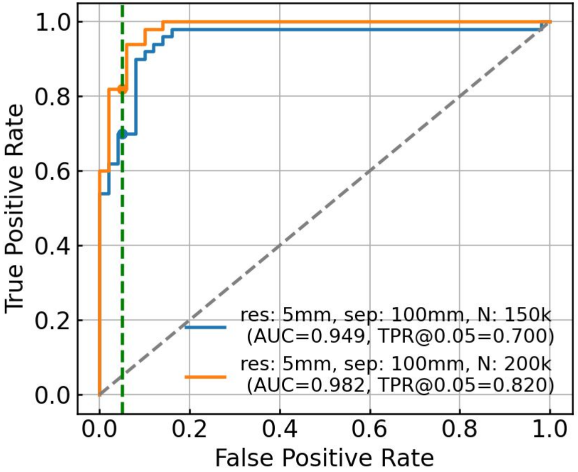

zahraa.zaher@uclouvain.be

Co-design optimization

After training

Differentiable hypothesis test:

$$D_{JS}(S_V, \bar{S}_{bkg}) = \frac{1}{2} D_{KL}(S_V \,\|\, M) + \frac{1}{2} D_{KL}(\bar{S}_{bkg} \,\|\, M),$$
$$M = \frac{1}{2}(S_V + \bar{S}_{bkg})$$

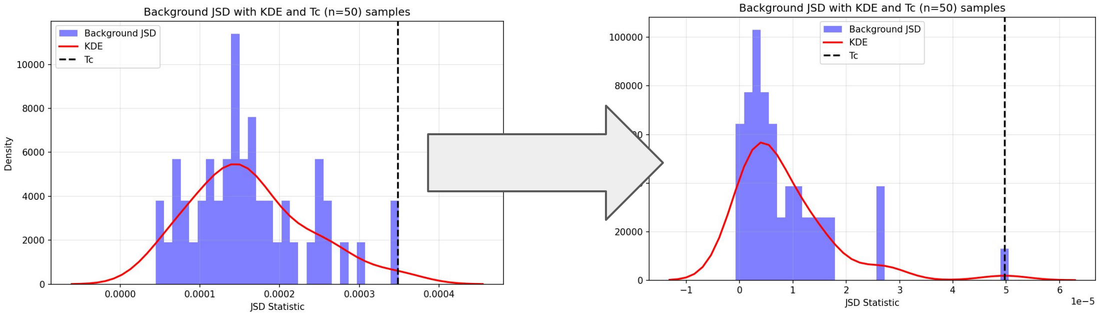

Threshold T_c at alpha=0.05 using bkg KDE

Detection sigmoid:

$$\hat{y}_i = \frac{1}{1 + \exp\left[-\tau_{\text{detect}}\left(D_{JS}(S^{\text{sig}}_{V_i}, \bar{S}_{bkg}) - T_c\right)\right]},$$

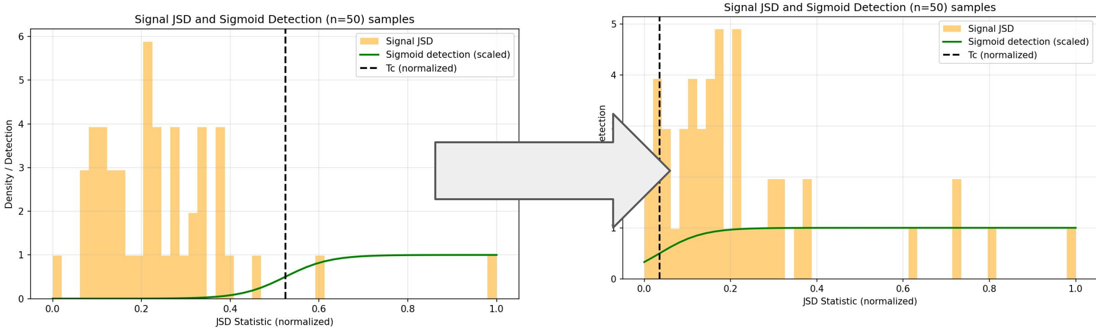

Detection power:

$$\Pi = \frac{1}{N_{\text{sig}}} \sum_{i=1}^{N_{\text{sig}}} \hat{y}_i.$$

Utility Loss: $$\mathcal{L}(\sigma_{hw}, \sigma_{sw}) = 1 - \Pi + \lambda \sigma_{hw},$$

Similar study

Nonparametric Dense-Object Detection Algorithm for Applications o
Cosmic-Ray Muon Tomography
(https://journals.aps.org/prapplied/pdf/10.1103/PhysRevApplied.14.
64032)