BERKELEY LAB

Machine Learning for Applied Nuclear Physics

Marco Salathe, March 30, 2022

Applied Nuclear Physics Program, LBNL

Graduate Seminar for Particle Physicist (Physics 290E)

BERKELEY LAB

About myself

- Scientific Engineering Associate
  - Applied Nuclear Physics Program
  - Nuclear Physics Division
  - Lawrence Berkeley National Laboratory
- Born and raised in Switzerland near Basel
- PostDoc in Radiation detectors for Nuclear Structure Physics at LBNL
- PhD in Particle Physics
  - Radiation detectors for neutrinoless double beta decay experiments
- Bachelors/Masters in Physics/Physical Engineering
  - Ecole Polytechnique Fédérale de Lausanne, Switzerland
- Work focus: Machine Learning as a tool to enhance and inform radiation detection in urban environments
  - Mostly Nonproliferation, National, and Homeland Security
  - Some basic science

Overview

- Introduction
- Machine Learning for detection and identification of radioactive material
- Semantic segmentation to support radiation detection
- Object detection and radiation attribution
- Machine Learning for beam line operation and optimization

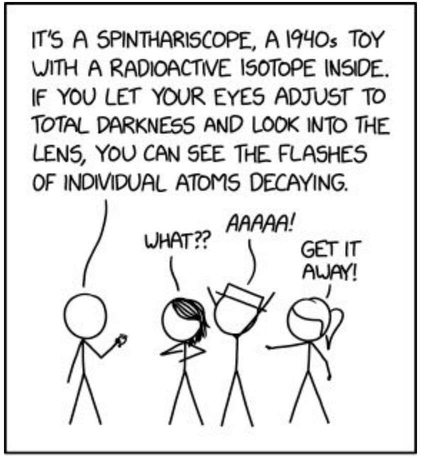

FUN FACT: SPINTHARISCOPES HAVE THE HIGHEST RATIO OF "THAT CAN'T POSSIBLY BE SAFE AND LEGAL" TO ACTUAL SAFETY AND LEGALITY OF ANY KNOWN TOY.

Introduction

to radiation detection in urban environments

Radiological source searches in urban environments

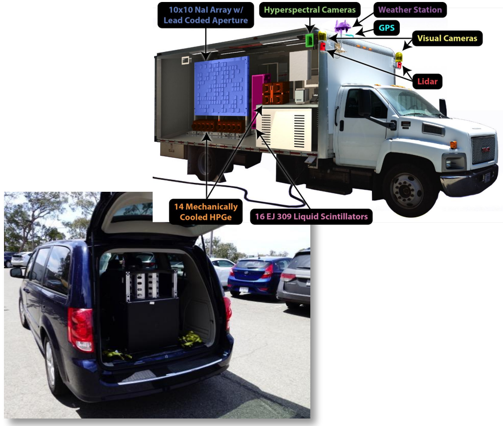

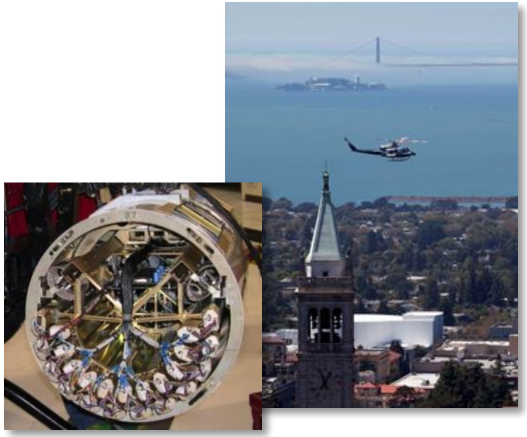

Street Scale

Block Scale

City Scale

How do we detect and identify radioactive material?

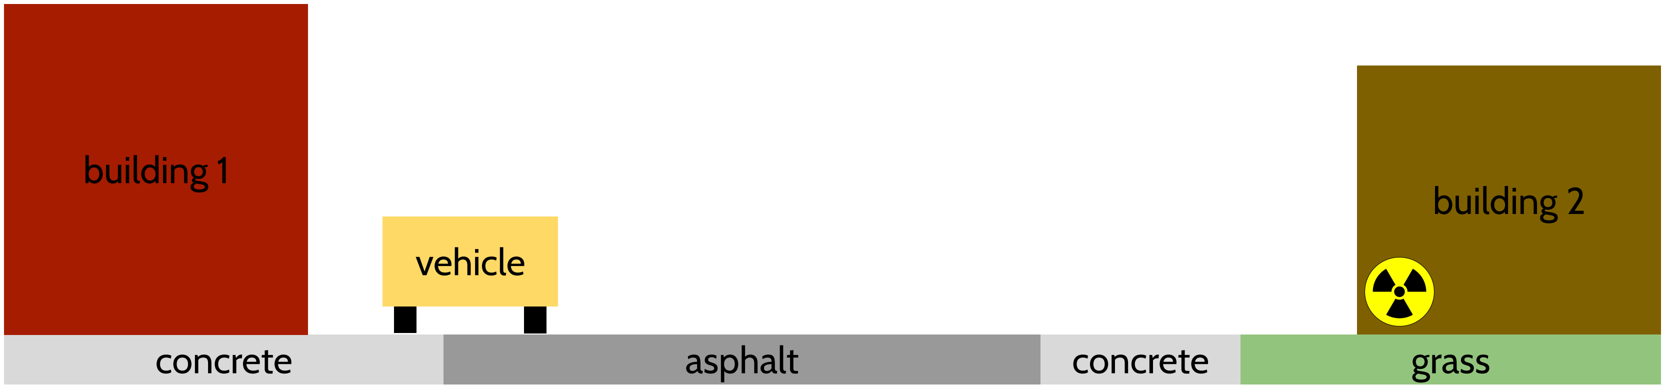

How do we detect and identify radioactive material?

- We need a signal:

Apha	Beta-minus	Beta-plus	gamma	neutron
- Very short-range - Characteristic energies	– Short-range – Energy continuum	– Short-range – Annihilation yields characteristic gamma-ray	– Penetrating – Characteristic energies – Scattering introduces noise	– Penetrating – Energy continuum – Scattering introduces noise but facilitates detection

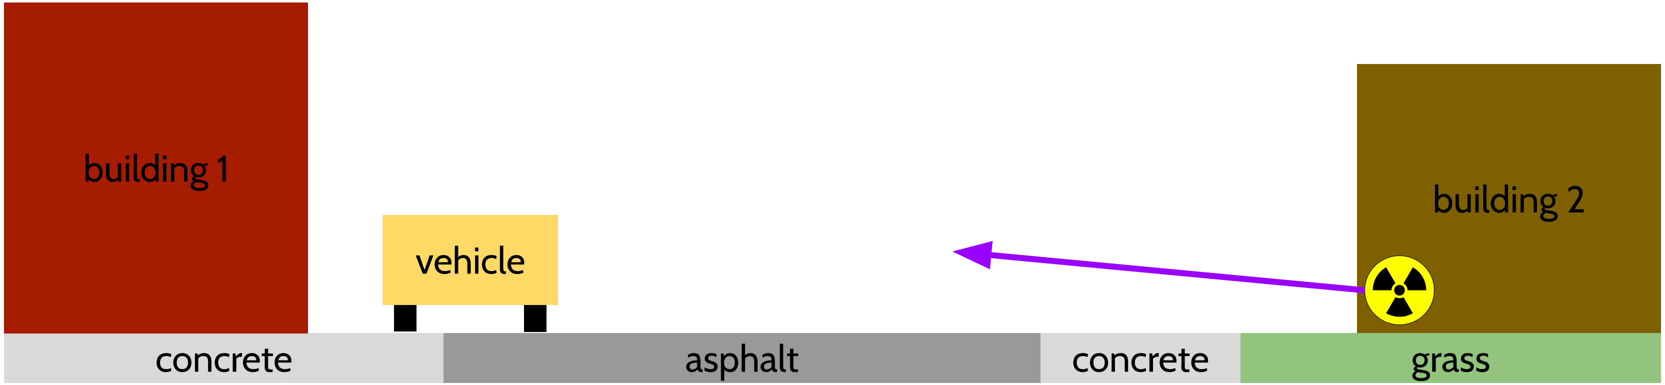

How do we detect and identify radioactive material?

- We need a signal:

Apha	Beta-minus	Beta-plus	gamma	neutron
- Very short-range - Characteristic energies	– Short-range – Energy continuum	– Short-range – Annihilation yields characteristic gamma-ray	– Penetrating – Characteristic energies – Scattering introduces noise	– Penetrating – Energy continuum – Scattering introduces noise but facilitates detection

- Emissions are a property of individual isotopes
  – Penetrating (allows detection) and characteristic (allows identification)

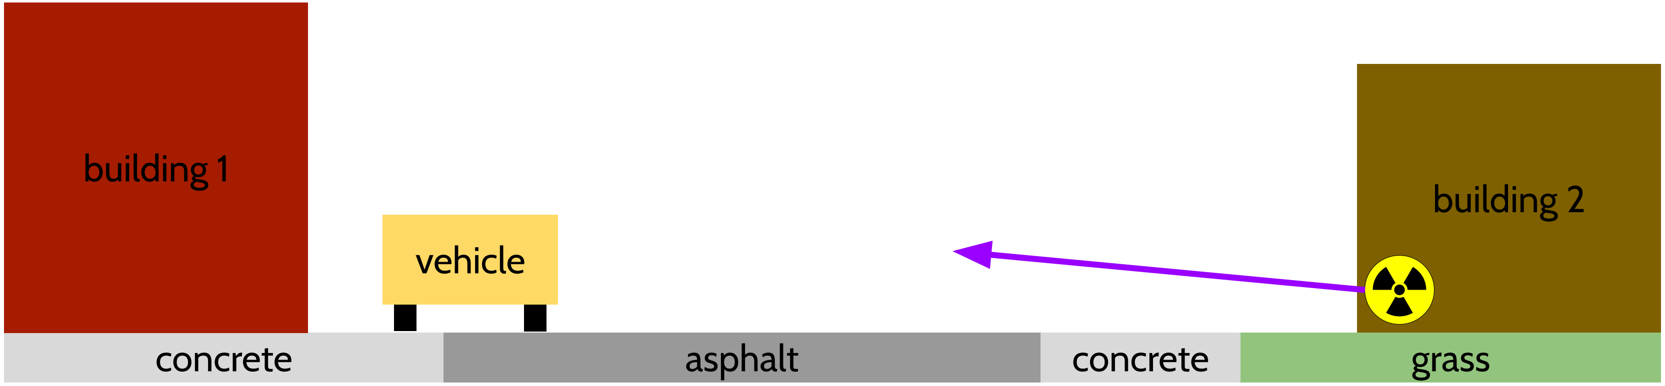

How do we detect and identify radioactive material?

- We need a signal:

Apha	Beta-minus	Beta-plus	gamma	neutron
- Very short-range - Characteristic energies	– Short-range – Energy continuum	– Short-range – Annihilation yields characteristic gamma-ray	– Penetrating – Characteristic energies – Scattering introduces noise	– Penetrating – Energy continuum – Scattering introduces noise but facilitates detection

- Emissions are a property of individual isotopes
  - Penetrating (allows detection) and characteristic (allows identification)
- And a detector: For example scintillator (NaI,CZT) or semiconductors (HPGe)

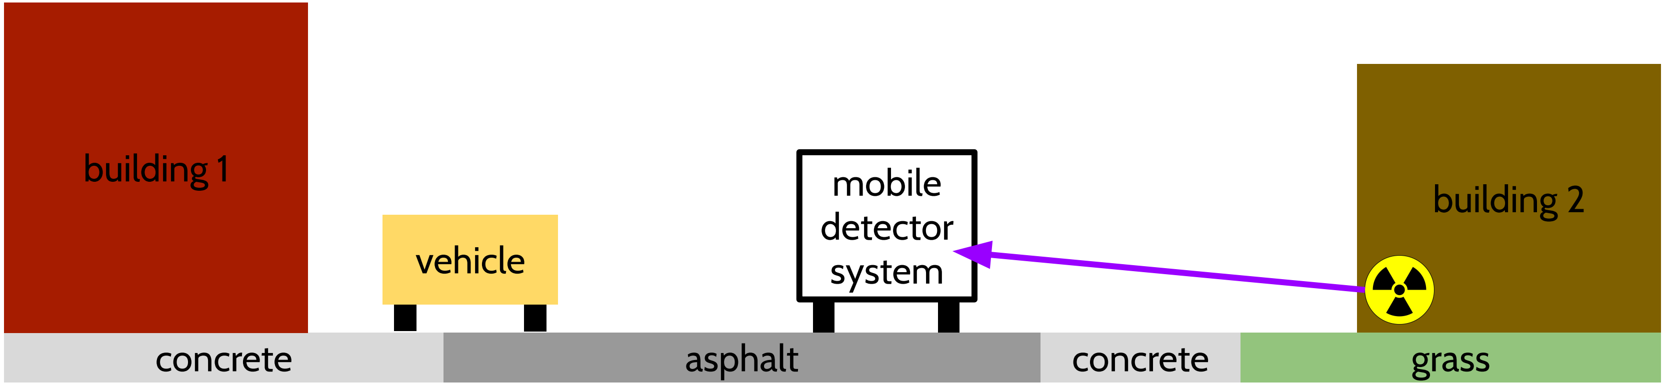

Lots of different detection system!

Hand-carried

Ground-based vehicle

Static

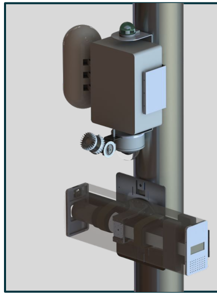

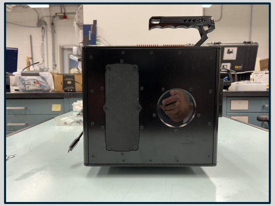

Airborne

Source localization

- To localize a source multiple measurements with varying configurations are required
- Static system only provide sufficient data if many detectors are used (masking)

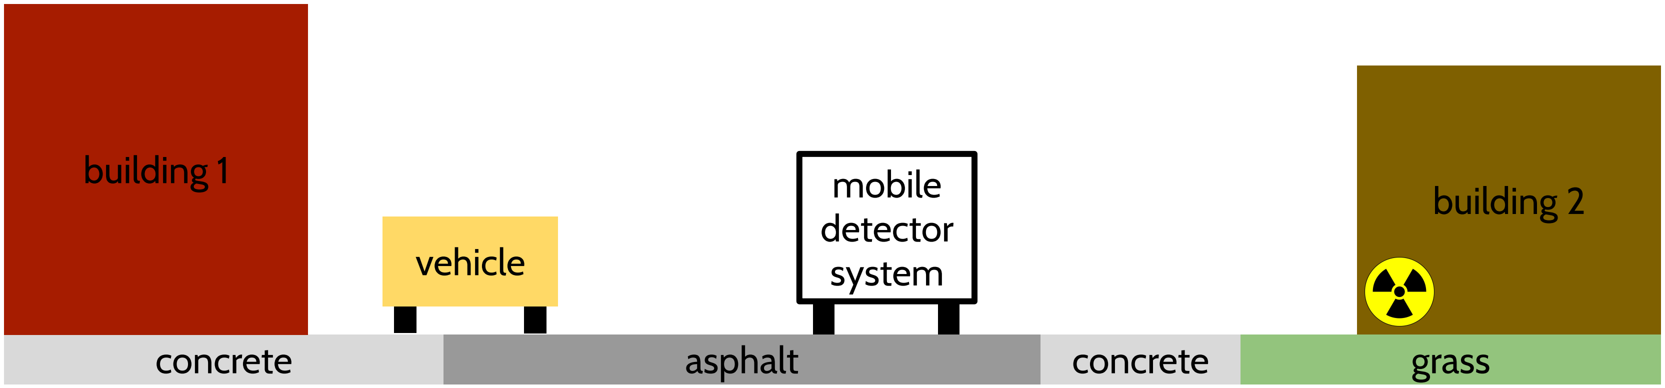

Source localization

- To localize a source multiple measurements with varying configurations are required
- Static system only provide sufficient data if many detectors are used (masking)
- Moving systems encounter different offsets and orientations to the source over time.
- This enables to formulate a linear system of equations and solve for the unknown source location and strength.
  - For Poisson data: Maximum Likelihood Estimation Maximization
  - Iterative update rule that converges towards a solution
- Requirements: Exact detector trajectory must be measured

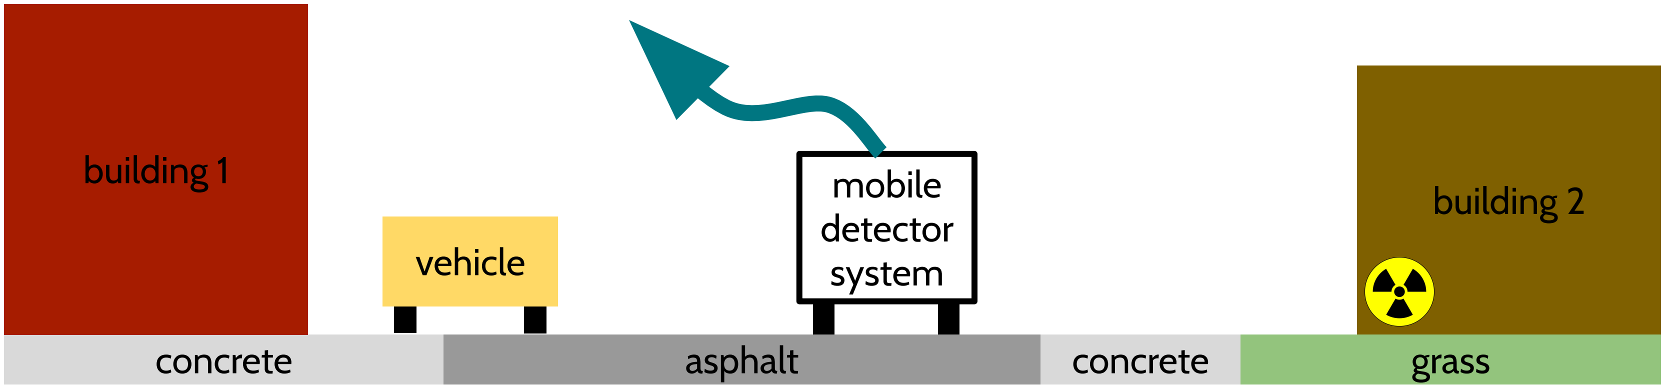

Contextual sensors

Leverage contextual sensors such as lidars, cameras, etc.

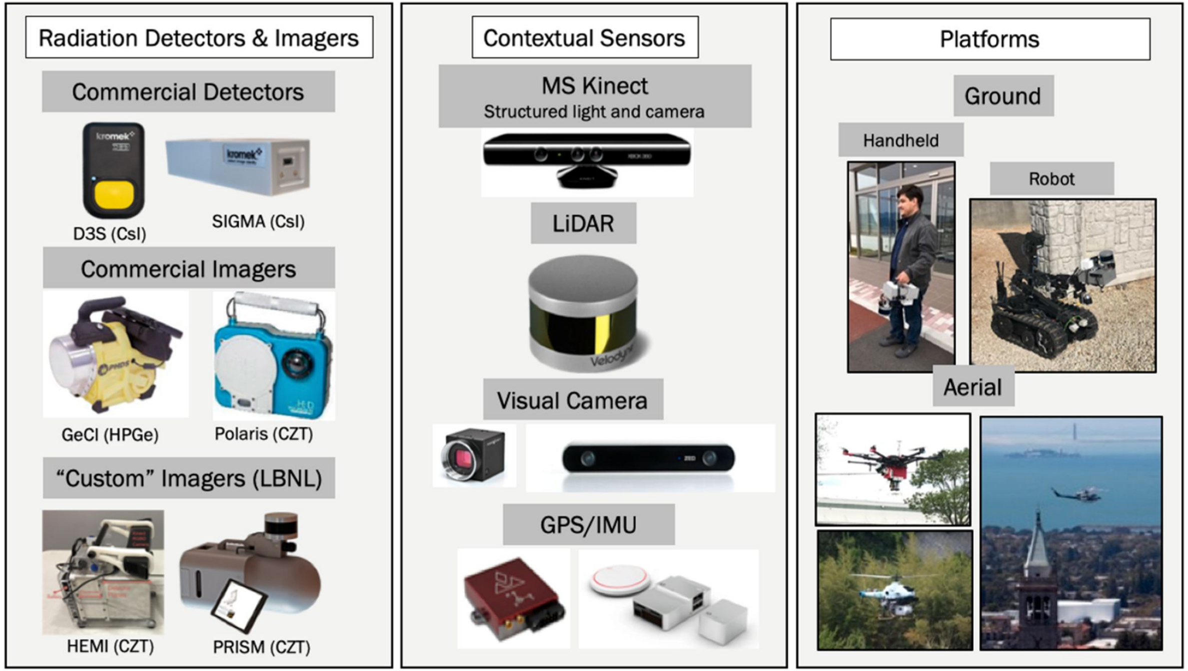

K. Vetter, et. al., Advances in Nuclear Radiation Sensing: Enabling 3-D Gamma-Ray Vision. Sensors 2019, 19, 2541. https://doi.org/10.3390/s19112541

Localization and mapping

- Simultaneous Localization and Mapping (SLAM) using Google Cartographer https://github.com/googlecartographer/cartographer
  - Minimize cost function between current LiDAR data and the reconstructed map from previous data
  - Rotation frequency of LiDAR 10Hz
- Requires LiDAR sensor and a Inertial Measurement Unit (IMU)
- Freebie: A map of the scene that can be leveraged to simplify the localization problem

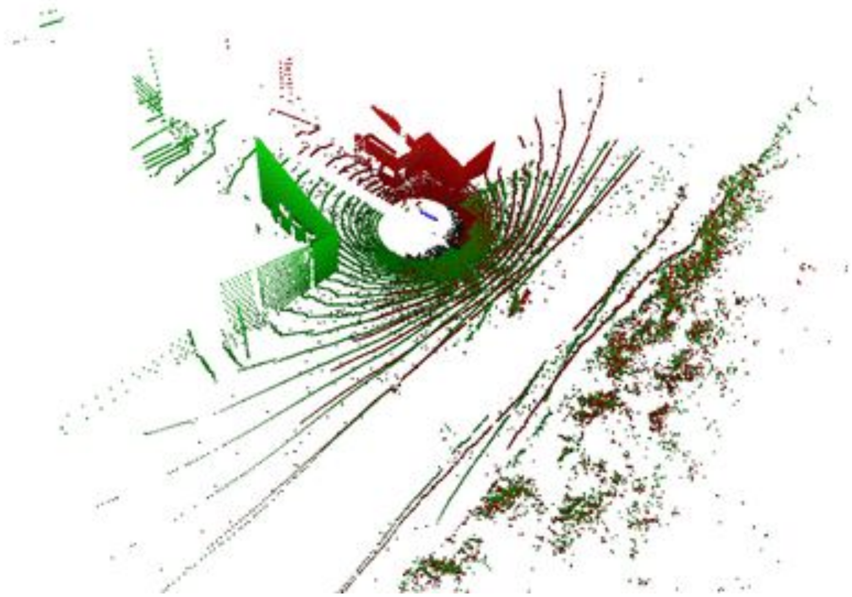

Nuclear Scene Data Fusion (SDF)

- Integrate auxiliary contextual sensors and onboard computer with radiation detection/imaging systems
- Perform SLAM in real-time (~10 Hz) to map the 3D scene and track the system pose as it moves freely through the environment
- Scene data fused with radiation data
  - Time synchronization
  - Coordinate transforms
- Real-time poses enable free-moving 3D imaging and real-time map provides a dynamic, and constrainable, image space
  - Map constraint can increase image accuracy, reduce noise, decrease reconstruction time

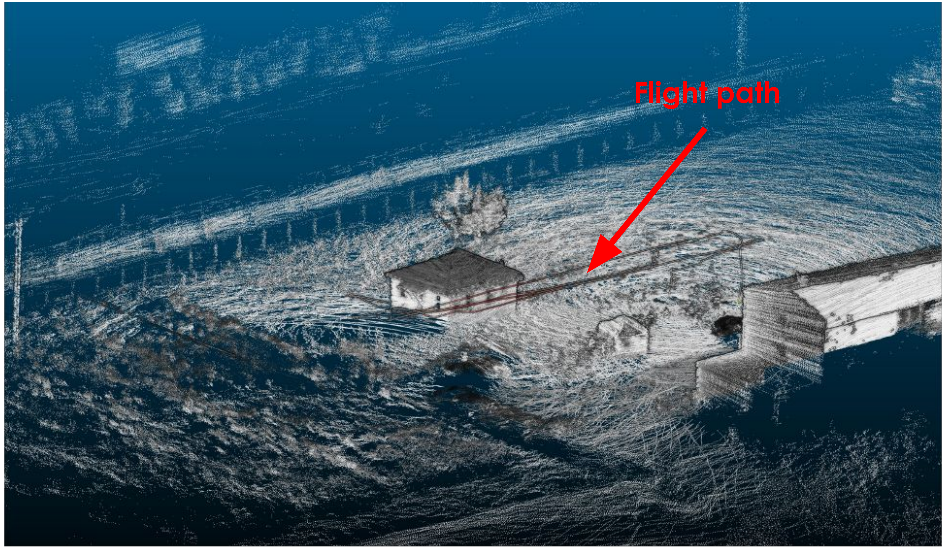

Nuclear Scene Data Fusion (SDF)

- Integrate auxiliary contextual sensors and onboard computer with radiation detection/imaging systems
- Perform SLAM in real-time (~10 Hz) to map the 3D scene and track the system pose as it moves freely through the environment
- Scene data fused with radiation data
  - Time synchronization
  - Coordinate transforms
- Real-time poses enable free-moving 3D imaging and real-time map provides a dynamic, and constrainable, image space
  - Map constraint can increase image accuracy, reduce noise, decrease reconstruction time

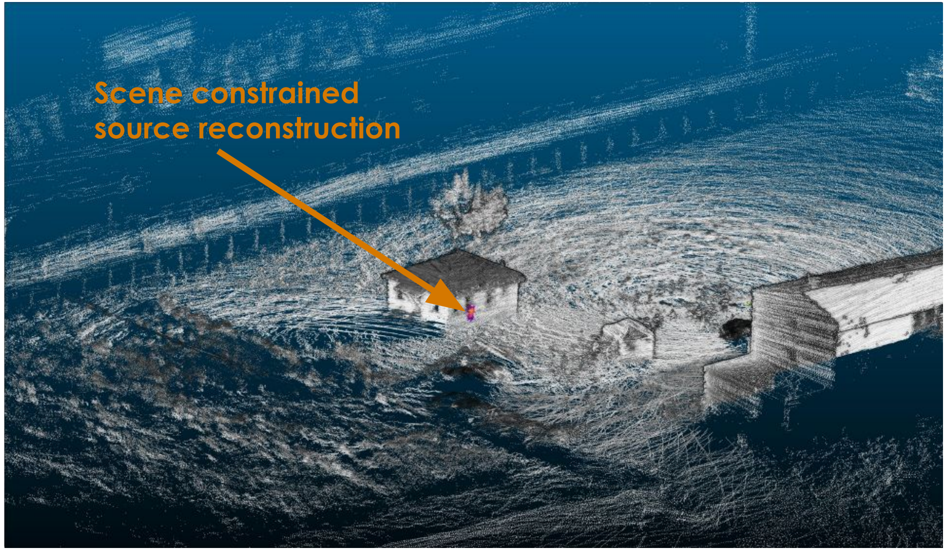

Nuclear Scene Data Fusion (SDF)

- Integrate auxiliary contextual sensors and onboard computer with radiation detection/imaging systems
- Perform SLAM in real-time (~10 Hz) to map the 3D scene and track the system pose as it moves freely through the environment
- Scene data fused with radiation data
  - Time synchronization
  - Coordinate transforms
- Real-time poses enable free-moving 3D imaging and real-time map provides a dynamic, and constrainable, image space
  - Map constraint can increase image accuracy, reduce noise, decrease reconstruction time

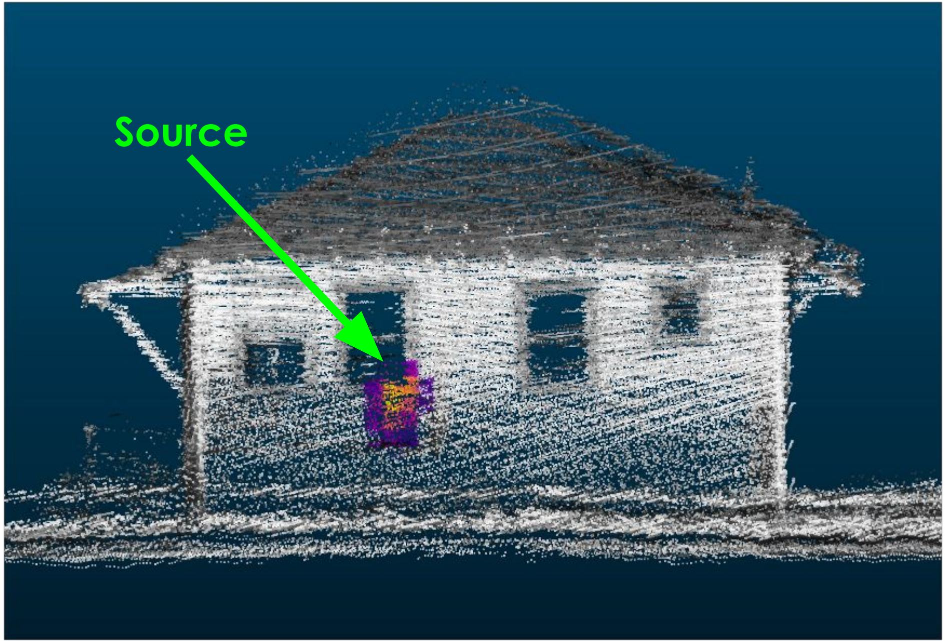

BERKELEY LAB

Nuclear Scene Data Fusion (SDF)

- Integrate auxiliary contextual sensors and onboard computer with radiation detection/imaging systems
- Perform SLAM in real-time (~10 Hz) to map the 3D scene and track the system pose as it moves freely through the environment
- Scene data fused with radiation data
  - Time synchronization
  - Coordinate transforms
- Real-time poses enable free-moving 3D imaging and real-time map provides a dynamic, and constrainable, image space
  - Map constraint can increase image accuracy, reduce noise, decrease reconstruction time

Part 1

Machine Learning for detection and identification of radioactive material

Biggest challenge: variable backgrounds

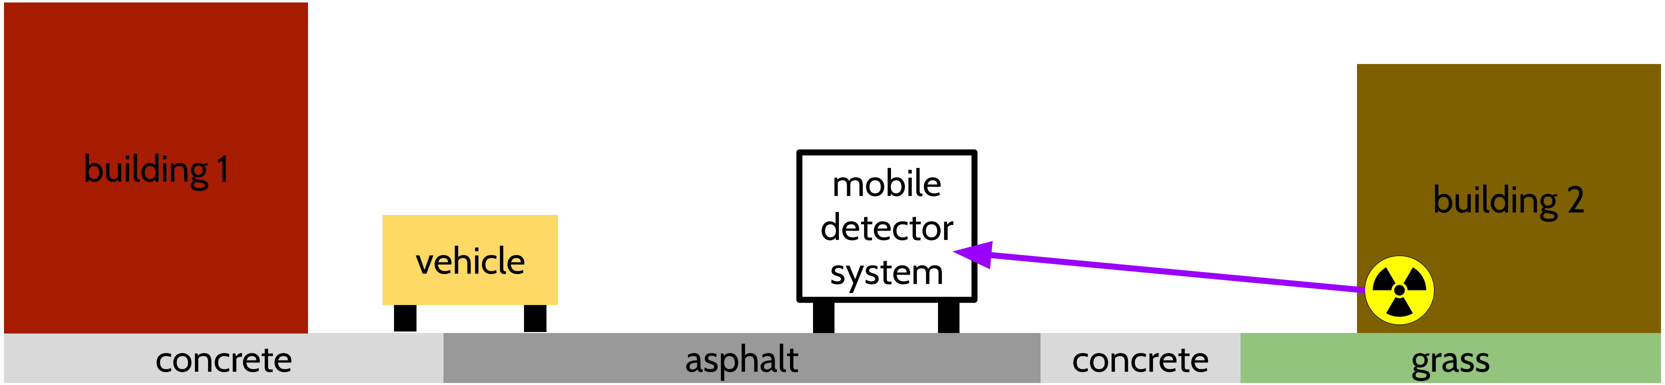

Biggest challenge: variable backgrounds

- Terrestrial – Trace, variable levels in building materials
  - ^{40}K, Uranium-series, Thorium-series
- Airborne – Modulated by weather and rain
  - ^{222}Rn and progeny (mostly from washdown during rain events)
- Cosmogenic – Modulated by space weather and buildings
  - Neutrons, Muons
- Artificial – Variable with human factors
  - Industrial, medical

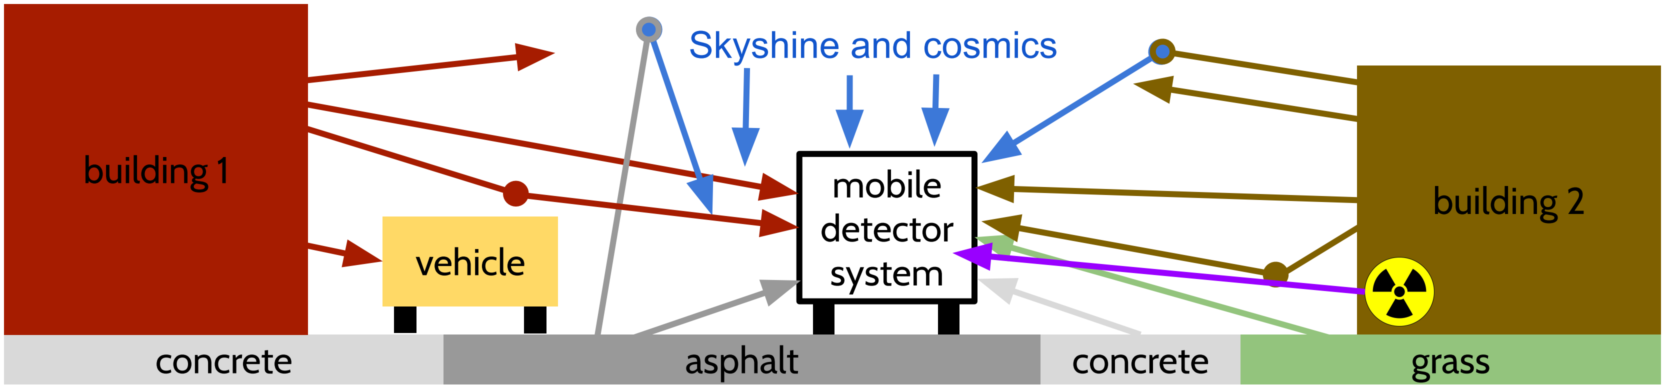

Signals in radiation detectors

^{137}Cs – Nuisance or threat

- Characteristic emissions
- Gamma-ray scattering in shielding
- Propagation through air
- Energy deposition in a detector
- Detector resolution

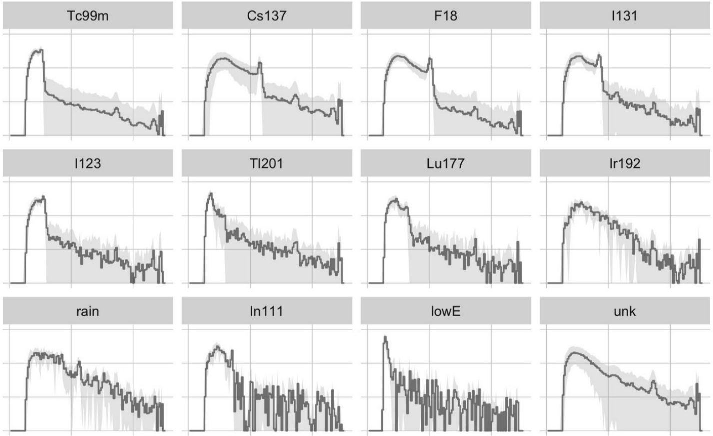

https://arxiv.org/abs/2104.04137

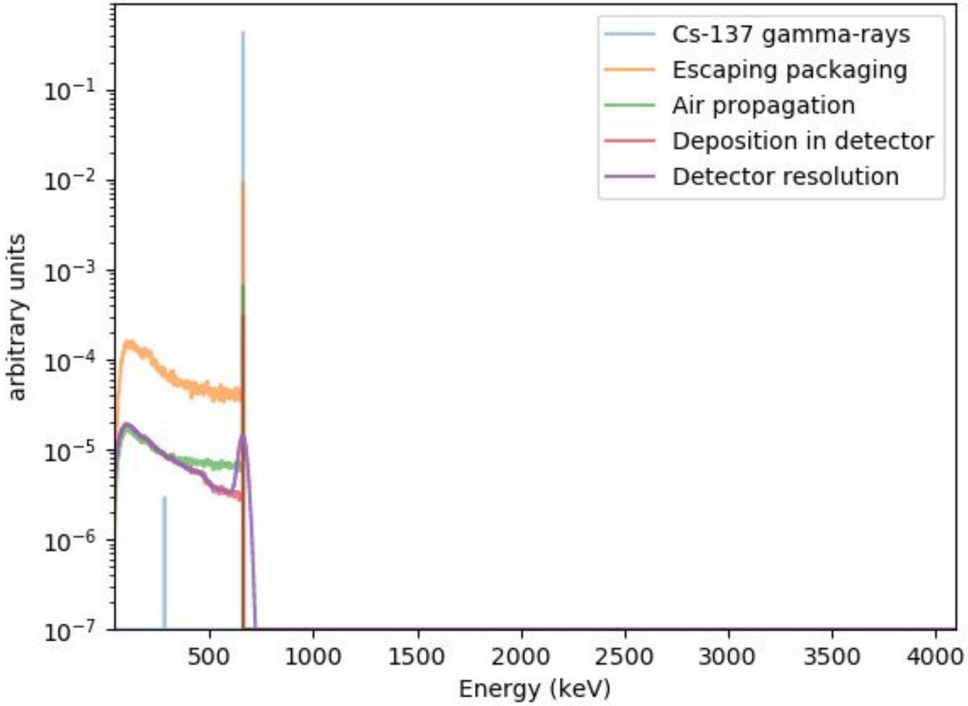

Signals in radiation detectors

- Also backgrounds undergo scattering in air and broadening of the resolution
- Scattering + diverse origin of backgrounds leads to a complex spectrum shape
- We are looking for a small signal at low energy…

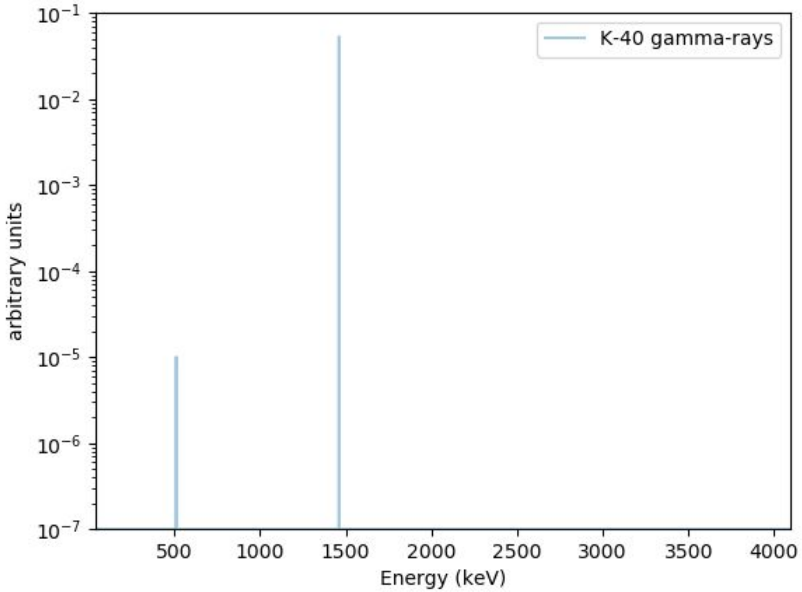

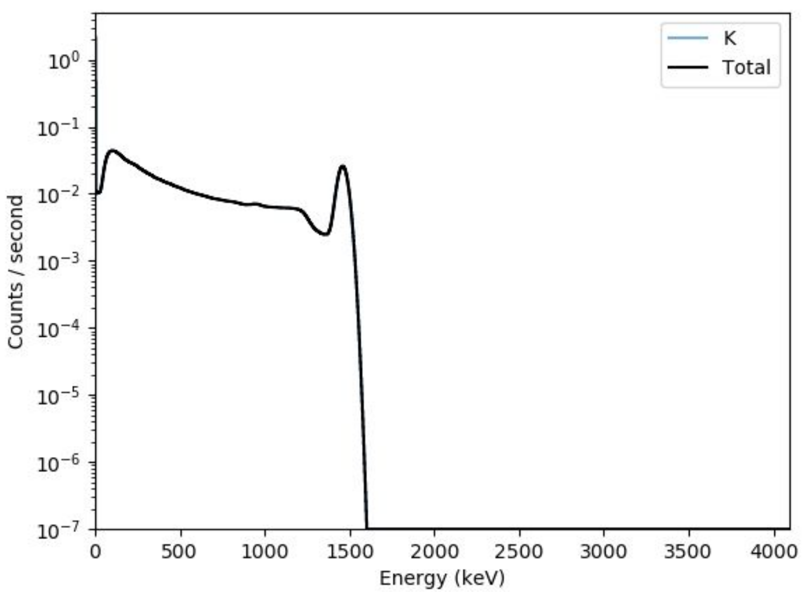

What do algorithmic approaches look like?

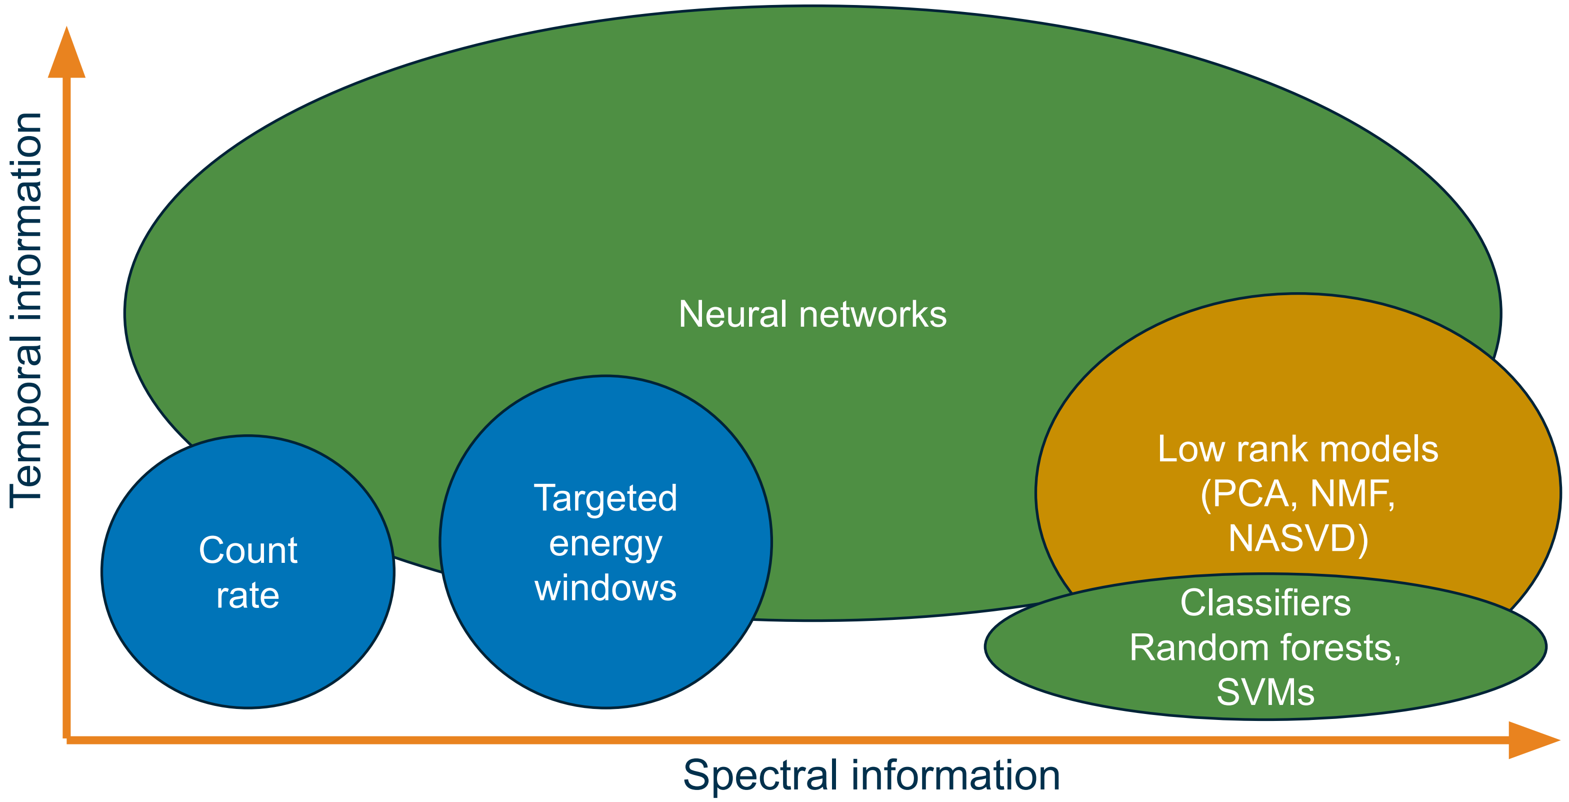

This is far from exhaustive

BERKELEY LAB

The Urban Radiological Search Competition

- Competitors are provided with
  - A training set of list mode data for ~10k runs, each run lasting 1-3 minutes.
  - A test set of ~16k runs: 43% public, 57% private.
  - Energy spectra for each source type.
- For each run in the test set, competitors must
  - Detect whether there is an anomalous source.
  - Identify the type of source.
  - Locate when the detector is closest to it.
- Two Competitions
  - .gov – National Laboratories
  - TopCoder – Open to the public with cash prizes
- Leaderboard scoring
  - Pre-specified scalar that combines all desired aspects of solution
- Data handling:
  - Training data: Input + answer provided
  - Public test data: Competitors receive feedback during competition
  - Private test data: Used to determine final standings

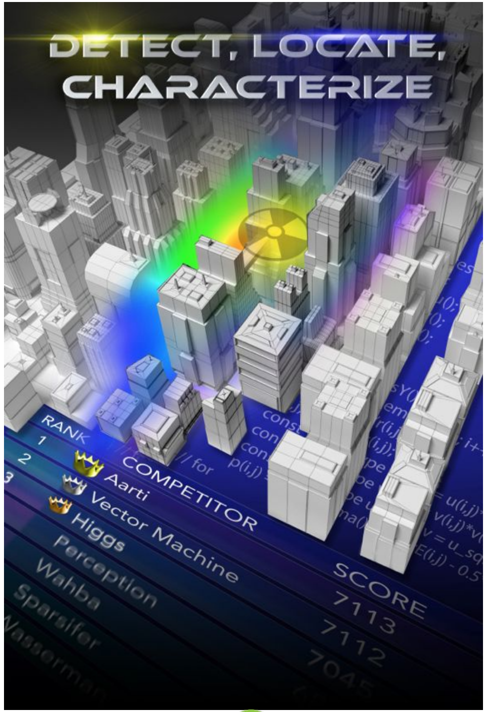

topcoder

https://www.topcoder.com/lp/detect-radiation

BERKELEY LAB

Data Generation

- Realistic background variations in K, Th, U
  - Developed a large 3D modular Monte Carlo model loosely based on Gay St., Knoxville, TN
  - 8 instances by arranging seven interchangeable blocks
  - Each block + material combo, was simulated independently so that material activates could be modified block by block
- Simulation of sources and backgrounds
  - Calculated background and source detector response functions for a 2”x4”x16” NaI(Tl) detector moving through one of four lanes of traffic (1 to 13.4 mph)
  - ORNL Monaco/MAVRIC Monte Carlo
  - 6 different sources at 15 locations with different offsets.
- Full knowledge of data set.

  No unknowns in the data.

1. HEU: Highly enriched uranium
2. WGPu: Weapons grade plutonium
3. ^{131}I: Iodine, a medical isotope
4. ^{60}Co: Cobalt, an industrial isotope
5. ^{99m}Tc: Technetium, a medical isotope
6. A combination of HEU and ^{99m}Tc

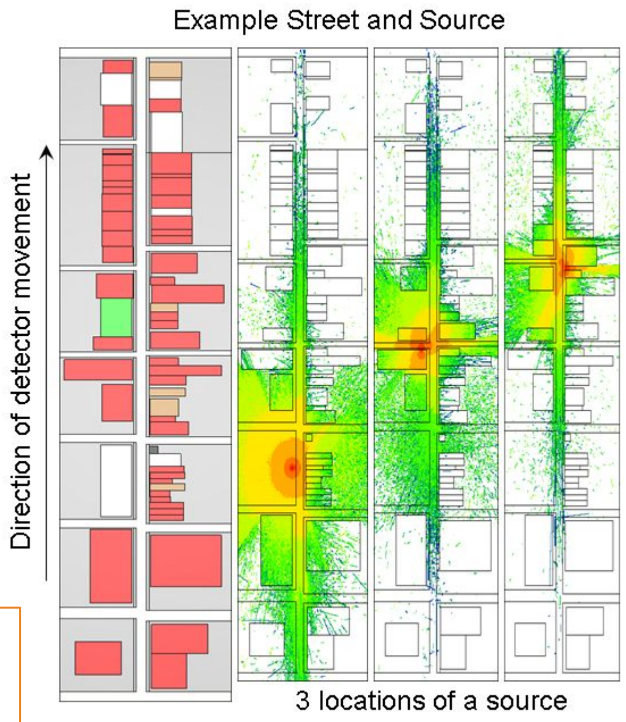

OAK RIDGE National Laboratory

Anomaly Identification with NMF (winner .gov)

- Factorize data X into weights A and components V

$$\boldsymbol{X} \approx \hat{\boldsymbol{X}} = \boldsymbol{A}\boldsymbol{V}$$

$$\underbrace{\overset{\text{Measured}}{\begin{bmatrix} \boldsymbol{x}_0^T \\ \boldsymbol{x}_1^T \\ \vdots \\ \boldsymbol{x}_n^T \end{bmatrix}}}_{\boldsymbol{X}\,(n\times m)} \approx \underbrace{\overset{\text{Weights}}{\begin{bmatrix} a_{00} & a_{01} & \dots & a_{0k} \\ a_{10} & a_{11} & \dots & a_{1k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n0} & a_{n1} & \dots & a_{nk} \end{bmatrix}}}_{\boldsymbol{A}(n\times k)} \cdot \underbrace{\overset{\text{Components}}{\begin{bmatrix} \boldsymbol{v}_0^T \\ \boldsymbol{v}_1^T \\ \vdots \\ \boldsymbol{v}_k^T \end{bmatrix}}}_{\boldsymbol{V}\,(k\times m)}$$

- Dimensionality reduction (ala PCA)
  - Non-negative
  - Components are additive
  - Components are not orthogonal
  - Number of components identified at outset
  - Consistent with Poisson statistics
- Lee and Seung, “Learning the parts of objects by non-negative matrix factorization,” Nature, 401: 788–791 (1999)

Anomaly Identification with NMF (winner .gov)

- Physics informed methodology
  - Learned templates from data
  - Experimented with multi-base templates to model shielding and scatter
- Empirical thresholding
- Limitations
  - Single integration time
  - Time independent implementation

- LBNL contribution to competition
- 1st place in the National Laboratory competition
- 13th place in the TopCoder competition
  - Outperformed by algorithms that included temporal information

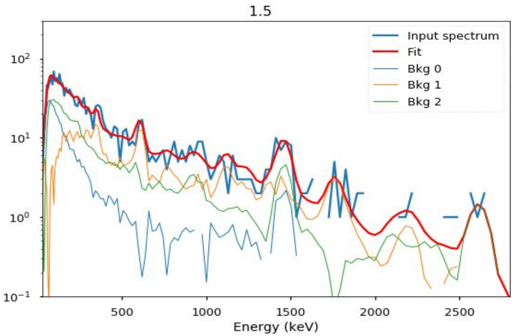

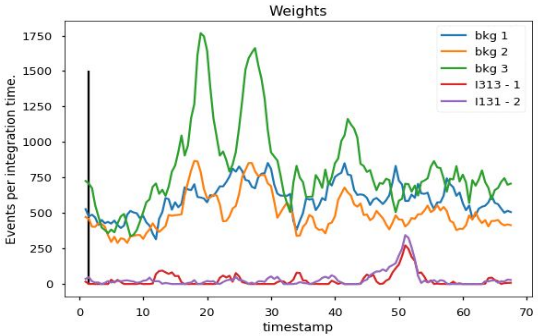

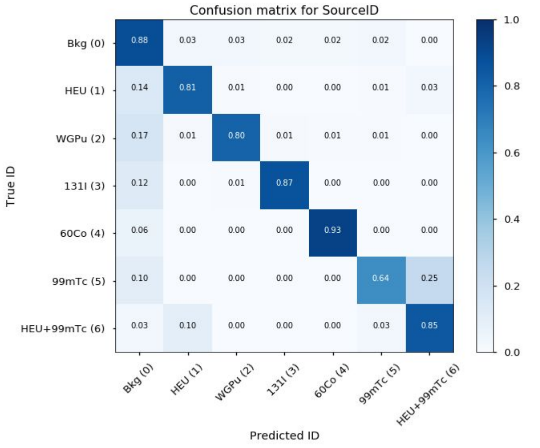

Anomaly detection with NNs (winner topcoder)

- Multi-point lead
- Binned energy approach, at different time scales
- Ensemble of neural networks
- Cross validation to avoid overfitting
- Augmented learning (add noise)
- Data driven training (no use of source templates)

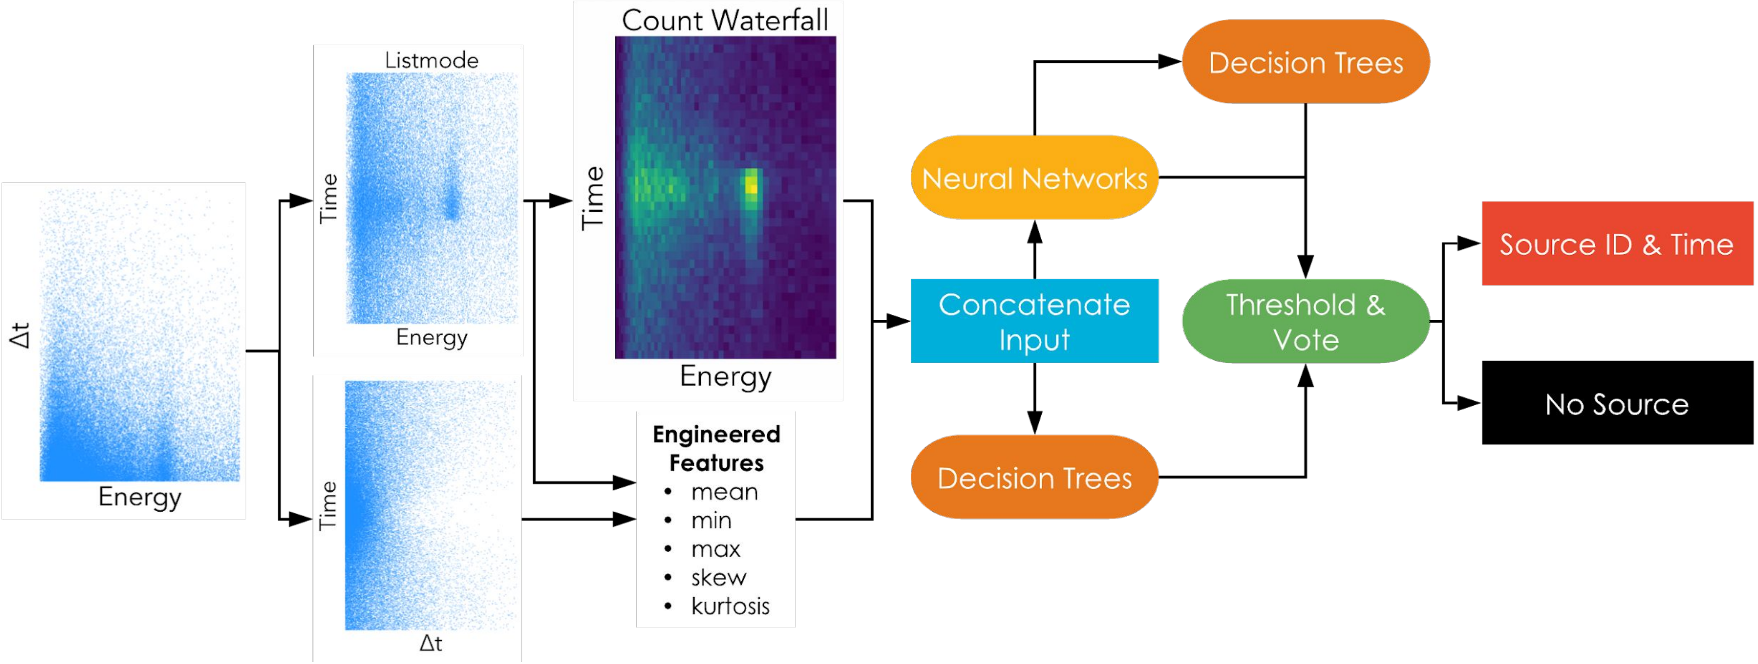

Anomaly detection with NNs (winner topcoder)

- Multiple time scales, passed through identical neural network
- Combine results with a Softmax
- Threshold to decide if source is present (source id + time) or not

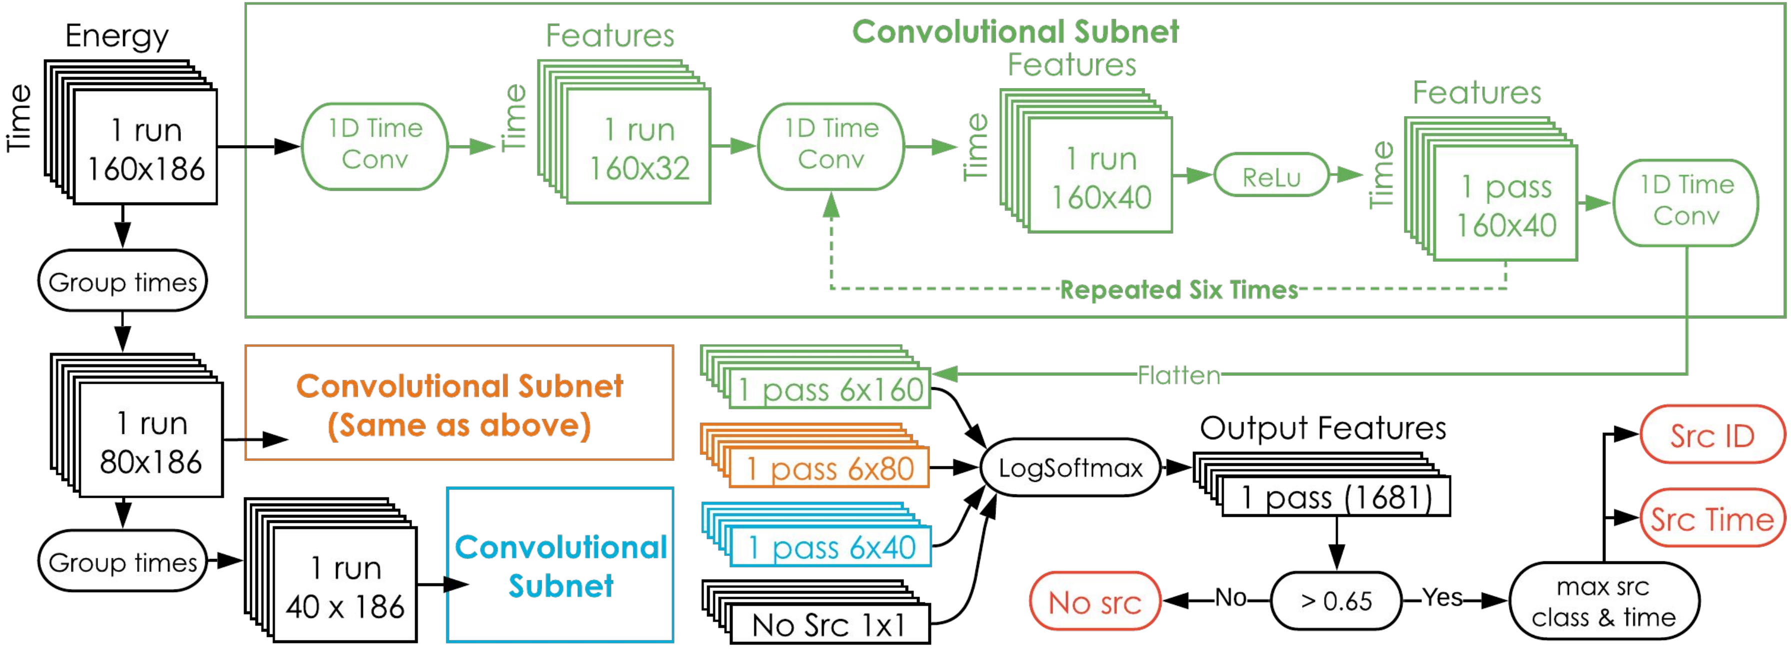

Anomaly detection with NNs (winner topcoder)

- Combine many identical networks (N experts are better than 1)
- Ensembling reduces variance without cost to bias
- Price: Increased computation / code complexity

- 10 training data splits (90% each)
- 3 repetitions per split
- Each network trained with down-sampled spectra which were randomly reversed in time

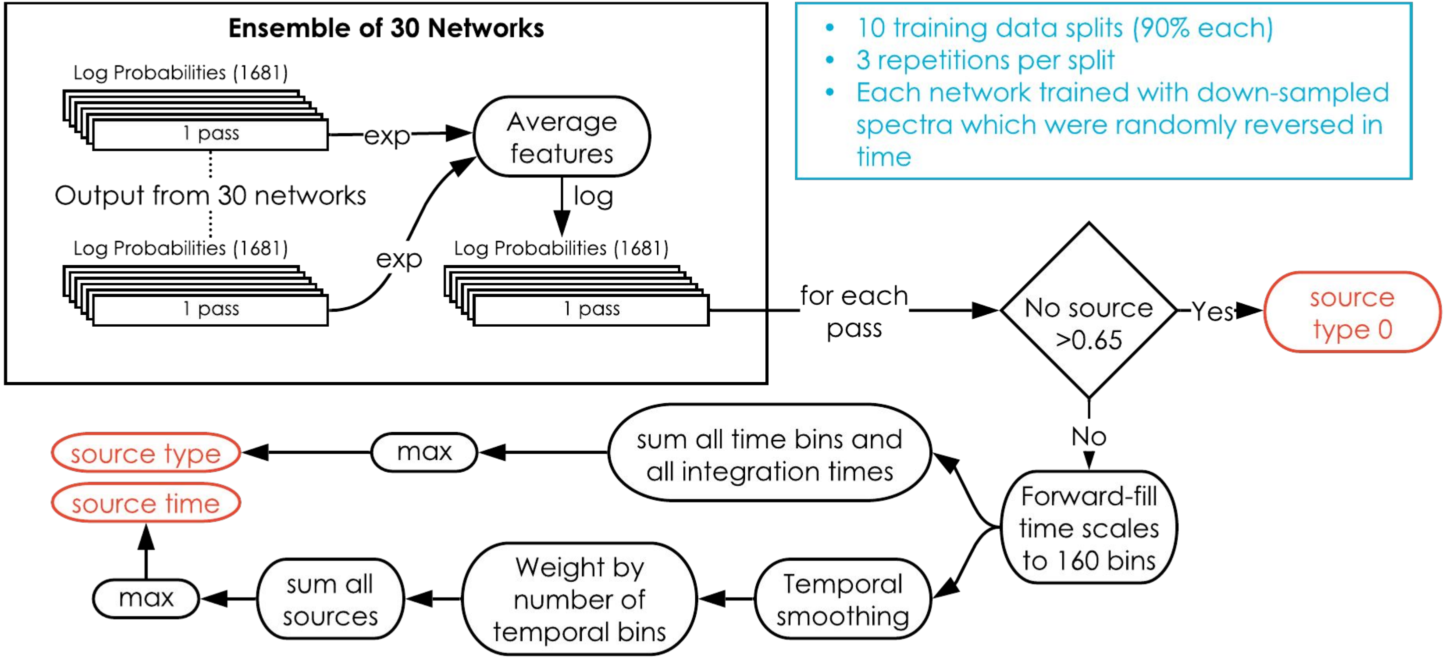

Conclusion and outlook

- Physics, instrumentation, and human considerations combine to make radiological search a challenging task
- Advanced algorithms have shown promise for improving sensitivity
  - Better performance with the same hardware
  - Same performance with cheaper hardware
  - Enable new concepts (networks, data fusion)
- Algorithm design must consider
  - Perform analysis in real time
  - Produce instantaneous results
  - Allow to be configured with meaningful quantities such as false positive rate
- Introspection and interpretable of NNs

Follow up project

- Goal: Build a community standard benchmark dataset for radiological search and use it to evaluate fieldable machine learning algorithms against literature benchmarks
  - More diversity (rural area, bridges over rivers, tunnels, etc)
  - More realistic (include environmental effects such as rain, etc.)
  - More sources
  - Different detector types
- Better understanding of algorithm and enable fieldability of algorithms
  - Develop an open source package: https://gitlab.com/lbl-anp/radai/radai

BERKELEY LAB

Team and references

- LBNL:
  - Tenzing H. Joshi, Joseph Curtis, Mark Bandstra,Reynold Cooper, Brian Quiter,Daniel Hellfeld
- UC Berkeley:
  - Kyle Bilton, Kai Vetter
- LANL:
  - Christine Anderson-Cook, Kary Myers
- ORNL:
  - Andrew Nicholson, Douglas Peplow, James Ghawley, Dan Archer

- Support:
  - NA-22 Data Science – Radiation Detection Data Competition
  - NA-22 Near-field -- PANDA

Sources:

- https://gitlab.com/lbl-anp/radai/radai
- Technical report: https://doi.org/10.2172/1778748
- Data generation: https://doi.org/10.1038/s41597-020-00672-2

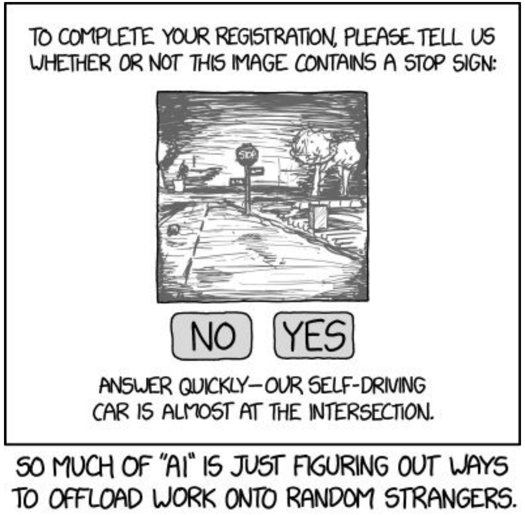

Part 2

Semantic segmentation to support radiation detection

Model naturally occurring materials in urban scenes

- Model radioactive background with contextual sensor data
  - Focus on a simple urban mock facility with known radioisotope composition
  - Build a three dimensional model of the surrounding that includes the most crucial features
  - Include energy-dependence though modeling radioisotope spectrum, providing access to activities
  - MLEM to attribute radiological measurements to surroundings

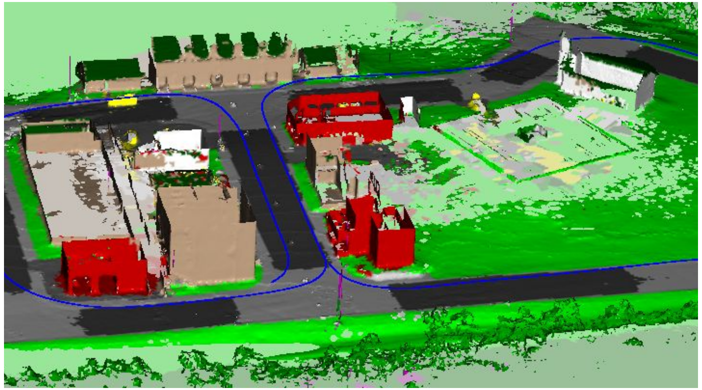

Linear system:

$$\boldsymbol{\lambda} = \boldsymbol{R}\boldsymbol{f}$$

Maximum Likelihood Maximization Estimation for solving system

Model naturally occurring materials in urban scenes

- Inversion problem
  - Predict radiation and it's transport by classifying visible surfaces as seen from the detector system
  - Build a system of linear equations (system response) to solve for the unknown gamma-ray flux from various surfaces
- System response
  - 3D description of the facility (distance and material class)
  - Effective area (detector efficiency and geometry) and description of gamma-ray transport in air
  - NORM modelling for complexity reduction originating from energy dependence of radiological data

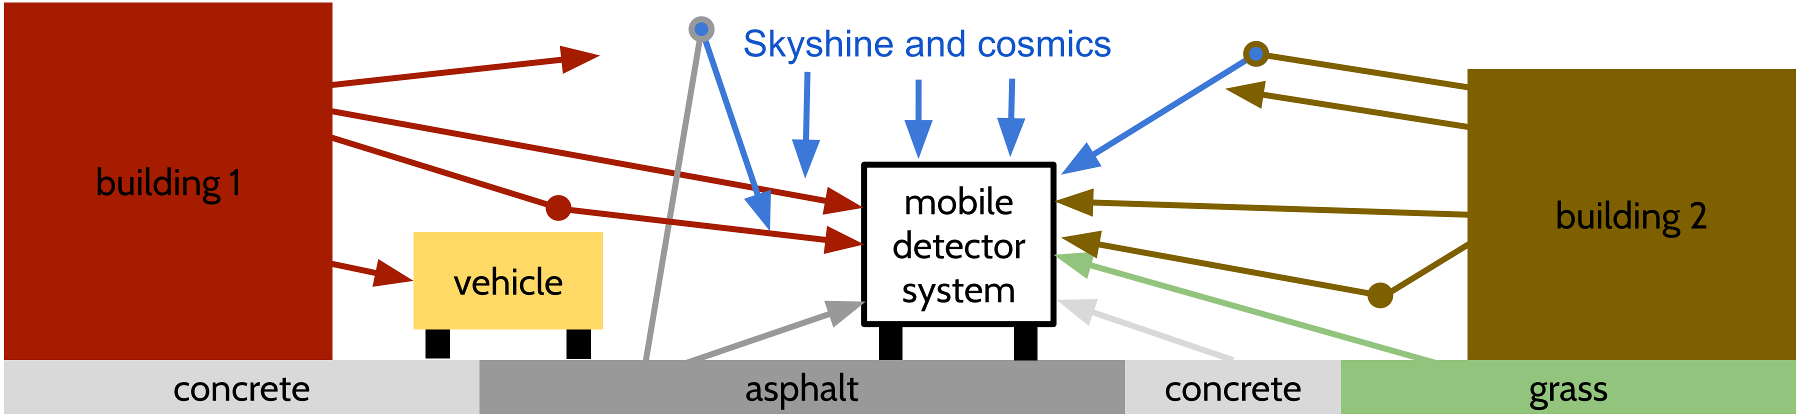

Model naturally occurring materials in urban scenes

- Inversion problem
  - Predict radiation and it's transport by classifying visible surfaces as seen from the detector system
  - Build a system of linear equations (system response) to solve for the unknown gamma-ray flux from various surfaces
- System response
  - 3D description of the facility (distance and material class)
  - Effective area (detector efficiency and geometry) and description of gamma-ray transport in air
  - NORM modelling for complexity reduction originating from energy dependence of radiological data

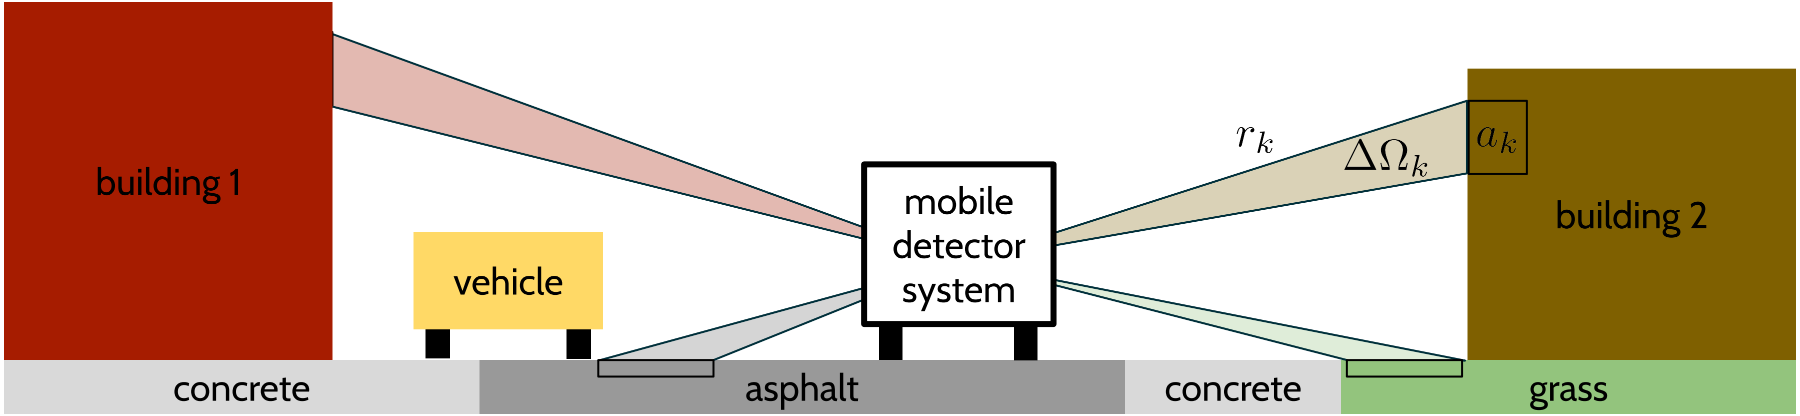

Building a 3D description of the facility

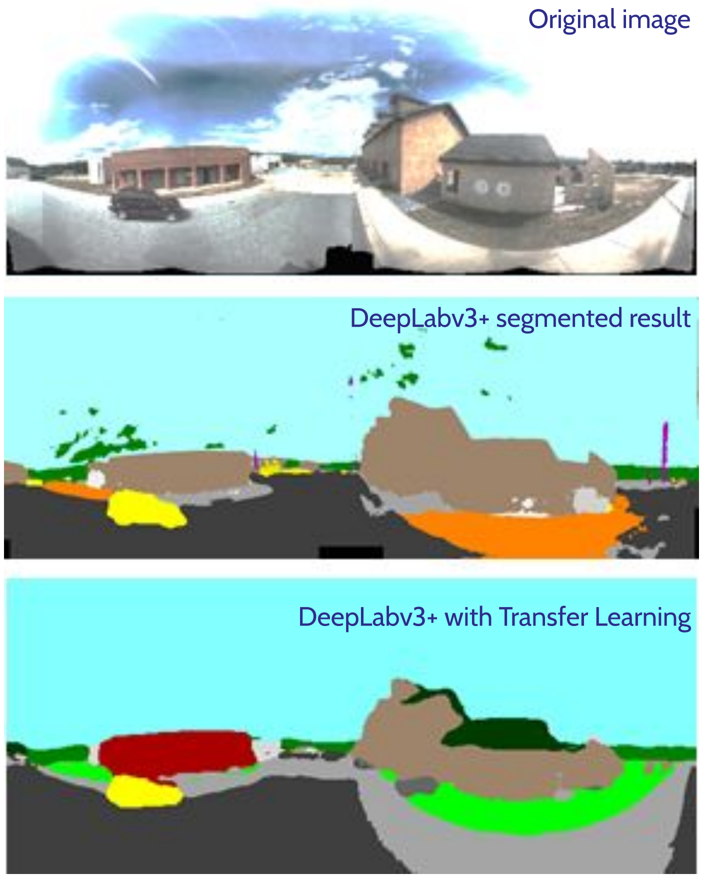

- Segmentation and classification of images
  - Used Google's Deep Labelling for Semantic Image Segmentation (DeepLabv3+) model on pre-trained Cityscapes1 dataset
- Applied transfer learning by retraining last, fully-connected neural layer with 45 hand-labeled images to be closer to ground truth labels:
  - Asphalt
  - Building red
  - Building brown
  - Building white
  - Building roof
  - Concrete
  - Forest
  - Grass
  - Gravel
  - Sky
  - Vehicle

Cityscapes dataset available at: https://www.cityscapes-dataset.com
L.Chen, et. al., Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation, ECCV, 2018, https://github.com/tensorflow/models/tree/master/research/deeplab

Building a 3D description of the facility

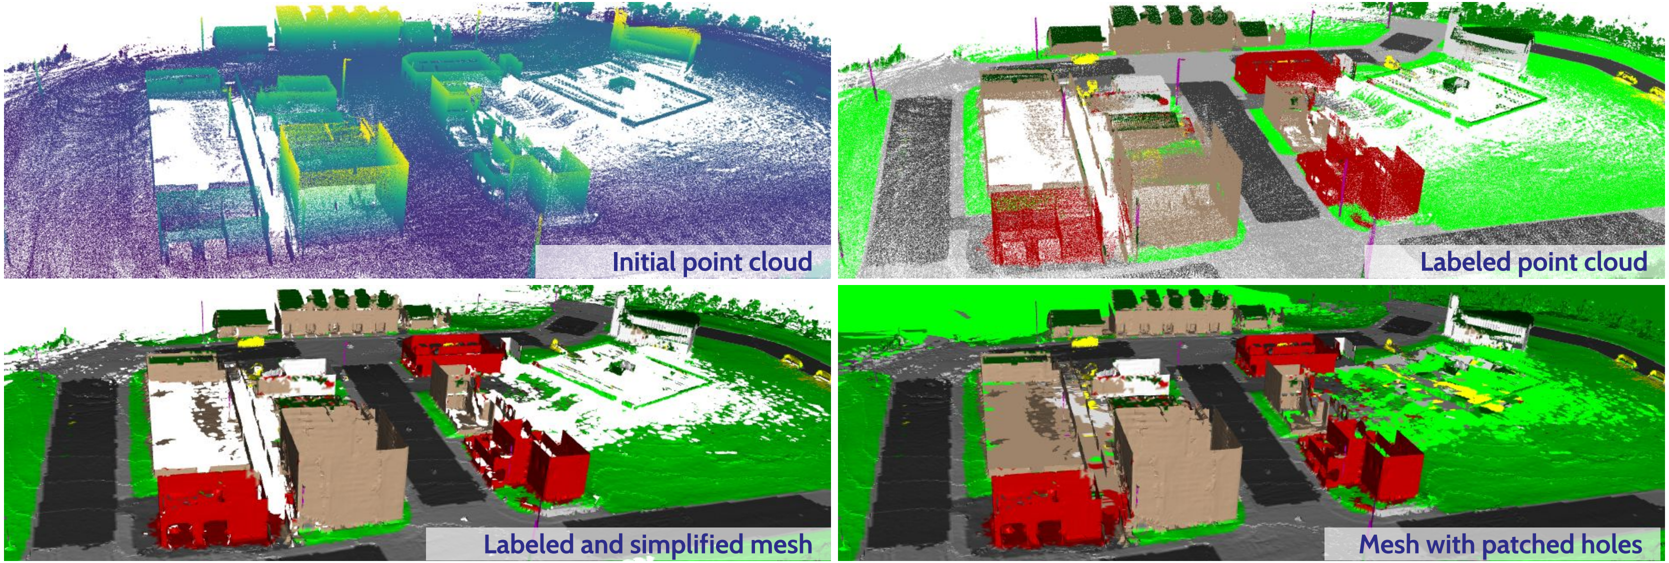

- Projecting labeled images back to point cloud and pick the label that is observed most often at each point.
- Convert labeled point cloud into a triangular mesh (based on ball pivoting algorithm with smart normal orientation algorithm)
- Simplify mesh to reduce number of vertices by a factor of ~10
- Remaining holes are patched using nearest neighbor interpolation and extending to a flat horizon

S. Katz, A.Tal, and R. Basri, Direct visibility of point sets, ACM Trans, Graph. 26, 3, Article 24, 2007
Q. Zhou, J. Park, V. Koltun, Open3D: A Modern Library for 3D Data Processing, arXiv:1801.09847, 2018, http://www.open3d.org

Building a 3D description of the facility

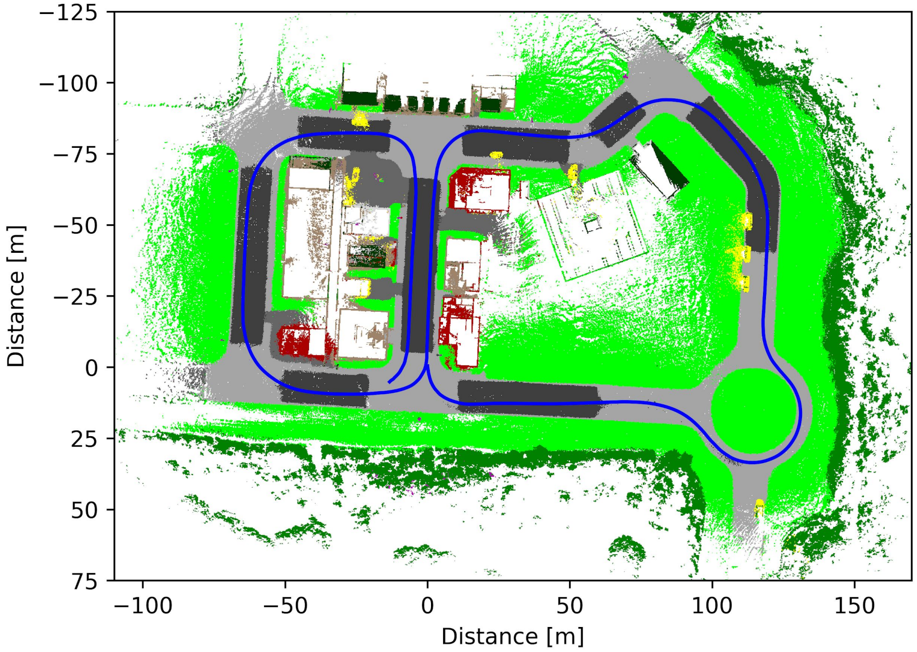

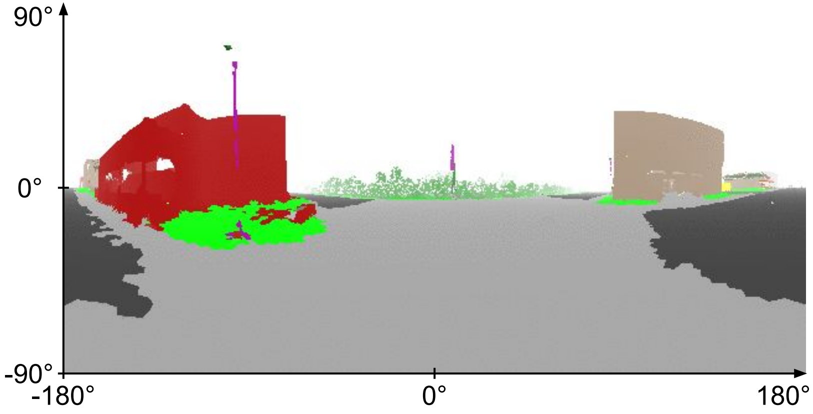

- The distance and material class of all the surfaces in the field of view of each detector can be calculated at every time step
- Visualization of panoramic view of mesh from detector array center
- Alpha channel is distance between 0 (transparent) and 80 meter (white)

Results of background modeling

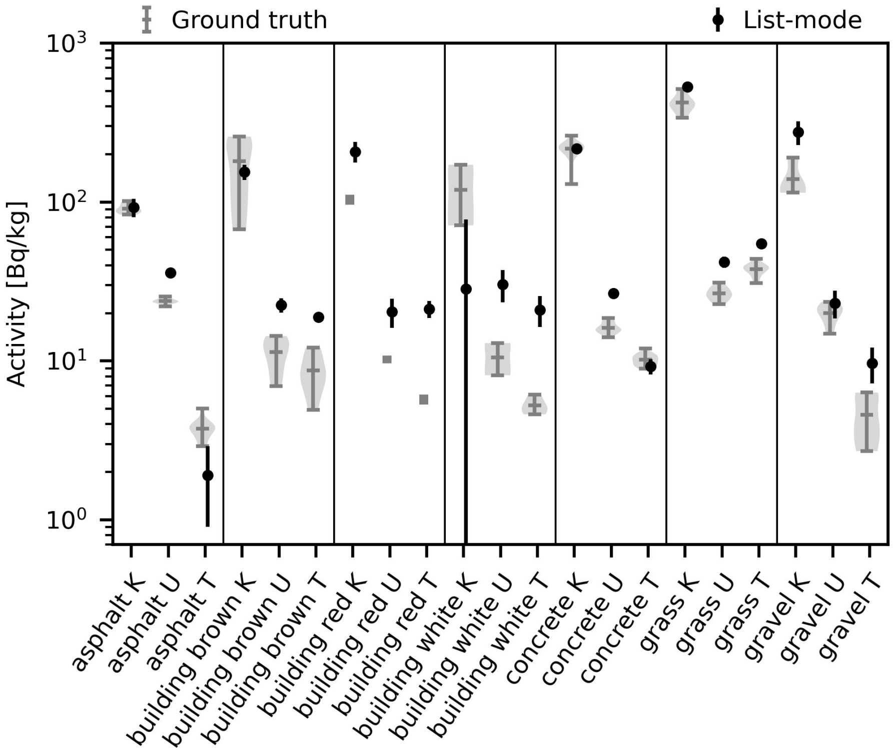

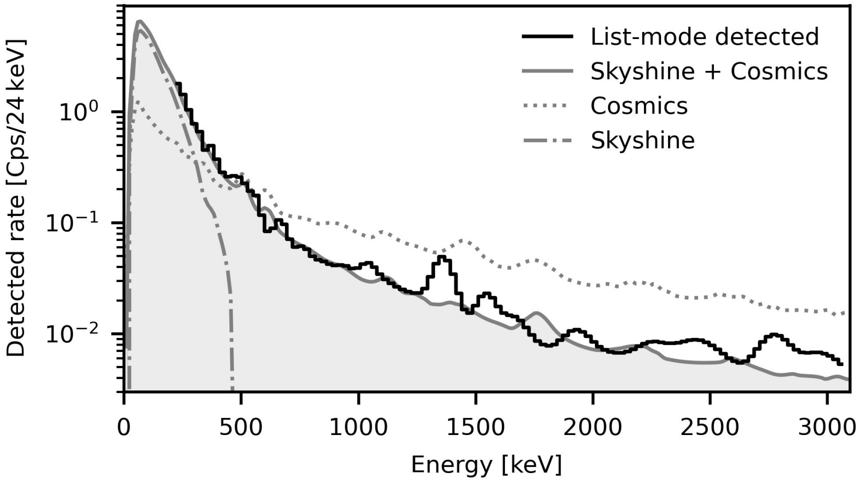

M. W. Swinney, et al., A methodology for determining the concentration of naturally occurring radioactive materials in an urban environment. Nuclear Technology, 203(3):325-335, 2018.

A. L. Mitchell, et al., Skyshine contribution to gamma ray background between 0 and 4 MeV. Technical report, Pacific Northwest National Lab. (PNNL), August 2009.

G. A. Sandness, et al., Accurate modeling of the terrestrial gamma-ray background for homeland security applications, 2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC), Orlando, FL, USA (IEEE, Piscataway, NJ, 2009), pp. 126–133.

Results of background modeling

- Our model describes the observed background fluctuations well
- Features seen in the exposure to distinct classes coincides with features seen in the background
- Searches for radioactive sources outside of regulatory control can benefit from background modeling and prediction based on contextual sensor data

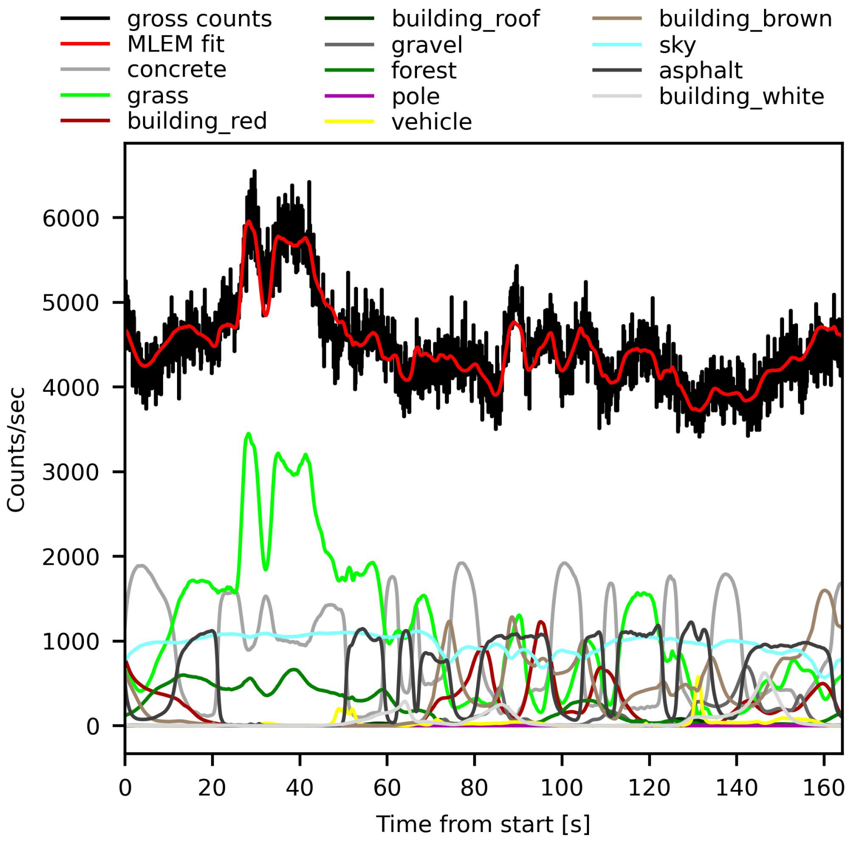

Team and references

- Team (all LBNL)
  - Marco Salathe
  - Brian J. Quiter
  - Mark S. Bandstra
  - Joseph C. Curtis
  - Ross Meyer
  - Chun Ho Chow
- References

  - M. Salathe, B. J. Quiter, M. S. Bandstra, J. C. Curtis, R. Meyer, and C. H.Chow, "Determining urban material activities with a vehicle-based multi-sensor system", Phys. Rev. Research 3, 023070, 2021
  - M. S. Bandstra, et al., Attribution of gamma-ray background collected by a mobile detector system to its surroundings using panoramic video, NIMA 954, 161126, 2020
  - M. S. Bandstra, et al., "Correlations between Panoramic Imagery and Gamma-Ray Background in an Urban Area", 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (2019), pp. 1–5
  - M. Salathe, et al., "Using 3D-Scene Data from a Mobile Detector System to Model Gamma-Ray Backgrounds", 2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (2019), pp. 1–4

Semi-automated Scene Generation for Diagnostics

- Scene capture and digitization for Nuclear Emergency Response missions
- Object of interest found, next step is answering “What is it?” and “What to do about it?”
  - γ and neutron detectors/imagers/ multiplicity counters & X-ray / API active interrogation systems deployed
- Scene digitization
  - Automatically provides record of what/where/when
  - Can automatically produce geometry portion of Monte Carlo input files
  - Helps with ensuring events (and exercises) are documented

Iphone 12 Max Pro

Semi-automated Scene Generation for Diagnostics

- Semantic segmentation of colorized point cloud
  - SparseConvNet trained on ScanNet
  - http://www.scan-net.org/ScanNet/
  - https://github.com/facebookresearch/SparseConvNet
- Semantic segmentation of images and projection to point cloud
  - DeepLab3+ trained on ADE20k
  - https://groups.csail.mit.edu/vision/datasets/ADE20K/
  - https://github.com/tensorflow/models/tree/master/research/deeplab
- Work in progress
  - Create CAD model for simulation
  - Tracking of objects and detectors during radiation measurements

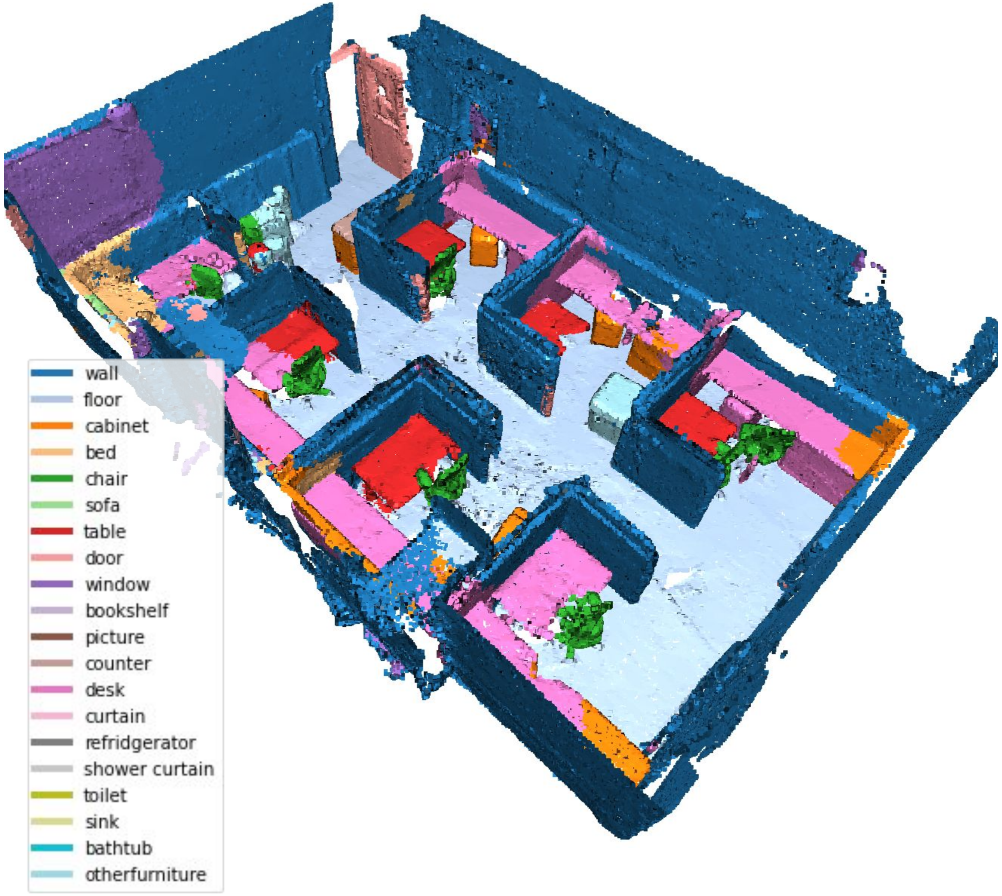

Semi-automated Scene Generation for Diagnostics

- Semantic segmentation of colorized point cloud
  - SparseConvNet trained on ScanNet
  - http://www.scan-net.org/ScanNet/
  - https://github.com/facebookresearch/SparseConvNet
- Semantic segmentation of images and projection to point cloud
  - DeepLab3+ trained on ADE20k
  - https://groups.csail.mit.edu/vision/datasets/ADE20K/
  - https://github.com/tensorflow/models/tree/master/research/deeplab
- Work in progress
  - Create CAD model for simulation
  - Tracking of objects and detectors during radiation measurements

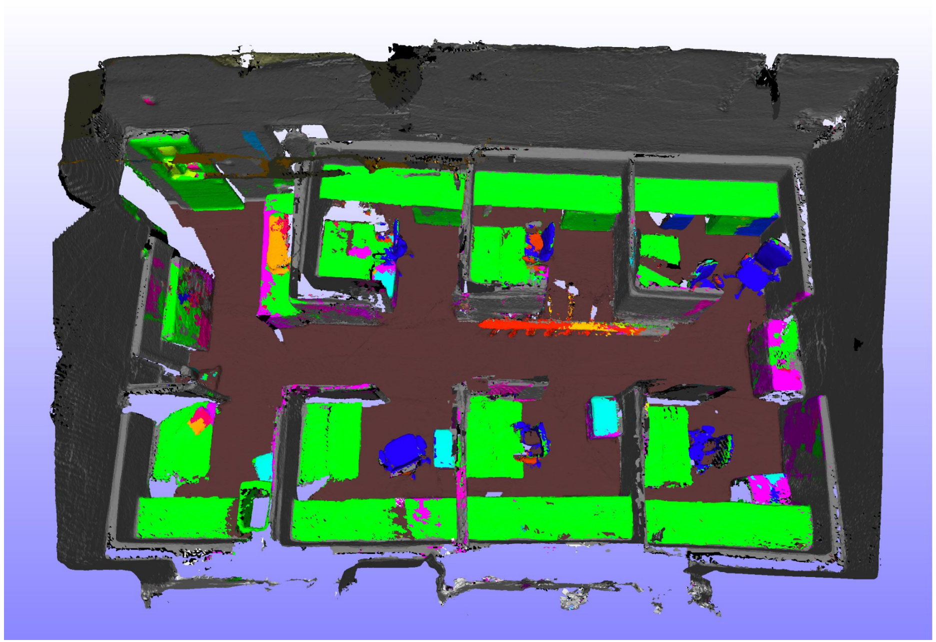

Team

LBNL

- Brian Quiter
- Marco Salathe
- Mark Bandstra
- Dan Hellfeld
- Xin Chen

ORNL

- Dan Archer
- Christopher Greulich
- Mathew Swinney

BERKELEY LAB

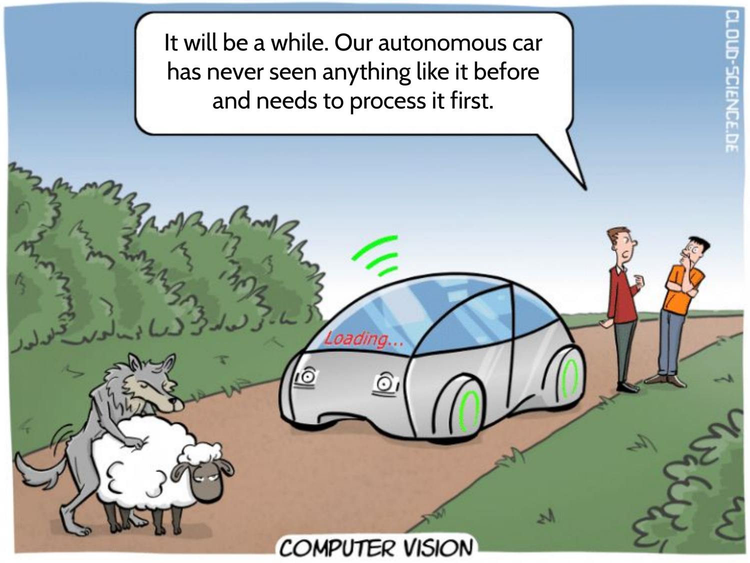

Part 3

Object detection and radiation attribution

Object detection, tracking and attribution

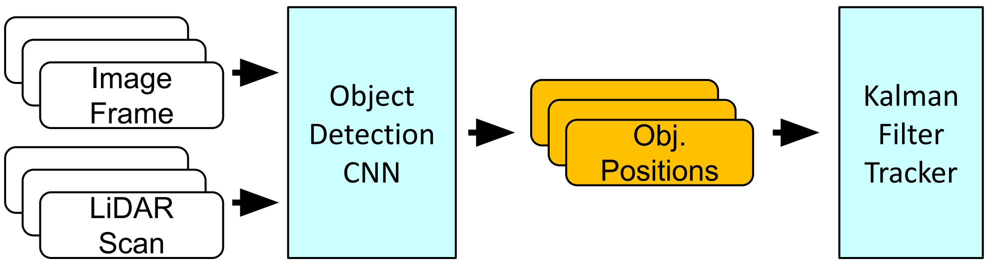

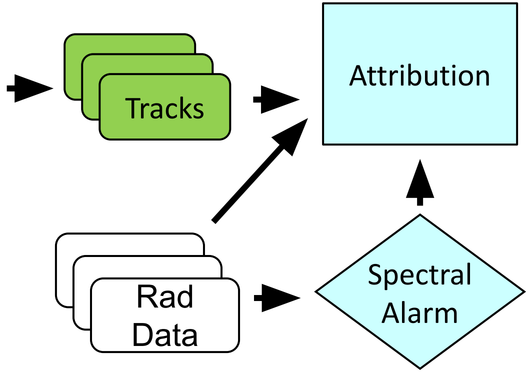

Detection and Tracking Pipeline:

1. Perform object detection on images or LiDAR with Convolutional Neural Networks
2. Match detections to those from earlier data through predictions from a Kalman Filter.
3. Search for tracks that describes radiation data best (Attribution)

Improved situational awareness and anomaly detection performance

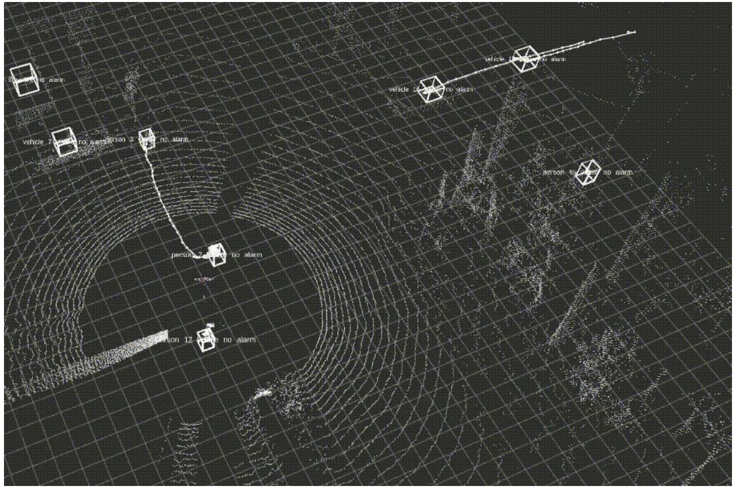

M. R. Marshall et al., "3-D Object Tracking in Panoramic Video and LiDAR for Radiological Source–Object Attribution and Improved Source Detection," doi: 10.1109/TNS.2020.3047646.

Object detection

Trade-off between performance and speed

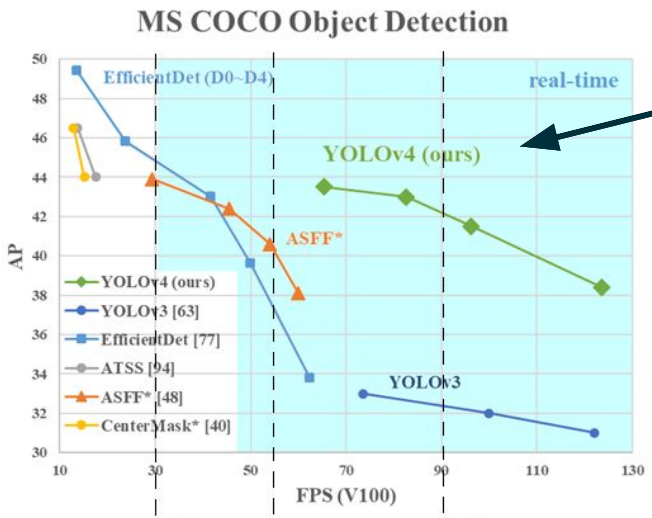

real-time (30 FPS) on
TitanV/TeslaV100 (server)
250 Watt
$2250 - $8000

real-time (30 FPS) on
RTX 2070 (desktop)
175 Watt
$450

real-time (30 FPS) on
RTX 2060 MXM (Mobile)
80 Watt
$700

9 FPS on
Jetson AGX Xavier (embedd
30 Watt
$700

Image based

A. Bochkovskiy, et al, "YOLOv4: Optimal Speed and Accuracy of Object Detection," arXiv:2004.10934v1

LiDAR based

A. H. Lang, et. al, "PointPillars: Fast Encoders for Object Detection From Point Clouds," doi:10.1109/CVPR.2019.01298

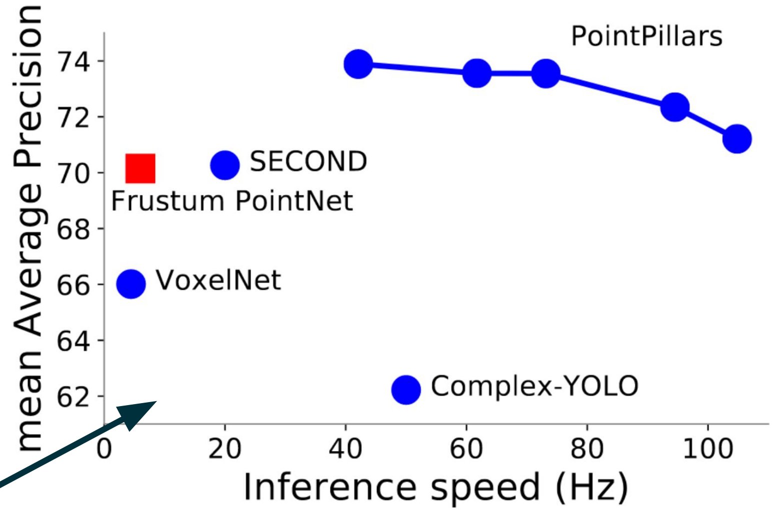

Object detection in real time on a NVIDIA Jetson NX

- Yolov4-tiny (2020) for image based detection
  - Use TensorRT for increased performance
  - ~45 images / second
- PointPillars (2018) for LiDAR based detection
  - PyTorch based implementation with speed tweaks
  - ~10 scans / second

Object detection

YOLOv4-tiny example at Richmond Field station

PointPillars example at Lawrence Berkeley National Laboratory

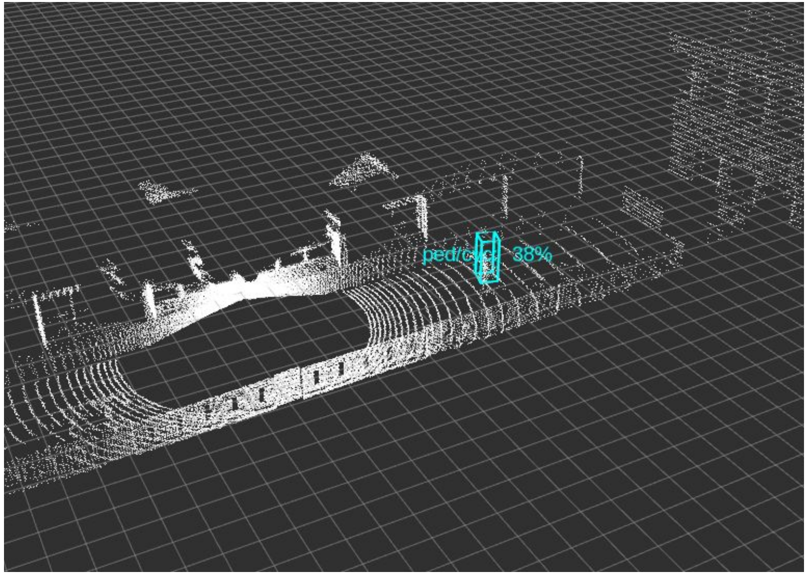

Object tracking

Simple Online and Realtime Tracking (SORT)

- Based on bounding boxes in image coordinates
- Kalman filter used to predict the next likely location of a given bounding box
- Intersection over Union comparison

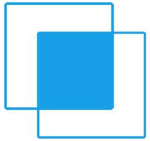

$$\text{IoU} = \frac{\text{Area of Overlap}}{\text{Area of Union}}$$

- Additional conversion to 3D required

A. Bewley, et. al, "Simple online and realtime tracking," doi: 10.1109/ICIP.2016.7533003.

Extensions

- 3D position (x, y, z) in Kalman Filter used
- Works with LiDAR
- Multivariate Normal: 3x3 matrix used to encoding x, y and z uncertainties and orientation
- Hellinger distance as matching criteria

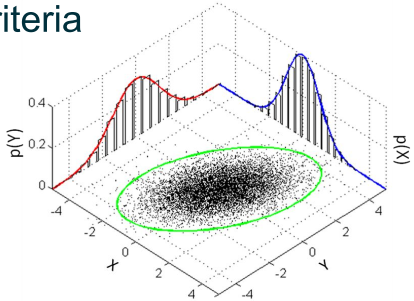

M. R. Marshall et al., "3-D Object Tracking in Panoramic Video and LiDAR for Radiological Source–Object Attribution and Improved Source Detection," doi:10.1109/TNS.2020.3047646.

FastMOT

- Omni-scale network (OSNet) feature vectors CNN for object re-identification
- Optical flow for tracking between detections (allows for slower more performant object detection network)
- Limited to images

Y. Yang, "FastMOT: High-Performance Multiple Object Tracking Based on Deep SORT and KLT", doi:10.5281/zenodo.4294717
Kaiyang Zhou, et. al, "Omni-Scale Feature Learning for Person Re-Identification", arXiv:1905.00953

Attribution

- Model the expected number of observed gamma rays for each trajectory (**r**):

$$c_i(E) = \frac{\epsilon\left(\hat{\Omega},\, E\right)\alpha e^{-\mu(E)\mathbf{r}_i}}{4\pi\,\mathbf{r}_i^2} \cdot \Delta t_i + b$$

- α and b are fit parameters that are optimized with respect to radiation data
- Calculate Poisson log-likelihood and p-value of observing data given the optimal model
- Use s-value ($-log_2[p\text{-}value]$) as exclusion metric

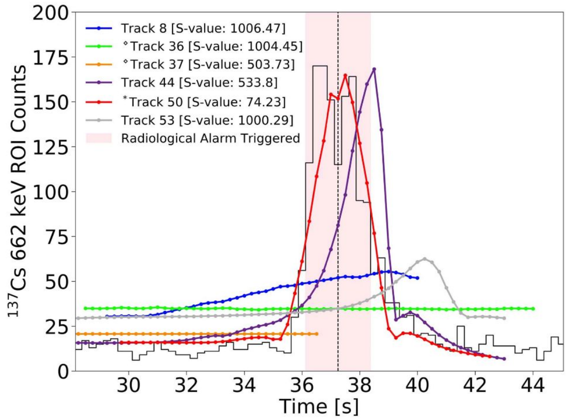

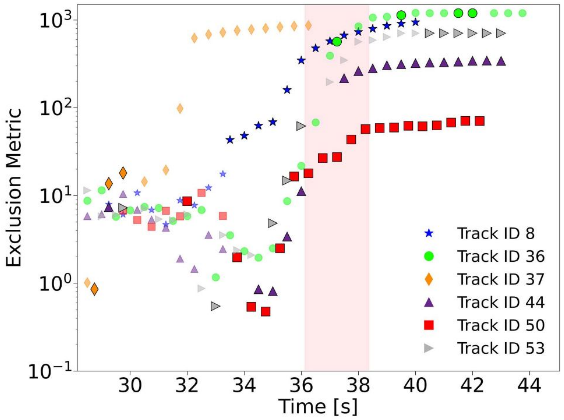

S. Greenland, “Valid *P*-Values Behave Exactly as They Should: Some Misleading Criticisms of *P*-Values and Their Resolution With *S*-Values”, doi:10.1080/00031305.2018.1529625

Experimental data

- Measurements performed with static contextual sensor system (Professor)
- Both LiDAR and video detection, tracking and attribution in real time (not simultaneously)
- Robotic operating system (https://www.ros.org) for real-time

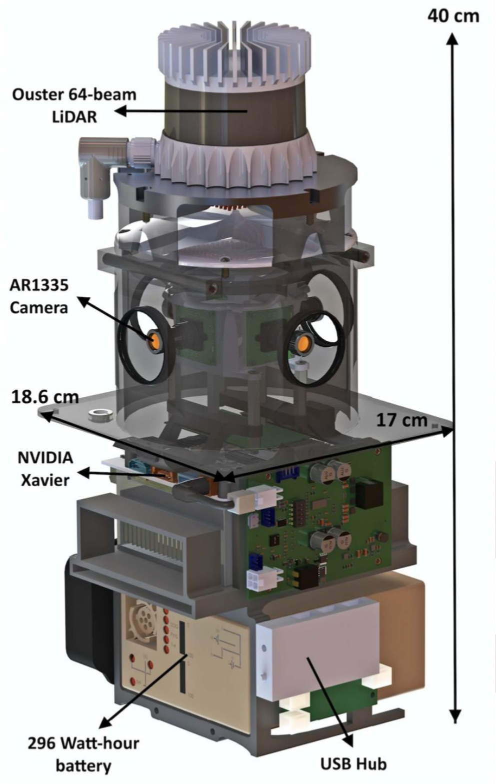

BERKELEY LAB

Experimental results

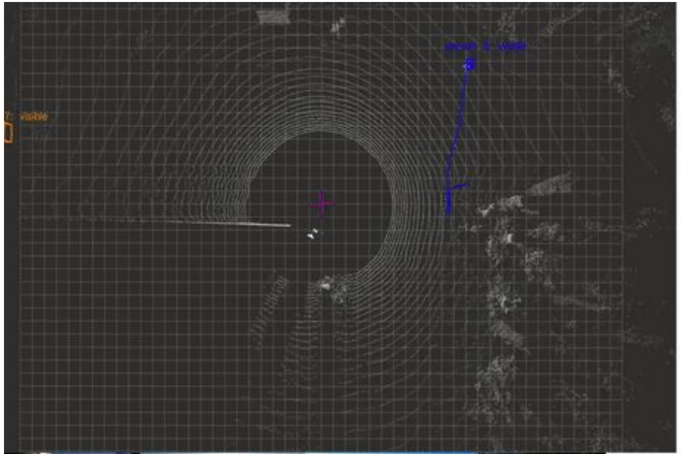

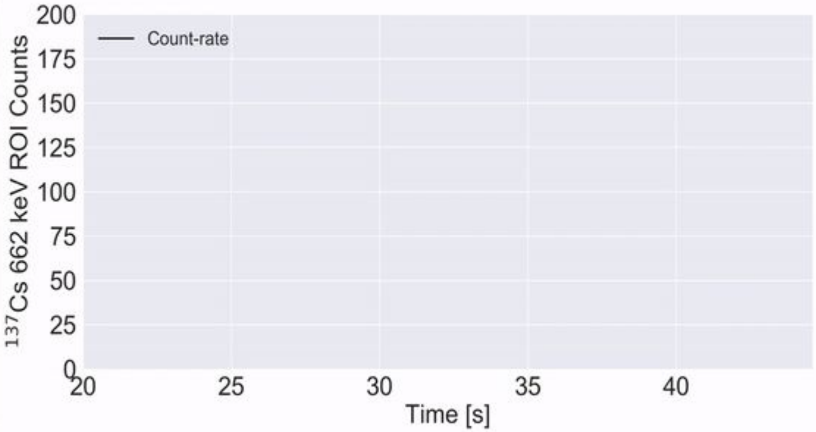

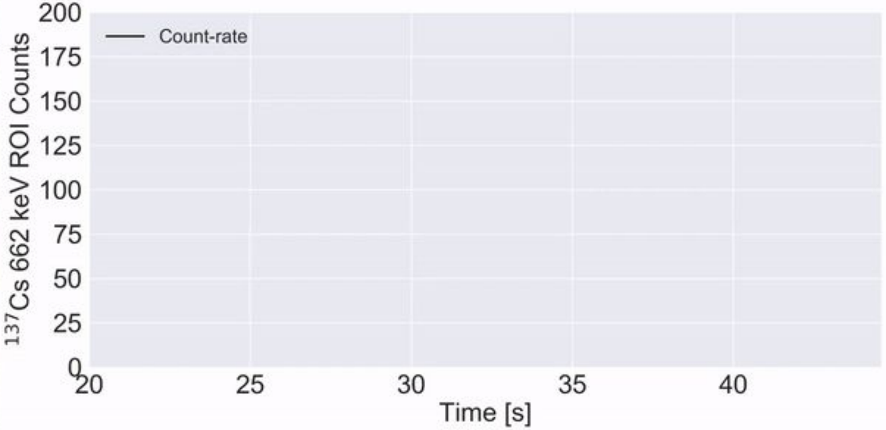

Team and references

- University of California, Berkeley
  - M. R. Marshall
- Lawrence Berkeley National Laboratory
  - M. Salathe
  - D. Hellfeld
  - T. H. Y. Joshi
  - M. S. Bandstra
  - K. J. Bilton
  - R. J. Cooper
  - J. C. Curtis
  - V. Negut
  - A. J. Shurley
  - K. Vetter

M. R. Marshall et al., "3-D Object Tracking in Panoramic Video and LiDAR for Radiological Source–Object Attribution and Improved Source Detection," doi: 10.1109/TNS.2020.3047646.

M. R. Marshall et al, “Mobile Object Tracking in Panoramic Video and LiDAR for Radiological Source-Object Attribution and Improved Source Detection”, submitted to TNS.

Container counting

- Count nuclear storage containers
  - System with LiDAR and Realsense camera
  - Object detection in images: https://github.com/Megvii-BaseDetection/YOLOX

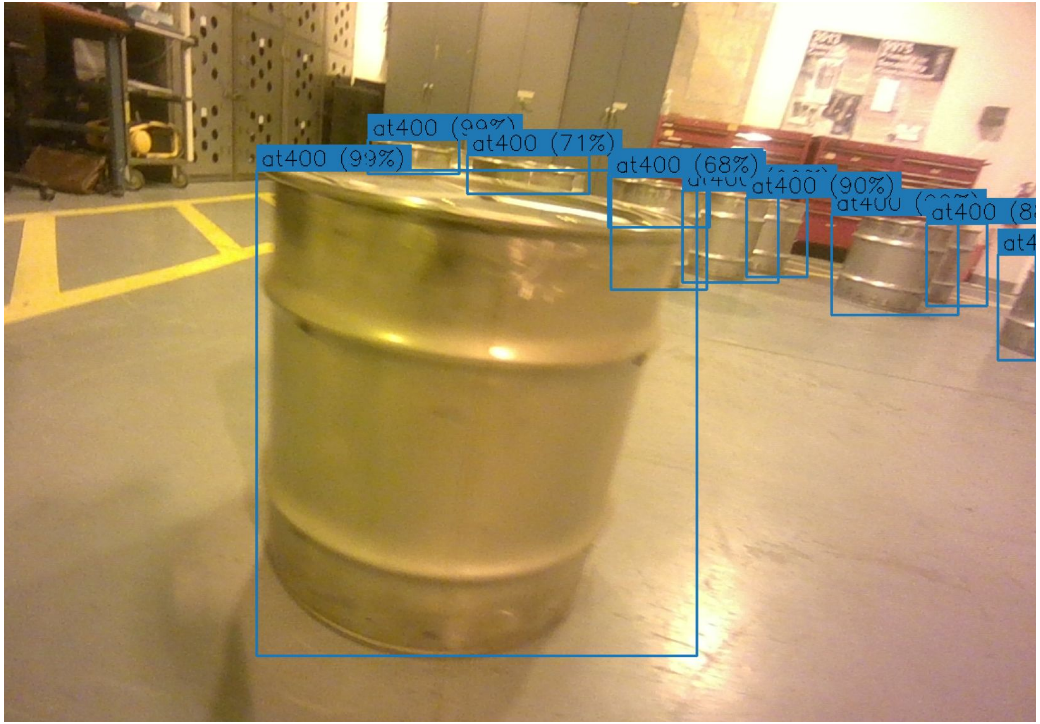

BERKELEY LAB

Combine SLAM result and object detection results

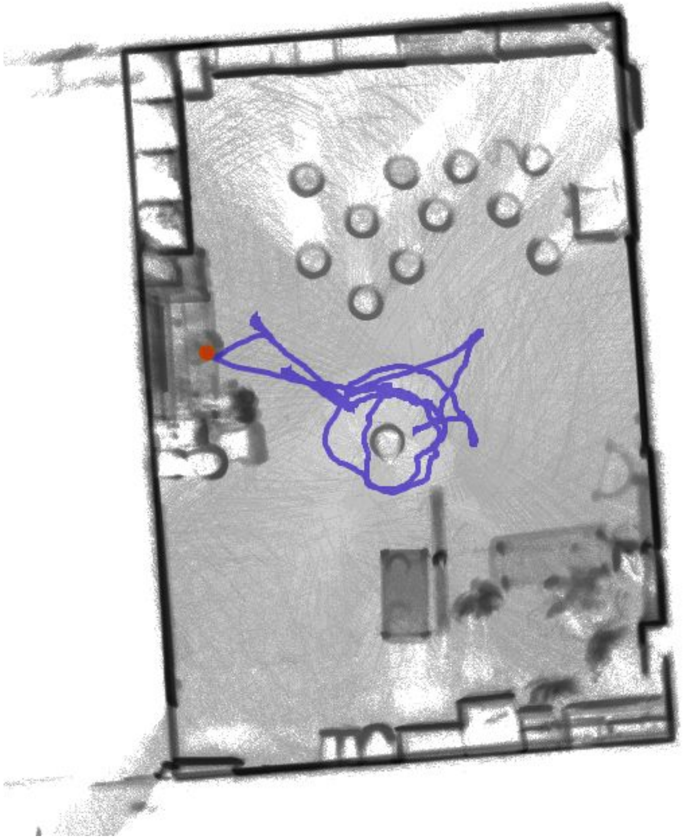

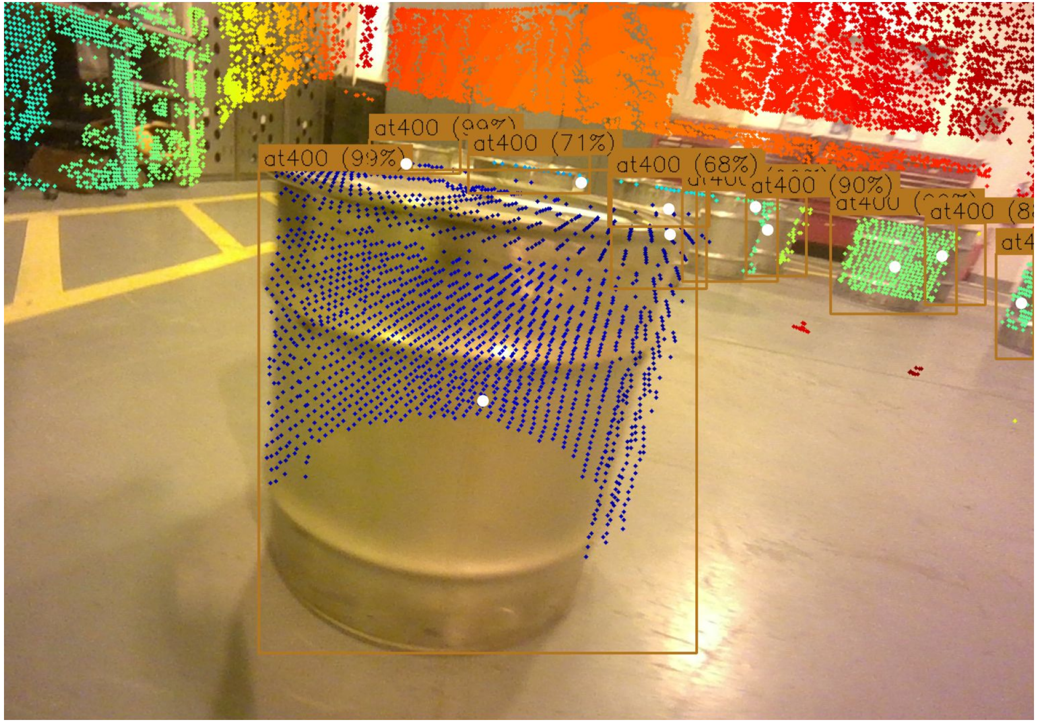

Filter detections

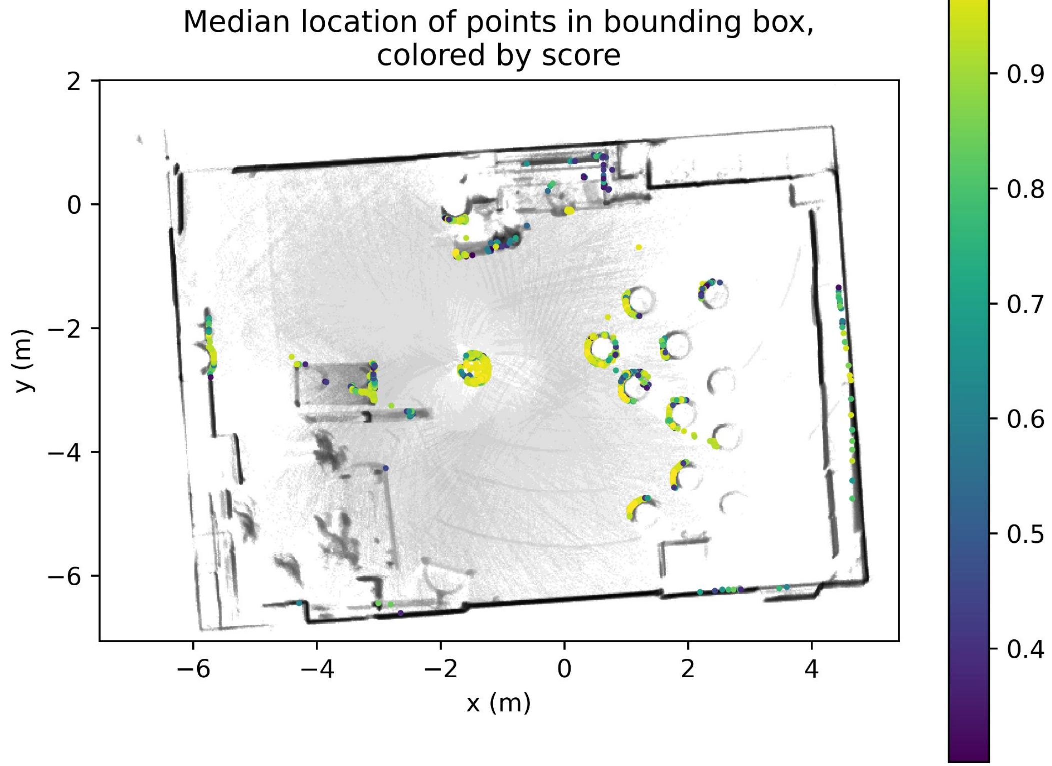

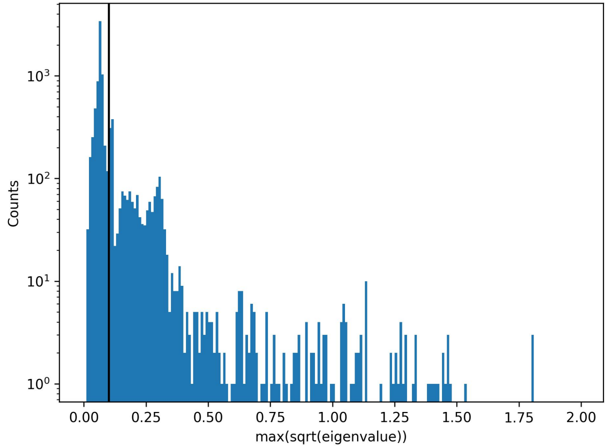

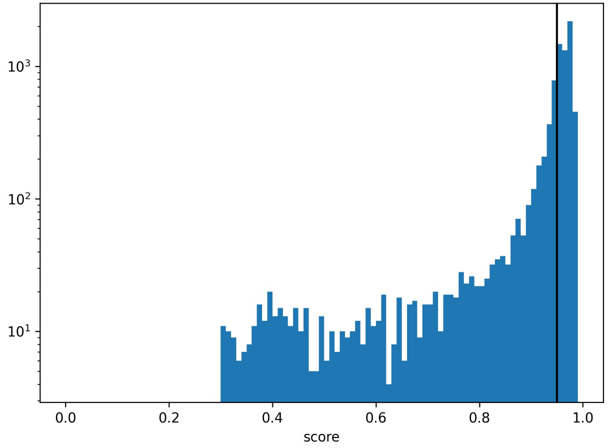

Filter detections

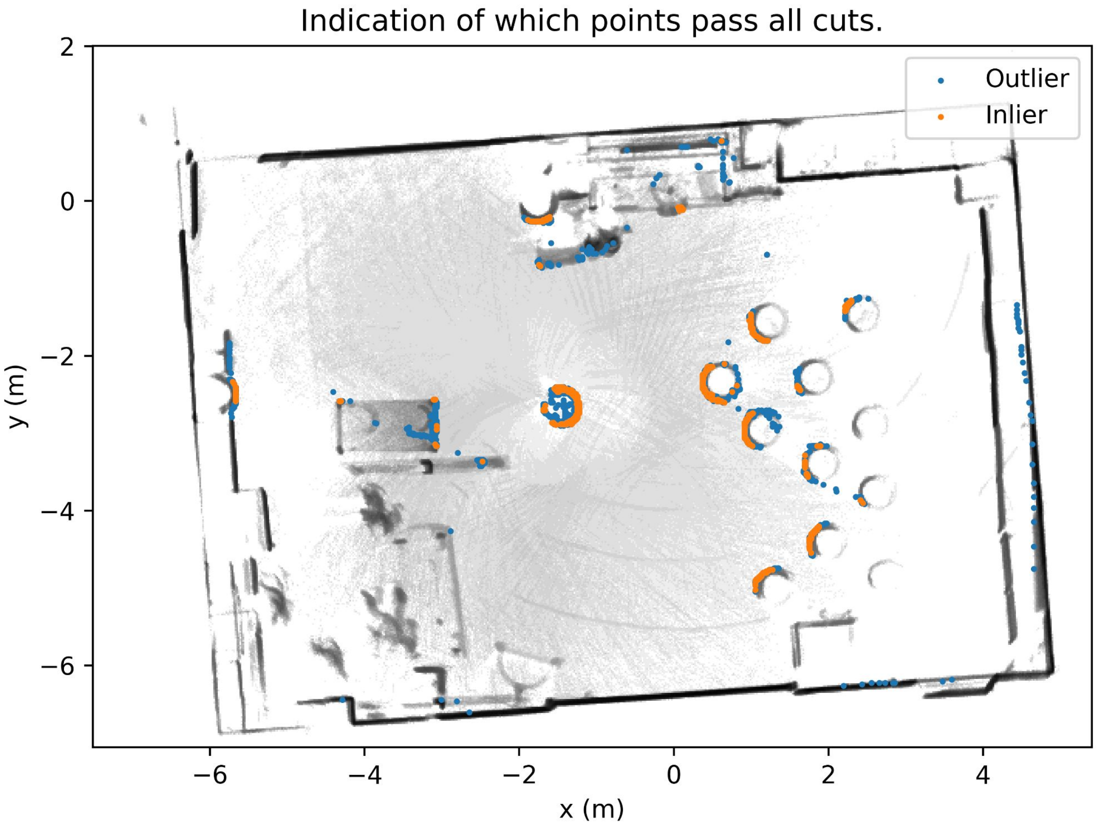

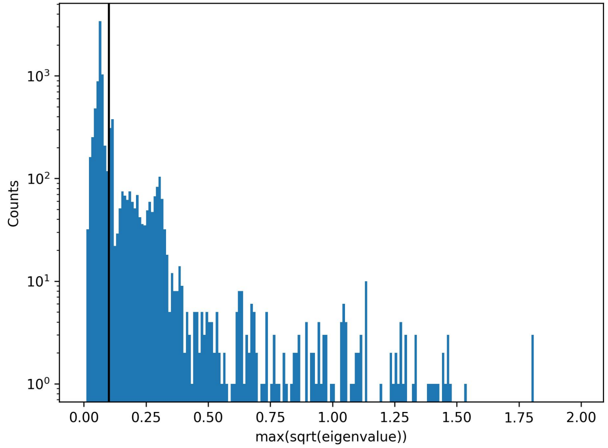

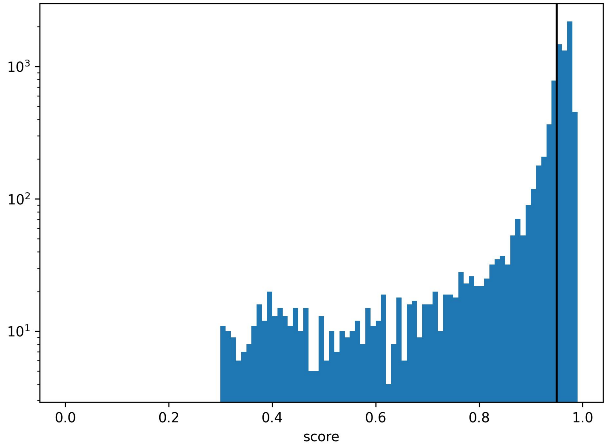

Clustering

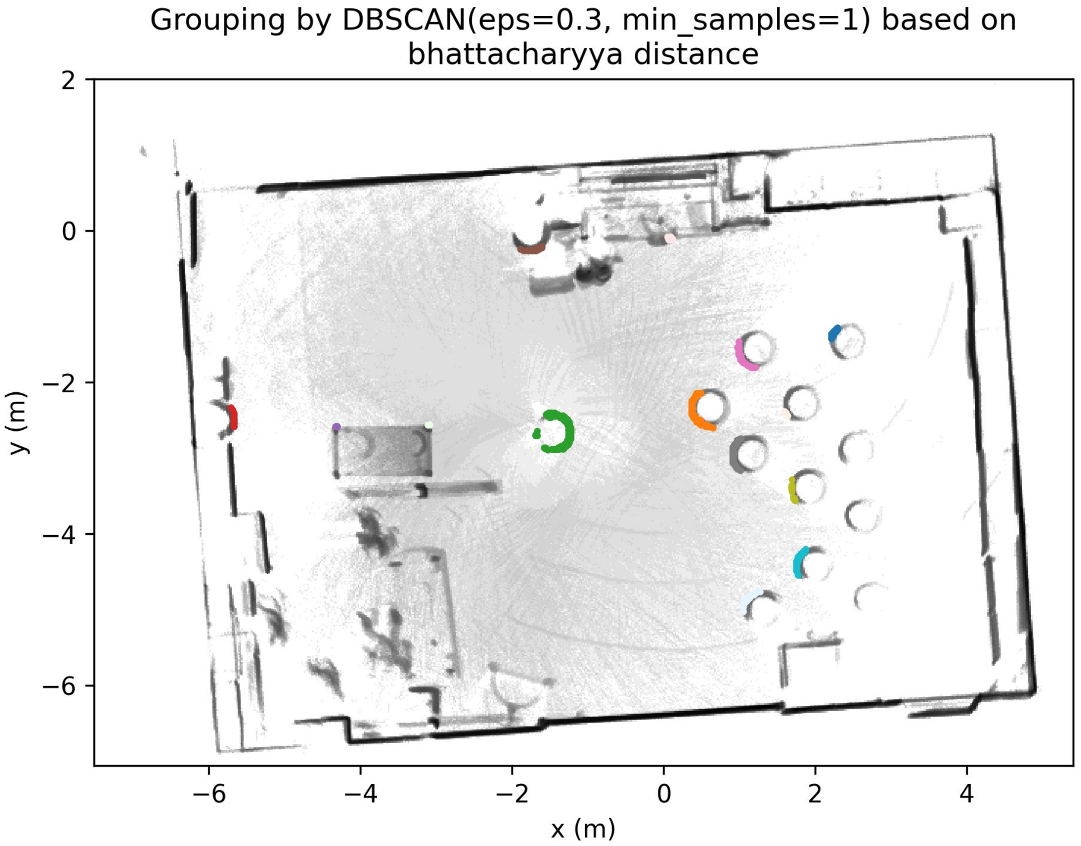

- Clustering with DBSCAN based on the Bhattacharyya distance (metric including uncertainties)
  - https://scikit-learn.org/stable/modules/clustering.html
  - https://en.wikipedia.org/wiki/Bhattacharyya_distance
- Good performance for clustering
- A bunch of false positives
- Possible improvements:
  - Use actual shape of containers during clustering
  - Tracking based approach

Team

LBNL

- Tenzing H. Joshi
- Reynold J. Cooper
- Marco Salathe
- Nicholas Parrilla
- Daniel Parker
- Victor Negut

Stanford - Charm Lab

- Allison Okamura and Team

- Vine Robots development
- http://charm.stanford.edu

BERKELEY LAB

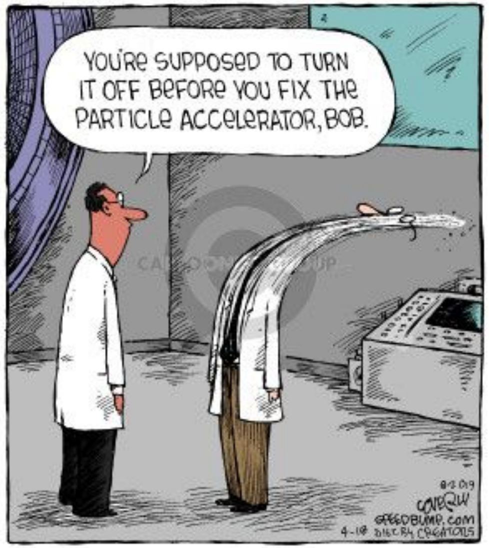

Part 4

ML for beam line operation and optimization

BERKELEY LAB

Machine learning for accelerator and detector control

- Study machine learning as a tool for controlling and optimizing accelerator and detector systems at LBNL
  - Gamma-ray tracking array (GRETA)
  - Venus ion source
- Both systems will be crucial for FRIB
  a new nuclear physics accelerator facility
  build at Michigan State University

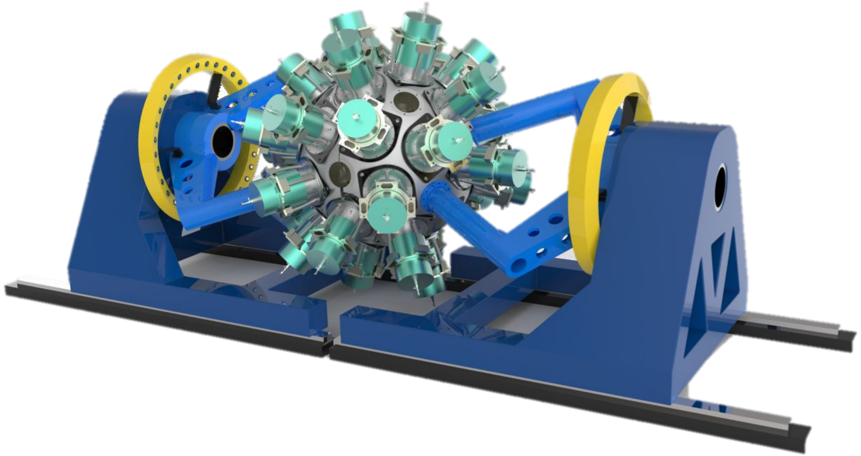

BERKELEY LAB

Integrating the Subsystems of GRETA

Detector Systems

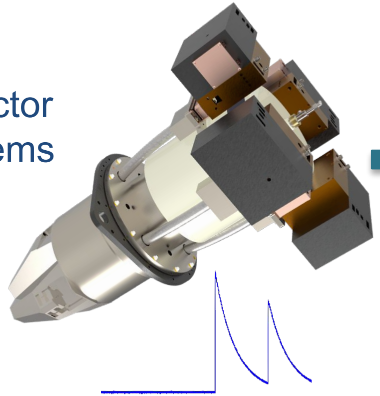

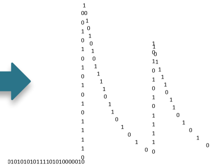

Continuous 100MHz digitization of **40** preamplifier signals **per crystal**

[E, t, 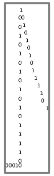] x 40

FPGA-based energy filters, event selection in response to physics triggers

Electronics Systems

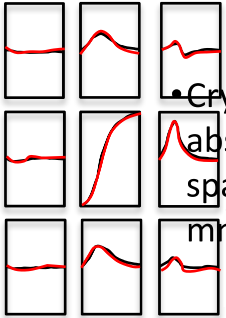

- Crystals located absolutely in space to +/- 0.4 mm

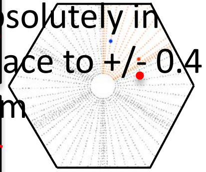

Computing Systems

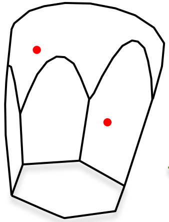

$\left[E, t, (x, y, z)_{crystal} \right]_n$

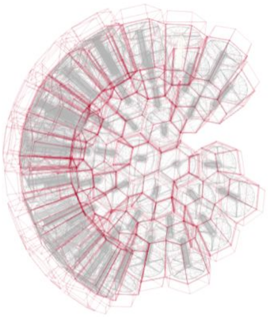

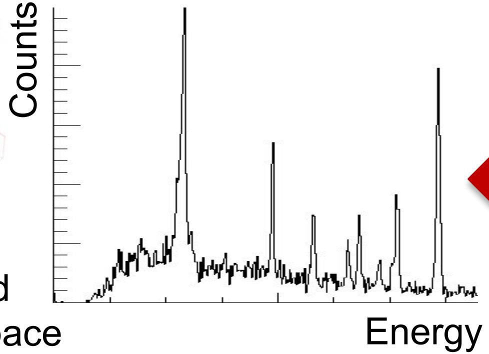

Crystals located absolutely in space to +/- 0.4 mm

Machine learning for GRETA

- 120 crystals with each 40 preamplifier signals
- Converted to energy, x, y, z and t for each detected gamma-ray
- Operational controls
  - Cooling systems (30 modules require liquid nitrogen controls and monitoring)
  - High voltage (120 crystals)
  - Filter parameters
  - Auxiliary devices
- Large amount of data (to much for a operator to handle)
  - Anomaly detection
  - Operational optimization (energy resolution, spatial resolution)

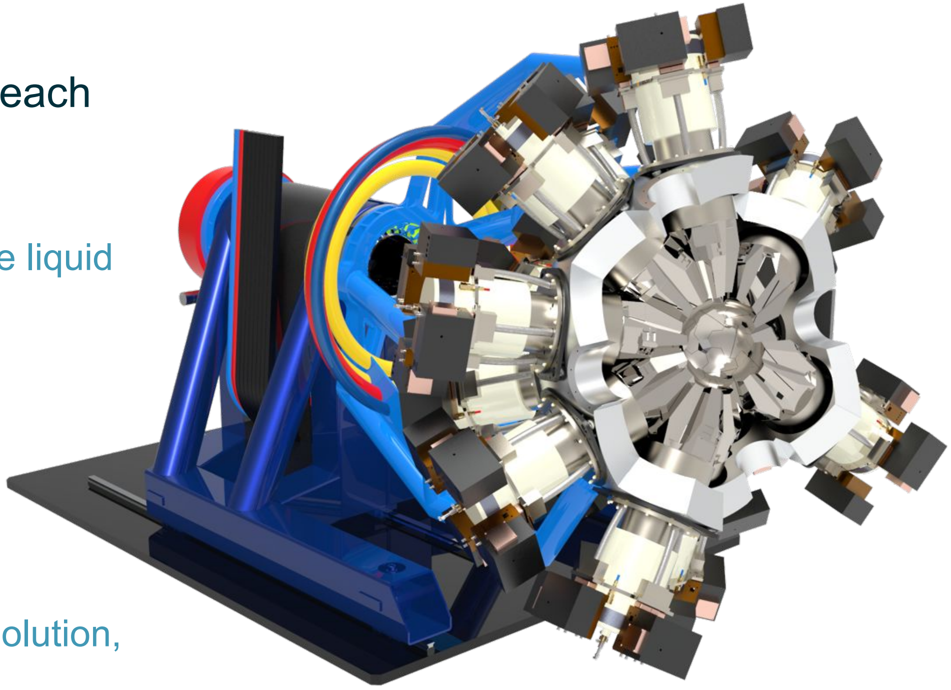

Venus ion source

- Plasma, contained by magnetic field created by a sextupole and 3 superconducting solenoids.
- Radio-frequency heating with two waveguides
- Isotope injection through back (oven, gas valves)

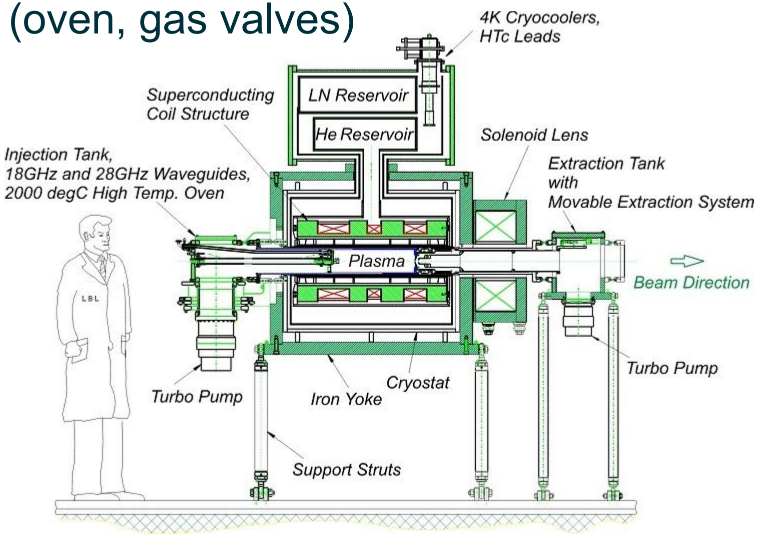

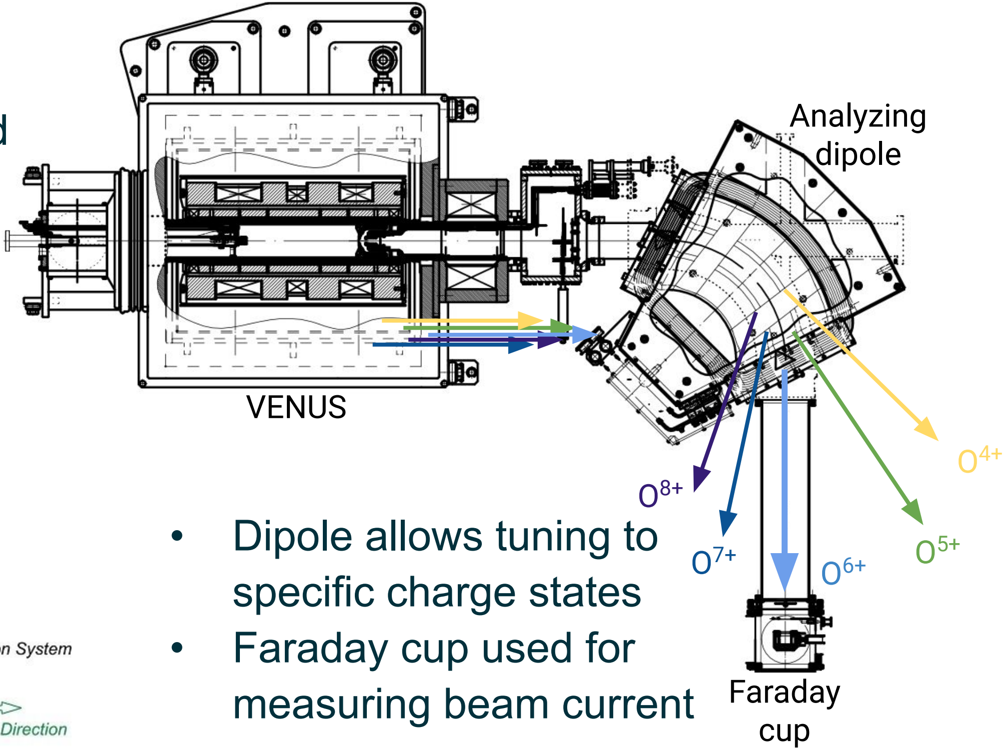

- Dipole allows tuning to specific charge states
- Faraday cup used for measuring beam current

Venus ion source

- Plasma, contained by magnetic field created by a sextupole and 3 superconducting solenoids.
- Radio-frequency heating with two waveguides
- Isotope injection through back (oven, gas valves)

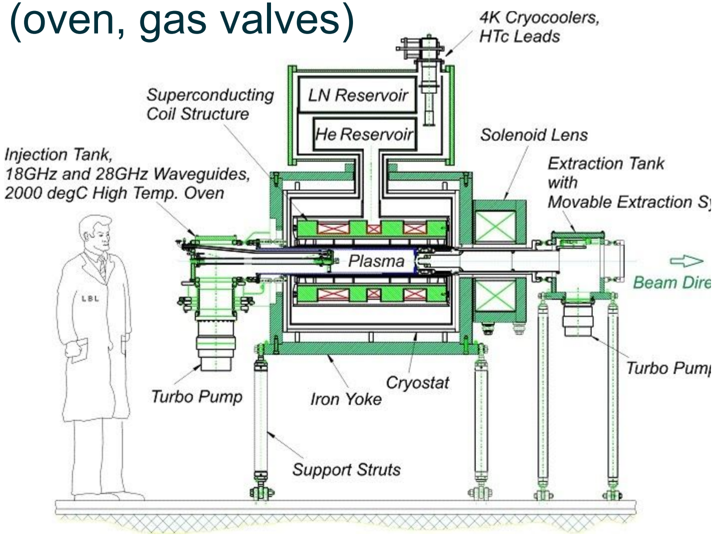

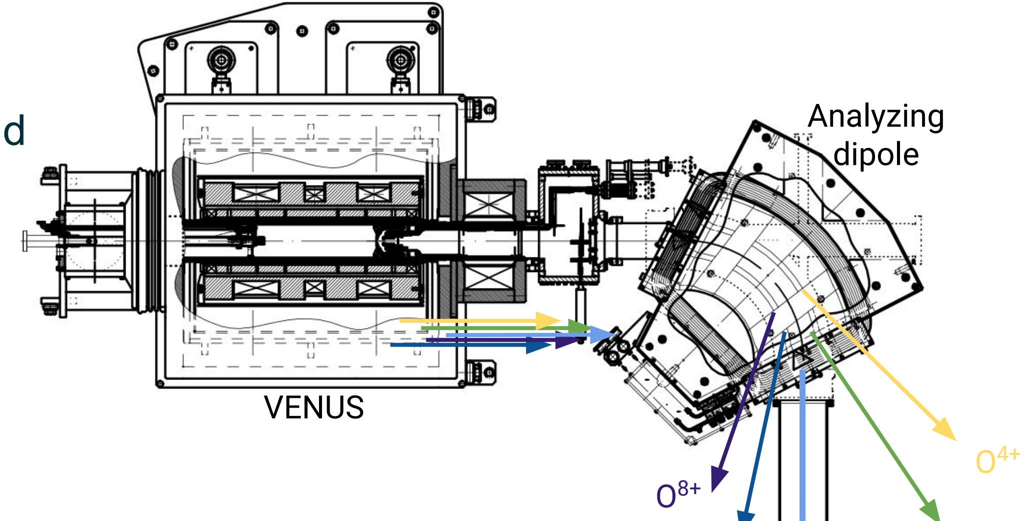

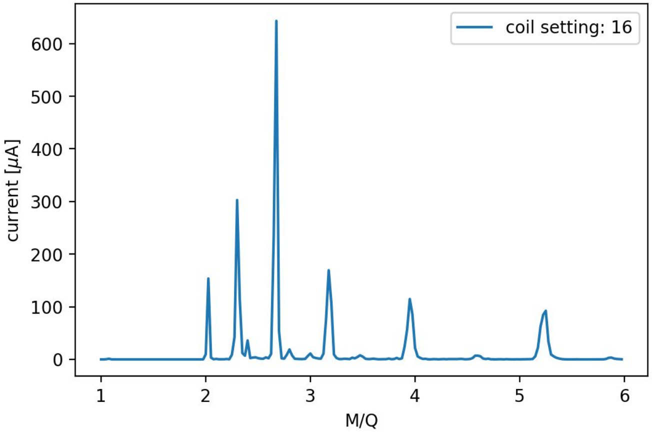

First results

- Automated tuning
  - Main parameters controlling beam current: 3 magnet currents
  - Beam requires about 5 min at each configuration to stabilize → Total tuning requires infrequent input from operator but will take hours to complete.
  - Idea: Use bayesian optimization to find configuration with best beam properties (for now current on the faraday cup)
  - Two weekend long runs to run different programs and understand parameter space

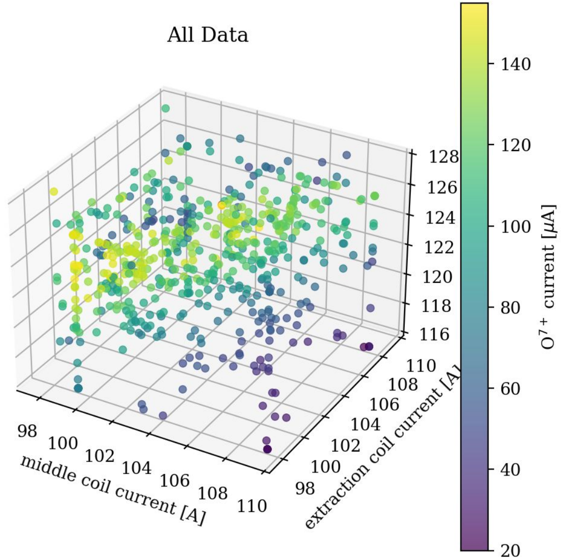

First results

- Automated tuning
  - Main parameters controlling beam current: 3 magnet currents
  - Beam requires about 5 min at each configuration to stabilize → Total tuning requires infrequent input from operator but will take hours to complete.
  - Idea: Use bayesian optimization to find configuration with best beam properties (for now current on the faraday cup)
  - Two weekend long runs to run different programs and understand parameter space
  - Gaussian Process regressor for parameter space modeling

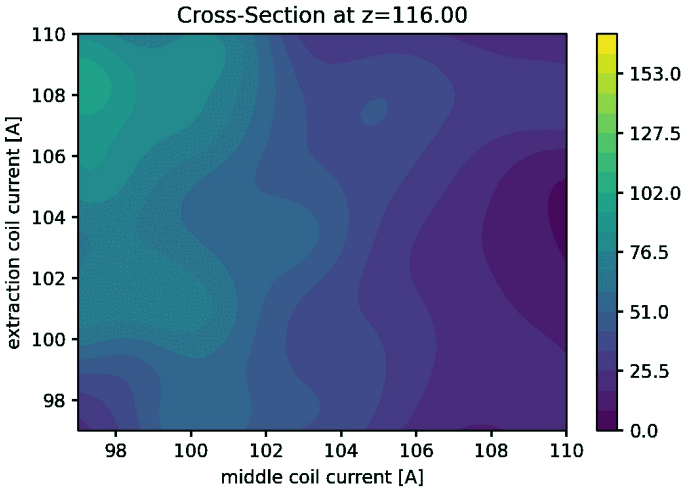

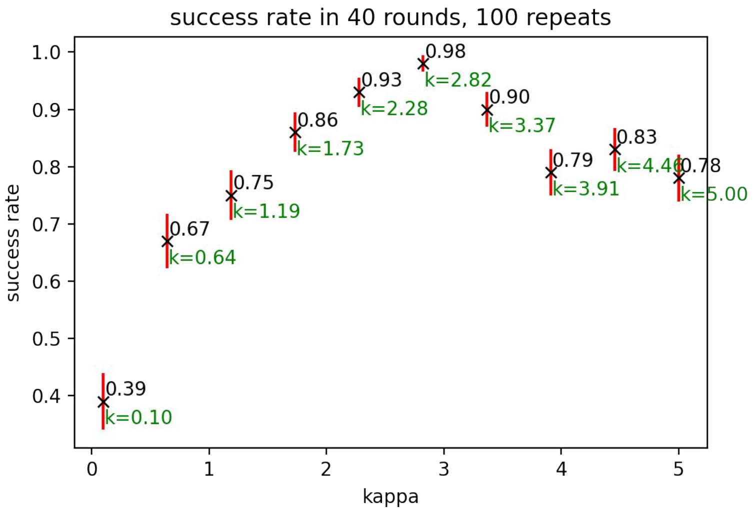

Team

- GRETA
  - Heather L. Crawford
  - Christoper M. Cambell
  - Paul Fallon
- VENUS
  - Damon S. Todd
  - Larry W. Phair
  - Janilee Benitez
- ANP
  - Marco Salathe
  - Brian J. Quiter
  - Reynold J. Cooper
- Undergraduates
  - Wenhan Sun, UC Berkeley
  - Yubin (Harvey) Hu, UC Berkeley
  - Alex Kireeff, Carnegie Mellon University

Thanks

Please feel free to contact me with questions and/or interesting ideas:

msalathe@lbl.gov

Detector response (Effective Area)

- Effective area A_i is product of efficiency and geometric area
- Simulated using a simple model of RadMAP in all 4π
- Folded with estimated detector energy resolution

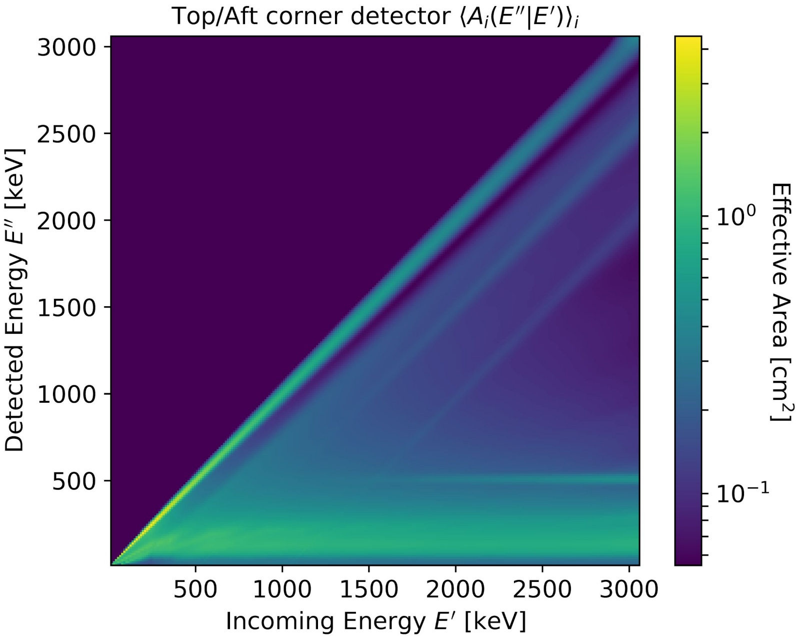

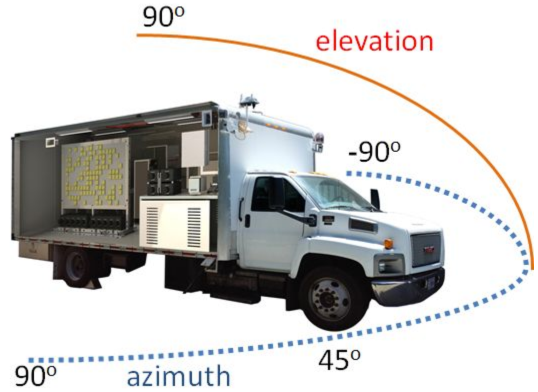

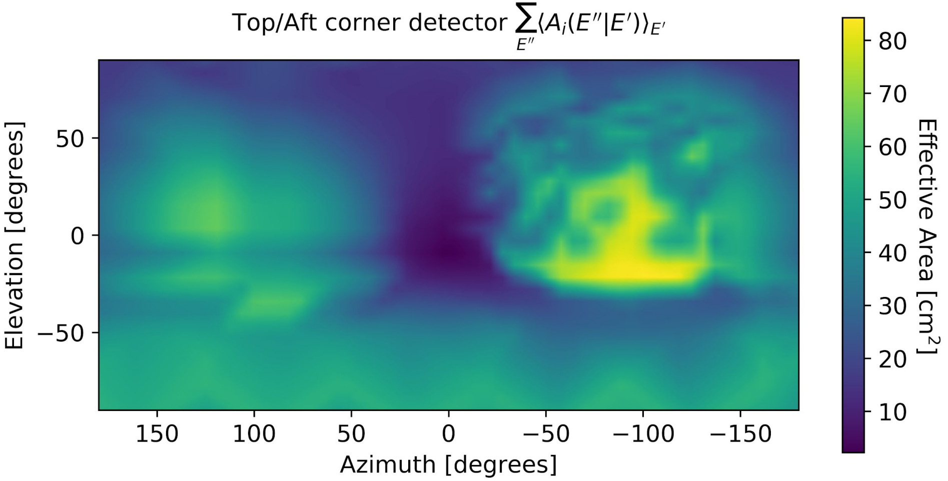

Down scattering in air

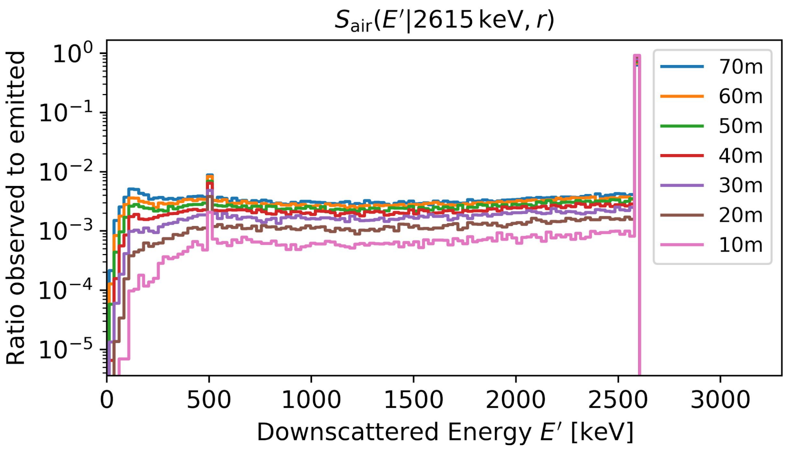

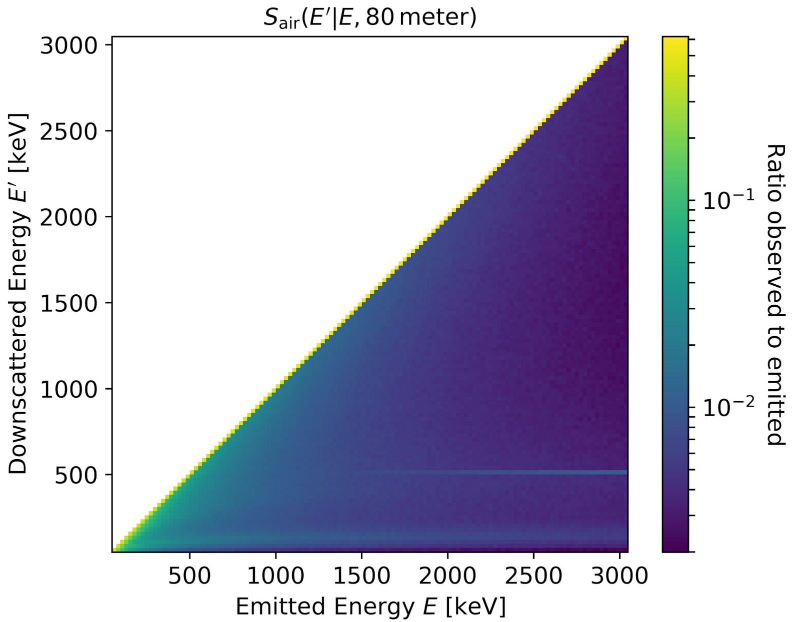

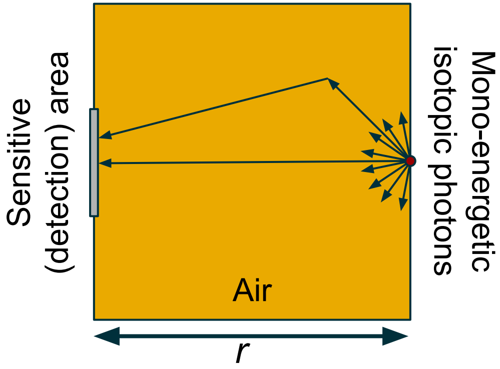

- Down scattering in air has been simulated with a tool developed by Mark S. Bandstra named *Ersatz* (not yet published)
  - A square box with equal sides was used as a simulation volume
  - Gamma-rays were emitted isotropically from a point-like mono-energetic source
  - The sensitive (detection) area covered 1/9th of the surface opposing the source

Complexity reduction with NORM modeling

Main sources of NORM:

- Terrestrial (KUT)
  - K-40,
  - U-238 series,
  - Th-232 series
- Airborne
  - Radon
  - skyshine
- Cosmic
  - Continuum
  - 511 keV from positrons

- K, U and T from simulation, leaving 3 free parameters for each label
- Modeling airborne and cosmic is hard, energy dependence was not enforced (~120 free parameters)
- About 155 free parameters in total, a factor of 10 improvement from an unconstrained fit

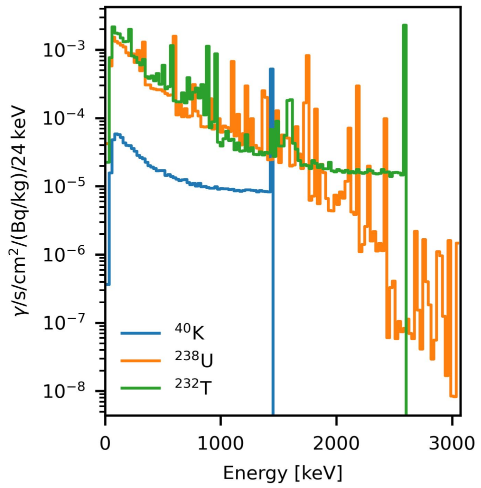

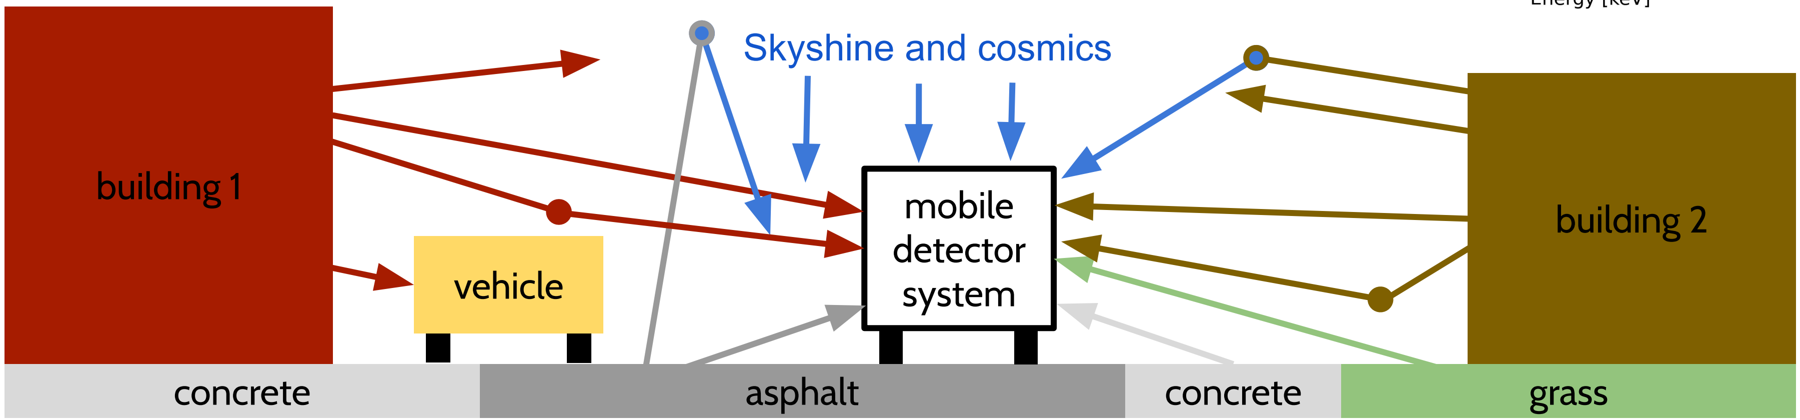

om 2D image to 3D representation

- Object's standoff required in attribution analysis (dominated by *$1/r^2$*)
- 3D position from LiDAR from object detection
- Not directly accessible in images based object detection
- Current method: Compare average size of object (persons height, car dimensions) with bounding box size and use camera intrinsics
    - https://docs.opencv.org/master/d9/d0c/group__calib3d.html

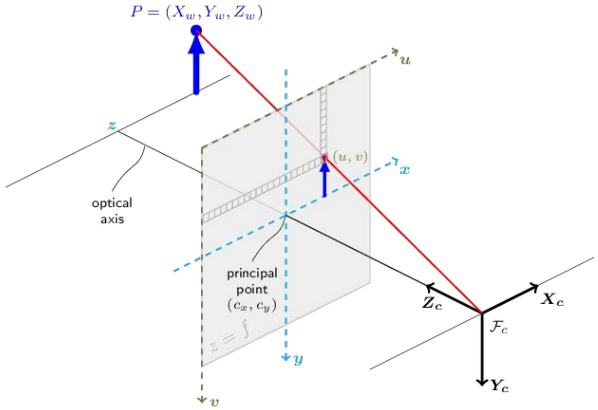

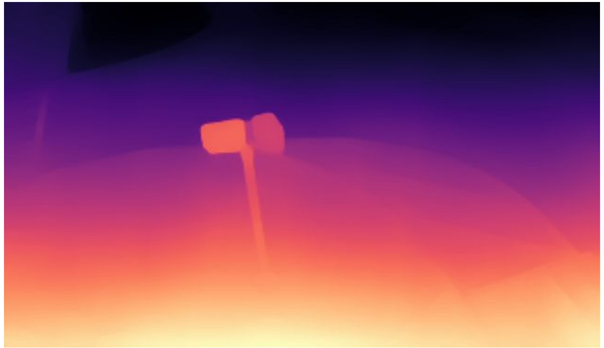

- Possible improvements: Neural network for depth extraction, ground plane estimation, etc.
    - https://filippoaleotti.github.io/demo_live/

Model introspection

Are common features learned in the ensemble? Yes.

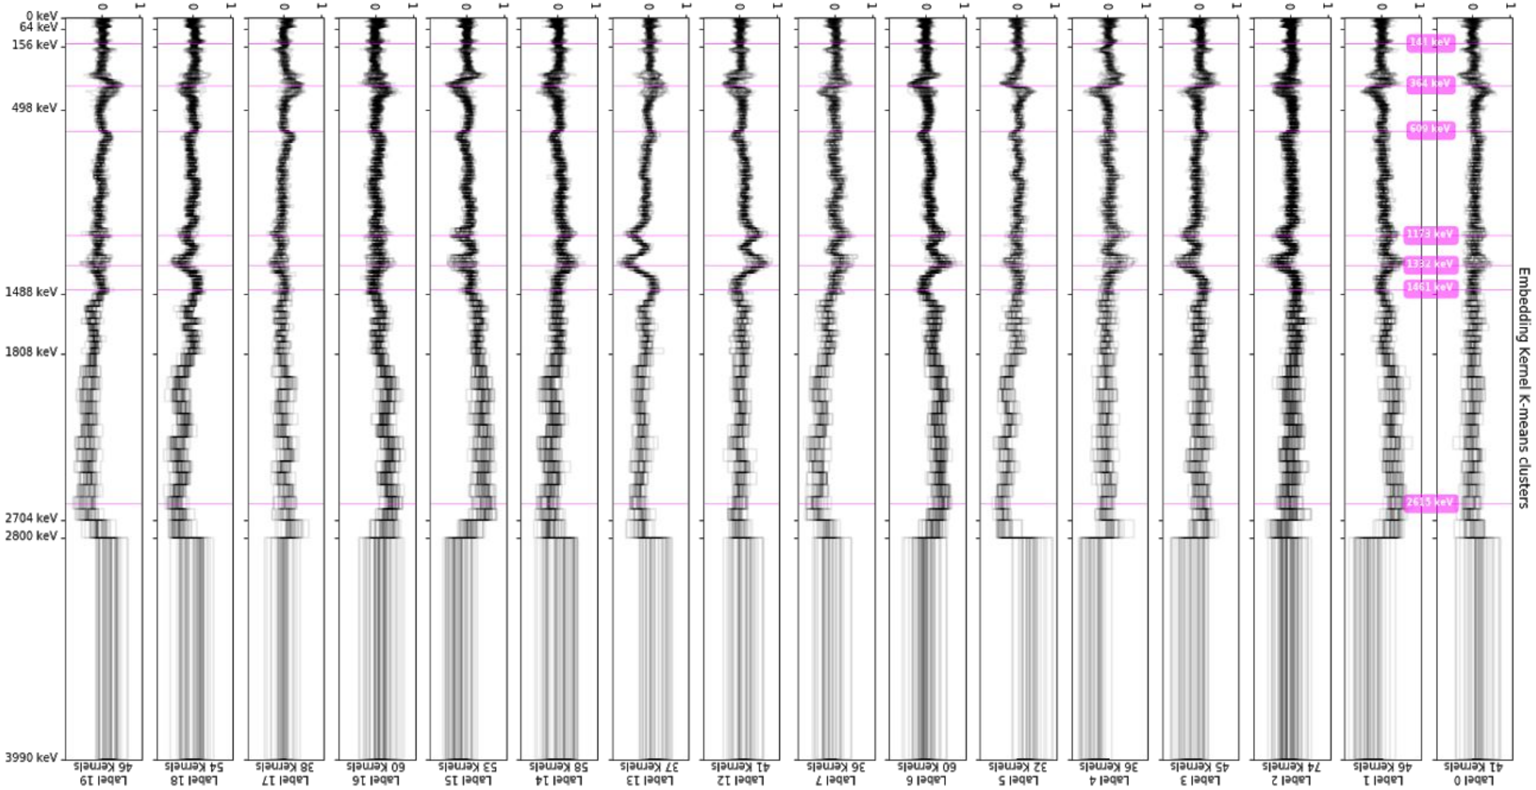

NMF as a physical model

- Airborne data is subject to large background fluctuations
- Aerial survey flying repeatedly over a land/water interface
- NMF finds structure with physical origins:
  - “Nearby KUT” decreases the most abruptly over water
  - “Distant KUT” has a slower ecrease
  - “Radon/cosmics/aircraft” component remains approximately constant

NaI detector data from Aerial Measurement System (AMS) at Lake Mohave, NV

Individual 2 Hz spectra

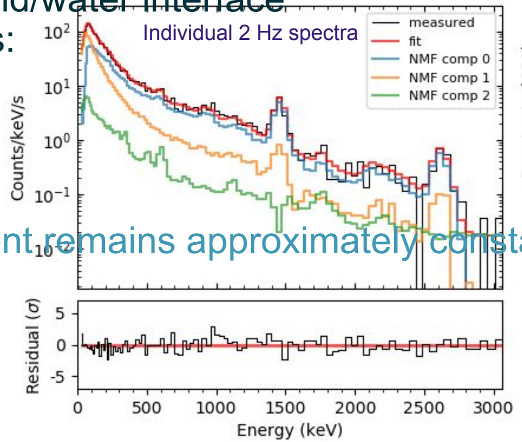

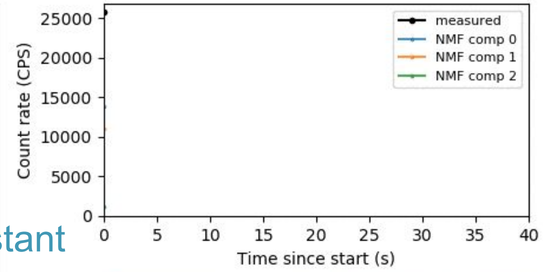

Mean NMF component spectra

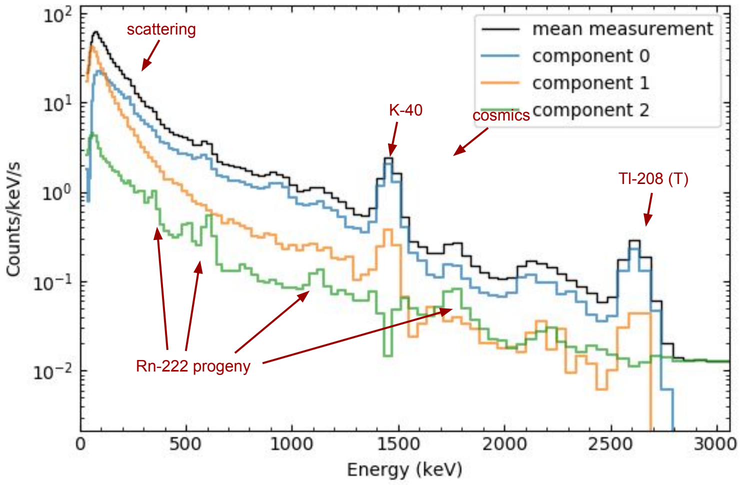

M.S. Bandstra, T.H.Y. Joshi, K.J. Bilton, A. Zoglauer, and B.J. Quiter, “Modeling Aerial Gamma-Ray Backgrounds using Non-negative Matrix Factorization,” IEEE Trans. Nucl. Sci., 67 (5), 2020.

BERKELEY LAB

Overview over submissions (Top Coder)

- Top 7 participants used a
  - Spectra rather than list mode
- Minimal emphasis on threshold setting
  - No mention of Receiver Operator Cost (ROC), False Alarm Rate (FAR), or False Positive Rate (FPR)
- 7 of 10 used neural networks (NN)
  - None in top 3 of .gov used NNs
- Each neural network approach:
  - was an ensemble (multiple nets)
  - contained at least one convolutional layer
- Most competitors used python
  - 2 Java (random forests)
  - 1 R (likelihood testing)

- Each top approach:
  - used some cross validation to avoid overfitting
  - added statistical noise to input before training
  - used labeled training data instead of source templates (data driven training)
  - calculated the closest approach time with a metric weighted average across temporally smoothed energy bins
- The first-place approach:
  - held a multi-point lead over competitors
  - used appropriate statistics during training
  - simultaneously analyzed all data from each run
  - made comments on augmentations for real-time operation

Model introspection

The first layer (1D conv) learns some spectral features

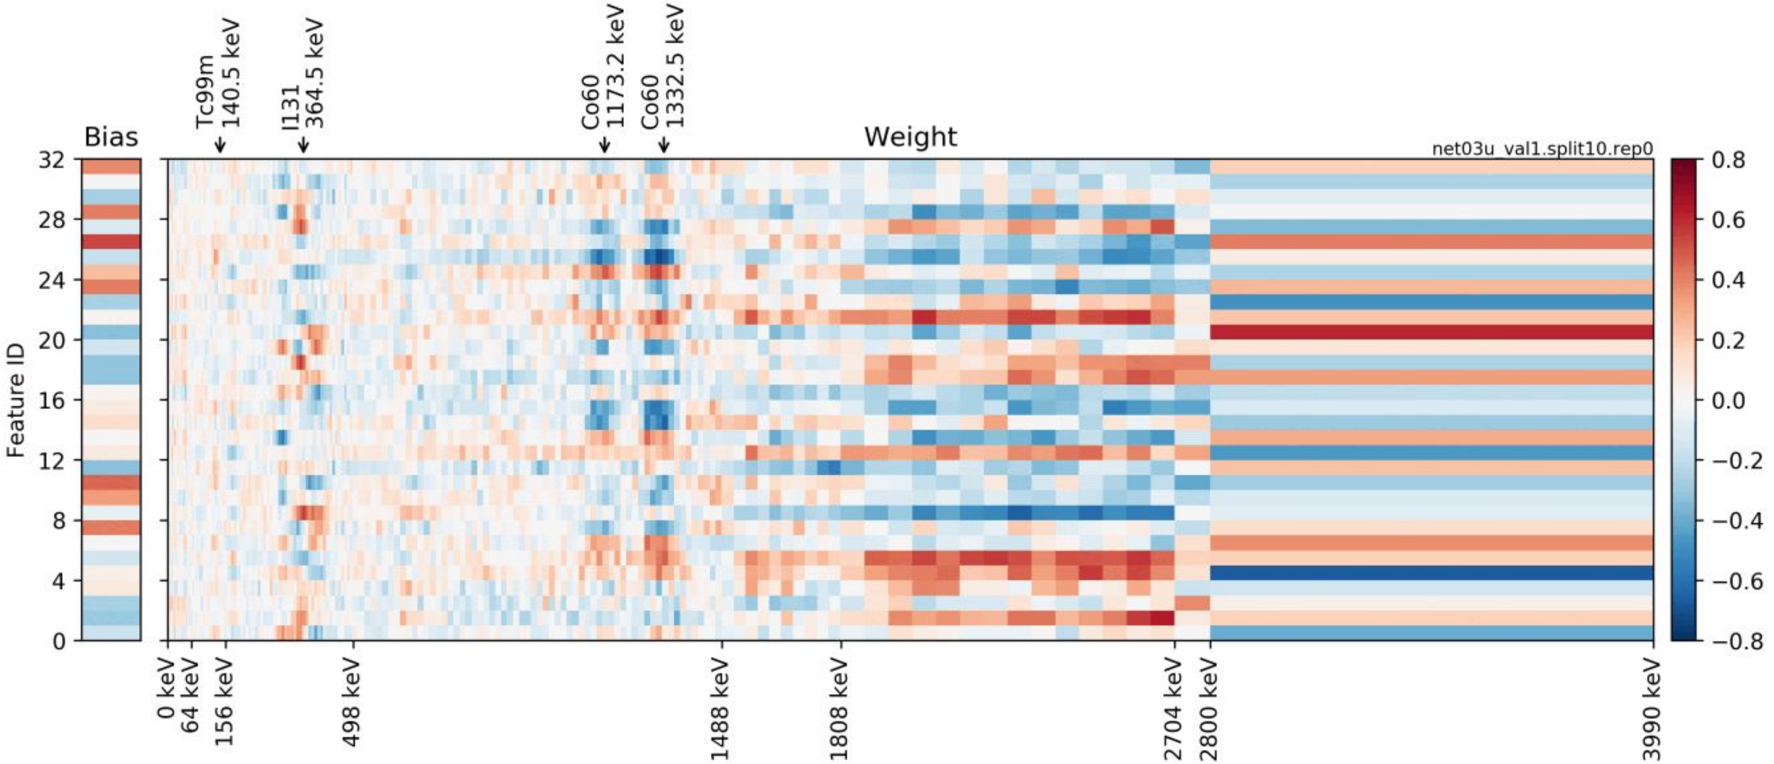

Computer vision challenge

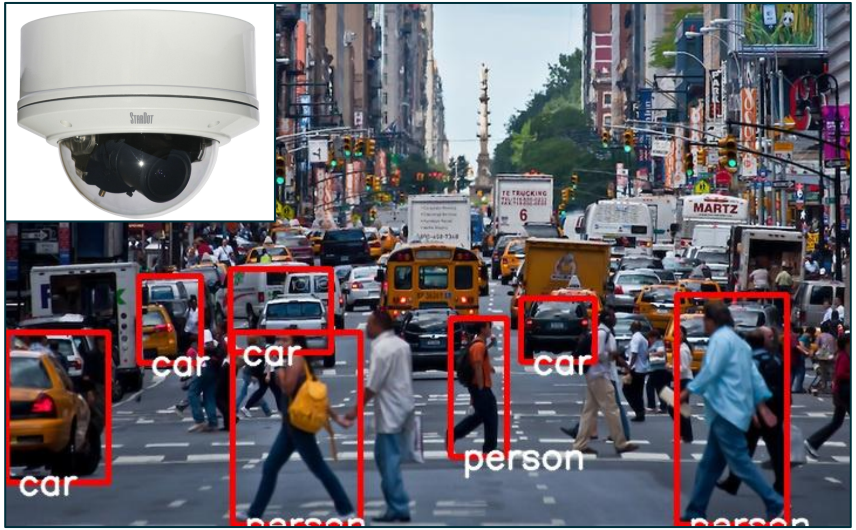

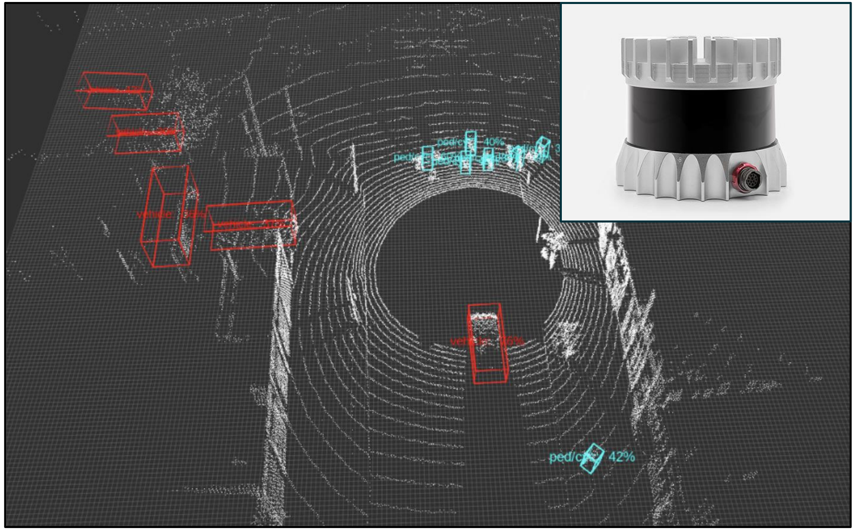

- Video and LiDAR both produce a rich representation of the nearby scene
    a. Finding specific objects in these data
    b. Tracking these objects as a function of time through the scene
    c. Use these data to improve anomaly detection performance
- Hardware limitation on the edge (NVIDIA Jetson NX)

Data handling for competitors

Available Data

Training Data:
Inputs and Answers provided

Test Data:
Inputs provided

Public Test Data:
Competitors receive feedback on this subset during competition

Private Test Data:
Used to determine final standings (no feedback during competition)

Other rules:
- Some constraint on # of submissions
- Competitors may have to specify which solution(s) they want considered for final standings
- Sometimes options for teams to combine/collaborate

Competitors don't know which runs are in each of the test data sets

Leaderboard scoring: pre-specified scalar that combines all desired aspects of solution

Courtesy of Christine Anderson-Cook

BERKELEY LAB

3D Radiation Mapping Applications

Contamination Avoidance

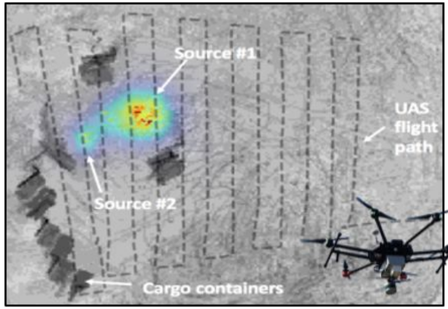

Consequence Management

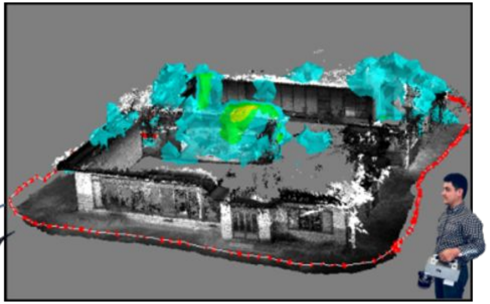

Directed Search

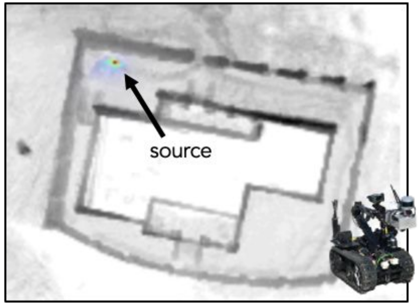

Nuclear Disablement

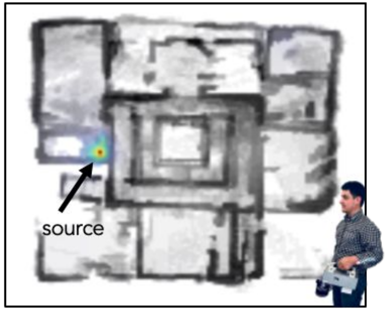

Damage Assessment

Decontamination Verification

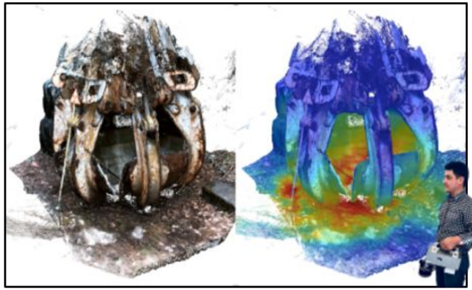

Sensitive Site Assessment

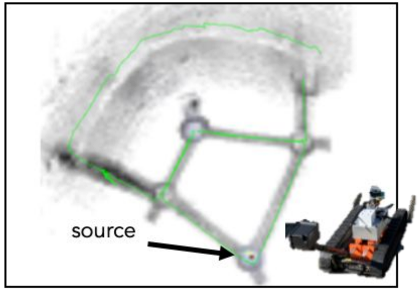

Accident Response

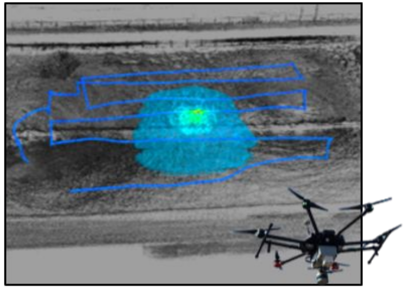

BERKELEY LAB

What about Machine Learning?

- Most systems have also cameras installed.
- Images provide more context and allow for semantic segmentation
- Combined with the point cloud these information can be used to build a simplified model of the detectors surrounding
  - Better understand/estimate backgrounds
  - Build more accurate simulations *to for example support diagnostic of emergency response operations (investigation of possible nuclear threats)*

Radkit (internal software suite)

- Collection of python libraries to process contextual and radiation data products
- Integrate with the Robotic Operating System for real time operation

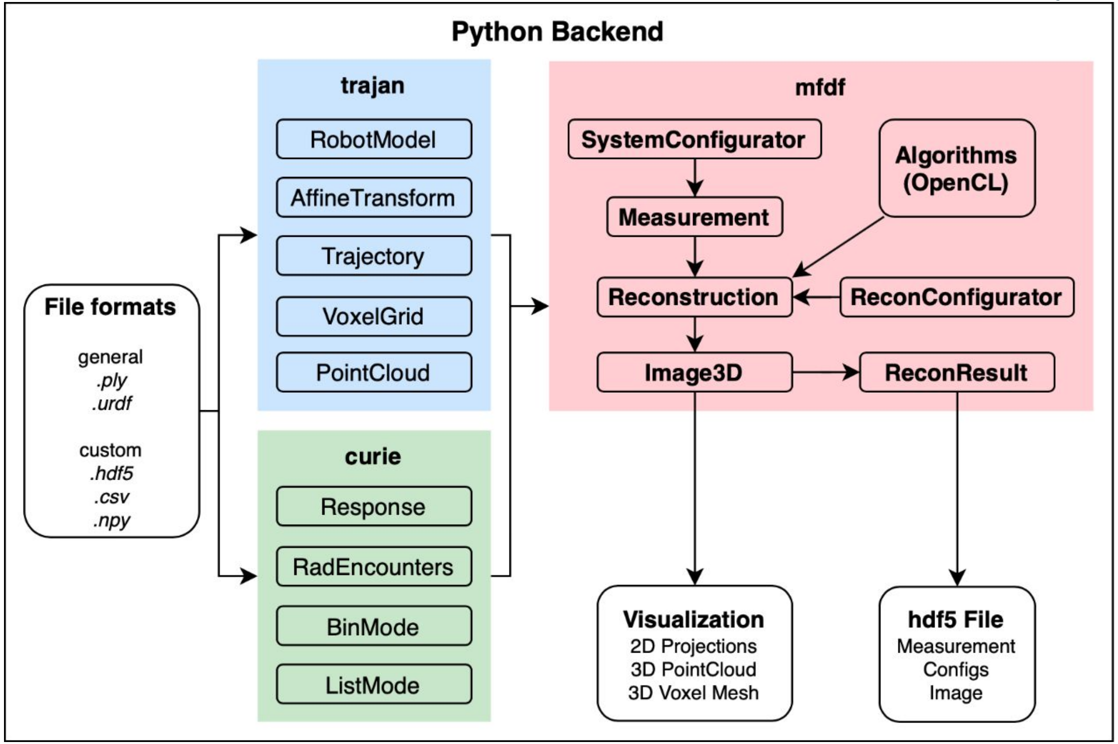

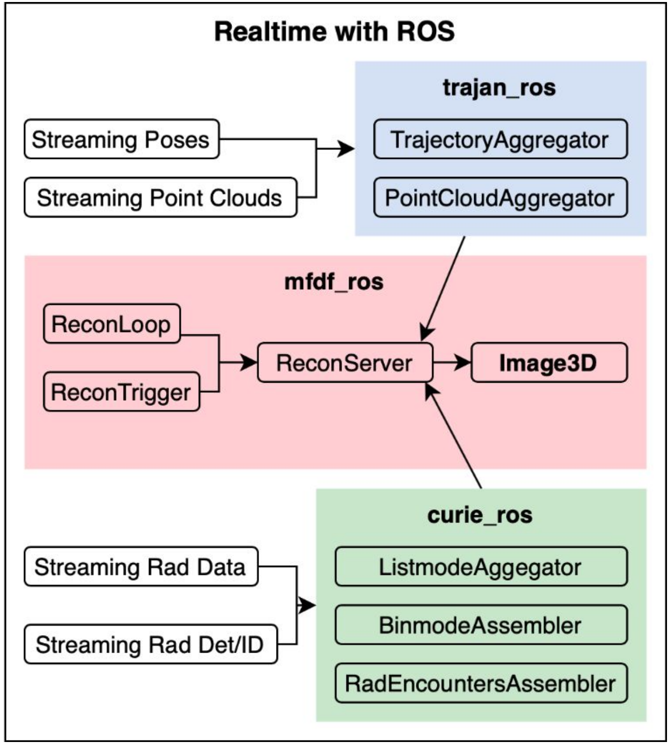

BERKELEY LAB

Summary of the top 5 competitors

User	Input Data Preprocessing	ML Approach	Cross Validation
pfr	• Fixed number of time bins (40, 80, 160) • Increasing size energy bins	• Ensemble of 1D CNN's at 3 time scales • Thresholded output	• 10 splits repeated 3 times • Driving direction reversed • Binomial downsampling
p_kuzmin	• 1 second bins • 50 keV bins • Listmode statistical features	• 3 1D CNN or MLP ensembles for determining det, ID and time • LGBM classifiers	• 5 splits
gardn999	• Square root energy bins • Listmode statistical features	• One random forest classifier per source	• 8 splits
rayvanve	• Multiple energy bin structures • Multiple constant time bin widths • 2D gaussian smoothing	• 2D CNN over sliding time window • Heuristic distance between spectral shape & source terms • Random forest classifiers	• Feature ranking and selection for RF's
cyril.v	• Constant 0.5 s time bins • Log(E) bin widths	• 1D CNN's passed into LSTM • 3 models for det/ID, coarse time and fine time	• Driving direction reversed

BERKELEY LAB

Keys to competition success

- Robust training
  - Input preparation (increasing energy bin widths, multiple time scales)
  - Data augmentation (add noise, leverage symmetry)
  - Cross-validation (split up data)
  - Appropriate labeling (binary vs continuous)
- ML classification methods
  - Neural networks
  - Decision trees
- Power in numbers
  - N experts are better than 1
  - Ensembling reduces variance without cost to bias
  - What is the price?
    - Computation
    - complex/poorly documented code submissions

BERKELEY LAB

Outlook

- Real-time
  - Reduce or remove ensembles
  - How deep is deep enough?
- Rolling analysis
  - "entire drive-by" based analyses deliver answers too late
  - Reshape networks to run on smaller blocks of waterfall data
  - Alarm time becomes implicit
- False positive rate
  - Need to be low and adjustable
  - Threshold perturbation rather than re-training
- Many more sources
  - Re-think network designs for generalization
  - Handling unknown source classes

Follow up project

- Goal: Build a community standard benchmark dataset for radiological search and use it to evaluate fieldable machine learning algorithms against literature benchmarks
  - More diversity (rural area, bridges over rivers, tunnels, etc)
  - More realistic (include environmental effects such as rain, etc.)
  - More sources
  - Different detector types
- Better understanding of algorithm and enable fieldability of algorithms
  - Develop an open source package: https://gitlab.com/lbl-anp/radai/radai