Title

Three-dimensional Object Tracking in Panoramic Video and LiDAR for Radiological Source-Object Attribution and Improved Source Detection

Author

Marshall, Matthew

Publication Date

2021

Three-dimensional Object Tracking in Panoramic Video and LiDAR for Radiological Source-Object Attribution and Improved Source Detection

by

Matthew Marshall

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Nuclear Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committe in charge:

Professor Kai Vetter, Chair
Assistant Professor Mark Mueller
Professor Per Peterson

Summer 2021

Three-dimensional Object Tracking in Panoramic Video and LiDAR for Radiological
Source-Object Attribution and Improved Source Detection

Abstract

Three-dimensional Object Tracking in Panoramic Video and LiDAR for Radiological Source-Object Attribution and Improved Source Detection

by

Matthew Marshall

Doctor of Philosophy in Nuclear Engineering

University of California, Berkeley

Professor Kai Vetter, Chair

The detection and localization of radiological and/or nuclear material remains a key challenge in homeland security, especially in urban environments. In an effort to improve detection and localization capabilities, networked detector systems can be deployed in urban environments to aid in the detection and localization of radiological and/or nuclear material. However, effectively responding to and interpreting a radiological alarm using spectroscopic data alone may be hampered by a lack of situational awareness, particularly in complex environments. This work investigates the use of LiDAR and streaming video to enable real-time object detection and tracking, and the fusion of this tracking information with radiological data for the purposes of enhanced situational awareness and increased detection sensitivity. This work presents an object detection, tracking, and novel source-object attribution analysis that is capable of operating in real-time. The analysis pipeline is implemented on a custom-developed system that comprises a static $2\ in. \times 4\ in. \times 16\ in.$ NaI(Tl) detector co-located with a 64-beam LiDAR and 4 monocular cameras. Using this analysis approach on the static system, physics-based models that describe the expected count rates from tracked objects are used to correlate vehicle and/or pedestrian trajectories to measured count-rate data through the use of Poisson maximum likelihood estimation and to discern between source-carrying and non-source-carrying objects. In this work, the source-object attribution approach is explained in detail and a quantitative performance assessment that characterizes the source-object attribution capabilities of both video and LiDAR is presented. Additionally, experimental results from a mock urban environment are shown using the contextual-radiological data fusion methodology. With this data, the ability to simultaneously track pedestrians and vehicles in a mock urban environment is demonstrated, and using this tracking information both detection sensitivity and situational awareness is improved.

The addition of contextual sensors to mobile radiation sensors provides valuable information about radiological source encounters that can assist in adjudication of alarms. This study explores how computer-vision based object detection and tracking analyses can be used to augment radiological data on a mobile detector system. Using these analyses on a mobile system, this work studies how contextual information (streaming video and LiDAR) can be used to associate dynamic pedestrians or vehicles with radiological alarms to enhance both situational awareness and detection sensitivity. To perform this study, data was collected in a mock urban environment where participants included pedestrians and vehicles moving in the vicinity of an intersection. Data was collected with a vehicle equipped with 6 NaI(Tl) $2\,in. \times 4\,in. \times 16\,in.$ detectors in a hexagonal arrangement and multiple cameras and LiDARs as well as an IMU and an INS. In this work, the source-object attribution approach as applied to a mobile system with multiple detectors in the presence of static and moving sources is demonstrated. The results show improved situational awareness and detection sensitivity using video and LiDAR-based trajectories. In addition, it is seen that LiDAR data produces more reliable position estimates of an object compared to using video data, which enables more effective object tracking and attribution, especially in scenarios with vehicle speeds of about 20 mph. Furthermore, with both video and LiDAR data, improved detection sensitivity is demonstrated using an optimal configuration of detectors within a detector array compared to the summed response of a detector array in a mock urban environment. Finally, by correlating vehicle and/or pedestrian trajectories to measured count-rate data, the source is inherently localized to an object, which enables a new paradigm for source localization and might reduce the complexity in detector array design. To test this concept, source-object attribution is performed using different detector array configurations with varying levels of complexity to understand the impact of the angular response on radiological source localization when tracking information is available. The findings demonstrate that using the source-object attribution analysis approach can enable simpler detector array designs to perform source localization in the investigated scenarios.

To Kenzie, Mila, and my parents.

Contents

List of Figures

List of Abbreviations

Alternate Nuclear Materials (ANM)
Berkeley Anomaly Detection (BAD)
Bayesian Information Criterion (BIC)
Contextual Sensor Package (CSP)
Convolutional Neural Network (CNN)
Dilution of Precision (DOP)
Field of View (FOV)
Full-width at half maximum (FWHM)
Global Positioning System (GPS)
Gridded Point-Source Likelihood (GPSL)
Hellinger Distance (HD)
High Purity Germanium (HPGe)
Horizontal Dilution of Precision (HDOP)
Inertial Measurement Unit (IMU)
Inertial Navigation System (INS)
Intersection-over-union (IOU)
Lawrence Berkeley National Laboratory (LBNL)
LiDAR Enhanced Radiation Urban Search (LEMURS)
Markov Chain Monte Carlo (MCMC)
Maximum Likelihood Expectation Maximization (ML-EM)
Miles per hour (mph)
Mobile Imaging and Spectroscopic Threat Identification (MISTI)
Multivariate Normal (MVN)
Non-negative Matrix Factorization (NMF)
Point Source Localization (PSL)
Photomultiplier Tube (PMT)
Region of Interest (ROI)
Sodium Iodide (NaI)
Simple Online Realtime Tracking (SORT)
Radiological Dispersion Device (RDD)
Radiological Multi-sensor Acquisition Platform (RadMAP)
Region of interest (ROI)
Robot Operating System (ROS)
Signal-to-Noise Ratio (SNR)
Simultaneous Localization and Mapping (SLAM)
Special Nuclear Material (SNM)
Universal Robotic Description File (URDF)
Vertical Dilution of Precision (VDOP)
You Only Look Once (YOLO)

Acknowledgements

As a PhD student, I never thought this day would come when I began writing my Acknowledgements section. Reflecting on my academic career, there are a number of people who have supported me emotionally and academically and I would not have reached this point without them. First and foremost, I want to thank my wife, Caroline, for all of her love and support throughout the years. Your patience and encouragement throughout my PhD has been incredible. Additionally, I want to thank my parents. Mom and Dad, thank you for always sending care packages and letters my way. Those little things made all the difference when I was knee deep in exams or papers. Finally, thanks to my siblings and in-laws for their love and support throughout this journey.

During my time at Berkeley I would not have been able to succeed without the help from Dr. Ren Cooper, Dr. Tenzing Joshi, Dr. Marco Salathe, Dr. Dan Hellfeld, and Joesph Curtis within the Applied Nuclear Physics (ANP) program at Lawrence Berkeley National Lab. Thank you all for your guidance and help with my research. You all made my PhD work and experience much more academically and personally fulfilling. I also want to thank Dr. Mark Bandstra and Victor Negut for their help and assistance with my research. In addition, none of this would have been possible without my academic advisor Dr. Kai Vetter. I truly appreciate all of your valuable insights and allowing me freedom with my research to explore different avenues and questions. The resources you made available to me throughout my PhD made a huge difference. I always appreciated your kindness and understanding throughout this process.

I also want to thank the Berkeley Applied Research on the Imaging of Neutrons and Gamma-rays (BeARING) group at UC Berkeley for participating in my research and providing valuable insights. In particular, thank you to Adam Glick, Joe Vanderlip, and Grey Batie for being people I could turn to for research support and to share a laugh or beer with. Also, Michael Bondin, Arthur Shurley, Ray Cheng, and Ivan Cho thank you for sacrificing your time in different data collection campaigns.

Thank you to my doctoral committee – Dr. Kai Vetter, Dr. Mark Mueller and Dr. Per Peterson – for agreeing to serve as committee members.

Within the Nuclear Engineering department, I want to thank Dr. Karl Van Bibber for his encouragement and always having his door open to me when I needed to chat. Also, thank you Alan Bolind for all your help when I was lab manager and helping me coordinate different research projects within Etcheverry.

Chapter 1

Introduction

This dissertation focuses on the improvement in detection, localization, and attribution of radiation detection systems composed of radiation detector(s) and contextual sensors in urban environments to improve both situational awareness and detection sensitivity. This topic is explored using both static and mobile radiation detection systems. The following sections introduce these topics and also discusses the motivation and challenge of detecting and localizing gamma-rays in urban environments.

1.1 Motivation for Gamma-ray Detection and Localization in Urban Environments

In an urban environment, the two main categories of domestic nuclear threats are special nuclear material (SNM) and a radiological dispersion device (RDD) or a "dirty bomb". The SNMs of main concern are highly enriched Uranium (HEU) and Pu (either weapons-grade or reactor-grade) because both of these nuclear materials could be used to construct an improvised nuclear explosive by subnational groups or a crude nuclear weapon by proliferant states [3, 4]. Due to easier access and a simpler device design, it is thought the crude nuclear weapon would most likely be constructed using HEU instead of Pu [5], while either material could be used to construct improvised nuclear explosives. It should be noted alternate nuclear materials (ANM) (i.e., ^{237}Np and Am) and other actinides could be used as well by proliferant states and subnational groups [4]. For additional information regarding how to assess the risk posed by different types of materials and different combinations of materials by proliferant states and subnational groups refer to [4, 6].

RDDs use conventional explosives to disperse radiological material. The isotopes most likely considered to be used in a RDD are ^{60}Co, ^{137}Cs, ^{192}Ir, and ^{241}Am [5]. These isotopes have medical or industrial purposes and could be more easily obtained compared to other isotopes, and most importantly, these isotopes can be obtained with high activities, which mainly increases the public perception of the risk.

Both the isotopes comprising SNM or considered in RDD emit characteristic gamma-rays, which can be used to detect and localize illicit nuclear and radiological material in urban environments. In the case of SNM, in particular HEU, the emission spectrum consists of low energy gamma-rays such as 186 keV for ^{235}U that require minimal shielding material (e.g. lead) around the SNM to reduce detection efficiency. HEU can have ^{232}U impurities which has a daughter isotope ^{208}Tl that emits a 2614 keV gamma-ray and is more penetrating compared to the 186 keV gamma-ray of ^{235}U. Depending on the shielding thickness, this line could be more easily detected, but it is not a reliable signature [3]. Since the isotopes used in RDD can be obtained with large activities, the gamma-ray signatures emitted from these devices would be harder to shield against detection. Generally, RDD are considered the more likely nuclear threat compared to SNM since the isotopes used in RDD are ubiquitous in medical or industrial settings and RDD require simpler dispersal designs [5]. Neutron signatures could also be used for detecting and localizing potential nuclear threats in urban environments. However, neutron emissions from HEU are weak [3] and none of the RDD isotopes have neutron signatures. For these reasons, this work focuses on detecting and localizing gamma-ray emitting isotopes.

1.2 Challenges for Detecting and Localizing Gamma-ray Sources in Urban Environments

The detection and localization of radiological and/or nuclear material remains a key challenge in homeland security, particularly in urban environments [3, 7]. These environments are dynamic and highly cluttered from both a physical and radiological perspective, making the detection and localization of radiological threats extremely challenging. The detection of radiological and/or nuclear material involves finding a potentially weak or shielded source in an environment where the background gamma ray rate is not well understood. Additionally, the presence of static and dynamic objects makes the attribution of a radiological signal to a particular physical object (e.g. a given vehicle or person) particularly difficult, and in most cases, attribution needs to be performed in a matter of a few seconds. Methods to overcome the source search problem in urban environments can involve using static or mobile detector systems.

Static networked detector systems for urban search and radiological surveillance can be deployed in transportation hubs, cities, and other high risk areas in an effort to improve detection and localization capabilities in urban environments. Network detector systems are capable of running autonomously and use an array of detectors that are usually separated by several meters to broaden the coverage of these sensors because urban environments are unconstrained in extent. Radiation detection is governed by the inverse square law and as the distance decreases between a sensor and a radiological source, the signal-to-noise (SNR) ratio increases for the detector. The SNR changes in a networked detector system can be used to track a radiological source over time [8, 9]. Mielke et al. [10] coupled contextual information (microphone, light sensor, thermistor, and two-axis micro electro mechanical system (MEMS) accelerometer) with radiation sensors to detect the presence of a vehicle and collect radiation data from the triggered sensor.

In comparison to static systems, mobile detector systems rely on the free-movement of the system to survey an urban environment. This enables more autonomy in decision-making when surveying for radiological and/or nuclear material in large-scale urban environments, such as city blocks. For example, an operator could decrease the distance between the detector system and a suspicious object to monitor how the SNR changes with time. Additionally, mobile detector systems can house large amounts of detectors allowing for detection of weaker sources that are at larger standoff distances (greater than 10 m) from the detector system [11, 12]. Penny et al. [11] reported the ability to detect and locate sources in 3D at standoff distances of over 100 m. However, unlike static systems, mobile systems need to be monitored continuously, have sparser coverage in space, and are more conspicuous. Thus, the use of static or mobile detection systems is largely dictated by the needs of the survey mission in urban environments.

Overall, for both static and mobile detector systems, when a radiological and/or nuclear alarm occurs, alarm adjudication for an operator is still difficult given the complex, cluttered nature of urban environments. Furthermore, with mobile systems, the motion of the mobile detector system relative to objects in the scene compounds the issue. The following two sections (Section 1.3 and Section 1.4) discuss a source-object attribution approach that is explored in this thesis to overcome the radiological source search problem in urban environments using static and mobile detector systems.

1.3 Source-Object Attribution on a Static System

Augmenting spectroscopic radiation data with contextual information, such as Light Detection and Ranging (LiDAR) or streaming video, can provide important situational awareness in urban environment scenarios and help inform the interpretation and adjudication of nuclear and radiological alarms. While contextual data can undoubtedly aid the analysis of radiological data by a human operator, recent advances in computer-vision based object detection [13, 14, 15] now make it possible to perform radiological-contextual data fusion in real-time, and provide automatic associations between radiological signals and the physical objects in a scene. This has the potential to improve the speed and efficiency with which radiological/nuclear sources can be detected and localized in complex environments while providing significantly enhanced situational awareness to a human operator.

Previous work has demonstrated the use of video and rudimentary vehicle tracking to enhance the performance of a large gamma-ray imaging system for portal monitoring [16, 17]. More recent advances in contextual sensing, and the processing of contextual data, have enabled the exploration of object tracking and the correlation of the trajectories of objects in a scene with data from isotropic radiation detectors. Such methods have the potential to provide source localization in poorly constrained measurement environments, and without the need for large, complex gamma-ray imagers.

Recently, object tracking and radiological attribution was demonstrated using a LiDAR point cloud projected onto the X-Y plane [18]. This top-down LiDAR projection was used to identify objects in the scene, the 2D trajectories associated with these 2D objects were then correlated with the data recorded by co-located radiation detectors. This correlation was performed using a non-negative least squares method and was used to attribute the radiological data to the person carrying the source. This approach neglects static sources when performing attribution, performs object tracking only in 2D, and is unable to handle dynamic occlusions. Therefore, while they may be adequate in certain situations, these specific methods are likely to offer limited utility in complex, real-world environments.

In this work, 3D object tracking and radiological data fusion using multicamera video and LiDAR is demonstrated. LiDAR and video are analyzed independently and the object tracking performance from these sensors are combined with radiological signals and their performance compared. Object detection is performed using lightweight neural networks. Lightweight neural networks have less layers compared to traditional neutral networks so they are less accurate at classifying objects; however, lightweight neural networks are capable of running in real-time on board a custom-developed system at Lawrence Berkeley National Laboratory (LBNL). Object tracking is performed in three dimensions using a class-specific (e.g. pedestrians, vehicles) Kalman filter. Tracking in a 360-degree field of view (FOV) is achieved by using four cameras and a 64-beam LiDAR. These tools are then used to evaluate the performance of object tracking on radiological alarm attribution and optimized spectroscopic search, with the objective of demonstrating these concepts and their performance in a challenging mock urban environment. It should be noted that given the scenario considered, a mock urban environment with minimal constraints in the 3D motion of people and vehicles and with various occluding objects, a system without video or LiDAR would be limited so no comparisons are drawn between the sensor system and one without video or LiDAR.

1.4 Source-Object Attribution on a Mobile Detector System

The addition of contextual sensors (streaming video and LiDAR) to free-moving detector systems can provide valuable information about radiological source encounters that can assist in adjudication of alarms. Wulf et al. [12] demonstrated the ability of simple contextual information (GPS and a camera) on a mobile detector system to improve situational awareness by overlaying a reconstructed gamma-ray image onto a camera image creating a 2D gamma-ray image. More recent work have explored and demonstrated 3D gamma-ray imaging with free-moving handheld devices [19, 20, 21] by leveraging advances in sensor and computational technology. These methods utilize a set of contextual sensors, such as LiDAR and/or streaming video, in conjunction with radiation sensors and algorithms that produce pose (position and orientation) estimates of the free-moving device in a consistent reference frame. All of the contextual sensor information, radiological data, and pose estimates are processed in real-time to produce 3D visualizations of both the scene and gamma-ray image as the device moves throughout a scene. While these methods can improve situational awareness by enabling 2D or 3D reconstructions of gamma-ray sources, these methods utilize maximum likelihood expectation maximization (ML-EM) to perform 2D or 3D gamma-ray image reconstructions, and ML-EM can suffer from overfitting the radiological data [22, 23] leading to distributed source intensities. In addition, these works focus on static sources in stationary environments and currently are not well suited for dynamic environments with moving sources. Thus, alarm adjudication by an operator in a cluttered environment with dynamic objects would still be difficult to perform quickly and efficiently.

Here the contextual-radiological data fusion concept introduced in Section 1.3 is applied to a free-moving ground vehicle to correlate trajectories from tracked objects with radiological data to better improve localization capabilities of moving sources compared to the conventional reconstruction approaches mentioned previously. This work builds upon the previous work outlined in Section 1.3 by applying the source-object attribution analysis concept on a mobile detector system equipped with video and LiDAR as well as six 2 *in.* × 4 *in.* × 16 *in.* NaI(Tl) detectors in a hexagonal arrangement. One of the objectives of this work is to demonstrate improved situational awareness on a mobile platform system in the presence of dynamic radiological sources in a mock urban environment by fusing contextual information with radiological data to perform real-time source-object attribution. Another objective of the work is to demonstrate that contextual information can be used to improve detection sensitivity of a detector array. These topics are explored using the LiDAR Enhanced Mobile Urban Radiation Search (LEMURS) vehicle.

Additionally, with source-object attribution the source is inherently localized to physical objects in a scene. Previous methods for source localization in urban environments involved using planar arrays of detectors with coded masks on large mobile systems [11, 12]. Cooper et al. [24] demonstrated the LEMURS vehicle (a mobile detector system) that uses a smaller vehicle with a less complicated detector array that relies on inter-detector modulation and contextual data for 3D radiation imaging. However, all of the detector arrays in these systems, along with others, are designed without considering the contextual information available. If the amount of source localization provided by source-object attribution is sufficient, a detection system that performs localization in the conventional sense (i.e., relying on the angular resolution of a detector array) might not be necessary. Thus, source-object attribution has the potential to reduce the cost and complexity of detector array designs. The work presented in this section also explores this topic and how source-object attribution might be considered in the design of detector arrays.

1.5 Dissertation Structure

The dissertation is organized as follows:

- Chapter 2 discusses how gamma-rays interact with matter, the gamma-ray detection process, the source-object attribution analysis pipeline, and a track-informed optimal integration window analysis.
- Chapter 3 demonstrates performing source-object attribution using a static contextual sensor system co-located with a single radiation detector as well as performs a spectroscopic analysis using track-informed optimal integration windows.
- Chapter 4 investigates performing source-object attribution and a spectroscopic analysis using track-informed optimal integration windows on a mobile platform system. The ability of contextual information to influence detector array design is also explored.
- Chapter 5 provides conclusions on the work and includes future outlooks.

1.6 Relevant Publications

Text and figures from the following paper, of which I was the primary author, are included in this dissertation with the permission of all authors:

- **M. R. Marshall**, D. Hellfeld, T. H. Y. Joshi, M. Salathe, M. S. Bandstra, K. J. Bilton, R. J. Cooper, J. C. Curtis, V. Negut, A. J. Shurley, and K. Vetter. 3-D Object Tracking in Panoramic Video and LiDAR for Radiological Source-object Attribution and Improved Source Detection. IEEE Transactions on Nuclear Science, 68(2):189202, 2021. [25]

Other text and figures presented will be included in a future publication.

Chapter 2

Methods for Attributing Radiological Sources to Objects

This chapter first presents how a gamma-rays interact in matter and the gamma-ray detection process in Section 2.1 and Section 2.2, respectively. The subsequent section discusses gamma-ray imaging concepts (Section 2.3). The following sections discuss each stage of the source-object attribution analysis pipeline starting with object detection (Section 2.4), object tracking (Section 2.5), modeling and fitting radiation data to trajectories (Section 2.6), and attributing trajectories to radiological data (Section 2.8). The final two sections – Section 2.9 and Section 2.10 – discuss performing source-object attribution when multiple sources of the same type are present in a scene and using track-informed signal-to-noise ratio (SNR) optimization to improve detection sensitivity, respectively.

2.1 Gamma-ray Interaction Mechanisms in Matter

Gamma-rays interact with matter primarily through three mechanisms: photoelectric absorption, Compton scattering, and/or pair production. This section will briefly describe all three interactions. For more information see [26].

Figure 2.1 displays the mass attenuation coefficients for gamma-ray interactions specifically in NaI(Tl). At low energies (1-300 keV), photoelectric absorption is the dominant gamma-ray interaction mechanism. Above 300 keV Compton scattering becomes the dominant interaction until about 6 MeV where pair production begins to dominate. The majority of gamma-ray energies that are of interest from a nuclear security standpoint are within the energy region where Compton scattering is dominant. It should be noted there are relevant gamma-ray energies below 300 keV; however, these isotopes can be more easily shielded.

Photoelectric absorption is the process where all the gamma-ray energy is transferred primarily to a bound k-shell electron. In photoelectric absorption, the energy is sufficient enough to liberate the photoelectron from the bound electron shell. Once the photoelectron is removed from the electron shell, a hole occupies the electron vacancy. Either the capture of a free electron in the surrounding medium or electrons from predominantly outer shell electrons will move to fill the vacancy. As electrons from outer shells move to fill the vacancy, either Auger electrons or X-ray photons will be emitted from the atom with an energy equal to the difference in the binding energy of the initial and final state. If photoelectric absorption

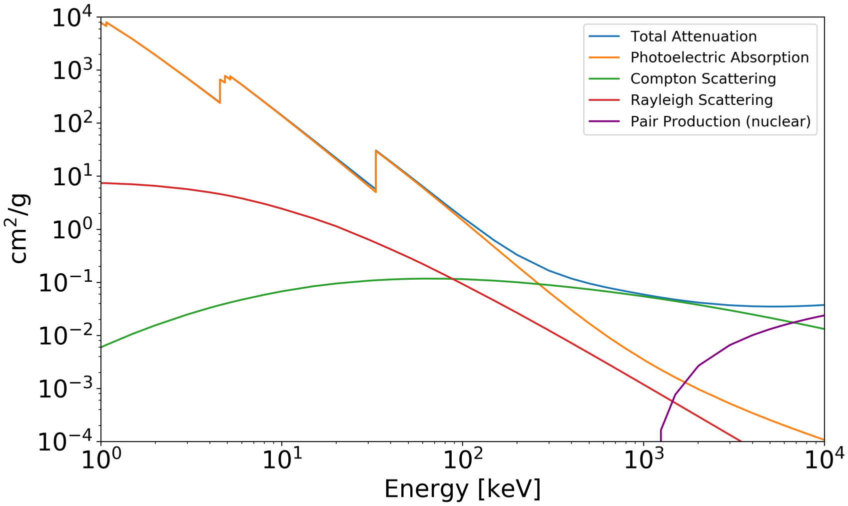

Figure 2.1: Mass attenuation coefficients for gamma-ray interactions specifically in NaI(Tl). Data taken from [1].

occurs near the surface of the detecting medium, the X-ray could escape the detecting medium without interacting resulting in an energy deposition that is less than the incident gamma-ray energy. This tends to occur more often with low energy gamma-ray isotopes such as ^{241}Am (59.54 keV) because the cross-section for photoelectric absorption is proportional to $Z^{4-5}/E^{3.5}$. As seen in Figure 2.1, at lower energies photoelectric absorption dominates compared to all the other interactions. For ^{241}Am, the mean free path in NaI(Tl) is ~0.4 mm so the gamma-ray has a high likelihood of interacting near the surface of the detector. As the energy increases, the probability of photoelectric absorption starts to decrease. With a 661.67 keV gamma-ray from ^{137}Cs, the mean free path in NaI(Tl) is ~3.7 cm; thus, the interaction has a higher probability of occurring deeper in the NaI(Tl) detector away from the detector edge. In this case, the characteristic X-ray will be absorbed more often resulting in an energy deposition that is equal to the incident gamma-ray energy. It is important to note that in order to know the full energy of the gamma-ray, a photoelectric absorption must take place.

In Compton scattering an incident gamma-ray with an initial energy E_{γ_i} interacts inelastically with a bound electron resulting in a scattered and lower energy gamma-ray E_{γ_f} and a recoil electron. The energy of the gamma-ray following the inelastic scatter depends on the incident angle of the gamma-ray (θ) described by

$$E_{\gamma_f} = \frac{E_{\gamma_i}}{1 + \frac{E_{\gamma_i}}{m_e c^2}(1 - cos(\theta))}, \tag{2.1}$$

where $m_e c^2$ is the rest mass energy of an electron (0.511 MeV). The energy of the scattered gamma-ray ranges from small energy transfers ($\theta = 0°$) up to the maximum amount of energy transfer ($\theta = 180°$). Following an inelastic scatter with an electron, the scattered gamma-ray can continue to Compton scatter within the detecting medium. Eventually, the scattered gamma-ray will either escape the detecting medium or will be photoelectrically absorbed. The cross-section for Compton scattering is proportional to Z/E, and as mentioned earlier, Compton scattering is the dominant interaction mechanism for most isotopes of importance to nuclear security.

Pair production occurs when an incident gamma-ray has an energy greater than twice the rest mass energy of an electron (1.022 MeV). The interaction likelihood increases as the incident gamma-ray energy increases as seen by Figure 2.1. In this interaction the gamma-ray energy is sufficient that when the gamma-ray interacts within the coulomb field of a nucleus an electron and positron pair are created. Once the positron slows down, the positron will annihilate with an electron creating two annihilation photons that have energies of 511 keV. The two annihilation photons are emitted approximately 180° from each other to conserve momentum.

2.2 Gamma-ray Detection Mechanisms

In order to attribute radiological signals to objects in a contextual sensors field of view (FOV) gamma-rays need to be detected. This section provides a brief overview on gamma-ray detection primarily focusing on NaI(Tl) since the majority of this work was performed with NaI(Tl) detectors. For more information please refer to [26].

The detection process starts with a gamma-ray interacting within the detector material and transferring some or all of its initial incident energy to an electron. The energy transferred to the electron is sufficient to liberate the electron so that it is free to move in the detecting medium, and a vacancy (hole) in the electron shell is created. Electron-hole pairs are produced along the electron's track as the electron deposits its energy. In NaI(Tl) to improve the scintillation process activator sites are added into the crystal. For NaI(Tl), the activator is thallium. Holes move to occupy and then ionize one of these activator sites, and the electrons will drift through the detecting medium until they encounter an ionized activator site. Once a drifting electron encounters an ionized activator site, the electron will be captured by the activator site creating a set of excited states. The de-excitation of these excited states can result in the emission of a photon in the visible light spectrum. The photons are emitted isotropically, and a fraction of them will make it to the photocathode. From here, the photocathode converts the photons to electrons, and these electrons are multiplied by the photomultiplier tube (PMT) to amplify the signal. The PMT amplifies the electrons such that the amplitude of the voltage pulse is still proportional to the charge deposited. The amplitude of the voltage pulses can be plotted using a histogram to create an energy spectrum (Figure 2.2).

During the scintillation process, it is possible for the electron and hole to drift together through the crystal. In this process, the electron-hole configuration is known as an *exciton*, and the *exciton* will migrate until it encounters an activator site. Similar to the aforementioned process, a set of excited states will be created, and a photon in the visible range can be emitted during the de-excitation process.

As outlined above, the scintillation process is an inefficient process. There are multiple stages throughout the detection process where valuable scintillation photons or photoelectrons can be lost. The following example illustrates the various inefficiencies and how these impact the photoelectron statistics and ultimately the energy resolution. With NaI(Tl), the light yield is about 38,000 photons/MeV and the scintillation efficiency is ~12%. Assuming a 662 keV gamma-ray deposits all of its energy through photoelectric absorption and the average energy of the photons produced through scintillation are 3 eV, ∼26,500 scintillation photons will be produced. Further assuming all of these 26,500 photons reach the photocathode, which will not be the case, a photocathode has a typical efficiency of ∼20% so 5,300 photoelectrons will be generated. The photoelectrons generated from the photocathode are proportional to the deposited charge so the amount of peak broadening is dependent on the photoelectrons leaving the photocathode. Finally, assuming Poisson statistics, the full-width at half maximum (FWHM) will be $2.355 \times \sqrt{5300}$ resulting in an energy resolution of 0.6%. However, clearly there is more that contributes to energy resolution than solely photoelectron statistics because the energy resolution for NaI(Tl) at 662 keV is ∼6-7%.

The other important factor that contributes to poor energy resolution in NaI(Tl) is nonproportionality of light yield. The emission of a visible photon is one of several competing processes (e.g., *quenching*) for an excited state to de-excite back to a stable configuration. This leads to a spectra of electron energies. Therefore, the total light yield can vary for each gamma-ray interaction if the detector does not have a linear response to electron energies.

Both photoelectron statistics and nonproportionality of light yield contribute to an energy resolution that is worse than some other – more expensive – scintillators such as LaBr or semiconductors such as high purity germanium (HPGe) or CZT; however, the advantage of NaI(Tl) detectors is that they are more affordable, available in large volumes, and more durable detectors compared to these other detectors.

2.2.1 Gamma-ray Background

With gamma-ray detection, there is an omnipresent background radiation (hereafter referred to as background) from a combination of terrestrial radionuclides and cosmic origins [27]. This is seen in the energy spectrum displayed in Figure 2.2. The energy spectrum is for ^{137}Cs, which only emits a 662 keV gamma-ray; however, more gamma-ray energies are observed throughout the spectrum, especially at 911.2, 1120.3, 1460, 1764.5, 2204.2, and 2614.5 keV. These energies arise from primordial radionuclides ^{40}K (1460 keV), ^{235}U, ^{238}U, and ^{232}Th and the daughter nuclides of the U – the 1120.3, 1764.5, and 2204.2 energy lines are from ^{214}Bi, respectively, a daughter nuclide of ^{238}U – and Th isotopes – the 911.2 and 2614.5 keV energy lines are from ^{228}Ac and ^{208}Tl, daughter nuclides of ^{232}Th, respectively. Additionally, the background contribution can fluctuate with time. ^{222}Rn is a gaseous isotope that is a daughter of ^{238}U, and the movement of ^{222}Rn can be influenced by meteorological conditions. All of these factors contribute to a background radiation that is not known *a priori* and thus can interfere with detecting radiological anomalies when surveying urban environments. Approaches can be made to minimize the background contribution to the gamma-ray spectrum (e.g., house the detector in lead shielding, *anticoincidence shield*); however, these methods sacrifice precious counts in radiation source search missions. In Section 2.6, an anomaly detection algorithm is briefly introduced that is used in this thesis to detect the presence of anomalies in unknown background situations.

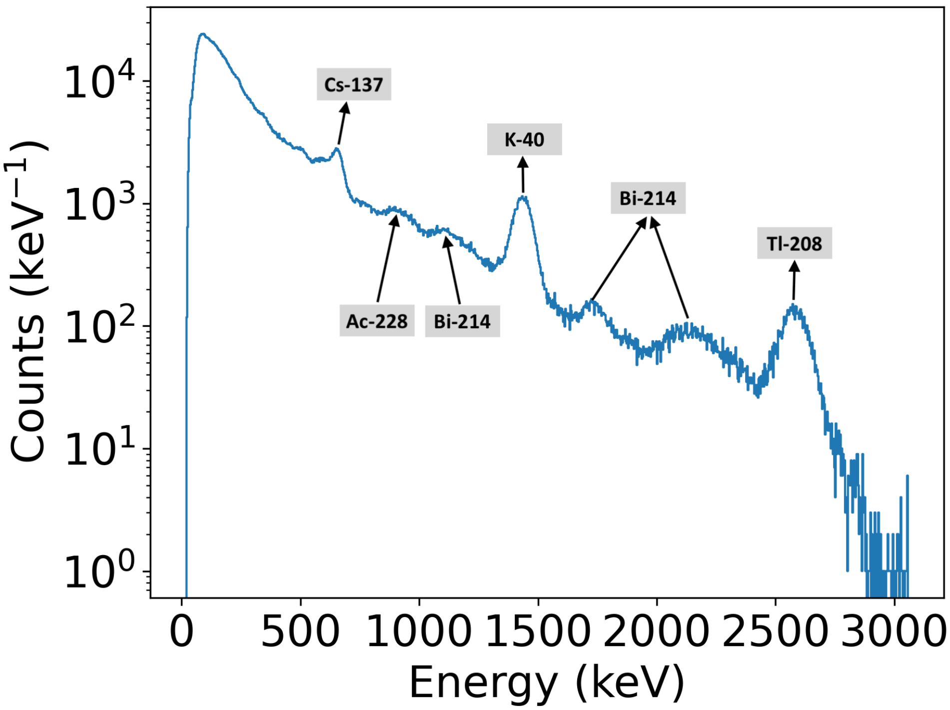

Figure 2.2: A ^{137}Cs gamma-ray spectrum using a NaI(Tl) detector. The ^{137}Cs, ^{228}Ac, ^{214}Bi, ^{40}K, and ^{208}Tl photopeaks are labeled.

2.3 Gamma-ray Imaging

The determination of a detected gamma-ray's origin (i.e. x, y, and z location in space) is referred to as gamma-ray imaging. For nondirectional detectors, source localization can be determined by monitoring count rates as either the detector is moved around an environment or as objects move in front of the detector because detection efficiency increase as the square of the source distance from the detector decreases. This method only works well for constrained environments, and has limited capacity in large, unconstrained environments such as city blocks. More complex methods of gamma-ray imaging use passive or active masks to infer gamma-ray directionality. In the coming sections, a brief overview on the more sophisticated gamma-ray imaging techniques will be given along with a contextual-radiological data approach (source-object attribution) to gamma-ray imaging.

2.3.1 Passive and Active Imaging

Passive gamma-ray imaging involves placing shielding material (i.e. Pb) in front of detectors to collimate the radiological source. In the case of multiple detectors, the shielding material can be arranged in unique patterns so that directionality can be inferred. Depending on the location of the radiological source, a certain *shadow* will be cast on the detectors and the source location can be found. Coded aperture imagers can provide accurate and precise source localization capabilities and have applications in nuclear security [12, 11, 28].

However, there are limitations with this approach. Coded aperture imagers are typically constructed using planar arrays of detectors with a coded mask, and this limits the FOV of the system. Also, the shielding material that is necessary for source localization reduces detection efficiency. In the case of weakly emitting sources or long-standoff distances (> 10 m), long dwell times are needed to get adequate counting statistics in the detectors. Thus, it is a trade-off between detection efficiency and localization with passive gamma-ray imaging systems.

More recent approaches utilize active imagers. With active imagers the detectors are arranged in order to increase detection efficiency while maintaining localization capabilities. In this case instead of using passive shielding material to cast a unique *shadow*, detectors are used as the occluding material and the inter-detector modulation enables source localization. Since detectors are the modulating material, active masks are lighter compared to passive systems and can be handheld. This concept of active mask imaging has been demonstrated using a spherical arrangement of detectors [29] and also a two-plane arrangement of detectors [20]. In addition, active mask imaging has been implemented on a mobile platform system [30]. While active mask imagers provide source localization without sacrificing detection efficiency, the systems typically involve complex arrangements of the detectors and can be expensive since detectors are the modulating material. This dissertation focuses on a new paradigm for source localization that also enables improved detection sensitivity which is described in the following section.

2.3.2 Source-Object Attribution

With recent advances in computer-vision-based object detection [14, 15], it now makes it possible to perform radiological-contextual data fusion in real-time and provide automatic associations between radiological signals and the physical objects in a scene, which enables a new paradigm for source localization and might reduce the complexity in detector array design.

In this work, source-object attribution is explored using panoramic video and LiDAR on either a static contextual sensor system or a mobile detection platform. Objects are detected in image frames or LiDAR scans using open-source object detection algorithms, and a class-specific (e.g. pedestrians, vehicles) Kalman filter is used to convert the series of detections into object tracks. The attribution analysis is triggered by spectroscopic alarming using a method based on non-negative matrix factorization [31]. During the attribution analysis, models that describe the time-dependent photopeak count-rate expected in a detector from each trajectory are calculated and compared with the time-series photopeak data. A model for each trajectory includes the $1/r^2$ profile of the trajectory (r is the distance from the detector to the source), the angular detector response, and gamma-ray attenuation in air and in intervening tracked objects. An estimate of goodness-of-fit is used as a discrimination metric to perform attribution (or exclusion) between the radiological data and the model from each trajectory. Optimized spectroscopic search is performed using the same models from trajectories to identify the temporal segments for each trajectory that provide and optimize the SNR, demonstrating enhanced detection sensitivity via data fusion. Sensor readout and synchronization was handled with the Robot Operating System (ROS).

Figure 2.7 shows an example output from running the source-object attribution analysis pipeline for both video and LiDAR on a static system. Each process of the pipeline will be explained in more detail in the subsequent sections for both the static and mobile system.

2.4 Object Detection

The following two sections provide a brief overview for the object detection convolutional neural networks (CNN) and open-source software used in this work. In addition, the process for converting 2D bounding boxes to 3D bounding boxes is discussed in Section 2.4.1. For more detail regarding object detection CNN please refer to current object detection literature [14, 32].

2.4.1 Video

Object detection from video frame images was performed using the open-source *you only look once* (YOLO) [13] detector and the ROS implementation, YOLO ROS [33], for the static system. In the static work, YOLOv3-tiny - a lightweight model pre-trained on the COCO dataset [34] capable of inference times of 66 ms on the the 4 camera system was used. It should be noted that while higher performing YOLO models exist, such as YOLOv3-608, they cannot be run in real-time on the 4 camera system due to the additional computational burden and limited resources.

For the mobile cases, YOLO ROS [33] was modified to handle YOLOv4-tiny [14]. YOLOv4-tiny was pretrained the on the COCO data set and runs at an inference of 66 ms similar to YOLOv3-tiny [34]. At the time of writing this dissertation, YOLOv4 was the most current yolo version. Additionally, YOLOv4-tiny was not available when the static analysis was performed.

YOLO ROS returns identification labels for the detected objects with confidence scores and 2D bounding boxes in image coordinates (see Figure 2.7a). The distance (depth) of the detected object to the camera is inferred by comparing the height of the 2D bounding box to the height of a person or a car. For a person, a nominal height of 1.75 meter (m) [35] is used and for a car, 1.43 m, which is what is considered to be an average size of a compact car in this study, is used. Heights of 1.80 m, 0.80 m, and 2.5 m are used for labels of truck, motorcycle, and bus, respectively. The 2D bounding boxes are mapped to 3D using

$$Z(depth) = \frac{H \cdot f_x}{H'}, \tag{2.2}$$

where f_x is one of the camera focal length parameters and H and H' are the nominal height of an object and the object height in pixels, respectively. The width in depth is given by

$$\Delta Z = \frac{W \cdot Z}{f_y}, \tag{2.3}$$

where W is the width of the object in pixels and f_y is a camera focal length parameter.

A separate object detection process is performed for each camera that is on the system. The camera frames are synchronized, and the object detections from all the cameras are collated. This is done to avoid double-counting detections that take place in regions where the camera images overlap.

2.4.2 LiDAR

The object detection in LiDAR point clouds was performed using the sparsely embedded CNN detection framework SECOND [15] with the PointPillars fast feature encoding [36]. A multi-class detector model was trained on the nuScenes dataset [37] (32-beam LiDAR spinning at 20 Hz) developed for autonomous driving research. In accordance with SECOND, the model was trained on point clouds generated from 10 LiDAR scans to achieve higher point density for inference. It is worth noting explicitly that because the LiDAR provides data which is inherently three-dimensional, the object detection algorithm is able to return 3D bounding boxes, removing the need to infer the distance to the object that is necessary in the case of video.

In the application presented here for the static system, two scans of the Ouster 64-beam LiDAR generated sufficiently dense point clouds without significant motion blurring from dynamic objects. The inference time using a NVIDIA Xavier GPU is sufficiently fast (~125 ms) to enable the object detection to keep up with the point cloud data (200 ms). The LiDAR object detection returns labelled 3D bounding boxes with associated confidence scores using the same ROS message format as the video results (see Figure 2.7b).

As mentioned, the mobile platform system used in this work was the LEMURS vehicle (discussed in more detail in Section 4.1). LEMURS has two separate LiDAR, and in this case, two scans from each LiDAR were performed to generate sufficiently dense point clouds. The point clouds were moved from each respective LiDAR reference frame into the body-fixed reference frame of LEMURS (hereby referred to as *base_link*) using a Universal Robotic Description File (URDF) file. A URDF is a file that describes the orientation of each component of a system. Figure 2.3 displays the LiDAR and *base_link* orientations from the LEMURS URDF in ROS. The labels correspond with the frame IDs for each component of the system, and *base_link* is the body-fixed reference frame. In this case, the origin and orientation of *base_link* corresponds with the onboard inertial measurement unit (IMU) sensor.

With the URDF, ROS is able to calculate a transformation matrix between any two components in the system. ROS uses the frame IDs from each sensor to lookup the transformation matrix between that sensor and the target coordinate frame. ROS provides the transformation matrix (T) to the target coordinate frame (G) from that of sensor frame (S), and with this information the data (x) is transformed using Eq. 2.4

$$[x]^G = [T]^{GS}[x]^S + [x]^G. \tag{2.4}$$

With the point clouds in *base_link*, the two point clouds can be aggregated into one point cloud when LEMURS is operated as a static system. However, when LEMURS is operated as a mobile system, the two point clouds need to be motion corrected before they are aggregated because they may not align in *base_link*. To do this, the point clouds are moved from *base_link* into a world-fixed frame. The transformation from *base_link* to the world-fixed frame was found using two different methods. In the first method, the pose (position and orientation) estimates of the LEMURS vehicle in a world-fixed frame are found by using the onboard inertial measurement unit (IMU) sensor and running the Google Cartographer simultaneous localization and mapping (SLAM) algorithm [38]. The other method involved using a global positioning system (GPS) stabilized by an Inertial Navigation System (INS) (SPATIAL [39]) data to track the pose estimates of the LEMURS vehicle in a world-fixed frame. The point clouds are aggregated in the world-fixed frame, but when SECOND performs its voxel grid on the aggregated point cloud it assumes the sensor is at the origin, which is *base_link*. If inference is performed in the world-fixed frame, SECOND will have no points to perform detection on. Thus, the aggregated point cloud is transformed back to *base_link* from the world-fixed frame, and the timestamp used to map the point cloud back to *base_link* is the average of all the timestamps from the point cloud messages that were used to generate the aggregated point cloud. Once in *base_link*, inference is performed on the aggregated point cloud for both situations where LEMURS is operated static or mobile.

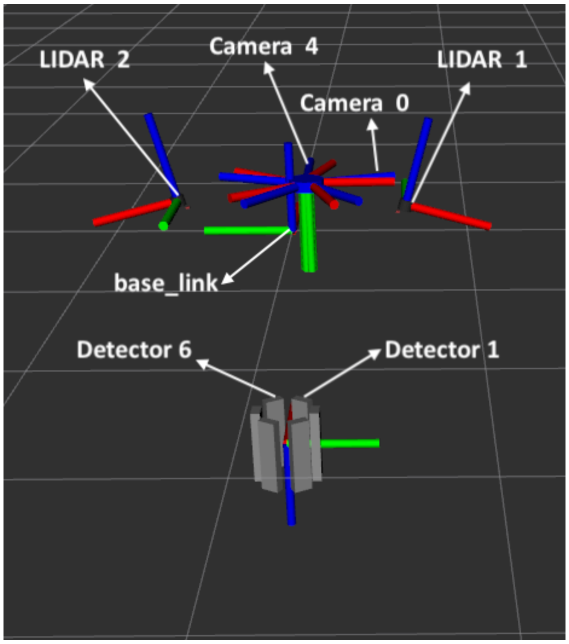

Figure 2.3: URDF file for LEMURS showing only orientations of the two LiDAR, *base_link*, Occam camera, and NaI(Tl) bar array. Certain components were omitted from the visualization. The orientation of each component in the system is indicated by the red (x-axis), green (y-axis), and blue (z-axis) lines. The frame ID of each component is labeled. *base_link* is the body-fixed reference coordinate system of LEMURS. Only two components of the Occam camera are labeled, but the frame IDs increase counter-clockwise from camera 0 to camera 4. The detector IDs increase clockwise from Detector 1 to Detector 6.

2.5 Object Tracking

2.5.1 Overview of Kalman Filter

This section provides an overview of the basic principles of a Kalman filter. The next section explains how the Kalman filter is applied to track detected objects between video frames or LiDAR scans. Refer to [2, 40] for more details.

Figure 2.4 shows the basic principles behind the Kalman filter. The first step in the Kalman filter is the prediction step. The goal of the prediction step is to use the prior state information (posterior) to predict the next state (prior). This is done using a set of linear, kinematic equations that model the expected behavior of each state variable (x) and predicts the next state at time k given the previous state at k-1. The state transition matrix F describes the kinematic equations for each state variable and is applied to the previous state at k-1

$$x_{k|k-1} = F x_{k-1|k-1}, \tag{2.5}$$

which gives the predicted next state. In Eq. 2.5, F is invariant and does not evolve with time. The covariance matrix in the the prediction step (P) is determined with

$$P_{k|k-1} = F P_{k-1|k-1} F^T + Q, \tag{2.6}$$

where Q is the process noise covariance matrix. The process noise tries to account for unknown factors that are not accounted for in the linear, kinematic model. It is assumed the process noise is a zero-mean Gaussian distribution with a known covariance matrix Q. Since the process noise has a zero-mean Gaussian distribution, it was omitted from Eq. 2.5.

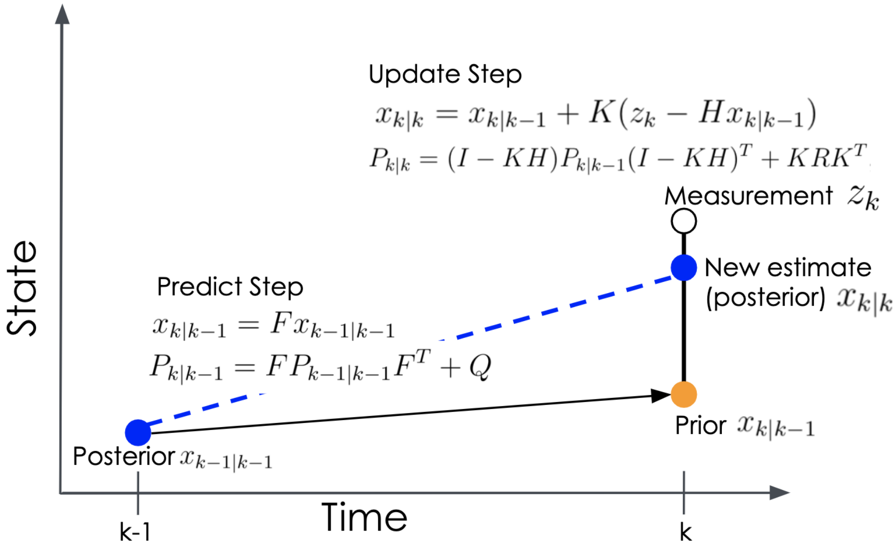

Figure 2.4: Modified from [2]. An example showing the Kalman filter equations and how an object's state is updated using a sensor measurement and an estimate of the object's state.

The update step for the Kalman filter weights the measurement z and the predicted state using the Kalman gain K and places the updated state somewhere between the measurement and predicted state. For example, if there is more confidence in the predicted state than the measurement, the updated state will be closer to the predicted state, and vice versa if the measurement has more confidence. The state is updated with

$$x_{k|k} = x_{k|k-1} + K(z_k - H x_{k|k-1}), \tag{2.7}$$

where H is the measurement function. The measurement function is a matrix that moves the predicted state into measurement space so the difference between the measurement and prediction can be calculated. This is commonly referred to as the residual. The uncertainty in the update step is given by

$$P_{k|k} = (I - KH)P_{k|k-1}(I - KH)^T + KRK^T, \tag{2.8}$$

where R is the measurement noise uncertainty, and I is the identity matrix. Similar to the process noise, the measurement noise is assumed to be a zero-mean Gaussian distribution with a known covariance matrix R - $\mathcal{N}$(0,R).

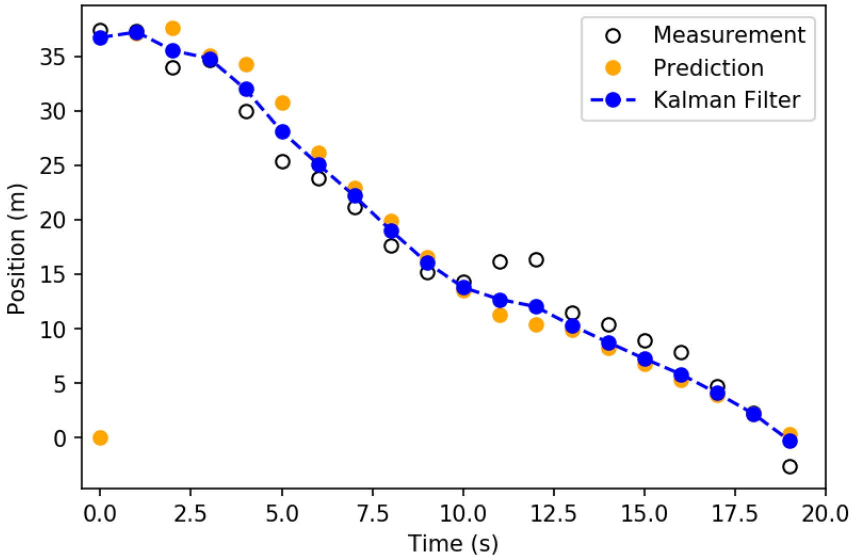

Figure 2.5: An example of a Kalman filter. The color-coding is consistent with Figure 2.4. The Kalman filter uses the sensor measurement and prediction to update the object's state over time.

Figure 2.5 shows an example of tracking an object over time using a Kalman filter and a noisy sensor. The colors and symbols are consistent with Figure 2.4. With the initial measurement, the object's position is uncertain so an initial guess is made and high uncertainty (P) is placed on this prediction. As seen, the Kalman filter gives more weight to the measurement in the initial measurement. With time, the state's covariance P starts to decrease as the Kalman filter becomes more confident in the object's position. This is seen around 11 - 12 seconds when the sensor makes a series of noisy measurements, and the Kalman filter gives more weight to the prediction than the measurement when updating the object's state. For each measurement, the Kalman filter will use the measurement and the prediction model to determine the best estimate for the object's state.

However, Kalman filters can be designed to have too much confidence in the prediction model. In these cases, the Kalman filter is *smug*, and the Kalman filter predictions can start to diverge from the actual measurements because more weight is consistently given to the predicted state than the measurement. To improve the performance of the Kalman filter, Q needs to be adjusted so the Kalman filter adapts to changes in the measurements and is not overly confident in the predictions. A small Q results in a *smug* Kalman filter that is too confident in the predicted model, but when Q is too large, the Kalman filter is slow to adapt to changes that occur. A value for Q is normally found through experimental data, and iterating through different values of Q until a desired Kalman filter performance is found. The process for determining Q in this work is explained in the following section.

An advantage of Kalman filters is the ability to estimate states that are not measured directly through the measurement sensor. These estimated states are called *hidden variables.* Both sensors, LiDAR and video, measure an object's position (x, y, z) but neither measure velocity (v_x, v_y, v_z). However, position and velocity are correlated, with velocity being the first derivative of position, and in this work, a multivariate Kalman filter was employed to track two state vectors: position and velocity. After a few position measurements over time, the Kalman filter can start to get a good estimate of the object's velocity. The covariance between position and velocity arises from $FP_{k-1|k-1}F^T + Q$ in Eq. 2.6. By taking advantage of the correlation between position and velocity, it improves the Kalman filters ability to estimate the object's state compared to relying solely on position.

2.5.2 Tracking Video or LiDAR Detected Objects

To convert sequences of detections into object tracks, the Kalman-filter-based Simple Online Realtime Tracking (SORT) algorithm [41] was used. SORT was extended to process both video and LiDAR detections in three-dimensions. SORT was further modified by replacing the method used to associate detections with the most likely corresponding track. The conventional intersection-over-union (IOU) method was replaced with a multivariate normal (MVN) representation of the 3-D bounding boxes and the Hellinger Distance (HD). The underlying representation of objects in SORT is a Kalman filter which uses an MVN as its state, so this was adopted for tracking, and MVNs are capable of accounting for variable uncertainties in object extent and enabling tracking across multiple cameras. The HD is a measure of similarity between two MVNs that is scaled between 0 and 1, where 0 represents high similarity [42]. It is important to note that other data association techniques exist such as the Mahalanobis distance. However, comparing the performance of different data association approaches in source-object attribution is outside the scope of this dissertation.

To create the MVN representation of a 3D bounding box, the center of the bounding box is converted to the mean of the MVN, and scaled dimensions of the bounding box populate the covariance matrix (P) for position uncertainty. Different scaled dimensions were used for video and LiDAR 3D bounding boxes. Since LiDAR has a relatively high precision of ±1.5-5 cm [43], the applied position uncertainty was 20% of the bounding box in all three dimensions. For video, the position uncertainty was 25% in the width of the bounding box, 33% in length, and 50% in depth. The largest uncertainty in the object's position is the depth because it is estimated using a nominal height.

Additionally, SECOND provides the heading (yaw) of bounding boxes. This information was used to transform the off-diagonal elements of the covariance matrix. For video, no heading information is provided and so the off-diagonal elements of the video covariance matrix are not transformed.

In the modified SORT algorithm, as mentioned earlier, a multivariate Kalman filter was employed to track two state vectors: position (x, y, z) and velocity (v_x, v_y, v_z). The diagonals of the velocity uncertainty in the covariance matrix are populated according to the detection label. For vehicles, 4.44 m^2/s^2 is used and 0.28 m^2/s^2 for pedestrians; thus, enabling the

Kalman filter to simultaneously track both pedestrians and vehicles. The determination of these values is explained later in this section.

Data association between the detected objects and the most likely corresponding track is performed using the HD. HD is calculated with the following equation

$$\text{HD} = \sqrt{1 - \frac{det(\Sigma_1)^{1/4} det(\Sigma_2)^{1/4}}{det(\frac{\Sigma_1+\Sigma_2}{2})^{1/2}} exp\Big(-\frac{1}{8}(\mu_1 - \mu_2)^T \Big(\frac{\Sigma_1 + \Sigma_2}{2}\Big)^{-1} (\mu_1 - \mu_2)\Big)}, \tag{2.9}$$

where μ_1 and μ_2 represent the MVN mean for the detected object and track, respectively, and Σ_1 and Σ_2 represent the MVN covariance matrix for the detected object and track, respectively. In Eq. 2.9, Σ_1 and Σ_2 are 3×3 matrices that have been moved into measurement space by

$$HP_{k|k-1}H^T + R, \tag{2.10}$$

which is often referred to as the system uncertainty. It is worth explicitly noting that in Eq. 2.9 the distance metric calculation is determined in the sensor frame so μ_1 and μ_2 are the center of the bounding boxes with Σ_1 and Σ_2 as the position covariance matrices.

Using Eq. 2.11, Eq. 2.9 is calculated for each possible detection and track pair creating a matrix that is $D \times T$, where D is the total number of detections present and T is the total number of available tracks:

$$\sum_D \sum_T \text{HD}_{DT}. \tag{2.11}$$

To assign the detected object with its most likely corresponding track, linear assignment (also known as the Hungarian Method) is applied to the matrix [44]. Subsequently, if a detection and a track have a calculated HD of less than 0.8 they are consolidated to a single track. The velocity uncertainties and HD thresholds for pedestrians and vehicles were found by running an optimization over those variables on a scene only containing either pedestrians or vehicles. The objective was to reduce the number of objects with disjoint tracks while also ensuring a low number of incorrect associations by the data association algorithm.

In order to track objects across multiple cameras, the MVNs are transformed from the image frame into either the *base_link* or world-fixed coordinate frame prior to computation of the HD for the static and mobile system, respectively. Additionally, the LiDAR-tracked objects are transformed into the *base_link* coordinate frame so all the trajectories are in one coordinate system. The data is transformed using Eq. 2.4, and the covariance matrix (c) is transformed by

$$[c]^G = [T]^{GS}[c]^S\overline{[T]}^{GS}. \tag{2.12}$$

With a mobile system, tracking and attribution needs to be performed in a world-fixed frame and not the body-fixed frame of LEMURS. Objects in the scene should be invariant to the motion of LEMURS which is not the case in a body-fixed frame. To generate pose estimates of LEMURS in a world-fixed frame, INS or SLAM information is processed, and pose information is produced at a rate of 10 Hz. It is thought that a system that has a navigational system that relies on GPS to produce pose estimates of the system's location in a global frame would have degraded tracking and attribution performance in an urban environment compared to applying SLAM. Urban environments are cluttered with buildings that can occlude or reflect signals from GPS satellites to the LEMURS system. To test this concept, the tracking and source-object attribution performance of using SLAM or INS to produce pose estimates of the LEMURS system in a world-fixed frame are compared in a mock urban environment using a mobile detection system.

Additionally, higher position uncertainty is assigned to objects near the edges of the camera images to improve tracking capabilities across multiple cameras. This is done because the cameras are not perfectly calibrated and the camera parameters can vary between cameras. Therefore, when an object passes from one camera to another, the object's estimated position can vary such that the calculated HD considers the new estimated position too dissimilar to the previous position and will spawn a new track. To determine the detected object's position in the image, the image width and center are calculated, and the center of the 2D bounding box in the image coordinate frame from the object detection CNN is determined. If the center of the bounding box is more than ∼33% away from the center of the image, the object is considered near the edge of the image, and an additional scaling factor is applied to the position uncertainty in addition to the aforementioned scaling factors.

Figure 2.6 shows an example of applying higher position uncertainty at the edge of images. The bounding boxes in the images are MVN representations of bounding boxes and the extent of the bounding boxes reflect the position uncertainty. When the object is near the edge of the image, their position uncertainty is much larger compared to when the object was near the center of the image.

(a)

(b)

Figure 2.6: Comparison of position uncertainties. While the object is near the edge of the image (a) the position uncertainty is higher, but once the object approaches the center of the image (b), the position uncertainty decreases. The bounding boxes are MVN representations of bounding boxes and the extent of the bounding boxes reflects the position uncertainty.

Using the object's label, different scaling factors are applied to pedestrians and vehicles. For people, the position uncertainty in width, which is the direction of travel between cameras, is scaled by 350%, and the uncertainty in depth is scaled 200%. For vehicles, the extent of the object is larger so the position uncertainty in width is scaled by 700% and depth is scaled 200%. However, with increased uncertainty at the edges this approach does result in more incorrect associations especially at the camera edges.

Small FOV overlaps exist between the cameras and therefore it is possible for multiple cameras to independently detect the same object. However, once the MVNs are transformed from the image coordinate system into the global frame, these bounding boxes will overlap. To prevent the same object from being tracked more than once, these overlapping bounding boxes are merged if the calculated HD is less than 0.6. This merging of bounding boxes supplements the nonmax suppression being performed with YOLOv3-tiny in each independent camera frame.

The tracker output is presented to an operator using the object's label (e.g., person or vehicle) and a track ID, where the track ID increases by 1 for each newly generated track. An example is shown in Figure 2.7c and Figure 2.7d for both video and LiDAR, respectively, using the static contextual sensor system. The bounding boxes displayed are MVN representations of bounding boxes and will no longer match the bounding boxes output from the object detection CNN. Also, it should be noted that an object's orientation is not considered so all the bounding boxes will have the same orientation.

Given the nature of the object detectors, false detections can occur so the track ID will normally be higher than the total number of objects that have been present in the system's FOV. Additionally, untracked objects and detections not being assigned right away to the correct track contribute to a higher total track ID number. To limit the number of tracked objects, a track is discarded from the Kalman filter if the track is not associated with a detection in 5 consecutive frames, and a new track will be spawned for the object if it is detected again.

2.6 Modeling and Fitting Radiation Data to Trajectories

The source-object attribution process is invoked when a radiological anomaly is observed. The presence of an anomaly is determined by the Berkeley Anomaly Detection (BAD) algorithm, a spectroscopic detection algorithm based on Non-negative Matrix Factorization (NMF) [31]. BAD provides high sensitivity detection, even in low signal-to-background regimes and serves as a robust trigger for the source-object attribution algorithm. This algorithm also provides an isotope ID which is used to define a spectral region of interest (ROI) for the extraction of time-series count-rate data. The source-object attribution analysis is performed on all trajectories that are within 7 seconds of the start and stop of the radiological alarm in the case of a static system. When source-object attribution is performed using a mobile platform system, trajectories within 2 seconds of the radiological alarm are modeled and fit to the count-rate data. For more information about BAD refer to [31].

In the case of multiple detectors, BAD is run independently on each detector, and if a radiological alarm is triggered for any detector, the attribution analysis starts for all detectors.

2.6.1 Generating Models based on Trajectories

The first step in the source-object attribution process is the generation of a set of models that describe the time-dependent count-rate that would be expected to be observed were each track to be associated with the detected anomaly. To model the expected counts in the detector, the following approach is used. For a given discrete time step, i, the expected number of detected events, c_i, within a spectral ROI, E, from a radioactive source with gamma-ray flux α in the presence of a constant background b can be described by

$$c_i(E) = \frac{\epsilon(\hat{\Omega}, E)\alpha e^{-\mu(E)\mathbf{r}_i}}{4\pi\mathbf{r}_i^2} \cdot \Delta t_i + b\,, \tag{2.13}$$

where ϵ is the effective area of the detector, $\mathbf{r}_i$ is the distance from the detector to the source, Δt_i is a given integration time, and μ is an energy and medium dependent linear attenuation coefficient. The effective area is a function of energy and the direction between the tracked object's position and the detector, $\hat{\Omega}$. In the current implementation, ϵ takes the form of a pre-computed response matrix which is used as a look-up table based on the spectral ROI and the direction between the object and the detector at any given time. Details of the process through which this response matrix was generated are provided later in Section 2.6.1.1.

The total attenuation coefficient, μ, describes the combination of attenuation in air and in occluding objects in the scene. Before this can be applied, it is necessary to determine whether a given object track is subject to occlusion from another object. To do this, the MVNs are converted back to 2D bounding boxes, and the IOU is calculated. If the IOU is greater than 90%, an object in the scene is considered to be obstructing another object. To determine which object is occluded, the distance from the system to each object is calculated using the estimated positions from the Kalman filter and the object farthest from the system is considered to be the one subject to occlusion. However, when objects pass within 1 m or less of each of other, their Kalman filter estimated positions can overlap, especially using video trajectories where depth is inferred, making it difficult to accurately determine which object is occluded in the scene. This is a limitation of the current attenuation approach.

The total attenuation imposed by the occluding object is then determined by its size and an estimated linear attenuation coefficient which is based on the label associated with the object. For example, pedestrians are modeled using the average elemental composition of a human (65% Oxygen, 18.5% Carbon, 9.5% Hydrogen, 3.2% Nitrogen, 1.5% Calcium, 1% Phosphorus) [45] while vehicles are modeled as hollow Aluminum boxes with 2 in thick sides. In the current implementation, additional attenuation from other components of the vehicle are not modeled and the orientation of the vehicle is not accounted for. This results in a slight underestimate of the total attenuation in some situations.

Once the attenuation has been calculated, the size of the bounding box is subtracted from the object's total distance from the system, and the remaining distance to the detector is considered to be air. This process is then repeated for each object in the image frame or LiDAR scan. This allows the model to account for cases in which a given object is subject to occlusion by multiple objects. It is important to note that if an object is occluded for more than 5 frames, a new track will be spawned for the object when it is detected again.

Accounting for the effect of occlusion when computing the best-fit models improves the source-object attribution performance by allowing tracks which would otherwise be degenerate, e.g., those which overlap closely in time and space, to be distinguished. A limitation of the current approach is that it accounts only for attenuation in tracked objects. Including attenuation in the environment would likely improve the source-object attenuation power in particularly complex scenes.

An additional attenuation effect that can occur is anisotropic shielding from the tracked object. To better handle anisotropic shielding from tracked objects, the following reasonable assumptions are made for both pedestrians and vehicles. With pedestrians, the radiological source is assumed to be in a backpack behind the object; whereas, with a vehicle, the source is assumed to be inside the vehicle as opposed to on top of it. In order to find how attenuation changes as a function of angle relative to LEMURS, simulations were run using the Monte Carlo simulator MEGAlib [46]. A pedestrian was modeled using the normal composition of a human, and a vehicle was modeled using a 1 m thickness of Aluminum (Al) for the hood of the vehicle and 5 cm thickness of Al for the doors. A ^{137}Cs source was placed either behind the pedestrian or inside the vehicle. A NaI(Tl) detector was moved in 5° azimuthal increments around either object at a constant elevation in line with the source. The amount of attenuation present at each azimuth was calculated and applied to Eq. 2.13 for a given angle relative to LEMURS. It should be noted that the heading of each object is necessary for this calculation and thus can only be performed with LIDAR detected objects.

2.6.1.1 Directional Response for the Static and Mobile System

For the static contextual system, the directional response matrix that describes the detector's energy and effective area was generated using a Geant4 simulation [47]. The matrix was populated by modeling a 2 *in.* × 4 *in.* × 16 *in.* NaI(Tl) detector in air and, for each of several discrete energies, placing a gamma-ray source at 3072 equal-area discretized pixels on a sphere [48] at a constant distance of 10 m (R^2_{sim}) from the detector. At each source position, the total number of counts within the relevant ROI (X_{cnts}) was computed. For a given energy and source position (i.e., direction), ϵ is given by

$$\epsilon = \frac{4\pi \cdot R^2_{\text{sim}} \cdot X_{\text{cnts}}}{N_{\text{SimParticles}}}, \tag{2.14}$$

where $N_{\text{SimParticles}}$ is the total number of particles emitted into 4π.

In order to validate these simulations, experimental measurements were made using a ^{137}Cs source placed at several positions around the detector. The validation process is discussed in more detail later in the dissertation (Section 3.2).

For mobile detector systems, different response functions with varying levels of fidelity were used. The first response function was generated using MEGAlib [46], a Monte Carlo method based on Geant4, which was modeled with a high fidelity. It should be noted the response function was not validated with experimental measurements. The MEGAlib model included the detectors, vehicle, operators, electronics, and mechanical supports [30]. The response function was generated by modeling a radiological source 10 m from the detector array center and 1.3 m off the ground, which corresponds with the center of the array in elevation. The source was moved in 10° increments in azimuth around LEMURS, and at each source position (i.e., direction), 5 × 10^{10} particles were simulated. Finally, to generate a 4π response function, the response function was cosine modulated. A detailed description and characterization of the response function for LEMURS is given by [30].

The other response function is a lower fidelity model of LEMURS and was generated using a zero-energy approach, which does not account for any gamma-ray physics (i.e. no scattering and infinite attenuation), with OpenGL [49]. OpenGL enables fast generation (≈ 30 sec) of 4π response functions. The 4π response function was generated placing a gamma-ray source at 3072 equal-area discretized pixels on a sphere. At each location, 2.25×10^6 uniformly spaced parallel rays were randomly generated along a circular plane and simulated. The visible surface area of each modeled surface at a given source location is determined by assigning a unique 8-b red value to each surface and histogramming the pixels by color. The generation of the OpenGL response functions are described in more detail in Section 4.5.

Finally, an isotropic directional response was generated by assigning the same detector response for all of the 3072 equal-area discretized pixels on a sphere, and this is the lowest fidelity model.

2.6.2 Fitting Models to Radiation Data

The best-fit model for each trajectory is found by extracting the time-series ROI count-rate in a time window around the radiological alarm, and maximizing the Poisson likelihood between Eq. 2.13 and the observed count-rate data with a maximum likelihood estimation algorithm [50], where α and b are free parameters. An overview of the maximum likelihood estimation algorithm is provided in the following section.

For the example in Figure 2.7, the images in panels Figure 2.7e and Figure 2.7f show the best-fit models for the tracks extracted from the video and LiDAR data, respectively, along with the measured count-rate. It is seen that the simple physics equation in Eq. 2.13 is quite effective at modeling the count-rate data from the detector because a clear association exists between the best-fit model for the object responsible for the radiological alarm (Track 5 in Figure 2.7e and Track 5 in Figure 2.7f) and the radiological data.

When modeling and fitting trajectories to multiple detectors, the radiation data is concatenated to a single array, and models for each detector for a given trajectory are generated and also concatenated to a single array. A global best-fit model is found by fitting the modeled trajectory array to the radiation data array.

2.6.3 Maximum Likelihood Estimation Maximization

The objective of Maximum Likelihood Estimation Maximization (ML-EM) is to find an estimated source distribution that best fits the measured radiological data. ML-EM is an iterative algorithm that converges to a solution, and it enforces that the estimated source distribution counts are non-negative. The following explanation gives an overview of the ML-EM algorithm.

To find the modeled time-dependent count-rate expected in the detector from each trajectory in Eq. 2.13, an initial estimate of the source distribution λ at a point in time i and iteration number n is made. The estimated source distribution is moved into detector-space by the system matrix A for each voxel k:

$$\sum_{k=1}^{K} A_{kl}\lambda_k. \tag{2.15}$$

The result from Eq. 2.15 is compared with the measured counts g indicating how the initial source distribution estimate should be updated for the next iteration:

$$\frac{g_i}{\sum_{k=1}^{K} A_{jk}\lambda_k}. \tag{2.16}$$

This value is projected back into the imaging space by A for each detector j and is sensitivity corrected. Finally, the initial estimate is updated. The full ML-EM equation is given by

$$\lambda_i^{n+1} = \frac{\lambda_i^n}{\sum_{j=1}^{J} A_{ji}} \sum_{j=1}^{J} \Big(\frac{g_i A_{ji}}{\sum_{k=1}^{K} A_{jk}\lambda_k} \Big). \tag{2.17}$$

The first iteration of ML-EM is a back-projection, and the choice of iteration number is arbitrary. Typically, the ML-EM iteration number is determined through several experiments with known outcomes. In this work, 100 iterations was chosen for fitting modeled data to experimental data.

2.7 Binomial Down-sampling Spectra

Most of the alarm encounters considered in this dissertation involve a 1.87 mCi ^{137}Cs source and close proximity between the source and radiation detector(s). This results in high count rates in the detector(s) even for transient alarm encounters. In order to better study weakly emitting sources, which is a goal in radiological source search, binomial down-sampling of the spectra was used. The counts in each spectrum are binomial sampled with a probability of acceptance which is the desired fraction of the original spectrum. This results in fewer counts in the new spectrum and preserves Poisson statistics.

For a given alarm encounter, the radiological data is a combination of source (S) and background (B) resulting in a total number of events $E = S+B$. To extract the source contribution during the alarm encounter, an estimation of the background contribution is needed. The estimated background contribution was found by averaging spectra over 10 second time windows 20 seconds before and after the alarm encounter. This estimated background contribution spectrum (B_E) was then Poisson sampled over time. Subsequently, the estimated source contribution (S_E) in a spectrum is found by $S_E = E - B_E$, and the estimated source contribution was binomial down-sampled (S_d). Thus, the background contribution during the alarm encounter is preserved. Finally, the down-sampled source contribution is added back to the spectrum $E_d = S_d + B_E$.

2.8 Attributing Trajectories to Radiological Data

The final step of the attribution analysis identifies the trajectories that are unlikely to be associated with a radiological alarm through the use of an estimated goodness-of-fit metric. The Poisson deviance or log-likelihood ratio statistic [51] was selected as the appropriate statistic for estimating goodness-of-fit. Its distribution is approximated by using the best-fit model to estimate its first three moments and adjusting the moments according to the number of model parameters (e.g., the mean is decreased by 1 for every model parameter). A shifted Gamma distribution that matches these moments was then used to calculate a p-value from the deviance statistic given this distribution. Finally, an S-value [52] was calculated from the p-value and used hereafter. The S-value is the $-log_2$ of the p-value, and it enables rejection of trajectories that are inconsistent with the data.

In the case of multiple detectors, an overall global goodness-of-fit metric is calculated from fitting the modeled trajectory array to the radiation data array.

The computation assumes that the source is located at the Kalman filter position estimate. If the source is not located at this position, such as in the case of a source located in the trunk of a vehicle when the position estimate is at the center of the vehicle, then an offset between the best-fit model and the photopeak count-rate could occur. To account for this potential offset between the assumed and true source location, the best-fit model is calculated multiple times over a two second window while the track is shifted in 0.1 second increments. The Poisson S-value is computed each time and the lowest value is used to determine the best fit. It should be noted that this only accounts for source offsets in the direction of travel, but does not account for potential source offsets in the other two dimensions that could exist.

With the attribution analysis, there are cases when an object is not tracked continuously during the radiological alarm. In these cases, tracks before or after the alarm are associated mostly with flat background, which, with α set to zero in Eq. 2.13, can always be well described by the model, and thus cannot be excluded from the analysis. To determine whether the best-fit model is responsible for the radiological alarm or if it is fit to background, a likelihood ratio test is performed between a background only model - α set to zero in Eq. 2.13 - and a source plus background model. The preferred model to describe the best-fit model, either a background only model or source plus background model, is selected using the Bayesian Information Criterion (BIC) [53]. The BIC is given by

$$\text{BIC} = log(n)k + 2\ell(\mathbf{x}|\boldsymbol{\lambda}), \tag{2.18}$$

where k is the estimated number of parameters in the model, n is the length of the modeled fit, and $\ell(\mathbf{x}|\boldsymbol{\lambda})$ is the negative log-likelihood of observing the counts within a spectral ROI $\mathbf{x}$ given the expected number of events within a spectral ROI $\boldsymbol{\lambda}$. In this case, $\boldsymbol{\lambda}$ is the best-fit model for a trajectory and $\boldsymbol{x}$ is the count-rate over the span of the trajectory. The negative log-likelihood is determined by

$$\ell(\mathbf{x}|\boldsymbol{\lambda}) = [\boldsymbol{\lambda} - \mathbf{x} \odot log\boldsymbol{\lambda} + log[\Gamma(\mathbf{x}+1)]]^T \cdot \mathbf{1}, \tag{2.19}$$

where $\odot$ denotes element-wise multiplication and $\Gamma(\cdot)$ is the gamma function.

It is determined whether the source term is merited by setting $k = 1$ or $k = 2$ in Eq. 2.18 for the background only model and the source plus background model, respectively. The model with the lower BIC better describes the best-fit model. In the cases where the source term is not merited and the best-fit model is consistent with a background only model, the best-fit model is flagged as background, and thus that track most likely is not responsible for the radiological alarm.

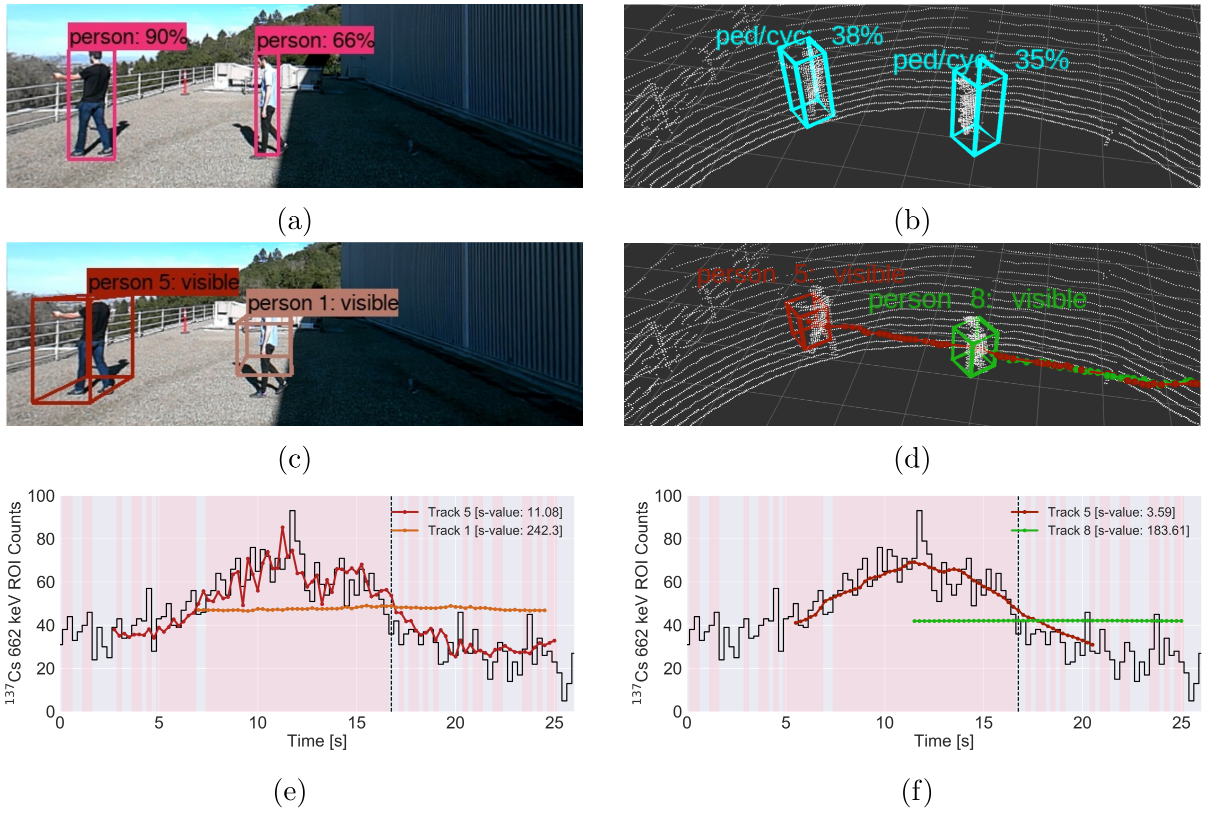

(a) (b)

(c) (d)

(e) (f)

Figure 2.7: Example showing the object detection, tracking, and attribution pipeline for both video [(a), (c), and (e)] and LiDAR [(b), (d), and (f)]. In the scene, two individuals are walking in a straight line past the detector at a standoff of 5 m at closest approach, and one of the individuals is carrying a 189 μCi ^{137}Cs source. The bounding boxes shown in (a) and (b) are the object detection CNN outputs for video and LiDAR, respectively, with the object label and confidence score. In (b) ”ped/cyc” refers to pedestrians or cyclists. The LiDAR CNN groups both labels into the same category. The bounding boxes displayed in (c) and (d) are the object tracking outputs, and the bounding boxes are MVNs converted back to bounding boxes and will not match the bounding boxes in (a) and (b). In (d) the object's trajectory up to that point is indicated by the line proceeding each respective bounding box. In (b) and (d) the white grid squares represent $1\,m^2$, and the area of no points exist from the field of view of the LiDAR. In (e) and (f) the best-fit models for each trajectory from (c) and (d), respectively, to the count-rate data are displayed. In both (c) and (d), Track 5 was carrying the ^{137}Cs source. The pink bars represent the time interval for a radiological alarm, and the count-rate from the $2\ in. \times 4\ in. \times 16\ in.$ NaI(Tl) detector is shown in black. Additionally, the dashed lines in (e) and (f) depict the moment in time these images were taken from.

2.9 Additive Track Source Localization

The formulation of the source-object attribution analysis mentioned in Section 2.6.1 assumes a single point source and does not account for multiple radiological sources of the same type. Recent work from [23, 22] demonstrates a mathematical framework for additive point source localization (APSL) that could be applied to scenarios where multiple trajectories in a scene are carrying the same source type. In [23, 22], a continuous 5D space consisting of source weights w_s, source background b, and position ($\vec{r_s}$ (x, y, and z)) is searched over to minimize the loss of Eq. 2.19. The source position is optimized over continuous space using a non-convex optimizer, and ML-EM is run at each trial source position for the optimal w_s and b. Additionally, the additive nature of Poisson random variables allows the inclusion of multiple, known source contributions into the formulation of the problem.

With tracking information, $\vec{r_s}$ is constrained to the track location so the space that needs to be optimized over is w_s and b. In this case, a track is modeled and fit to the radiation data, and the negative log-likelihood is calculated both for the modeled fit and background only. If the calculated BIC value is lower for the modeled fit then the background only, the track is accepted. Subsequently, the w_s for the track is fixed and another track is added. The new track is fit to the radiation data given the fixed source weight of the first track. Using Eq. 2.18, the BIC is calculated with the addition of the new track. BIC penalizes the addition of new trajectories by setting $k = N + 1$ in Eq. 2.18 where N is the number of parameters estimated by the model (source intensity and background). If the BIC for a two-track-fit is lower than the single track, the new source is accepted. The algorithm then performs a check to determine whether one of the trajectories is weakly contributing or low intensity. The trajectory is excluded from the analysis if this is the case. Finally, if a new track was accepted as being responsible for the radiological signal, the source weights for each trajectory are re-optimized. This process is continued for the number of trajectories present.

If more than one track describes the data, all the tracks that best describe the data are flagged, and the source-object attribution for those tracks is performed using the summed fit from all the contributing trajectories. Thus, all the trajectories that describe the data will have the same S-value score.

2.10 Track-informed Signal-to-Noise Ratio Optimization to Improve Detection Sensitivity

With knowledge of an object's trajectory, it is possible to identify a combination of time-segments that, when combined, maximize the expected signal-to-noise ratio (SNR) for spectral analysis of that object. For this analysis the SNR for a trajectory is considered, present across N discrete time windows, as,

$$\mathrm{SNR} = \left(\sum_{i}^{N} s_i \Delta t_i\right)\left(\sum_{i}^{N} b_i \Delta t_i\right)^{-1/2} \tag{2.20}$$

where $s_i = \epsilon(\theta, E)\alpha e^{-\mu(E)r_i}/(4\pi r_i^2)$ is the photopeak count-rate within a fixed integration window Δt at a point in time i, and b$_i$ is the background rate within Δt_i. From this equation it can be seen that SNR may be maximized by considering only a subset of time-segments ($T \in [1, N]$). Since the radiological sources of interest are long-lived (i.e., half-lives on the order of hours or greater), the source strength is considered constant. The mean background rate is also considered constant in time. With these assumptions, the sensitivity ($\$_T$) which is proportional to the SNR_T is described by

$$\mathrm{SNR}_T \propto \$_T = \frac{\displaystyle\sum_{i \in T} \frac{\epsilon(\hat{\Omega}, E)\Delta t_i}{4\pi r_i^2} e^{-\mu(E) r_i}}{\sqrt{\displaystyle\sum_{i \in T} \Delta t_i}}. \tag{2.21}$$

Eq. 2.21 may then be used to solve for the subset of measurements T, that maximizes \$ and, in turn, maximizes SNR for the trajectory in question. Two observations about this formulation are noted. First, the optimal time windows do not need to be contiguous. Second, under the assumption of constant source strength and background, the optimal integration windows are independent of source strength and background and are entirely determined using tracking information.

To account for the position uncertainties of each track in the optimal integration window formulation, Markov Chain Monte Carlo (MCMC) [54] was applied to the data to appropriately sample from the position uncertainties to better determine the optimal integration window. MCMC is a method that draws samples directly from the posterior probability density function (PDF) distribution [55]. MCMC does this by creating M walkers that explore the parameter space and generate models of the data at each position. The walker vector is defined by

$$\theta_i = \begin{pmatrix} x_i \\ y_i \\ z_i \end{pmatrix}, \tag{2.22}$$

where θ_i is the position at a discrete time step i for a given trajectory. The priors for each walker are the position uncertainty in x_i, y_i, and z_i around each respective mean value. The θ_i parameters are modeled to the count-rate data using Eq. 2.13, and source activity α and background are extracted from the best-fit model for the track.

The initial guess for θ_i is the best-fit model for the track. A total of 400 iterations with 600 walkers are run. However, the first 100 iterations are discarded because the walkers start close to the initial guess before fully exploring the parameter space. This results in 180,000 samples. The optimal integration window is calculated for a subset of the 180,000 samples along with the negative log-likelihood using Eq. 2.19. The model that produces the largest integration window but lowest negative log-likelihood within the subset is chosen as the optimal window.

The formulation in Eq. 2.21 is for a single radiation detector; however, it can easily be extended to account for multiple detectors. With multiple detectors, the model from each detector along with the detector ID associated with the model are concatenated together. The subset of measurements T, that maximizes \$, is calculated using the concatenated data, and the signal from each detectors respective optimal window is combined together to produce an optimal spectrum.

In addition, track uncertainty can be accounted for with multiple detectors. A random subset of samples is chosen from the 180,000 samples. The optimal integration window for each detector is summed together to produce an overall integration window. Similar to above, the largest summed integration window with the lowest negative log-likelihood within the subset is chosen as the optimal window.

To test the optimal integration window concept for either a single or multiple radiation detectors, a spectroscopic analysis to either fixed integration windows (1.0, 2.0, 3.0, 4.0, and 5.0 seconds), a common analysis approach, or to each trajectory's optimized integration window is applied. The spectroscopic analysis computes an anomaly value, the Poisson deviance between the observed data and a mean background spectrum scaled to match the observed counts. Subsequently, the magnitude of the anomaly value for the different integration windows are compared, a larger anomaly value for the optimized integration window suggests improved detection sensitivity through this track-informed analysis. This comparison is presented as a proof-of-concept, and it should be noted that quantitative evaluation of increased detection sensitivity would require appropriate handling of anomaly thresholds, which must vary with the number of statistical tests performed to maintain a fixed False Alarm Rate (FAR). Such a quantitative analysis would, in turn, require definition of strategies for when to used fixed integration windows (because tracking is not perfect) and when to use track-optimized windows.

Chapter 3

Static system: Contextual Sensor Package

In the following sections (Section 3.3 – Section 3.6), the methods for object detection, tracking, and attribution on a static system, previously described, are demonstrated. First, the static system is described in Section 3.1, and in Section 3.2 the generation of the effective area for the NaI(Tl) detector is detailed. Then the ability of the attribution analysis to discriminate degenerate tracks is evaluated by determining the minimum separation between two tracks (in the worst case scenario of identical time of closest approach and velocity) required to allow correct identification of the trajectory associated with the radiological alarm (Section 3.3). Subsequently, the performance of the source-object attribution analysis is demonstrated in a mock urban environment (Section 3.4). Using this experimental data, the potential for tracking-enhanced detection sensitivity is demonstrated in Section 3.5 by showing spectroscopic anomaly values can be increased using track-informed integration windows. Finally, in Section 3.6 using synthetic data multiple radiological sources of the same type are simulated in a scene and source-object attribution is performed.

3.1 Static system: Contextual Sensor Package

The source-object attribution analysis analysis pipeline was implemented on a custom developed platform which is shown in Figure 3.1. This system includes an Ouster 64-beam LiDAR unit [43] and four AR1335 monocular cameras [56] arranged to provide a 360-degree panoramic view around the system. A NVIDIA Jetson AGX Xavier computing platform with 512-core Volta GPU [57] is used to readout the camera and LiDAR sensor and perform on-board data processing. The system is designed as a standalone contextual sensor package (CSP) and can be used to augment a range of radiation detectors. In this work, a 2 *in.* × 4 *in.* × 16 *in.* NaI(Tl) gamma-ray detector was co-located with the system. The detector was instrumented with an Ortec DigiBASE [58] multi-channel analyzer and configured to publish gamma-ray spectra at 20 Hz. The system was powered by a 296 Watt-hour (24 V) battery. Sensor readout and synchronization was handled with ROS [59].

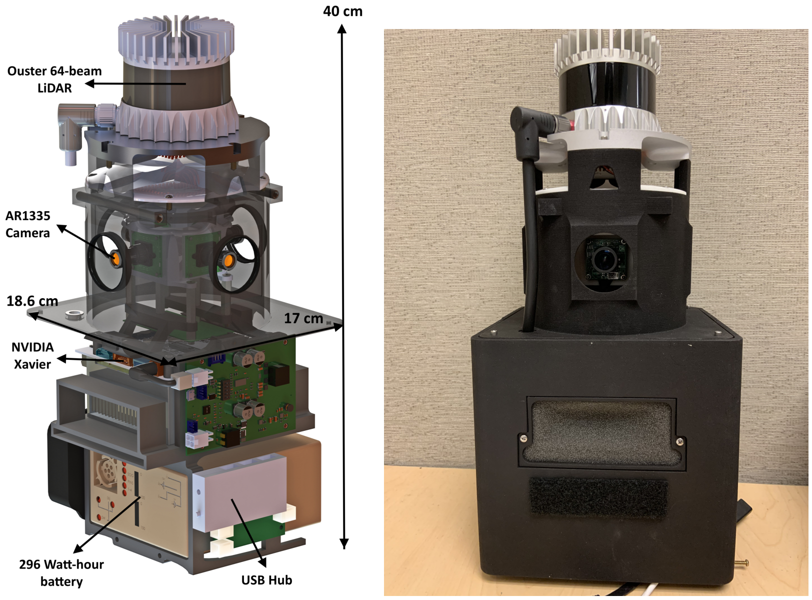

Figure 3.1: Object detection, tracking, and source-object attribution system. The system consists of an Ouster 64-beam LiDAR that is mounted above 4 AR1335 monocular cameras to provide a full panoramic field of view. A NVIDIA Xavier is located in the base of the setup for on-board processing. The system is powered by a 296 Watt-hour (24V) battery.

3.2 Geant4 Modeling and Benchmarking of Directional Response

A directional response matrix describes the expected detector response and efficiency to a particular radiological source from a given location. This information was found by modeling different radiological sources at different source locations around the NaI(Tl) detector in air using Geant4. The simulated responses were generated for photon energies at 122 keV (^{57}Co), 662 keV (^{137}Cs), 1173 keV (^{60}Co), and 1332 keV (^{60}Co). For each photon energy, the source was placed a constant 182 cm from the edge of the NaI(Tl) detector and 3072 different source locations around the NaI(Tl) detector were modeled. At each source location, 1,000,000 gamma-rays were simulated using a directional cone-beam directed at the NaI(Tl) detector to increase the efficiency of the simulations. The number of gamma-rays that would have been emitted isotropically into 4π was found by determining the fraction the cone subtends of the 182 cm source distance sphere.

The simulations were validated using the experimental setup shown in Figure 3.2. A 10μCi ^{137}Cs source was placed 182 cm from the detector edge and was elevated to be 20 cm above the benchtop. The NaI(Tl) detector was then rotated along its azimuth from the 4 in. × 16 in. detector-face side (0°) to the 2 in. × 16 in. detector-face side (90°) in 30° increments. At each source location data was taken for 5 minutes. The measurements (i.e., rotating the detector from 0° to 90°) were repeated multiple times to account for variations in moving the NaI(Tl) detector. Before and after the measurements, 30 minutes of background data was collected.

An experimental and simulated spectrum generated from when the source was in front of the 2 in. × 16 in. side of the detector for ^{137}Cs is shown in Figure 3.3a. The experimental

Figure 3.2: Experimental setup of a NaI(Tl) bar and ^{137}Cs source to find the photopeak efficiency of the detector at different source angles.

spectrum is background-subtracted using the 30 minutes of background data collected and is normalized by the emission rate of the radiological source. The simulated spectrum is also normalized by emission rate and was re-sampled using a normal distribution to agree with the experimental data. For a given energy location, the mean of the normal distribution was the energy, and the standard deviation at that particular energy was found from an energy resolution model that was generated from performing the experimental measurements with ^{57}Co, ^{137}Cs, and ^{60}Co. It is seen there is good agreement between the two spectrum; however, there is a discrepancy in the the magnitude of the photopeak between the simulated and experimental spectrum.

In order to benchmark the simulations and experiments, the photopeak efficiency was compared. The photopeak efficiency was found by fitting a Gaussian model to the photopeak and using the area of the fit. Figure 3.3b shows an example of fitting a Gaussian model to the simulated photopeak. In addition to the Gaussian model, a line is fit to the photopeak to remove excess background from the photopeak. This is more applicable for the experimental data.

The results of the photopeak efficiency between the experimental and simulated data for the different source locations is shown in Figure 3.4. There is agreement between the experimental data at each source location. However, there is a discrepancy between the Geant4 simulations and experiments. To ensure the discrepancy was not due to a modeling error, the simulations were repeated using MEGAlib. Both the Geant4 and MEGAlib simulations are within a two standard deviation statistical agreement indicating the discrepancy is not a modeling error. Understanding the discrepancy between the experimental and simulated photopeak efficiency is outside the scope of this dissertation, but surface roughening by the crystal manufacturer is believed to be the cause of the issue. Curtis et al. [30] provides a more detailed explanation. In order to get the experimental and simulated data to agree, an energy dependent correction factor was applied to each directional response to account for

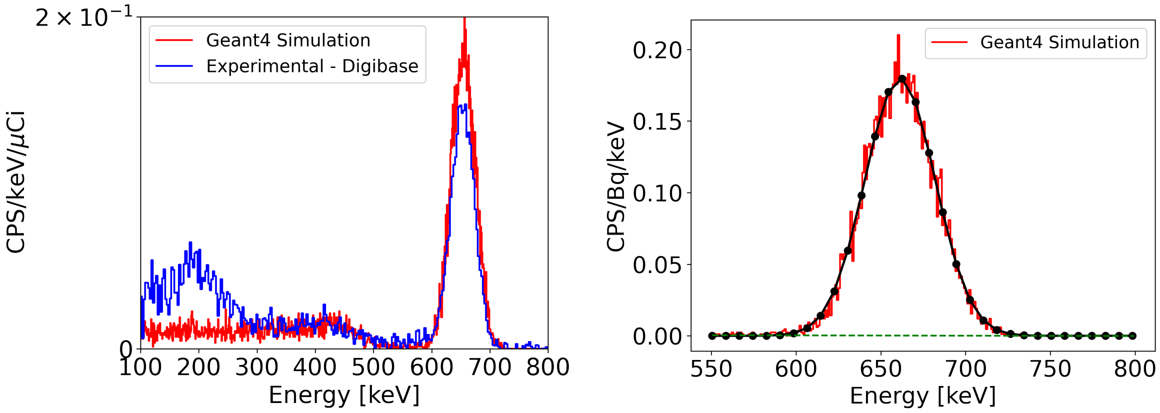

(a) (b)

Figure 3.3: Comparing an experimental (blue) and simulated (red) spectrum both generated with a NaI(Tl) detector and a ^{137}Cs oriented in the middle of the 2 in. */times* 16 in. side of the detector shown in (a). The spectra are normalized to the emission rate of the ^{137}Cs source, and the experimental spectrum is background subtracted. In (b), a Gaussian model plus a line model was used to fit the photopeak of the simulated ^{137}Cs source.

the offsets.

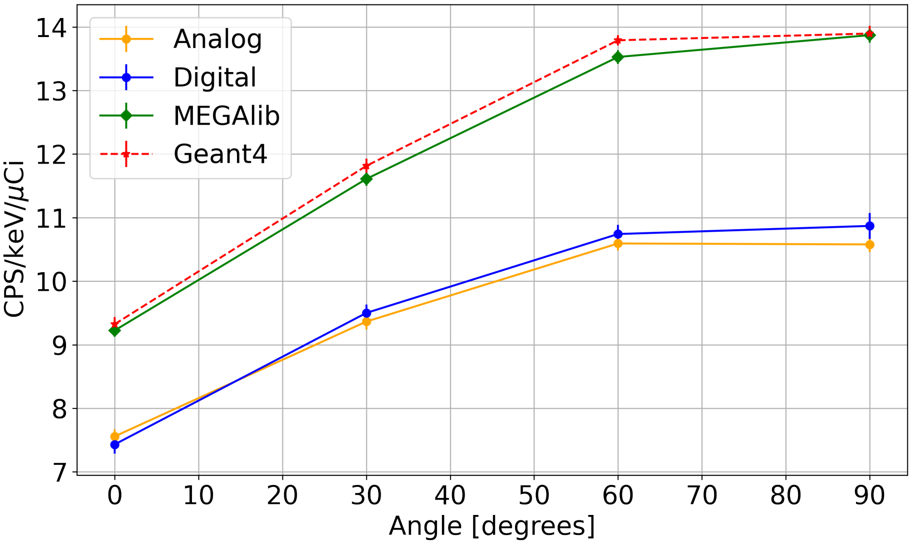

Figure 3.4: Comparison of photopeak efficiencey between experimental measurements (Digibase and analog) and simulation (Geant4 and MEGAlib). The experimental data is within statistical variation at all source angles, and the simulated data is within statistical variation at all source angles except for at 60°. The photopeak efficiency is within two standard deviations at this particular location but not when experiment and simulation are compared.

3.3 Track Discrimination in Source-Object Attribution

Most encounters with a radioactive source traveling along a straight trajectory will result in count-rate distributions as a function of time that mimic $1/r^2$, the falloff of detection efficiency with the square of the source distance from the system, and will exhibit the highest count-rate at closest approach. The worst case situation from an attribution perspective are two objects traveling at identical speeds and reaching the closest approach distance simultaneously. However, even slight variations in closest approach distance and the angular sensitivity of the detector prevent full degeneracy.

To judge the performance of the attribution pipeline in this type of scenario, a systematic study was performed using a 189 μCi (6.993 MBq) ^{137}Cs source to determine the distance needed between two nearly degenerate tracks to correctly identify the trajectory associated with the radiological alarm. In this analysis occlusion/attenuation from objects was not present/considered because the trajectories for this study were created by a single person walking in front of the system (Figure 3.5). Radiation data and object trajectories were collected from 3-8 m distance of closest approach in 1 m increments. For each source distance, the ^{137}Cs source was walked past the system at a walking speed of approximately 0.30 m/s, and at each distance, the person carrying the source walked in front of the system 8 times. This analysis was done by using the photopeak ROI (614 keV - 685 keV) for ^{137}Cs.

Figure 3.5: A person walking 5 m at closest approach in front of the static contextual system carrying a ^{137}Cs source.

The radiation data and trajectory from the 5 m source distance at closest approach were chosen as reference. The trajectory at 5 m was replaced with the trajectories from 3, 4, 6, 7, and 8 m to determine how well these trajectories would correlate with the radiological data produced by the 5 m trajectory. The peak count-rate from each trajectory was aligned temporally with the radiation data at 5 m to ensure the trajectories were compared in the most degenerate scenario. Figure 3.6 shows an example of fitting a trajectory 3 m at closest approach to radiation data generated from a source passing by the detector at 5 m at closest approach. Track 1 was generated using LiDAR. The peak is aligned temporally with the radiation data and there is some correlation between the best-fit model for Track 1 and the count-rate data. However, comparing the best-fit model with a 5 m best-fit model (Track 5 [source carrier] in Figure 2.7f), there is reduced quality in the best-fit model for Track 1 indicated by a S-value of 21.8 compared to 3.59 for 5 m.

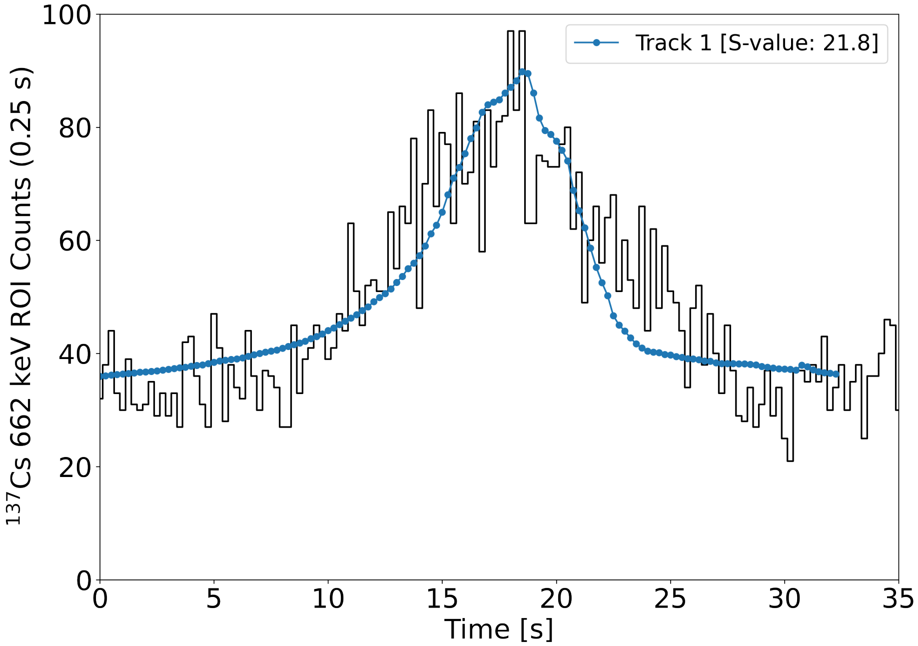

Figure 3.6: A trajectory at 3 m aligned temporally and fit to radiation data generated from a source passing by the detector at 5 m closest approach.

To supplement the limited statistics of these measurements, synthetic count-rate data and trajectories were generated. At each distance from the system, the mean and standard deviation of the position uncertainty for each trajectory was calculated as a function of both distance and angle from the system. These experimentally measured position uncertainties were used to noise simulated trajectories. Count-rate data were sampled according to Poisson statistics from the expected number of detected events calculated through Eq. 2.13. 1000 trials of this random sampling process were then subjected to the same analysis as the experimental data.

The Poisson S-values from analysis of these experiments, comparing 5 m count-rate data with different trajectories, are shown in Figure 3.7. The results show the experimental trajectories from video and LiDAR are both within one standard deviation of each other and their synthetic trials.

S-value thresholds, 11.75 for video and 5.25 for LiDAR, were defined based on a 95% true positive rate for 5 m synthetic trajectories with the 5 m synthetic count-rate data. Under these criteria, degeneracy is successfully broken 45% (58%) of the time for video (LiDAR) trajectories with 1 m difference from the true trajectory. When there is a 2 m separation, degeneracy is successfully broken 99% (99%) of the time for video (LiDAR) trajectories. These results demonstrate the ability of an object tracking and alarm attribution analysis, using either video or LiDAR, to be quite robust against track degeneracy within 2 m for a 5 m distance of closest approach.

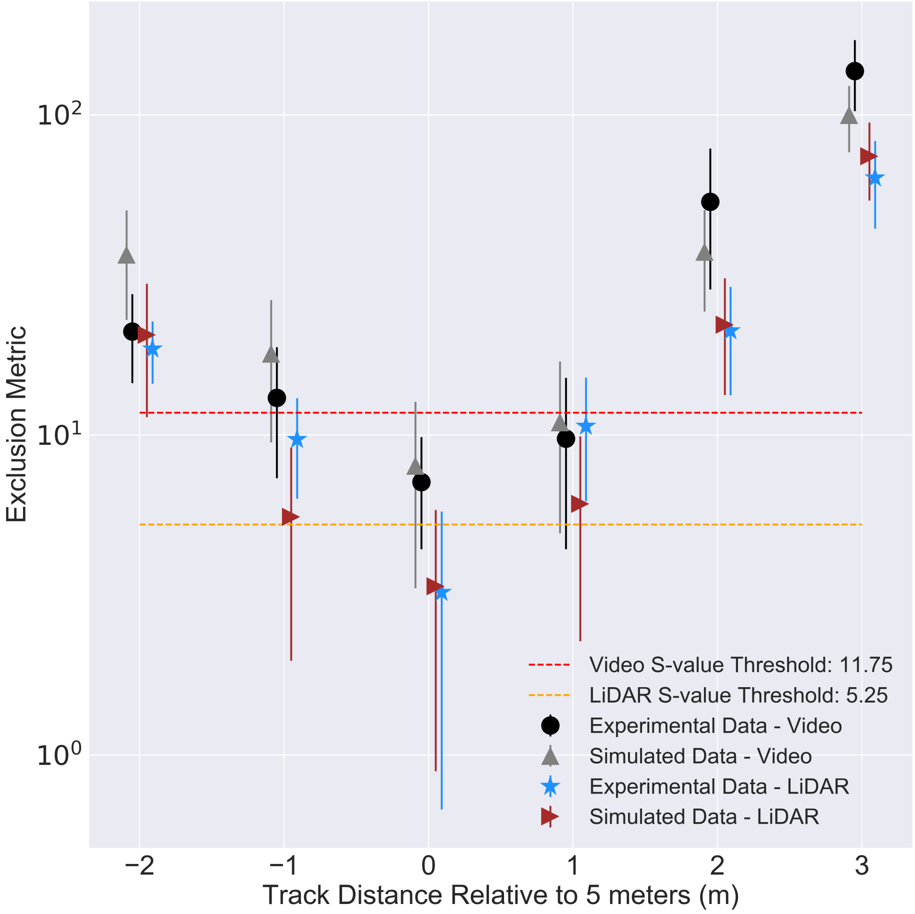

Figure 3.7: Discrimination of orthogonal trajectories in the alarm attribution analysis relative to 5 m with both LiDAR and video trajectories. The simulated data represents synthetic trajectories and count-rate data that was randomly sampled for 1000 trials and subjected to the same analysis as the experimental data. The large uncertainties in the synthetic trials reflects the sensitivity of the goodness-of-fit metric. The exclusion metric is the S-value.

An additional phenomena that modulates the radiological signal is attenuation, and in dynamic scenes one would expect objects occluding a particular object to also attenuate a radiological signal associated with that object. This could help to break the degeneracy in tracks, particularly, when two objects cross in front of the detector system. Synthetic data were used to study how occlusion, in conjunction with 3D object tracking, might further improve attribution in scenarios where degenerate tracks are present.

For this analysis, two objects with similar stand-off distances and walking velocities were considered, such that the two objects crossed paths at the distance of closest approach. Figure 3.8 demonstrates the experimental setup this analysis is capturing with synthetic trajectories. Two sets of trials were performed to test how occlusion might further improve attribution. In the first set of trials, the source-carrier was 5 m from the system at closest approach, and the non-source-carrier was either 6 m or 7 m from the system. For the second set, the source-carrier was again 5 m from the system at closest approach while the non-source-carrier was either 3 m or 4 m from the system. To simulate two objects crossing paths, synthetic detections were generated, and random noise based on the derived track uncertainties from LiDAR or videos was applied. The detections were then fed into the MVN SORT algorithm to create tracks. This procedure accounts for the data association algorithm potentially switching track IDs when the two objects cross paths. The count-rate data was simulated based on the object at 5 m using the same approach as outlined in the previous section. 1000 trials were simulated for each case within each trial scenario.

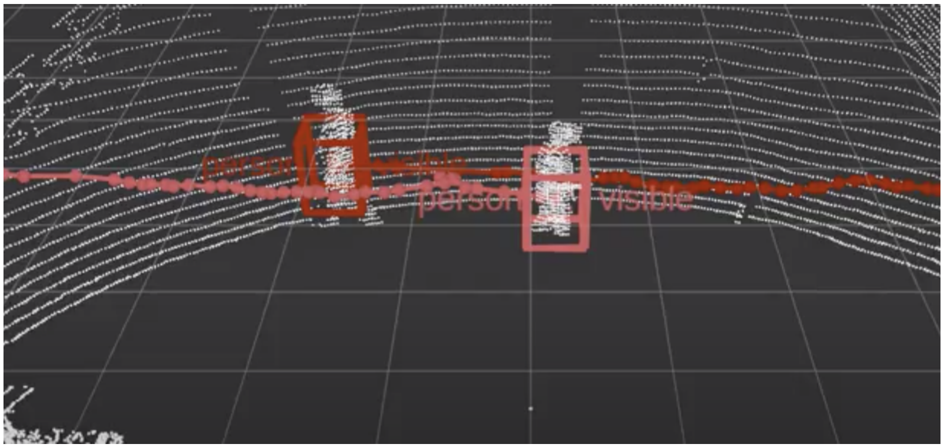

Figure 3.8: An example of two trajectories walking degenerate paths and crossing in the middle of the scene. Synthetic trajectories were created to replicate the experimental setup shown in this figure.

The attribution analysis was again applied to this synthetic data, with the inclusion of an attenuation factor in the response calculation. The applied attenuation factor accounted for either air or people. The attenuation factor for people was applied for portions of trajectories where objects were being tracked but occlusion was occurring. In the first set of trials, using the S-value thresholds of 11.75 (5.25) for video (LiDAR), degeneracy was considered successfully broken if the source-carrier at 5 m had a S-value score below the S-value threshold while the other trajectory at 6 m or 7 m was above the same threshold amount. The results of this analysis show that when the two objects are separated by 1 m, degeneracy is successfully broken 39% (69%) of the time for synthetic data with track uncertainties derived from video (LiDAR). When the distance between the two objects is 2 m, degeneracy is broken 82% (96%) of the time for video (LiDAR). For the second case when the source was occluded by the non-source-carrier, the results show when the two objects are separated by 1 m, degeneracy is successfully broken 33% (63%) of the time for synthetic data with track uncertainties derived from video (LiDAR). With a separation of 2 m between the two objects, degeneracy is broken 74% (84%) of the time for video (LiDAR).

In theory, by accounting for the effect of occlusion in the attribution analysis, the ability to break degeneracy should increase. However, here it is observed that the results are only better for LiDAR at 1 m separation in both cases where the source was and was not occluded. In the other cases, by including MVN tracking through SORT, the position uncertainties sometimes lead to wrong data associations or disjoint tracks, which negates the gains from occlusion and results in worse performance. Additionally, when the source-carrier was occluded, the ability to break degeneracy is diminished compared with the results where the source-carrier was not occluded. In this scenario, the best-fit model underestimated the amount of attenuation present in some of the trials resulting in a S-value that was higher than the threshold.

While the results, including attenuation, are not better compared to the earlier results, they provide a more realistic picture of crossing scenarios, where often, it is challenging to track through an occlusion without losing or misidentifying the two tracks, relying on only HD as a metric. Without this problem, occlusion should be an effective tool to deal with track degeneracy. This topic was going to be explored in greater detail but due to the pandemic caused by COVID-19, two people could not get close enough to truly disentangle the two effects of attenuation and tracking through temporary occlusions.

3.4 Source-Object Attribution in a Mock Urban Environment

To evaluate the system's source attribution performance for encounters with multiple vehicles and pedestrians simultaneously, data was collected at the University of California's Richmond Field Station (RFS) in a mock urban environment. The contextual system, co-located with the detector, were placed at one corner of an intersection containing traffic lights and crosswalks (Figure 3.9). The detector was oriented vertically with the 4 *in.* × 16 *in.* face towards the center of the intersection. Long-dwell measurements, spanning ∼30 minutes, were performed with dynamic pedestrian and vehicle traffic with vehicle speeds ranging from 0-9 m/s. During data collection a 1.87 mCi (69.19 MBq) ^{137}Cs source was placed in the trunk of a vehicle (white station wagon) and was shielded by 1.5 cm of lead reducing the activity of the ^{137}Cs source to 318 μCi (11.78 MBq). The following analyses are associated with individual source encounters triggered by the spectroscopic alarm.

An example of a single alarm encounter, showing multi-camera object tracking and attribution analysis results, is shown in Figure 3.10. The colors used for each trajectory are consistent across the sub-figures. Figure 3.10a shows images from the four cameras (at different times throughout the alarm encounter), overlaid with the 3D bounding boxes of the labeled object trajectories, illustrating the consistent tracking of the white vehicle through the encounter. A top-down view of the object trajectories during this encounter are shown in Figure 3.10b, with the FOV of the cameras also shown. The estimated track positions have large variations especially for objects more than 10 m away from the system. Both Track 5 and Track 19 have particularly large variations at an X position of 10-20 m and a Y position of 5-10 m. These tracks were stationary vehicles at the beginning of the alarm encounter. However, the noise in the trajectories is expected given the heuristic approach to depth estimation used for visual tracking.

The best-fit results from the attribution analysis and the conclusion of the encounter are

Figure 3.9: The contextual system, co-located with a NaI(Tl) detector, at the corner of an intersection at the University of California's Richmond Field Station (RFS).

shown in the top plot of Figure 3.10c. The majority of tracks are clearly excluded as having an association with the alarm, but three tracks (Track 7 (white vehicle), Track 20 (pedestrian in the background of Figure 3.10a image B), and Track 5 (black car)) show consistency with the radiological signals. Of these, the source-carrier (Track 7) is tracked for an extended period of time, while the other trajectories were not continuously tracked throughout the alarm encounter and cannot be excluded.

The temporal evolution of the attribution analysis scores (S-values) are shown in the bottom plot of Figure 3.10c. Short trajectories often have little modulation of $1/r^2$ and are thus likely to result in a model that fits the radiological data. Furthermore, before and after the alarm window, tracks are associated mostly with flat background, which, with α set to zero in Eq. 2.13, can always be well described by the model. However, continuous and consistent tracking of objects throughout an alarm encounter allows for effective attribution (or exclusion).

In this encounter, it is seen that the total number of objects present (6) does not correspond with the total track ID number, which is 20 by the end of the alarm encounter. This is caused from both a high number of untracked objects and detections not being assigned right away to the correct track. Short trajectories, which is defined in this dissertation as 5 or less Kalman filter poses, are more easily fit to the count-rate data and therefore, cannot be excluded from the analysis. In an effort to reduce false positives in the attribution analysis, the attribution analysis is only performed on tracks that have 3 or more Kalman filter poses.

In Figure 3.11a the results from the same source-encounter, but now using LiDAR tracking are shown. The color-coding of the tracked objects (but not their labels) has been kept consistent with Figure 3.10. The accurate depth estimates from the LiDAR result in substantially smoother object trajectories. In addition, it is seen again that the track ID's are higher than the total number of objects present in the scene, and this is due to false detections from the object detection CNN and detections not being assigned right away to the

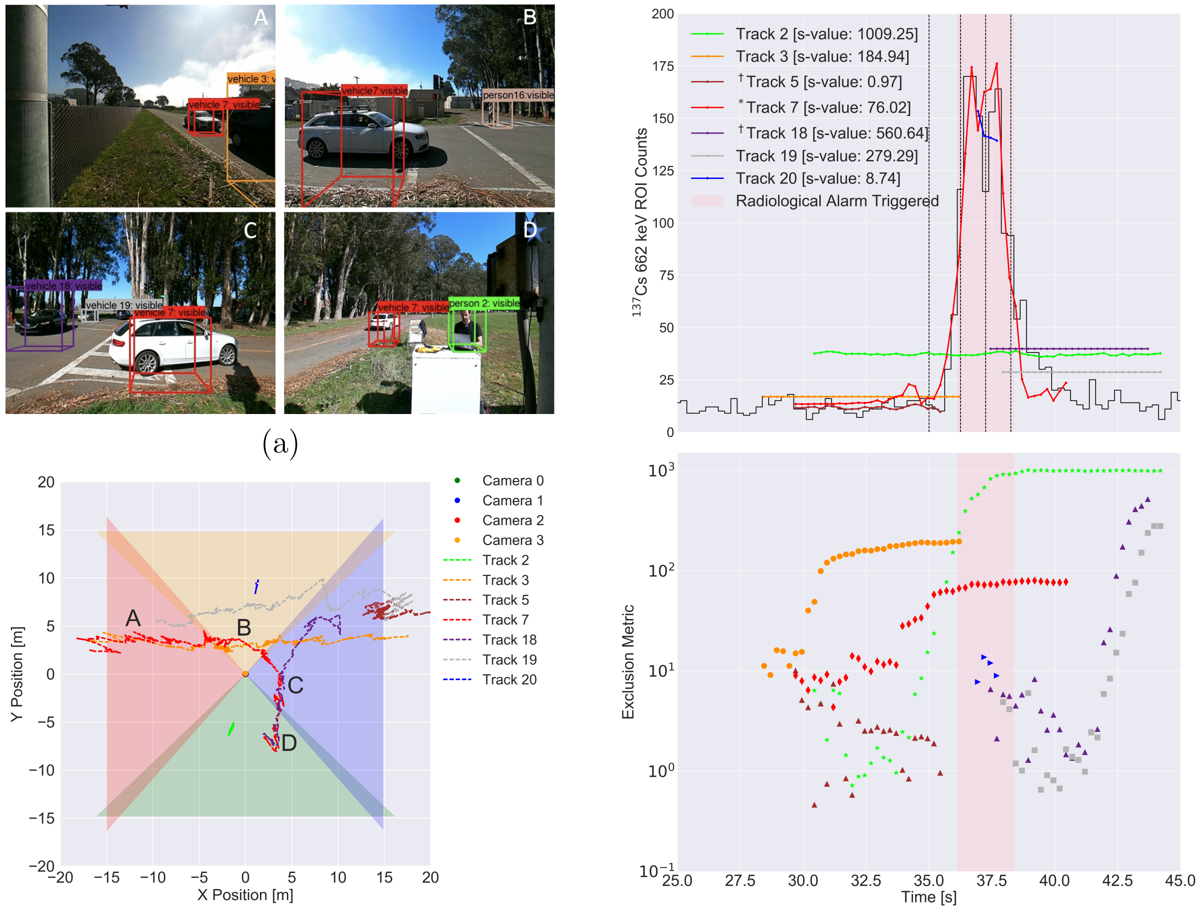

(a)

(b)

(c)

Figure 3.10: Source-object attribution in a mock urban scene with both pedestrians and vehicles using camera data. The progression of the scene as the vehicle - Track 7 (the white station wagon) - carrying the radiological source moves past the system is shown in (a). The associated trajectories for each object in the scene with each camera's FOV overlaid on top of it is displayed in (b), and the different letters correspond with (a). In (b), Track 20 corresponds with the pedestrian in the background of (a) image B. Track 5 and Track 18 are examples of disjoint tracks, and both tracks correspond with the black car in (a) image C. The black car is not tracked continuously through the occlusion temporarily caused by Track 3. The top plot in (c) is the associated best-fit models of each object in the scene to the ROI count-rate data from the 2 $in. \times 4$ $in. \times 16$ $in.$ NaI(Tl) detector, which is shown in black, and the dashed lines correspond with (a) starting with the upper left image, then upper right image, then lower left image, and finally the lower right image. The * indicates the source-carrier, and the † indicates the tracks were generated by the same object. The bottom plot in (c) displays the exclusion metric as a function of time for each trajectory.

correct track.

Furthermore, the LiDAR object detector is more robust at detecting objects further away from the system and tracking through temporary occlusions compared to the visual object detector. In particular, Track 8 was continuously detected and tracked throughout the scene with the LiDAR; whereas, the same pedestrian using video data (Track 20, the pedestrian from the background of Figure 3.10a image B) was not continuously detected or tracked. Additionally, Track 44 was tracked through a temporary occlusion from Track 3 driving straight, but with the video-based trajectories, this same object (Track 5 and Track 18 which are both associated with the black car from Figure 3.10a image C) has disjoint tracks from the temporary occlusion. The improved tracking from LiDAR allowed both this pedestrian and vehicle to be rejected as source-carriers, which was not the case in video. As seen in Figure 3.11b, by continuously tracking objects, the source-object attribution analysis is able to take advantage of the full dynamics in the scene to exclude trajectories that do not correlate with the count-rate data.

In this scenario, the best-fit model (Track 50 - white station wagon from Figure 3.10a) clearly follows the count-rate data observed in the detector. This attribution is correct as the ^{137}Cs source was located in the trunk of Track 50. The time-dependence of the source-object attribution metric is shown in Figure 3.11c. The continuous tracking of object through LiDAR allows rapid exclusion of all tracks except Track 50 by the time of closest approach (38 sec).

Though there is a clear correlation for both Track 7 and Track 50 in video and LiDAR, respectively, the S-values are larger than the statistically motivated thresholds for association defined in Sec. 3.3. This result is most likely driven by using a simple physics model within a scenario that contains anisotropic shielding/attenuation, which is not included in the model.

Figure 3.12 shows the results from all of the alarm encounters at RFS for both video and LiDAR. The exclusion metric for both video and LiDAR show large variability between trials mainly due to the use of the simple physics model. Nonetheless, with LiDAR trajectories, the source-carrier was still assigned the lowest S-value score in 21 out of 26 alarm encounters. In the 5 encounters where the source-carrier could potentially be excluded as the object responsible for the radiological alarm, there is a clear correlation between the count-rate data and the best-fit model for the source-carrier, and an operator monitoring the system in real-time would be able to sift through the trajectories present to identify the most likely trajectory for alarm adjudication.

Using video trajectories, the results show a large number of tracks that have 5 or fewer poses in a majority of the alarm encounters. These short tracks are caused by YOLOv3-tiny's inability to consistently detect objects, especially pedestrians, more than about 10 m from the system. This results in short trajectories that are more easily fit to the count-rate data and therefore, cannot be excluded from the analysis. In addition, inconsistent tracking allows tracks to only associate mostly with flat background, which will be well described by the model since α will be set to zero in Eq. 2.13; thus, preventing these tracks from being excluded from the analysis. However, in 16 out of the 26 alarm encounters an operator monitoring the system would be able to correctly identify the trajectory responsible for the radiological alarm. In these 16 cases, there was a clear correlation between the source-carrier and the radiological data. Of these 16, there were 5 alarm encounters where the source-carrier had both a clear correlation and the lowest S-value score. There were 10 alarm encounters out of the 26 where a clear correlation did not exist because the source-carrier had disjoint tracks from either inconsistent tracking or the data association incorrectly assigning the wrong track ID to the source-carrier when it passed closely by another vehicle.

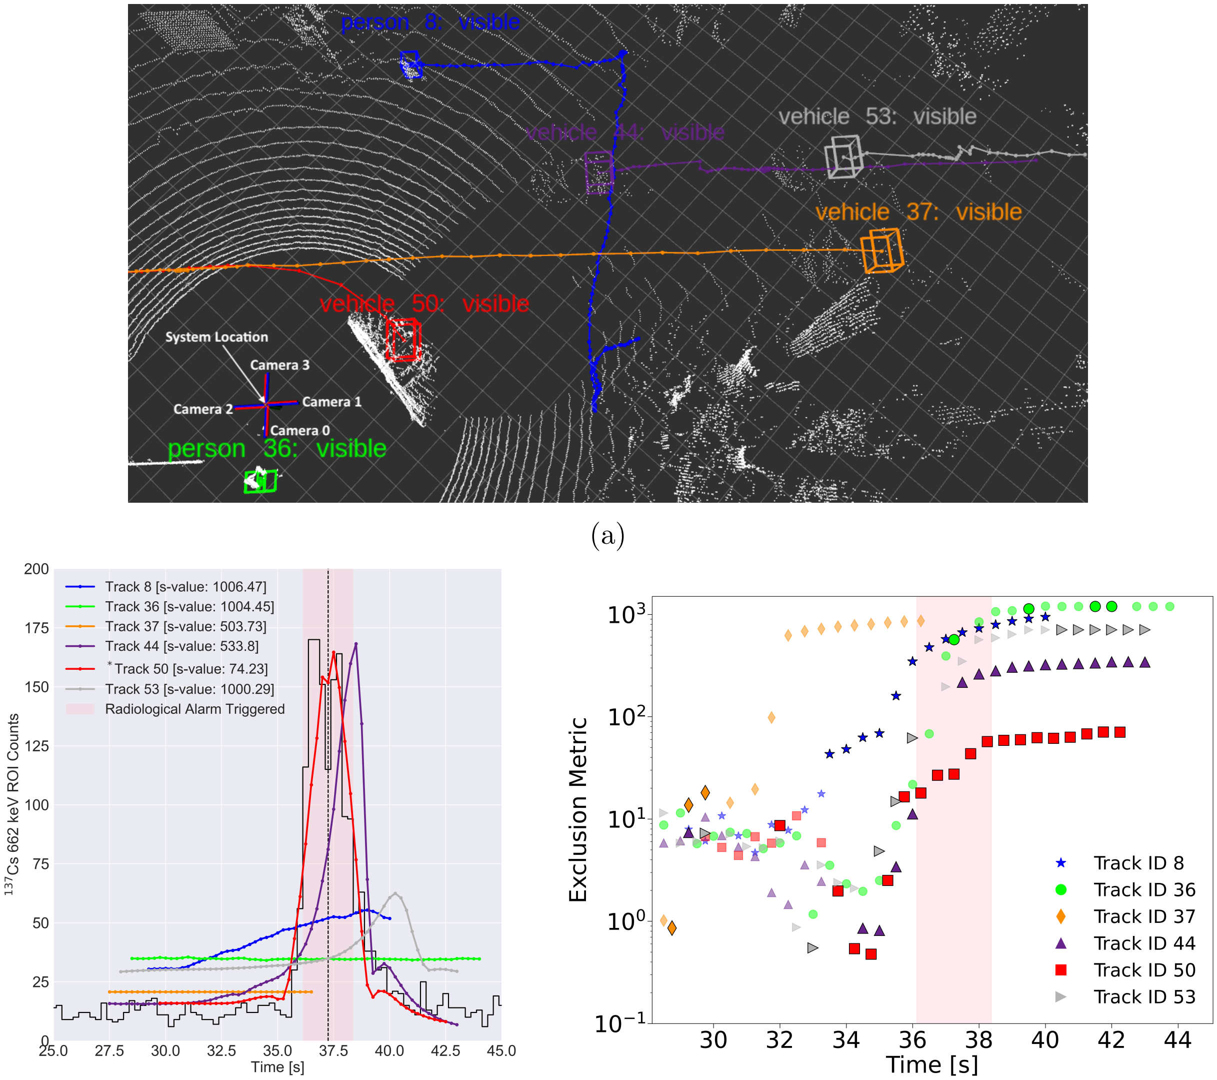

(a)

(b)

(c)

Figure 3.11: Source-object attribution in a mock urban scene with both pedestrians and vehicles using LiDAR data. The same source-encounter from Figure 3.10a, where a vehicle - Track 50 (white station wagon from Figure 3.10a) - carrying a ^{137}Cs source drove past the system, is shown in (a), and the * in (a) indicates the source-carrier. (a) corresponds with camera image C in Figure 3.10a. The bounding box colors (but not the labels) for each object are consistent with Figure 3.10a. In (a), the trajectory of each object in the scene up to that point in time is indicated by the lines proceeding each respective bounding box. The white grid lines represent 1 m^2, and the area in the middle without points is caused by the field of view of the LiDAR. The count-rate data from the 2 *in.* × 4 *in.* × 16 *in.* NaI(Tl) detector (black line) and the best-fit models to the ROI count-rate data is displayed in (b). The dashed line corresponds with the image from Figure 3.11a, and (c) displays the evolution of exclusion metric as a function of time for each trajectory, and he pink bars indicate the time interval for the radiological alarm.

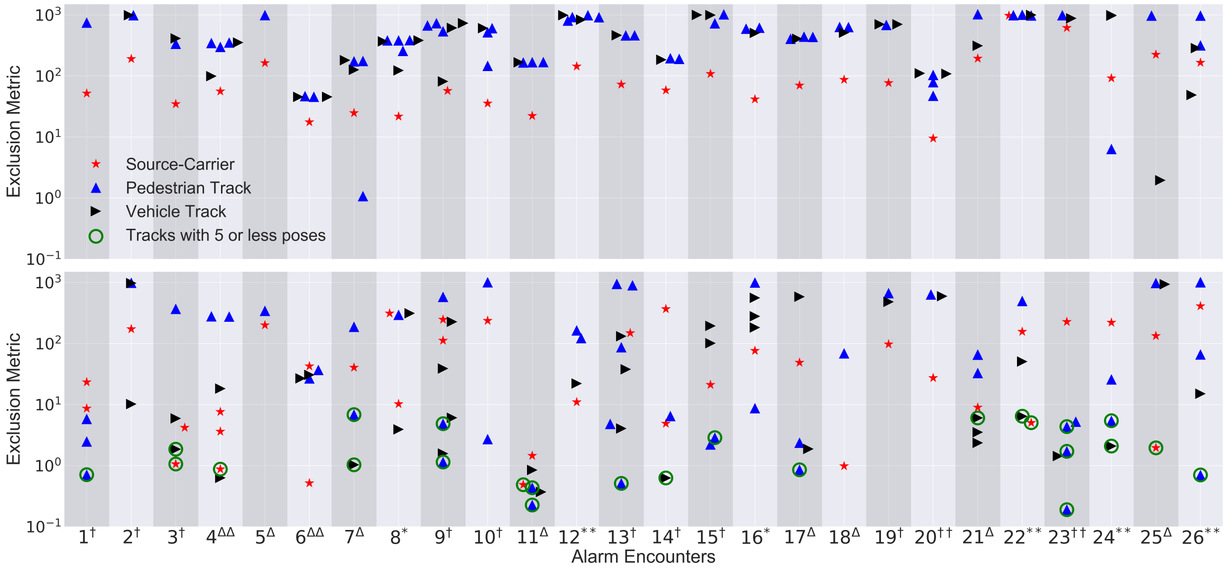

Figure 3.12: Source-object attribution analysis results from 26 alarm encounters in a mock urban environment. The total number of tracks generated during each alarm encounter for pedestrians, vehicles, and the source-carrier are presented. The top (bottom) plot was performed with LiDAR (video) trajectories. † indicates that the source carrier went straight past the operator during the trial. ‡ indicates that the source carrier was first stopped at a red light then proceeded straight past the operator during the trial. $*$ indicates that the source carrier turned right during the trial, and $**$ indicates that the source carrier was stopped at a red light and then turned right. Lastly, Δ indicates a left turn by the source carrier during the trial, and $\Delta\Delta$ indicates the source carrier was stopped at a red light and then performed a left-hand turn.

These results, collected in the mock urban environment, demonstrate that situational awareness can be improved for both LiDAR and video using the source-object attribution analysis. In the majority of alarm encounters for both LiDAR and video, a connection between the object carrying the radiological source and the radiological data could be made.

3.5 Improved Detection Sensitivity with Track-informed Optimized Integration Windows

In addition to enabling automated identification of pedestrians and/or vehicles correlated with radiological alarms, object tracking has the potential to enable increased detection sensitivity through track-specific time integration windows. As discussed in Section 2.10, using a model encapsulating the geometric and detector response associated with a trajectory, one may identify the selection of time-segments that should optimize SNR under the assumption of isotropic source emission and constant background. This optimum integration window analysis was performed on the same alarm encounters in the mock urban environment as the previous analyses and compared with fixed integration times.

In Figure 3.13a and Figure 3.13c, the calculated time segments that should optimize SNR when combined for the source carrier using video-based trajectories and LiDAR-based trajectories are shown, respectively. The source carrier trajectory used for this analysis correspond with the source carrier trajectory from Figure 3.10 and Figure 3.11, and the color coding of the source carrier is consistent with each respective figure. It should be noted this analysis was performed without using MCMC. The optimal integration for both video and LiDAR occurs during the period of time while sensitivity, which is proportional to SNR (Section 2.10), is increasing. For both video and LiDAR, the calculated track-informed contextual integration window is 2.25 seconds. The optimal integration window overlaid on the radiological data for both video and LiDAR is displayed in Figure 3.13b and Figure 3.13d, respectively. The optimal window for both video and LiDAR corresponds with the time of closest approach for the source carrier. However, due to the anisotropic attenuation of the source carrier, which is not accounted for in the physics equation model, the optimal time window does not fully encapsulate the period of time in which the source carrier was closest to the system.

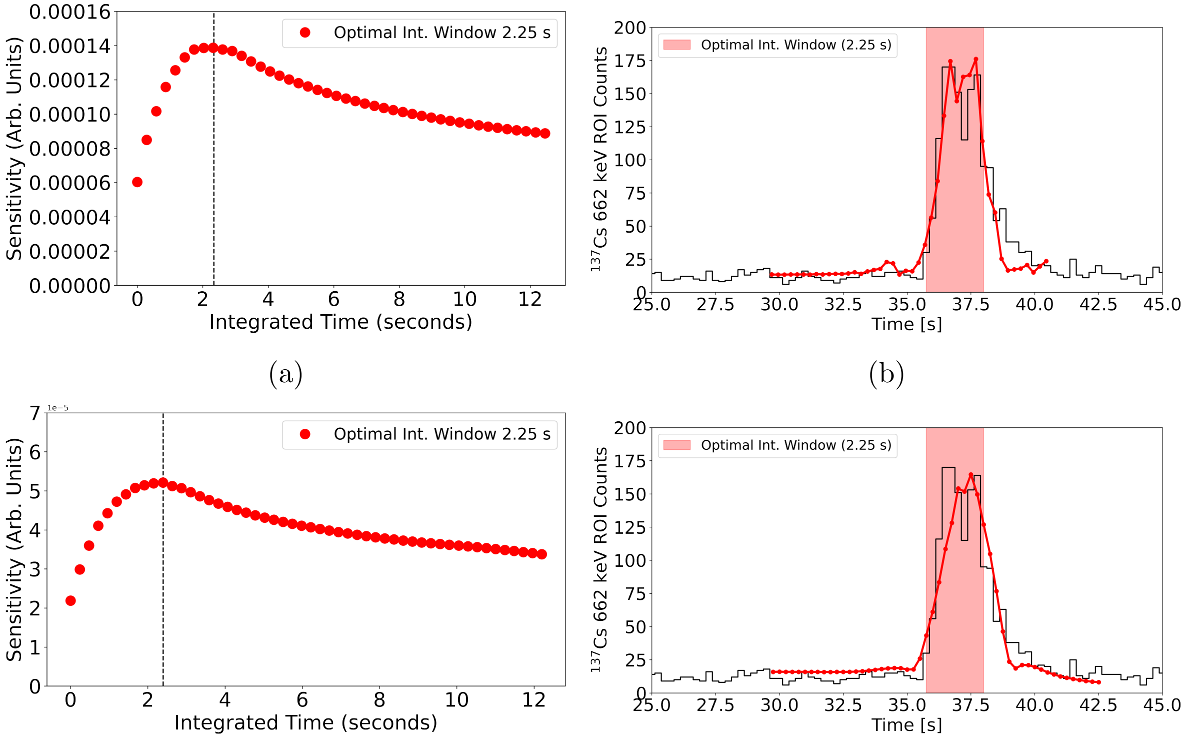

(a) (b)

(c) (d)

Figure 3.13: Track-informed integration windows using video-based trajectories (a-b) and LiDAR-based trajectories (c-d) for the source carrier (Track 7 and Track 50) from the alarm encounter in Figure 3.10 and Figure 3.11, respectively. In (a) and (c), the calculated time segments that should optimize SNR for the source carrier from video and LiDAR, respectively. The dashed line indicate the time period in which sensitivity was increasing. In (b) and (d), the time segments that should maximize SNR overlaid on the radiological data.

Figure 3.14 shows the results for a single alarm encounter, showing the maximum relative anomaly value, normalized against the maximum anomaly value with a 1.0 second integration time, for all tracks in the scene as well as different fixed integration times. Again, the trajectories used for this analysis correspond with the trajectories from Figure 3.10 and Figure 3.11, and the color-coding for the tracked objects is consistent with the color-coding from each respective figure. In this encounter results are consistent with the hypothesis that the optimized time-window for the source-carrying trajectory should yield a larger anomaly value. However, in both video and LiDAR, the source-carrying trajectory did not have the maximum anomaly value amongst the different integration times which was produced using a 4 second integration time. In this case, the optimal window does not produce the highest anomaly value due to anisotropic shielding from the vehicle, which is significant enough such that the assumptions of the analysis fail. Also observed is an elevated anomaly value for another track which corresponds to the researcher who was operating the system, who was nearby and stationary during this particular alarm encounter.

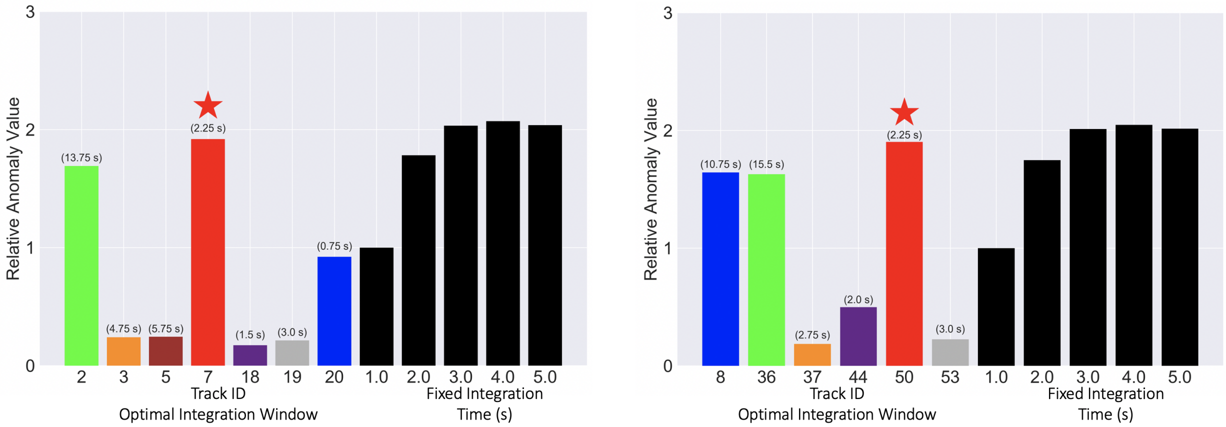

(a) (b)

Figure 3.14: Relative anomaly values from track-informed optimized integration windows and fixed integration window of 1, 2, 3, 4, and 5 seconds for both video (a) and LiDAR (b) data. The anomaly values are normalized to the 1.0 second fixed integration window. These results are from the alarm encounter shown in Figure 3.10 and Figure 3.11. The star indicates the track carrying the source in each scenario, and the time duration for each optimal window is provided in the parentheses.

The results from applying this analysis to the 26 different alarm encounters in a mock urban environment are shown in Figure 3.15 for both video and LiDAR.

The results show that using a tracking-informed integration window can improve the anomaly value, a proxy for detection sensitivity, compared to various fixed integration windows. Two alarm encounters (22** and 24**), in particular, produced large relative anomaly values compared to fixed integration windows for both video and LiDAR. In both of these cases, the vehicle carrying the source was stopped at a red light next to the system for an extended period of time before the vehicle proceeded making a right-hand turn past the system.

From the results in Figure 3.15, the LiDAR-generated trajectories are more consistent than the video-based trajectories in yielding this enhanced sensitivity. This is consistent with the previous results showing that the LiDAR detection and tracking is more effective than video at consistently tracking objects through the scene. Even in the case of LiDAR tracking, there are alarm encounters where a fixed integration window of 1, 2, 3, 4, or 5 seconds yields larger anomaly values. These cases were often driven by the object not being detected or fully tracked across the system's entire FOV. In several cases, however, the object was fully tracked and again it is hypothesized that the anisotropic shielding from the vehicle was sufficiently significant that the assumptions of the analysis fail. On the whole, this analysis suggests that contextual information can be used to improve the detection sensitivity of a static spectrometer.

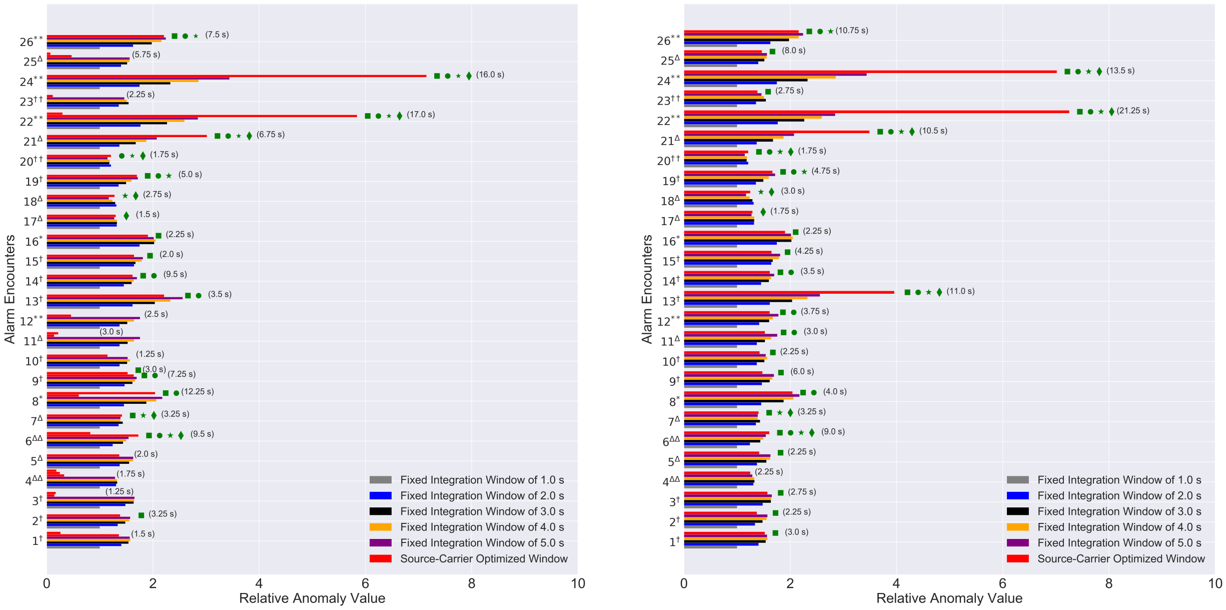

(a)

(b)

Figure 3.15: Maximum relative anomaly values from each alarm encounter using either track-optimized time-windows or fixed integration windows from 26 alarm encounters in a mock urban environment. The anomaly values are normalized to the 1.0 second fixed integration window. Results from the video and LiDAR are shown in sub-figure (a) and (b), respectively. The green squares, circles, stars, or diamonds indicate the encounters where the optimal integration window yielded a higher anomaly value than a 2, 3, 4, or 5 second integration window, respectively. The time duration for each optimal window is provided in the parentheses, and in the case for (a) the time duration with the longest window is provided in certain alarm encounters where the source-carrier had disjoint tracks. † indicates that the source carrier went straight past the operator during the trial. ‡ indicates that the source carrier was first stopped at a red light then proceeded straight past the operator during the trial. $*$ indicates that the source carrier turned right during the trial, and $**$ indicates that the source carrier was stopped at a red light and then turned right. Lastly, Δ indicates a left turn by the source carrier during the trial, and $\Delta\Delta$ indicates the source carrier was stopped at a red light and then performed a left-hand turn.

3.5.1 MCMC Track-informed Optimized Integration Windows

The above optimal window analysis does not account for the uncertainty in the object's position when determining the optimal integration time. In order to account for this, MCMC can be applied to the optimal integration window formulation for the source carrier in Figure 3.13(a–b) (video) and Figure 3.13(c–d) (LiDAR). As mentioned in Section 2.10, the position of the object is varied and the count-rate data is modeled and fit for that trajectory at each varied position. The position of the object is constrained to the position uncertainty in x, y, and z around each respective mean value. The optimal integration window is calculated from a randomly chosen subset object positions, and the largest integration window with the lowest negative log-likelihood within these random object positions is chosen as the optimal window.

The results from applying MCMC to the optimal integration window analysis are shown in Figure 3.16 for video and LiDAR. The color coding is still consistent with Figure 3.13. By accounting for position uncertainty, the time duration of the optimal integration for both video (Figure 3.16a) and LiDAR (Figure 3.16b) is longer relative to Figure 3.13(a–b) and Figure 3.13(c–d), respectively. With the longer time lengths for both video and LiDAR, the signal caused by the anisotropic shielding from the source carrier's vehicle is almost completely integrated.

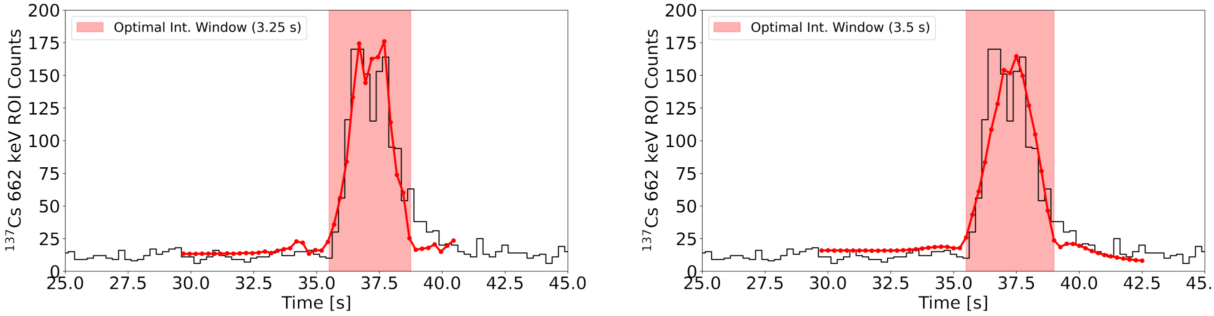

(a)

(b)

Figure 3.16: Accounting for the position uncertainty of an object in the optimized integration window calculation using video (a) and LiDAR-based trajectories (b). The trajectories for this example are for the source carrier (Track 7 and Track 50) from the alarm encounter in Figure 3.10 and Figure 3.11 for video and LiDAR, respectively. The optimized integration window overlaid on the count-rate data for video (a) and LiDAR (b) is shown.

A spectroscopic analysis was applied to the MCMC optimized integration windows and was compared to either the non-MCMC optimized integration window or fixed integration windows. The results of this analysis are shown in Figure 3.17. The relative anomaly values for the MCMC optimized integration windows offer improved detection sensitivity compared non-MCMC optimized integration windows for both video (Figure 3.17a) and LiDAR (Figure 3.17b). In Figure 3.17b, the MCMC optimized integration window has an anomaly value that is equal to a fixed integration window of 4 seconds which is the optimal integration window.

Subsequently, this analysis was applied to all of the alarm encounters for LiDAR (Fig-

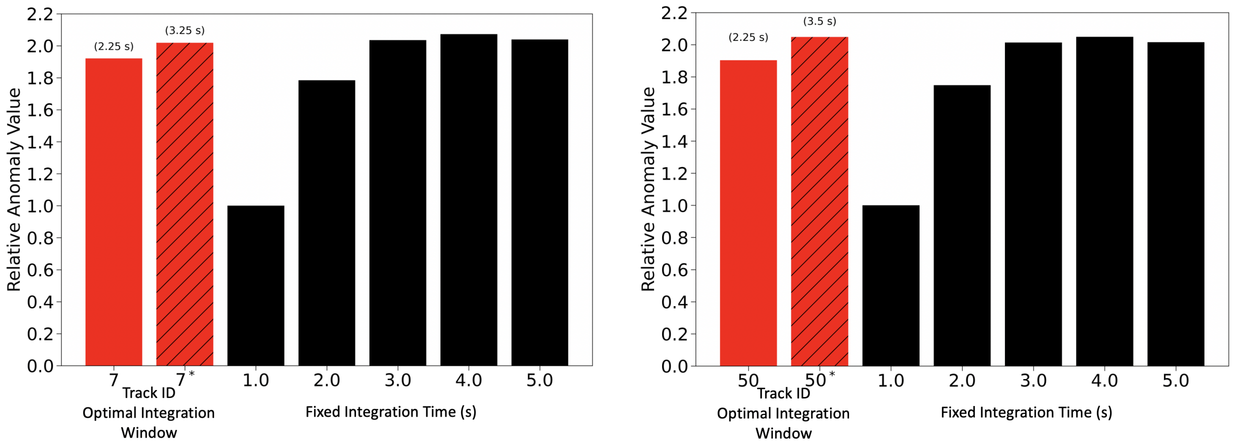

(a) (b)

Figure 3.17: Relative anomaly values from track-informed optimized integration windows and fixed integration window of 1, 2, 3, 4, and 5 seconds for both video (a) and LiDAR (b) data, respectively. The anomaly values are normalized to the 1.0 second fixed integration window. These results are from the alarm encounter shown in Figure 3.10 and Figure 3.11. The star indicates the track-informed optimized integration window analysis was performed with MCMC, and the time duration for each optimal window is provided in the parentheses.

ure 3.18) – video results are not shown because of inconsistent tracking in multiple alarm encounters. It is seen that accounting for the position uncertainty in the optimal window formulation can improve the detection sensitivity compared to not accounting for position uncertainty. In 25 out of 26 trials, the MCMC optimized integration window offers improvement, and in the remaining trial, the MCMC optimized integration window offered similar detection sensitivity to the non-MCMC optimized integration window. These results demonstrate that accounting for an object's uncertainty in the optimized integration window can improve detection sensitivity relative to not accounting for the object's uncertainty.

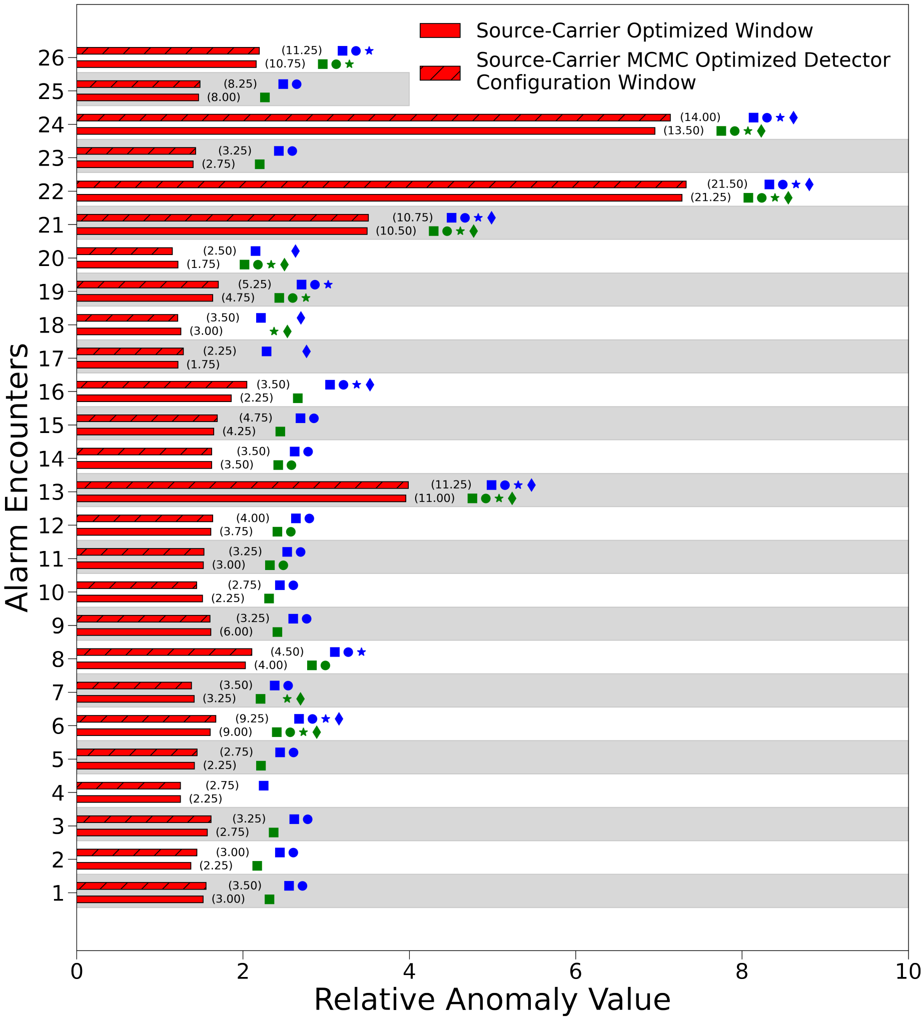

Figure 3.18: Maximum relative anomaly values from each alarm encounter using either track-optimized time-windows, MCMC track-optimized time-windows, or fixed integration windows from 26 alarm encounters in a mock urban environment using LiDAR-based trajectories. The anomaly values are normalized to the 1.0 second fixed integration window. The hatch marks represent the MCMC track-optimized time-windows. The blue (green) squares, circles, stars, or diamonds indicate the encounters where the MCMC optimal integration window (non-MCMC optimal integration window) yielded a higher anomaly value than a 2, 3, 4, or 5 second integration window, respectively. The time duration for each optimal window is provided in the parentheses.

3.6 Additive Tracking Source Localization and Source-Object Attribution

The previous sections (Section 3.3 - 3.5) investigated scenarios in which only a single point source was present at a time. However, in urban environments, multiple trajectories in a scene could be carrying the same source type. The mathematical framework of APSL outlined earlier in Section 2.9 can be directly applied to this problem. Since each track is considered to be carrying a source the APSL formulation can be changed from searching x, y, and z as well as the source weights to only optimizing the source weights of each trajectory. This adjusted formulation is hereby referred to as Additive Track Source Localization (ATSL).

To demonstrate the capabilities of ATSL, synthetic trajectories and count-rate data was generated. Three random trajectories all with the same track length and moving with a constant speed of 3 m/s past a detector were generated (Figure 3.19), and the number of trajectories carrying a ^{137}Cs source past the system was varied. The trajectories were shifted in time so that each trajectory passed in front of the detector at different times. In Figure 3.20, three different scenarios are presented.

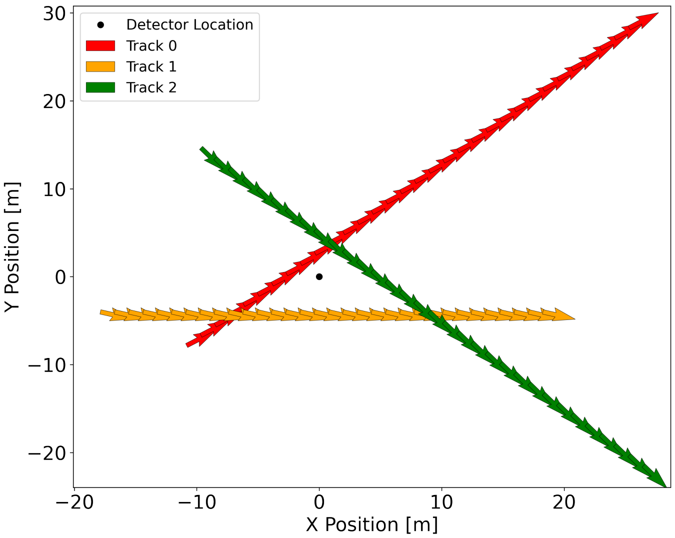

Figure 3.19: Three random trajectories walking past a detector (a). In any given scenario Track 0, Track 1, or Track 2 could be carrying a 120 μCi, 140 μCi, or 189 μCi ^{137}Cs source, respectively.

In the first scenario, Track 1 is carrying a 140 μCi ^{137}Cs source past the system. Track 1 is 4.02 m from the system at closest approach. Figure 3.20a and Figure 3.20b shows the results of using ATSL and source-object attribution in the case of a single source. ATSL searches through the present trajectories in sequential order searching for the trajectory or trajectories that best describe the data. Both Track 0 and Track 2 were correctly rejected as weakly contributing to the source, and ATSL correctly determined Track 1 was carrying the source. Not surprisingly the best-fit model for Track 1 correlates well with the radiological data.

To test how ATSL performs when multiple trajectories are carrying sources of the same source type, Track 0 and Track 1 both were given a source. The result of this encounter using ATSL is shown in Figure 3.20c. The two trajectories have different times of closest approach which results in a bump in the count-rate data around the 6 second mark when Track 0 passes in front of the detector first before Track 1. Track 0 and Track 1 are correctly identified as the trajectories responsible for the radiological signal while Track 2 is rejected. The source-object attribution of this encounter are shown in Figure 3.20d. The best-fit models for both the single source approach and ATSL are presented. Using the single source approach, all the trajectories could be excluded from the analysis because none correlate with the radiological data. However, by using the summed contribution of Track 0 and Track 1 there is a clear correlation and neither trajectory can be excluded.

In Figure 3.20e and Figure 3.20f the same random trajectories from the previous two source encounters are used but now all three trajectories are carrying a source. In this scenario, all three trajectories are correctly identified as contributing to the source, and the summed best-fit model for all three trajectories has a clear correlation with the radiological data.

The ATSL scenarios above demonstrate the capabilities of performing source-object attribution, when multiple sources of the same type are present, in relatively high SNR environments. Using the same scenario, the source activities were lowered and the track lengths of each object are varied reflecting more realistic situations.

Figure 3.21a demonstrates that the trajectory responsible for the radiological data can still be correctly identified in low SNR situations. Track 1 was carrying the source, but performing attribution none of trajectories can be excluded (Figure 3.21b). In low SNR situations, the model for a trajectory can be more easily fit to the data, which improves the S-value. This is a limitation of this approach. In Figure 3.21c, ATSL correctly identifies both trajectories contributing to the count-rate data when both Track 0 and Track 1 are carrying a radiological source, and the correct trajectories are attributed to the radiological data (Figure 3.21d). Without using ATSL, Track 1 could have potentially been missed. In Figure 3.21e, the addition of Track 2 did not improve the likelihood enough to warrant the addition of a third trajectory and was incorrectly rejected. However, performing attribution in this scenario (Figure 3.21f), Track 2 correlates with the radiological data and would probably be flagged by an operator.

The results of this analysis demonstrate that attribution can be performed when multiple sources of the same type are present in a scene by applying the ATSL formulation.

Due to restrictions placed by Coronavirus Disease 2019, experimental results demonstrating these capabilities were not able to be performed. However, using synthetic data, this work was able to show the effectiveness of ATSL in performing source-object attribution when multiple sources of the same type are present.

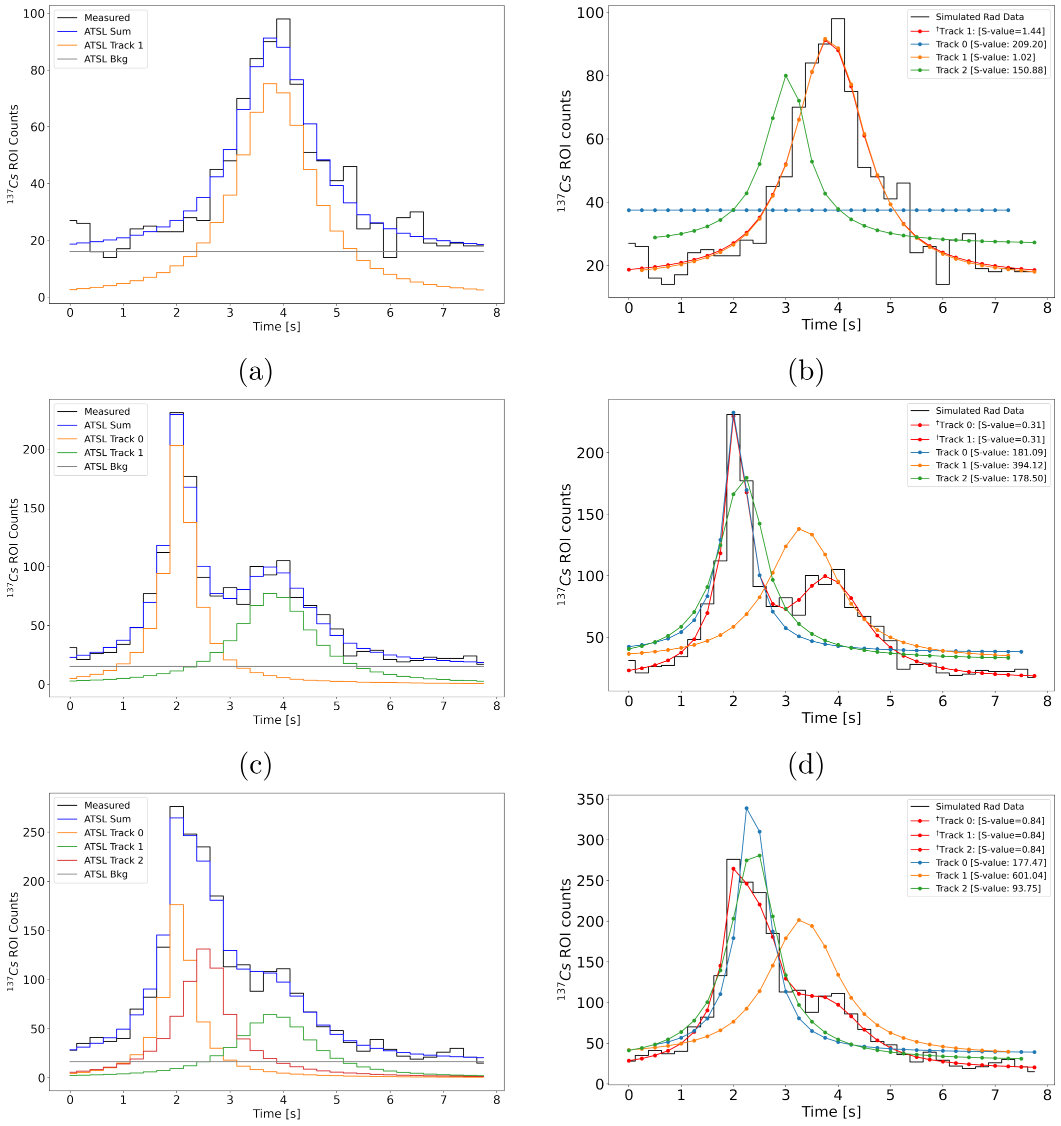

(a) (b)

(c) (d)

(e) (f)

Figure 3.20: Performing source-object attribution when multiple trajectories are carrying a radiological source of the same source type. The trajectories presented correspond with Figure 3.19. In (a) and (b), Track 1 is carrying a 140 μCi ^{137}Cs is shown. In (a), the expected contribution of the source is shown along with the summed contribution, and (b) is the best-fit models for each trajectory to the count-rate data. The † indicates that both trajectories contributed to the best-fit model. The results in (c) and (d) are from Track 0 and Track 1 carrying a 120 μCi ^{137}Cs or 140 μCi ^{137}Cs source, respectively. The results in (e) and (f) are from Track 0, Track 1, and Track 2 carrying either a 120 μCi ^{137}Cs, 140 μCi ^{137}Cs source, or a 189 μCi ^{137}Cs, respectively.

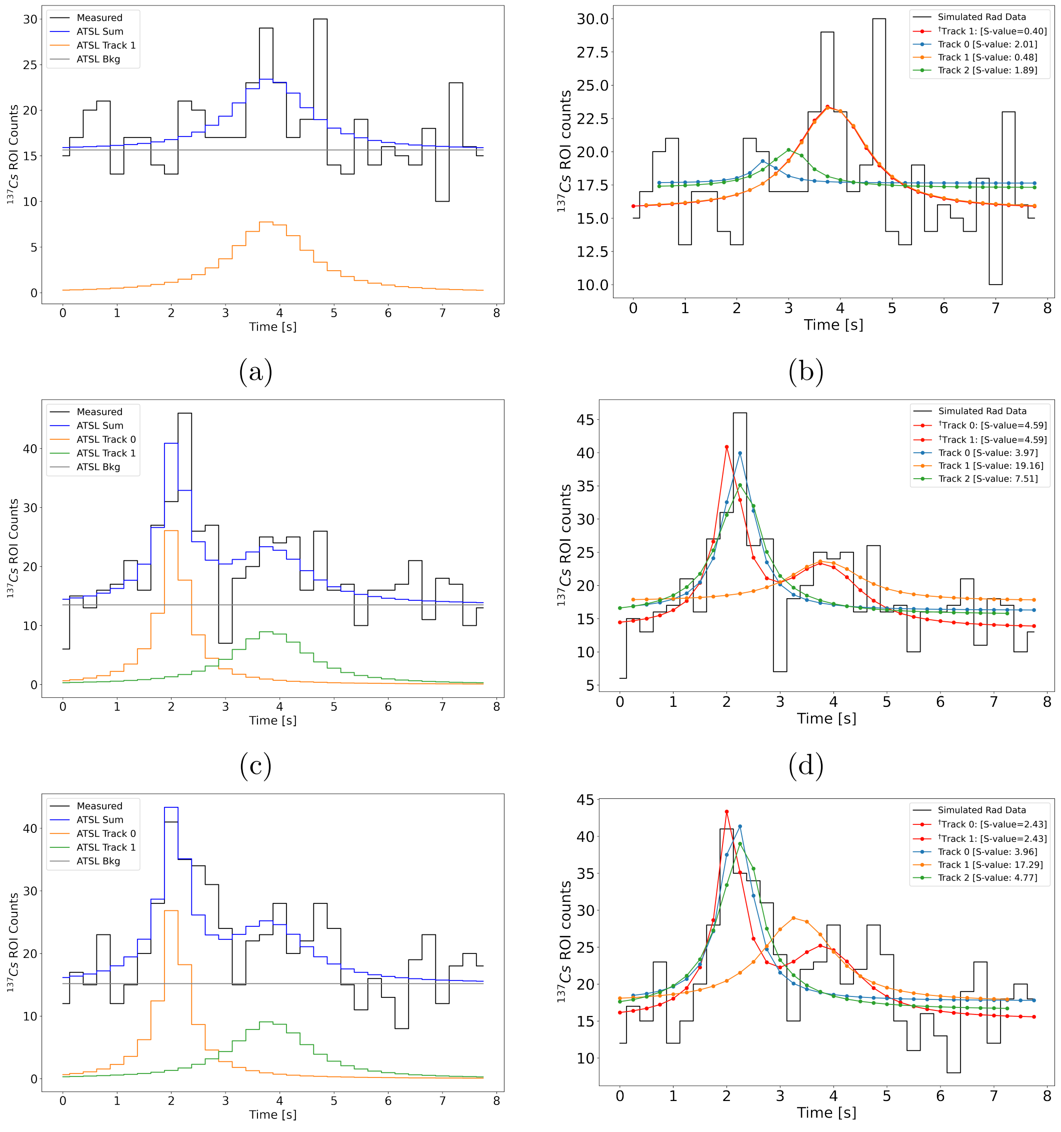

(a) (b) (c) (d) (e) (f)

Figure 3.21: Performing source-object attribution when multiple trajectories are carrying a radiological source of the same source type. The trajectories presented correspond with Figure 3.19. In (a) and (b), Track 1 is carrying a 12 μCi ^{137}Cs is shown. In (a), the expected contribution of the source is shown along with the summed contribution, and (b) is the best-fit models for each trajectory to the count-rate data. The † indicates that both trajectories contributed to the best-fit model. The results in (c) and (d) are from Track 0 and Track 1 carrying a 10 μCi ^{137}Cs or 12 μCi ^{137}Cs source, respectively. The results in (e) and (f) are from Track 0, Track 1, and Track 2 carrying either a 10 μCi ^{137}Cs, 12 μCi ^{137}Cs source, or a 17 μCi ^{137}Cs, respectively.

3.7 Conclusions & Future Directions

A contextual sensor system has been developed and demonstrated that performs multi-class object detection, 3-D tracking, and source-object attribution in real time using either LiDAR or visual cameras. This system, and the methods it implements, can improve the localization of radiological/nuclear materials in urban environments by enhancing situational awareness and allowing nondirectional detectors to provide localization information. Furthermore, in urban environments, this system has the potential to improve detection of radiological/nuclear materials by increasing detection sensitivity.

The findings from this study demonstrate that radiological sources can be successfully attributed to objects derived from video and LiDAR in a mock urban environment. The findings further show that LiDAR offers superior tracking performance when compared to video, and this improves the ability of the source-object attribution analysis to reject trajectories that are inconsistent with the data, while providing a more apparent correlation between the object responsible for the radiological alarm and the count-rate data. With LiDAR, the source carrier had a clear correlation with the radiological data in all 26 of the alarm encounters considered; whereas, using video-based trajectories an operator monitoring in real time would only be able to identify the source carrier in 16 alarm encounters.

Using track-specific integration windows it was demonstrated that object tracking can improve the anomaly value, a proxy for detection sensitivity, compared to fixed integration windows for both video and LiDAR. Again, since LiDAR offers superior tracking performance compared to video, more enhanced sensitivity was seen with LiDAR-based trajectories. Furthermore, accounting for uncertainty in an objects extent when performing the optimal integration window analysis improves the anomaly value relative to not accounting for uncertainty.

Finally, through simulation it was shown that multiple sources of the same type within a scene can be correctly attributed to the objects responsible for the radiological alarm using source-object attribution in conjunction with ATSL. ATSL was able to correctly reject and identify tracks that did or did not contribute to the radiological alarm even in low SNR regimes.

From these results, it can be concluded that using video and LiDAR to augment a static radiation detector does enhance situational awareness and strongly suggests that contextual information can be used to improve the detection sensitivity of a static radiation detector.

Future work is needed to explore improving tracking capabilities for video to limit track switches and disjoint tracks. One area of improvement is investigating different lightweight neural networks that offer better performing object detection algorithms. While attribution could be performed with visual imagery, the results were limited compared to LiDAR. The current object detection network is not built to deal with small objects in images that are far away from the system. At distances larger than 10 m pedestrians cannot reliably be tracked anymore. This leads to short or split tracks that often score better on the attribution metric than the true source-carrier's trajectory. New developments, particularly feature pyramid networks [32], are trying to address these issues and might result in lightweight networks with better detection performance. Recent tests with yolov4 [14] show very promising results.

The ATSL simulations demonstrated the capabilities of using ATSL in conjunction with source-object attribution; however, this needs to be further investigated using experimental data.

Additionally, the results presented in this section were limited to a static system co-located with a single radiation detector. Section 4 will investigate the application of tracking and attribution performance for mobile systems where both sensors and objects are in motion. Furthermore, performing source-object attribution using a detector array will be explored in Section 4.

Chapter 4

Mobile system: LiDAR Enhanced Mobile Urban Radiation Search Vehicle

The following sections (Sections 4.2 – 4.6) detail the source-object attribution capabilities on a mobile detector system. Section 4.2 performs a quantitative study with LEMURS in the presence of mobile and static sources. Following, the performance of the source-object attribution approach is tested in a mock urban environment with a mobile source (Section 4.3). Within Section 4.3, an environment with source attenuation present and a mobile source is also tested (Section 4.3.3). In Section 4.4, optimized spectroscopic search using track-informed integration windows is applied to the data from Section 4.3 to improve detection sensitivity. Additionally, the importance of the detector response and the impact of the angular response on radiological source localization when tracking information is available is evaluated using source-object attribution in Sections 4.5 and 4.6, respectively. Finally, the localization performance of source-object attribution is compared to different reconstruction algorithms in Section 4.7.

4.1 LiDAR Enhanced Radiation Urban Search (LEMURS) Vehicle

For this analysis, the LiDAR Enhanced Mobile Urban Radiation Search (LEMURS) vehicle [24] (Figure 4.1) was used. LEMURS consists of 2 16-beam LiDAR, multiple IMU and INS (SPATIAL [39]) devices, a 360° panoramic camera, and 6 NaI(Tl) detectors arranged in a hexagonal array. Each detector is equipped with an Ortec DigiBASE multichannel analyzer [58] and configured to publish list-mode gamma-ray interaction data packets at $\approx 20\,Hz$. The contextual sensor data and gamma-ray events are acquired within the ROS [59] across multiple single-board computers with clocks synchronized by the network time protocol (NTP) [60].

Figure 4.1: The LEMURS system which consists of a panoramic Occam Omni 60 camera, 2 LiDAR, multiple IMU and INS devices, and 6 2 *in.* × 4 *in.* × 16 *in.* NaI(Tl) detectors.

4.2 Quantitive Assessment of Attribution Analysis

To assess the performance of the source-object attribution analysis, LEMURS did straight line passes past moving and static sources. The following sections (Section 4.2.1 – 4.2.2) outline the source-object attribution capabilities using the LEMURS vehicle. In both sections, SLAM was used to produce pose estimates of LEMURS in the world-fixed frame.

4.2.1 Degenerate Scenario

An important aspect of the attribution analysis is correctly performing attribution when static objects are closely spaced together, such as a parking lot, and LEMURS is mobile. To explore this scenario, a degenerate scenario was created. The degenerate scenario consisted of vehicles parked on either side of the road such that the time of closest approach and distance from LEMURS was identical as LEMURS drove by the vehicles. This scenario represents a worst-case scenario for LEMURS. Using the inter-detector attenuation effects of the detector array could add the necessary constraints to break degeneracies present.

In the scenario, LEMURS drove straight past static cars on either side of the road going 10 mph. The cars were aligned such that cars on opposite sides of the road had similar time of closest approach (less than 0.1 seconds) and the rear of the vehicles were 4.1 ± 0.1 m from the edge of LEMURS as shown in Figure 4.2. In addition, there was a 2 meter separation between the cars on the same side of the street. A 1.87 mCi ^{137}Cs source inside of 2 cm Pb shielding was placed in the trunk of a vehicle. For each trial, LEMURS performed 20 source drive-bys. Attribution was performed using the photopeak ROI (600 keV - 725 keV) for ^{137}Cs.

Figure 4.2: A degenerate scenario where LEMURS drove straight down the center of the road with cars parked on either side of the road. A 1.87 mCi ^{137}Cs source inside of 2 cm Pb shielding was placed in the trunk of a vehicle.

Figure 4.3 shows the results from one source encounter using LiDAR-based trajectories and a LEMURS speed of 10 mph in the degenerate scenario. The source carrier and surrounding objects are continuously tracked throughout the alarm encounter. The Kalman filter's localization prediction is based on the center of an objects bounding box and thus deviates continuously by about 2 meters from the actual source location, in the trunk of the car.

Figure 4.4 shows the result of performing source-object attribution on the above source encounter using all 6 NaI(Tl) detectors. The variation in the Kalman filter positions and the position-source offset cause the best-fit models to underestimate the magnitude of the source activity. However, using the lowest exclusion metric as a metric for localization performance, the source is correctly attributed and localized to the source carrier (Track 39) while all the other trajectories can be rejected from the analysis.

Synthetic count-rate data and LiDAR trajectories were created to model the degenerate scenario. LEMURS was modeled traveling 10 mph past the source carrier. The source was offset from the source carrier's position by 2 m to match experiment. The expected count rate from the source carrier was generated using Eq. 2.13 and then was Poisson sampled. At each position of LEMURS, the angle between LEMURS and the objects present is calculated. An attenuation parameter is then applied to the count-rate data before Poisson sampling to account for attenuation changing over time as LEMURS moves past the source carrier. Also, the position and distance of LEMURS from each object is utilized to randomly sample the trajectories according to experimental trajectories. Using the synthetic data, 1000 trials were created for each scenario and attribution was performed similar to the experimental data.

The results from all the source encounters and the synthetic data are shown in Figure 4.5 for LiDAR (Figure 4.5a) and video-based trajectories (Figure 4.5b). In both Figure 4.5a and Figure 4.5b, the color coding from the experimental data is consistent with the objects shown

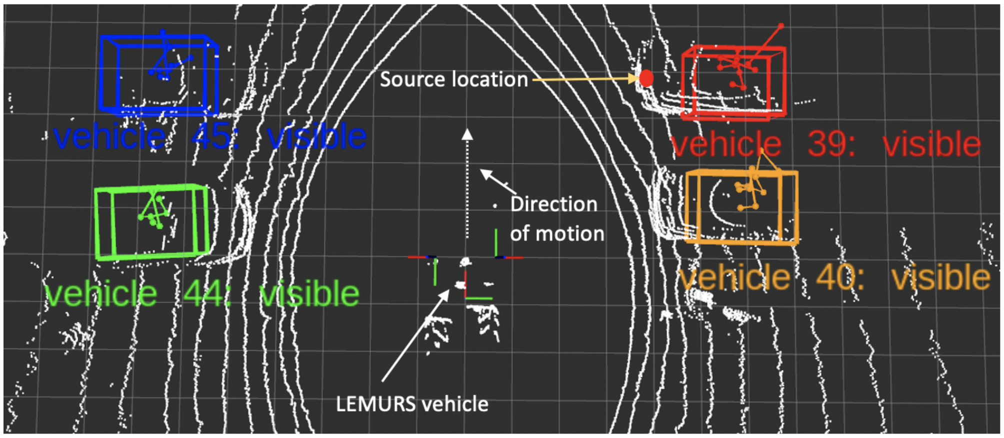

Figure 4.3: LEMURS drives past 4 stationary cars, where a vehicle – Vehicle 39 (red bounding box) – has a 1.87 mCi ^{137}Cs source inside of 2 cm Pb in the trunk of its vehicle. The trajectory of each object in the scene up to that point in time is indicated by the lines associated with each respective bounding box. The white grid lines represent 1 m^2, and the area in the middle without points is caused by the field of view of the LiDARs.

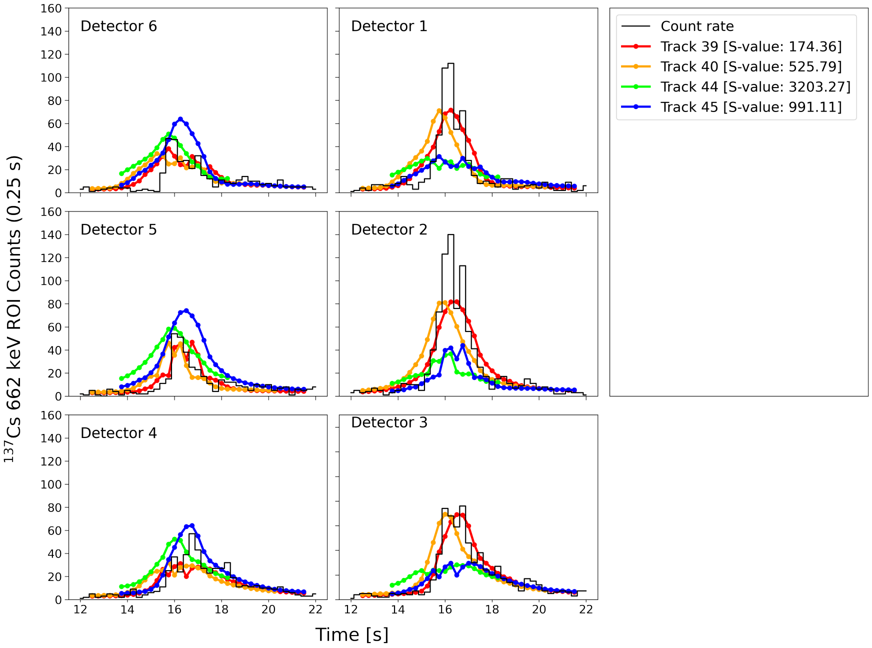

Figure 4.4: Performing source-object attribution with LiDAR-based trajectories and using all 6 NaI(Tl) detectors in the array. The trajectories are from Figure 4.3 and the color coding of the objects is consistent. The count-rate data from the full 6 NaI(Tl) detector array and the best-fit models to the count-rate data are displayed.

in Figure 4.3. The results of this analysis show that the experimental and simulated data are within a standard deviation for each respective object in the scene for both LiDAR and video. A S-value threshold of 363.5 (444.5) was set to define a 95% true positive rate for the source carrier for LiDAR (video). With a S-value threshold of 363.5 (444.5), the left-right degeneracy that exists between the source carrier and the vehicle directly across the road is successfully broken 100% (100%) of the time using LiDAR-based trajectories (video-based trajectories). This is due to the angular resolution of the LEMURS array. When static vehicles are separated by 2 meters, the source carrier is correctly identified 67% (16%) of the time with LiDAR (video).

Even though a 1000 trials were run using the synthetic data, the statistical variation is large, in particular for the source carrier and the adjacent vehicle. This represents the sensitivity of the goodness-of-fit metric. Variations in the estimated Kalman filter track position can influence the quality of fit, and the large statistical variation would be expected in the experimental data as well with more trials. With video, the variations for the source carrier are due to the inferred distance estimate which causes large uncertainties.

The results of this analysis demonstrate the ability of the source-object attribution analysis along with using the full NaI(Tl) array to model and fit trajectories to the count-rate data to be quite robust against track degeneracy for LiDAR. Using LiDAR-based trajectories, attribution was successfully performed in 67% of the trials. With video, the large uncertainties primarily due to the inferred distance estimate reduces the ability to distinguish cars parked within 2 m of each other.

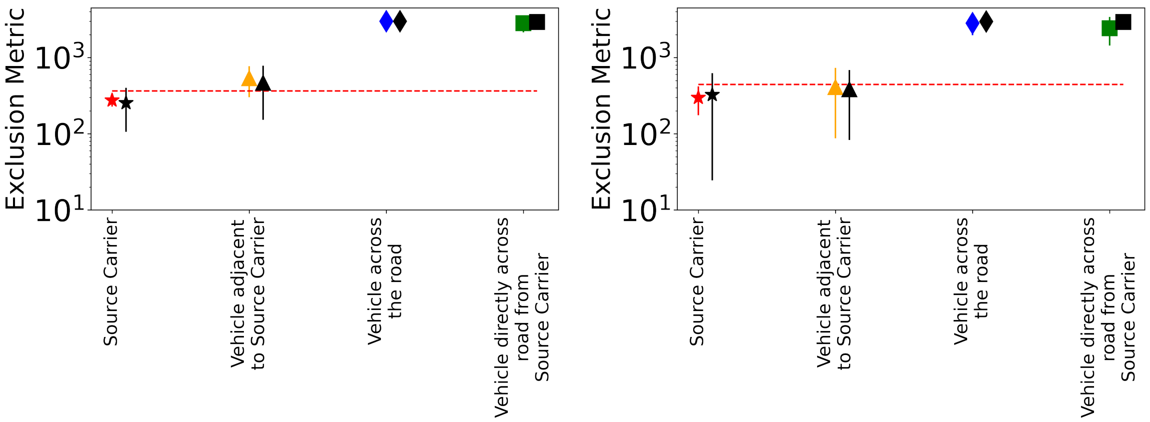

(a)

(b)

Figure 4.5: Results from all 20 alarm encounters in the degenerate scenario using source-object attribution using LiDAR (a) and video-based trajectories (b), respectively. The color coding in (a) and (b) is consistent with Figure 4.3, and source-object attribution was performed using the 6 NaI(Tl) detectors independently. In (a) and (b), the red dashed line represents a S-value threshold of 363.5 or 444.5, which defines a 95% true positive rate for the source carrier for LiDAR or video-based trajectories, respectively. The black markers indicate the attribution analysis was performed using synthetic data in (a) and (b).

4.2.2 Mobile Source

In urban environments, radiological sources can also be mobile. To understand attribution performance in the presence of mobile sources, the source carrier from Section 4.2.1 was used.

LEMURS and the source carrier performed straight line passes at 10 mph on a straight two-lane road. In this case, three different scenarios were explored: source carrier only, source carrier followed by a vehicle, and source carrier between two vehicles. For each scenario, a total of 10 source passes was conducted.

Similar to Section 4.2.1, synthetic count-rate data and trajectories were generated. Again, the expected count rate from the encounters was calculated using Eq. 2.13 and was Poisson sampled. In this case, the ^{137}Cs source was displaced from the center of the source carrier vehicle by 1 m corresponding with the trunk of the vehicle. The synthetic trajectories were randomly sampled using experimental data and the angle and distance of the objects from LEMURS. 1000 trials were conducted using the synthetic data and was subjected to the same analysis as the experimental data. However, in this case, synthetic count-rate data and trajectories could not be generated for the non-source-carrying vehicles because there was too much variation in each alarm encounter. For each trial, the distance between the source carrier and non-source-carriers was about 3 m. However, for a given trial, deviations in the distance between the cars could affect the best-fit model for that particular object. Too much time or distance between the vehicles resulted in a best-fit model for the non-source-carrier that was outside the alarm window and therefore was better described by a background only model.

The experimental results from the scenario where the source carrier is between two vehicles is shown in Figure 4.6 for video and LiDAR. In this case, the attribution analysis is able to correctly localize the source to the source carrier in all 10 trials using both video and LiDAR while the vehicles surrounding the source carrier can be rejected. These results, while limited in statistics, demonstrate that typical car lengths between moving vehicles is sufficient to effectively perform attribution.

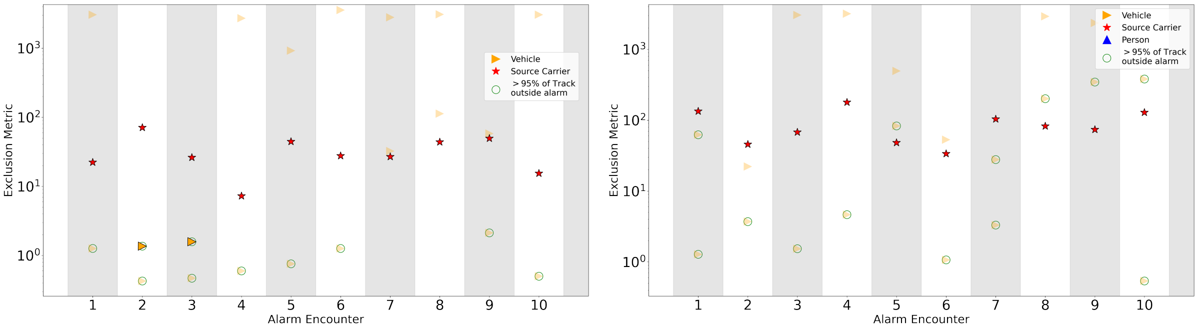

(a)

(b)

Figure 4.6: Results from performing source-object attribution for 10 alarm encounters with LEMURS in a mobile scenario using video (a) and LiDAR (b). In the scenario, LEMURS drives straight past three cars in a row where one of the vehicles is carrying a 1.87 mCi ^{137}Cs source inside of 2 cm Pb. In both (a) and (b), the faded (outlined) points indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

The experimental results of the 30 alarm encounters for the source carrier along with the synthetic results for video and LiDAR are shown in Figure 4.7. The experimental results for the source carrier from all the three scenarios are presented together since all the alarm encounters involved straight line passes with LEMURS. For both video and LiDAR, the experimental and synthetic data are within a standard deviation. A threshold value of 261.78 (118.36) was set to define a 95% true positive rate for video (LiDAR). The frame-to-frame variation in the distance estimate for video causes the large statistical variation in the exclusion metric results. Additionally, the inferred distance is the reason the threshold value for video is higher compared to LiDAR.

Comparing the threshold value of 118.36 for LiDAR with the threshold value of 363.5 from Section 4.2.1 for LiDAR demonstrates the large statistical variations that can exist when performing source-object attribution. For a given alarm encounter, the duration an object is tracked, the distance of the source from the detectors, the complexity of the scene (i.e. attenuation and scattering from untracked objects), and the location of the source in the object can all influence the goodness-of-fit metric. By setting a threshold value for video or LiDAR, this could cause a high number of false negatives, which cannot be tolerated in radiological source search in urban environments. Thus, determining the object responsible for the radiological alarm in a particular encounter varies from encounter to encounter, and it is dependent on identifying the object(s) with the lowest S-value(s) present that are within the alarm window and have best-fit models that are not best described by a background-only model. Subsequently, the object(s) can be flagged as potentially being responsible for the radiological alarm. By not using a threshold value, a low number of false positives is achieved by only performing attribution when BAD detects the presence of an anomaly.

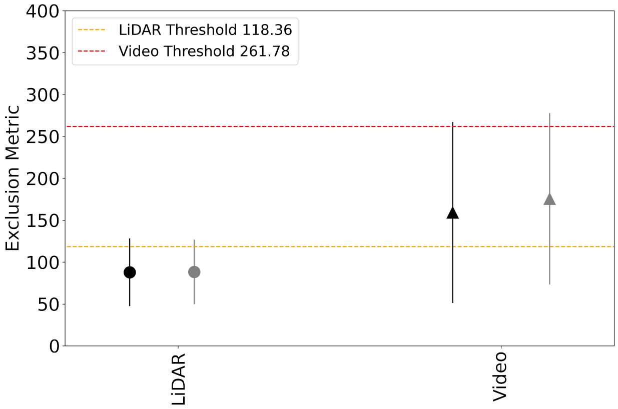

Figure 4.7: Exclusion metrics from 30 different experimental alarm encounters performing source-object attribution with the LEMURS system driving straight past a mobile vehicle carrying a 1.87 mCi ^{137}Cs source inside of 2 cm Pb shielding using LiDAR (a) and video-based trajectories (b), respectively. The simulated data represents synthetic trajectories and count-rate data that was randomly sampled for 1000 trials and subjected to the same analysis as the experimental data. The sensitivity of the goodness-of-fit metric is indicated by the large uncertainties in the synthetic trials. The exclusion metric is the S-value.

4.3 Source-Object Attribution in a Mock Urban Environment

After characterizing the performance of the source-object attribution analysis in the presence of static and mobile sources, a mock urban environment at RFS was created to further characterize the attribution analysis approach. A 1.87 mCi ^{137}Cs source inside of 2 cm of lead-shielding was placed in the trunk of a vehicle. LEMURS and the source carrier performed straight line drive-bys going either 10 or 20 mph. Both LEMURS and the source carrier drove straight for 15 m before passing in the middle of an intersection that had pedestrians on either side of the intersection walking parallel to LEMURS and the source carrier. Two stationary cars were on both sides of the intersection and were perpendicular to the direction of motion for LEMURS. Additionally, the source carrier was followed by a car traveling 10 or 20 mph depending on the trial. This scenario was repeated more than 18 times for both speeds.

In the below Sections 4.3.1–4.3.2, attribution was performed using the individual detectors from the detector array. Comparisons are made between using LiDAR and video for source-object attribution. In addition, the tracking and source-object attribution performance of using SLAM or INS to produce pose estimates of the LEMURS system in a world-fixed frame are also compared.

In Section 4.3.3, source attenuation is present and attribution is performed. Attribution is performed using the individual detectors from the detector array and SLAM is used to produce pose estimates.

4.3.1 10 mph Scenario

Figure 4.8 shows the results from a single alarm encounter using video data and SLAM to create pose estimates of LEMURS in a world-fixed frame. Figure 4.8a shows the progression of the alarm encounter as the source carrier (Track 4 - white vehicle) drives past LEMURS. The source carrier and surrounding objects are continuously tracked throughout the alarm encounter. The trajectory for each object is depicted in Figure 4.8b with the FOV of each LEMURS camera overlaid on top. The inferred depth estimation for each object causes the noise in the trajectories. Both Track 6 and Track 2 were stationary vehicles throughout the encounter, but due to frame-to-frame uncertainty in the distance estimate for video-base trajectories, the Kalman filter pose estimation for both objects varies.

The results of the alarm encounter in Figure 4.8c show that Track 4 (white vehicle) has a clear correlation with the radiological data. This is the correct attribution as the white vehicle was responsible for the radiological alarm. Track 5 has some correlation with Detectors 1-3, but with the angular response of the detector array, the trajectory does not follow the count-rate data in Detectors 4-6 and can be excluded. The remaining trajectories can also all be excluded from the analysis.

In Figure 4.9, the same alarm encounter as Figure 4.8 is shown but now using LiDAR trajectories and SLAM. The color-coding of each object (but not the labels) has been kept consistent with Figure 4.8a. The scene shown in Figure 4.9b corresponds with Figure 4.8a image C. The trajectories of each object are smoother compared to video data (Figure 4.8b), and the variability of Track 11 and Track 13 (Track 2 and Track 13 from Figure 4.8a image B and A, respectively) is more concentrated around each object's respective position. The

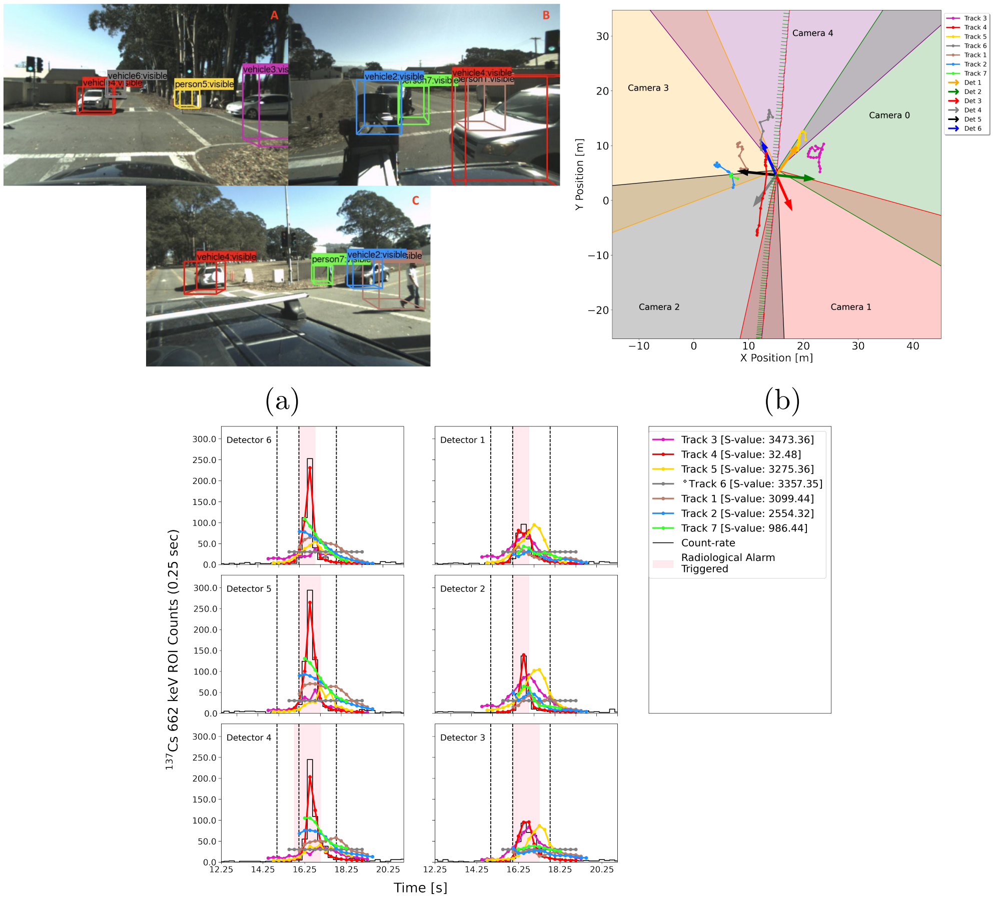

(a) (b)

(c)

Figure 4.8: Performing source-object attribution in a mock urban environment using video data and SLAM to produce pose estimates in a world-fixed frame with the LEMURS system driving 10 mph straight past a mobile vehicle (Vehicle 4 - white SUV) traveling 10 mph and carrying a 1.87 mCi ^{137}Cs source inside of 2 cm Pb shielding. In (a), three different moments of the alarm encounter are depicted, and (b) shows the trajectories for each object with the FOV of each camera overlaid on top of it. Additionally, the orientation of LEMURS is depicted by the red (x-axis), green (y-axis), and blue (z-axis) axes, and each detector's orientation is represented by the arrows. The camera FOV and detector positions correspond with image C in (a). In (c), the result of the alarm encounter is shown. The individual count-rate data from each 6 NaI(Tl) detector is displayed, and the pink bar indicates the period of time the radiological alarm was triggered. Also, the dashed lines from left to right correspond with images a–c in (a).

inherent depth information extracted from LiDAR enables more reliable position estimation.

The results from the alarm encounter are shown in Figure 4.9b. All of the trajectories can be excluded from the analysis except for Track 17, which correlates with the radiological data and is the correct attribution. With LiDAR, there are more trajectories than objects present in the scene. The LiDAR detection CNN has a higher number of false positive detections (i.e., detecting an object when in fact no object is present at the given location) compared

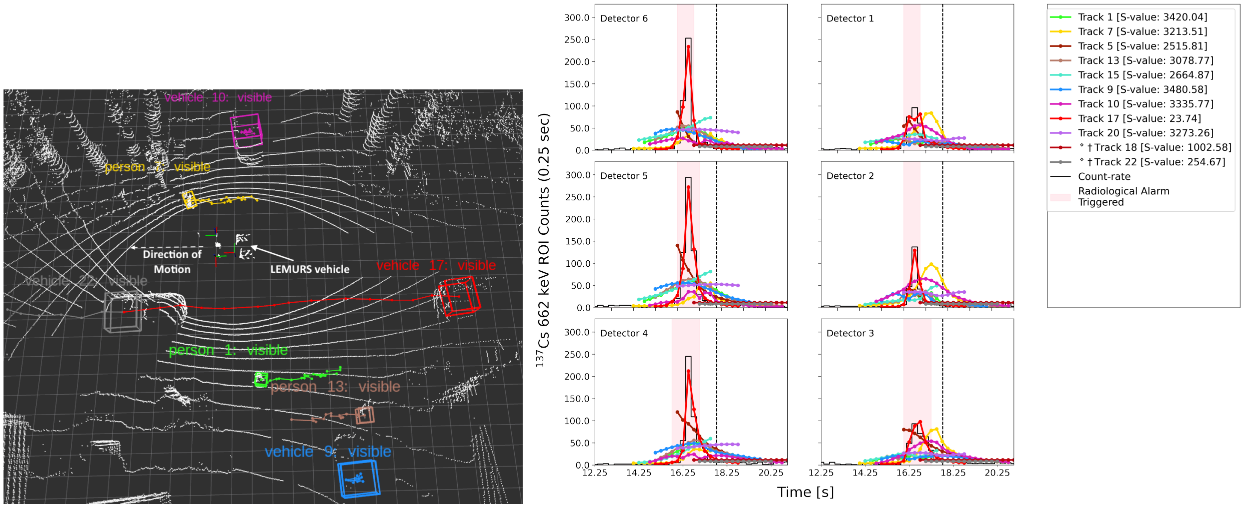

(a) (b)

Figure 4.9: Performing source-object attribution in a mock urban environment using LiDAR data with the LEMURS system driving straight past a mobile vehicle (Vehicle 4 - white SUV from Figure 4.8a) carrying a 1.87 mCi ^{137}Cs source inside of 2 cm Pb shielding. SLAM was used to generate pose estimates of the LEMURS system in a world-fixed frame. The alarm encounter is the same as shown in Figure 4.8. The bounding box colors in (a) of the objects (but not the labels) are consistent with Figure 4.8. In (a), the trajectory of each object to that point is shown. The white grid lines represent $1\,\mathrm{m}^2$, and the orientation of the two LiDARs and detector array on LEMURS are indicated by the red (x-axis), green (y-axis), and blue (z-axis) axes. (b) shows the trajectories for each object with the FOV of each camera overlaid on top of it. Additionally, the orientation of LEMURS is depicted by the red (x-axis), green (y-axis), and blue (z-axis) axes, and each detector's orientation is represented by the arrows. In (c), the result of the alarm encounter is shown. The count-rate data is the summed response from all 6 NaI(Tl) detectors, and the pink bar indicates the period of time the radiological alarm was triggered. Also, the dashed line corresponds with image c in Figure 4.8a. The diamond indicates the best-fit model was better described by a background only model, and the dagger indicates more than 95% of the trajectory was outside of the radiological alarm.

to Yolov4-tiny. These artifacts could be limited by increasing the minimum confidence score needed to track an object. However, increasing the minimum confidence score could decrease tracking performance for certain objects. The false positive detections would need to be discarded by an operator in real time.

The results from analyzing the above alarm encounters using INS to produce pose estimates in a world-fixed frame for video are shown in Figure 4.10a-b. The color-coding is consistent with Figure 4.8a. In Figure 4.10a, the trajectories of each object in the scene are overlaid on the FOV of the cameras. The trajectories created from tracking using the INS frame are comparable to the trajectories shown in Figure 4.8b when tracking with SLAM. In addition, the trajectory of LEMURS is continuous and smooth similar to the trajectory using SLAM. The result of this alarm encounter is seen in Figure 4.10b. Similar to Figure 4.8c, there is a clear correlation between the source carrier (Track 5 – white vehicle) and the count-rate data and all other trajectories can be rejected.

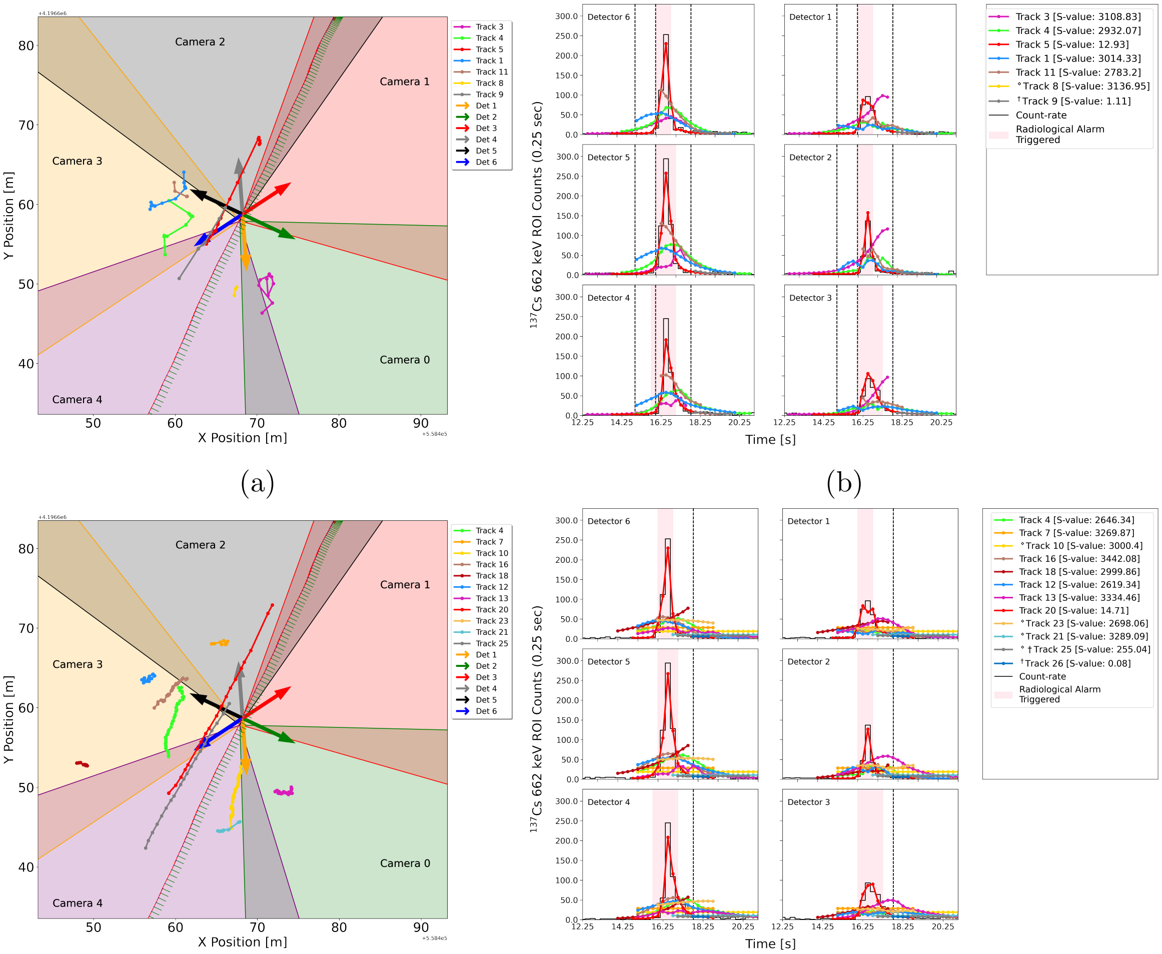

(a) (b)

(c) (d)

Figure 4.10: Performing source-object attribution in a mock urban environment using video (a-b) and LiDAR (c-d) data and INS to produce pose estimates in a world-fixed frame with the LEMURS system driving 10 mph straight past a mobile vehicle (Vehicle 5 (a) - white SUV from Figure 4.8a - and Vehicle 17 (Figure 4.9b)) carrying a 1.87 mCi ^{137}Cs source inside of 2 cm Pb shielding. The alarm encounter corresponds with the alarm encounter shown in Figure 4.8. (a) and (c) show the trajectories for each object with the FOV of each camera overlaid on top of it for video and LiDAR, respectively. Additionally, the orientation of LEMURS is depicted by the red (x-axis), green (y-axis), and blue (z-axis) axes, and each detector's orientation is represented by the arrows. In (b) and (d), the result of the alarm encounter is shown for video and LiDAR, respectively. The individual count-rate data from each 6 NaI(Tl) detector is displayed, and the pink bar indicates the period of time the radiological alarm was triggered.

Using LiDAR-based trajectories and tracking with the INS enables smooth trajectories and consistent tracking of objects similar to Figure 4.9a (Figure 4.10c). The color-coding of each object has been kept consistent Figure 4.8a. The results of this analysis are similar to Figure 4.10b. The source carrier (Track 20) has a clear correlation with the count-rate data. In this case, the quality of the fit is improved relative to performing source-object attribution with SLAM trajectories. Track 20 has an S-value of 14.71 compared to 23.74 for Track 17.

In Figure 4.11, all of the alarm encounters at 10 mph are shown using SLAM for tracking. For both video and LiDAR, the source carrier was detected and tracked throughout all the alarm encounters. Using LiDAR, the source carrier has the lowest exclusion metric in 19 out of 19 trials, and for video-based trajectories, the source carrier has the lowest exclusion metric in 18 out of 19 trials. In some of these trials, best-fit models for other trajectories were best described by a background only model or more than 95% of the trajectory was outside of the radiological alarm so the trajectories can be excluded from the analysis. The LiDAR source carrier best-fit model has better correlation than the video source carrier best-fit model due to the inherent depth from LiDAR. The frame-to-frame depth estimation can vary significantly which can affect the attribution analysis. In the one trial where the source carrier did not have the lowest exclusion metric for video, an obvious correlation between the best-fit model and radiological data existed. An operator monitoring in real time would be able to correctly attribute the radiological alarm to the source carrier.

While direct comparisons cannot be made between Section 3.4 and the above video-based results since the alarm encounters considered vary, it is seen that using YOLOv4-tiny allows for more consistent object detection compared to YOLOv3-tiny. As a result, this enables better tracking and source-object attribution.

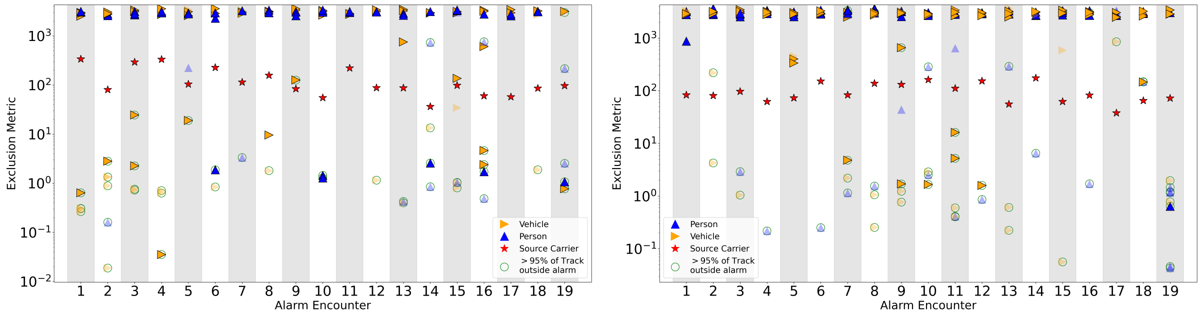

(a)

(b)

Figure 4.11: Source-object attribution on a mobile detector system in a mock urban environment using video (a) and LiDAR (b) and SLAM to produce a world-fixed frame. In both (a) and (b), the faded (outlined) points indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

The results of tracking using pose estimates produced with the INS for all the alarm encounters for video and LiDAR are shown in Figure 4.12a and Figure 4.12b, respectively. The results are similar to tracking and performing source-object attribution using SLAM. In this case, the source carrier has the lowest exclusion metric in 19 out of 19 trials using both LiDAR and video-based trajectories. The results from this analysis do not match the hypothesis that INS should produce degraded pose estimates in an urban environment, which would adversely impact the source-object attribution analysis. In a typical urban environment, buildings will reflect and block satellite signals causing signal interference reducing the position accuracy in the pose estimate. With navigational systems, this loss of position accuracy from obstructions is expressed as dilution of precision (DOP). The fewer satellites available for the pose estimate, the higher the DOP value. In the mock urban environment considered here the vertical DOP (VDOP) and horizontal DOP (HDOP) values range from 1–20, respectively, where values ranging from 5–20 indicate moderate to low confidence levels in the pose estimates due to high satellite interference. The larger DOP values are due to a high number of tall trees in the mock environment that occlude the satellite signal. While the INS does filter between GPS coordinates at ∼1 Hz to improve GPS accuracy and reduce jitter in the pose estimates, in these environments the position accuracy is still reduced. When the INS pose information for the alarm encounters is overlaid onto a map there is an obvious drift over time of the pose estimates (Figure 4.13). However, for a given alarm encounter, all the objects in the scene are relative to this drift since the objects are transformed into the world-fixed frame. Thus, tracking and attribution can still be performed effectively.

Overall, this analysis demonstrates that source-object attribution can improve situational awareness in a mock urban environment using either video or LiDAR as well as INS or SLAM. In all of the alarm encounters, an obvious correlation existed between the best-fit model for the source carrier and the count-rate data.

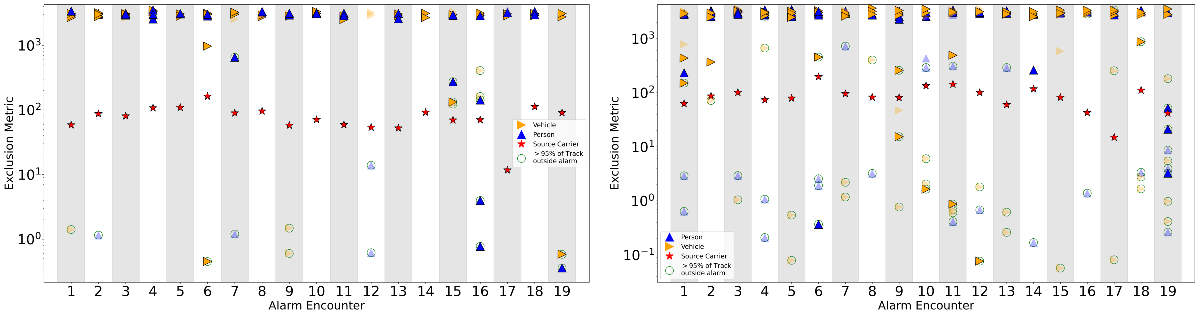

(a) (b)

Figure 4.12: Source-object attribution on a mobile detector system in a mock urban environment using video (a) and LiDAR (b) and INS to create the world-fixed frame. In both (a) and (b), the faded (outlined) points indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

Figure 4.13: Top-down view of the INS pose estimates overlaid onto a map for all the alarm encounters. The blue lines should be within the bounds of the road, Owl Way, but there is a drift of the estimates over time.

4.3.2 20 mph Scenario

The same scenario presented in Section 4.3.1 was repeated for a LEMURS and source carrier speed of 20 mph. These trials were performed to better understand tracking and attribution performance at higher vehicle speeds that are more relevant for urban environments. In this case the scenario was repeated 23 times.

Figure 4.14 shows the results from one of the alarm encounters using video-based trajectories. Figure 4.14a shows the progression of the alarm encounter as the source carrier (Track 15 - white SUV) drives past LEMURS. The source carrier is tracked across three different cameras during this alarm encounter. The large bounding boxes in image B for Track 11, Track 15, and Track 17 reflects the high position uncertainty assigned to objects near the edge of camera images. Once the objects are closer to the center of the image frame in image C, the bounding boxes are much more concentrated around the objects. This indicates a higher certainty in the object's position.

The source carrier and surrounding objects are continuously tracked throughout the alarm encounter. The results of the alarm encounter in Figure 4.14b show that the source carrier has a clear correlation with the radiological data, which is the correct attribution, and Track 15 has the highest association between the physics-based model for the trajectory and radiological data indicated by a S-value of 201.10. Both Track 14 and Track 17 correlate well with the count-rate data; however, both can be excluded from the analysis. Track 14 has a best-fit model that is better described by background, and more than 95% of Track 17 is outside of the radiological alarm.

In Figure 4.15, the results of LiDAR tracking using SLAM (Figure 4.15a) and performing attribution (Figure 4.15b) are shown. The color-coding is consistent with Figure 4.14. It is seen that the objects are consistently tracked throughout the alarm encounter in Figure 4.15a. With this tracking information, a clear correlation between the source carrier (Track 33) and the count-rate data from all 6 NaI(Tl) detectors exists. In this case, Track 17 and Track 38

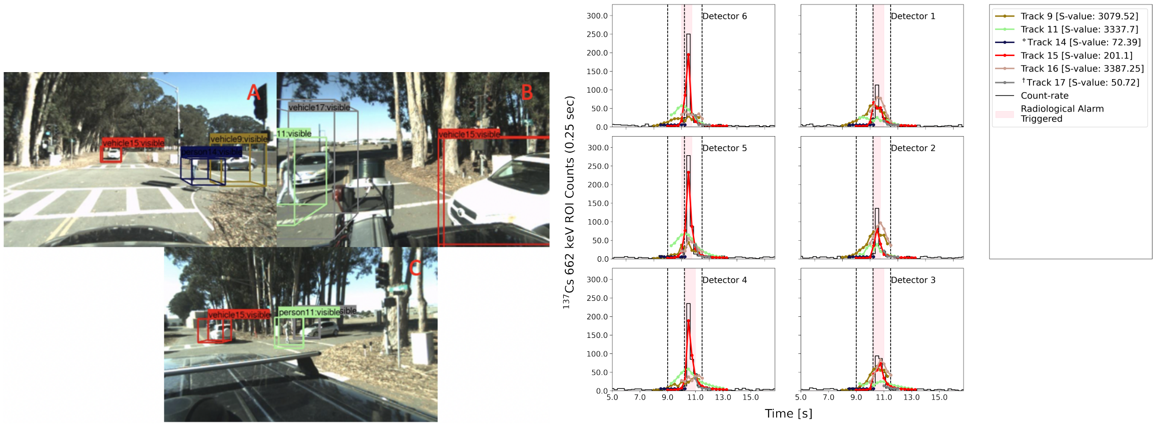

(a) (b)

Figure 4.14: Performing source-object attribution in a mock urban environment using video data with the LEMURS system driving straight past a mobile vehicle (Vehicle 15 - white SUV) carrying a 1.8 mCi ^{137}Cs source inside of 2 cm Pb shielding. In (a), three different moments of the alarm encounter are depicted. The bounding boxes are MVN representations and the size of the bounding boxes reflects the certainty in an object's position. In (b), the attribution analysis of this alarm encounter is shown. The count-rate data from each detector is labeled with the detector ID. The pink bar indicates the period of time the radiological alarm was triggered. The dashed vertical lines from left to right correspond with images a–c in (a). The diamond indicates the best-fit model was better described by a background only model.

correlate with the count-rate data, but both tracks can be excluded from the analysis because a background only model better describes the best-fit models for both tracks.

The results from these individual encounters demonstrate the flexibility of a Kalman filter and the advantage of using MVN tracking. In both cases, objects were only detected and tracked for a short period of time; however, for both video and LiDAR, all the objects were continuously tracked throughout the transient alarm encounter. This enabled more effective alarm attribution in these cases.

Using INS to track and perform source-object attribution is shown in Figure 4.16. Similar to the results shown in Figure 4.15b, the source carrier (Track 34) has a clear correlation with the count-rate data. The best-fit model for Track 34 has the highest association between the physics-based model for the trajectory and the radiological data indicated by a S-value of 20.47, which is the correct attribution. Best-fit models with lower S-values can be excluded because the best-fit models for those trajectories are better described by a background only model.

The INS analysis was also extended to video-based trajectories; however, due to inconsistent tracking, attribution could not be performed for this particular alarm encounter (results not shown).

The results from all the alarm encounters for the 20 mph scenario using SLAM are shown in Figure 4.17 for both video and LiDAR. In Figure 4.17a, 20 out of 23 trials the source carrier had the lowest exclusion metric. In one of the trials, the source carrier was inconsistently

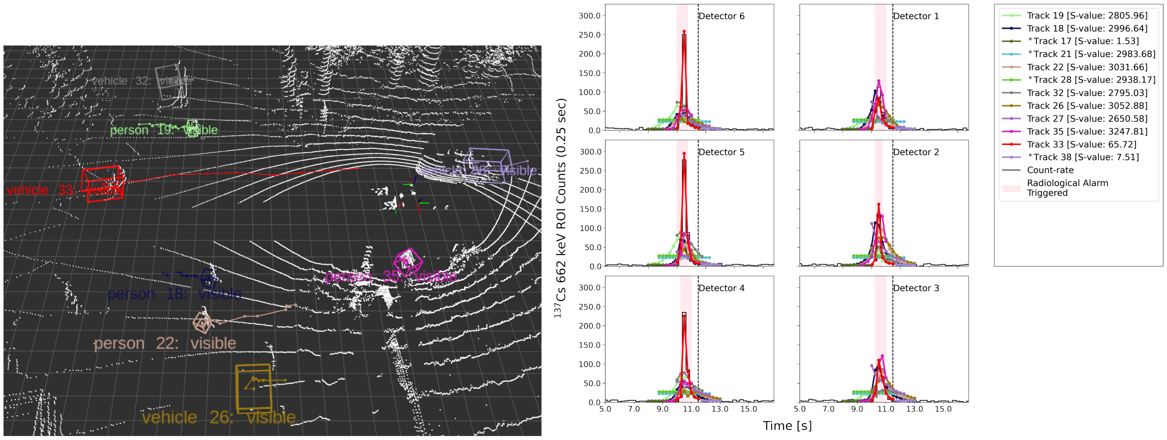

(a)

(b)

Figure 4.15: Performing source-object attribution in a mock urban environment using LiDAR data with the LEMURS system driving straight past a mobile vehicle (Vehicle 33 - white SUV from Figure 4.14) carrying a 1.8 mCi ^{137}Cs source inside of 2 cm Pb shielding. In this example, tracking and attribution was performed using SLAM. In (a), a LiDAR point cloud of the alarm encounter is depicted and corresponds with the image C from Figure 4.14b. In (b), the attribution analysis of this alarm encounter is shown. The count-rate data from each detector is labeled with the detector ID. The pink bar indicates the period of time the radiological alarm was triggered. The dashed vertical line corresponds with the moment in time in (a). The diamond indicates the best-fit model was better described by a background only model.

tracked and no attribution between the object responsible for the radiological alarm and the count-rate data was made. For the remaining two trials, a clear correlation existed and an operator monitoring in real time would be able to identify the object responsible for the radiological data. Using LiDAR-based trajectories (Figure 4.17b), the source carrier was assigned the lowest exclusion metric in 23 out of 23 trials, but a clear correlation existed in all of the trials.

Using the source-object attribution analysis with SLAM, these results demonstrate that situational awareness can be improved for both LiDAR and video with increased source carrier and LEMURS vehicle speeds. In these transient encounters, the object responsible for the radiological alarm was correctly attributed to the count-rate data in a majority of the alarm encounters.

Tracking using INS, the source carrier has the lowest exclusion metric in 10 out of 23 alarm encounters with video-based trajectories (Figure 4.18a) and 22 out of 23 alarm encounters with LiDAR-based trajectories (Figure 4.18b). With video, there were 13 alarm encounter where alarm attribution could be not performed due to inconsistent tracking. This result is in contrast with the video SLAM results from Figure 4.17a where the source carrier had the lowest exclusion metric in 20 out of 23 alarm encounters. Also, a clear correlation existed in 2 of the 3 remaining trials between the source carrier and the count-rate data. It is thought the discrepancy in the results is due to the noise from the INS pose information. While all the objects in the scene are relative to the INS drift, the INS noise can cause position

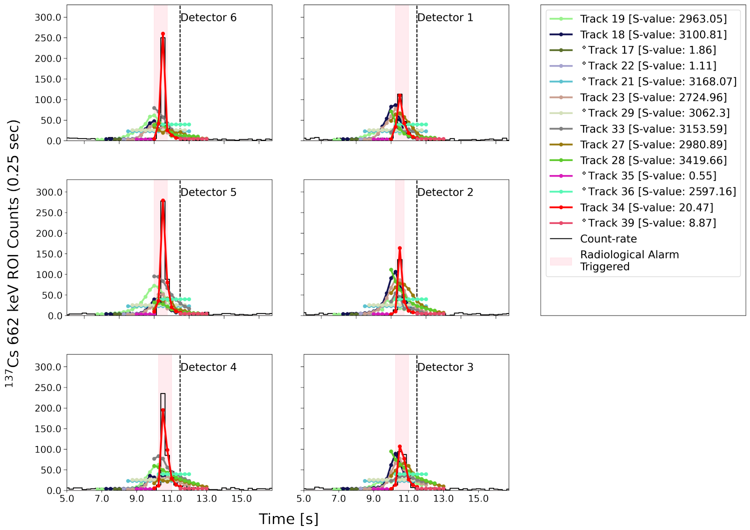

Figure 4.16: LiDAR-based source-object attribution results from Figure 4.15 using INS information for tracking. The color-coding is consistent with Figure 4.15. The count-rate data from each detector is labeled with the detector ID. The pink bar indicates the period of time the radiological alarm was triggered. The dashed vertical line corresponds with the moment in time in (a) from Figure 4.15. The diamond indicates the best-fit model was better described by a background only model.

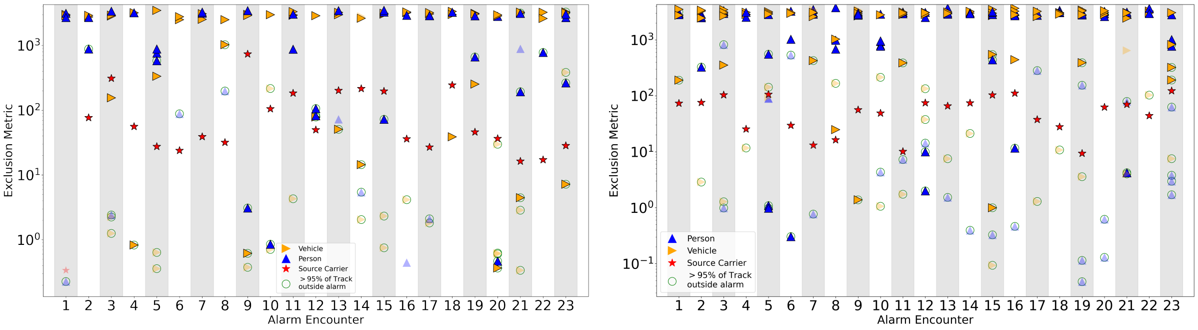

(a) (b)

Figure 4.17: Results from performing source-object attribution on all radiological alarm encounters when both LEMURS and the source carrier were traveling 20 mph relative to each other in the mock urban environment using video (a) and LiDAR (b). SLAM was used in (a) and (b) to generate pose estimates in a world-fixed frame. In both (a) and (b), the faded (outlined) points indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

variations ranging from 25 to 50 cm between data packets for a given alarm encounter. In the 20 mph case, this noise is sufficient enough along with the inferred distance information and increased speed to cause degraded tracking and attribution performance for video. The LiDAR results are comparable to Figure 4.17b due to the inherent depth information.

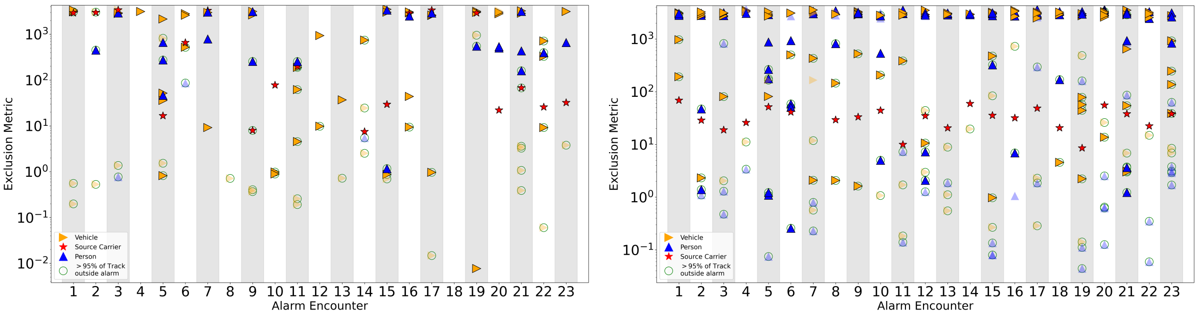

(a)

(b)

Figure 4.18: Source-object attribution on a mobile detector system in a mock urban environment using video (a) and LiDAR (b) and INS to create the world-fixed frame. In both (a) and (b), the faded (outlined) points indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

The findings from this analysis demonstrate that situational awareness can be improved using LiDAR and INS; whereas with video, there was a reduction in attribution performance compared to using SLAM. However, it is expected performance with an INS to worsen in an environment subject to degraded GPS performance (e.g. an urban canyon).

4.3.3 Source Attenuation Scenario

The above two sections (Section 4.3.1 and Section 4.3.2) involve alarm encounters where the source carrier is never occluded from LEMURS. Additionally, the scenarios involve localizing the source to a vehicle and not a pedestrian. To explore performing source-object attribution with attenuation present, a mock environment was created using the intersection at RFS. A 1.87 mCi ^{137}Cs source in 2 cm of lead-shielding was put in a briefcase and carried by a pedestrian. LEMURS traveled through the intersection at a speed of 10 mph in the same direction as the source carrier. As LEMURS approached the source carrier, vehicles traveling in the opposite direction of LEMURS and the source carrier temporarily occluded the source carrier from LEMURS causing modulation in the count-rate data at different moments in time. This scenario was repeated 10 times.

An example of one alarm encounter using LiDAR-based trajectories is shown in Figure 4.19. Figure 4.19a shows the LiDAR point cloud with the tracking bounding boxes and each object's trajectory to that point in time overlaid on top of it. All the objects are continuously tracked throughout the alarm encounter. In particular, the pedestrian carrying the source (Track 16) is effectively tracked through temporary occlusions caused by Track 28 and Track 41 as both objects drive through the intersection. Track 38, Track 39, and Track 46 are false detections. No objects were present at those positions.

In Figure 4.19b, the result of the alarm encounter is shown. From about 18.5-19.25 seconds and 20-21 seconds in the alarm encounter, the source attenuation from Track 28 and Track 41, respectively, occluding the pedestrian (Track 16) carrying the radiological source is clearly present in the count-rate data. With the continuous tracking in particular through the temporary occlusion, the attenuation that is present is captured in the best-fit model for Track 16, and the best-fit model for Track 16 clearly follows the radiological data, especially the modulation from the source attenuation. In this case, Track 16 was responsible for the radiological alarm and has the lowest S-value. Track 38, 50, and 52 have lower S-values but can be excluded from the analysis because the best-fit models for these objects are either better described by a background only model or are more than 95% of the track is outside the alarm window. The remaining trajectories can be excluded from the analysis and do not follow the radiological data.

To determine if accounting for attenuation improves attribution performance relative to neglecting attenuation effects in the best-fit model, the above analysis was repeated (Figure 4.20). However, this time attenuation was not accounted for in the best-fit model. The best-fit model for Track 16 follows the count-rate data and has the lowest S-value among all the objects present, but the modulation in the count rate is not captured by the best-fit model, in particular in Detectors 4-6. Interestingly, the best-fit model in this case produces a lower S-value compared to the best-fit model from Figure 4.19b (550.36) even though the best-fit model does not follow the source attenuation in the count-rate data. This result suggests that the goodness-of-fit metric lacks sensitivity to transient attenuation effects in the scene.

In Figure 4.21, the same alarm encounter as Figure 4.19 is shown but now using video-based trajectories. The color coding has been kept consistent with Figure 4.19. Figure 4.21a shows three different camera images as LEMURS drives past the source carrier. The tracks for each object throughout the alarm encounter overlaid on each cameras FOV is shown in Figure 4.21b. The source carrier is continuously tracked through the temporary occlusions caused by Vehicles 18 and 14 as they pass through the intersection. The large variation in the source carrier's position around an X-position of 14 m is due to the temporary occlusion from Vehicle 18. Also, both Track 9 and Track 16 (Track 24 and Track 26 in Figure 4.19a) were stationary objects throughout the alarm encounter; however, there are large variations in the estimated position for these objects due to the distance estimation.

The result of this alarm encounter is seen in Figure 4.21c. The source carrier (Track 10) clearly follows the count-rate data, and the source attenuation caused by the temporary occlusion from Vehicle 14 and Vehicle 18 is reflected in the best-fit model. All of the trajectories can be excluded from the analysis except for Tracks 14, 9, and 20. Of these, only Track 9 follows the count-rate data; however, the model for Track 9 misses the attenuation that is present in the count-rate data. Track 14 and 20 are tracks generated from the same object due to inconsistent tracking between cameras. Due to this, Track 14 and 20 are shorter tracks and thus are more easily fit to the radiological data, but an operator monitoring would be able to determine that the trajectories are most likely not responsible for the radiological alarm.

The video results were repeated for this alarm encounter without accounting for attenuation in the model generation for trajectories. Similar to the results for LiDAR without accounting for attenuation, the best-fit model for the source carrier had a lower S-value (559.61 – result now shown) compared to the S-value generated by accounting for attenuation. This result along with the S-value generated for Track 9 from Figure 4.21c demonstrate that the goodness-of-fit metric is not sensitive enough to attenuation effects.

The results of all of the alarm encounters with attenuation present are seen in Figure 4.22

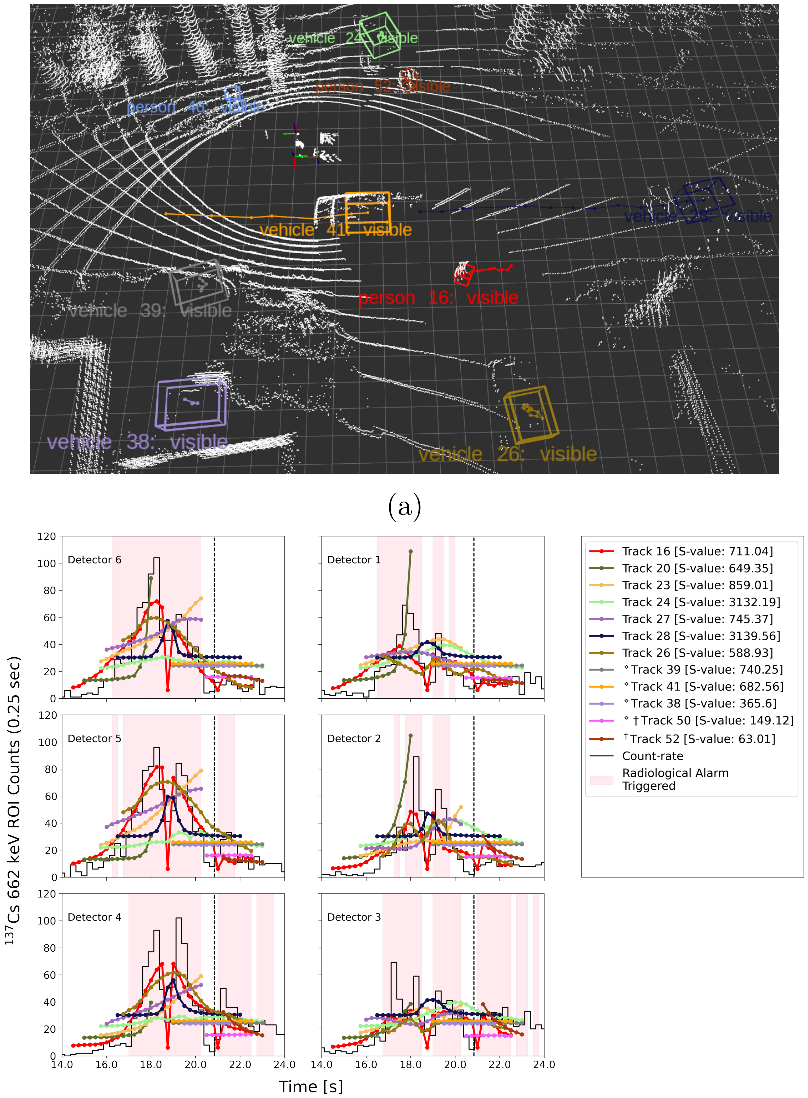

(a)

(b)

Figure 4.19: Performing source-object attribution in a mock urban environment using LiDAR. In this example, the LEMURS system drives straight past a mobile pedestrian (Person 16 - red bounding box) carrying a 1.8 mCi ^{137}Cs source inside of 2 cm Pb shielding, and when LEMURS passes Person 16, Vehicle 28 and Vehicle 41 occlude the source carrier from LEMURS at different points in time throughout the alarm encounter. In (a), a LiDAR point cloud of the alarm encounter is depicted. In (b), the attribution analysis of this alarm encounter is shown. The count-rate data from each detector is labeled with the detector ID. The pink bar indicates the period of time the radiological alarm was triggered. The dashed vertical line corresponds with the moment in time in (a). The diamond indicates the best-fit model was better described by a background only model, and the dagger represents the track was more than 95% outside of the alarm window.

for both video and LiDAR. In Figure 4.22a, the results show that the source carrier has the lowest exclusion metric in 3 out of 10 trials. In 6 of the remaining trials, inconsistent

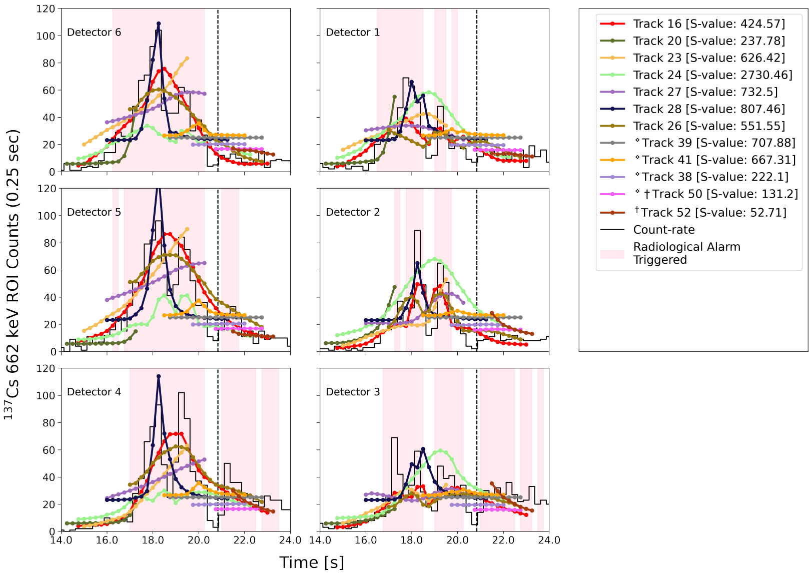

Figure 4.20: Performing source-object attribution without including attenuation effects in the best-fit model calculation in a mock urban environment using LiDAR. The alarm encounter considered corresponds with Figure 4.19. The count-rate data from each detector is labeled with the detector ID. The pink bar indicates the period of time the radiological alarm was triggered. The dashed vertical line corresponds with the moment in time in Figure 4.19a. The diamond indicates the best-fit model was better described by a background only model, and the dagger represents the track was more than 95% outside of the alarm window.

tracking of objects allowed best-fit models to have lower exclusion metrics than the source carrier. However, there was 1 trial (Trial Number 6) where adjudicating the alarm could not be performed. In this case, the tracking of the object was consistent, but there was too much noise in the source carrier's trajectory preventing correct attribution. When the above analysis was performed without accounting for attenuation in the physics model, the best-fit model for the source carrier had the lowest exclusion metric in 3 out of 10 trials, which matches the results when attenuation is accounted for in the physics model. However, the S-values were lower when attenuation was not included in the model generation.

With LiDAR (Figure 4.22b), the source carrier is correctly assigned the lowest exclusion metric in 1 out of 10 alarm encounters. As mentioned earlier in Section 4.3.1, the threshold to track a detected object using SECOND is low to enable more consistent tracking of objects; however, this allows for more false positive object detections. In 6 of the remaining trials, false positives prevented the source carrier from having the lowest exclusion metric. Repeating the analysis without accounting for attenuation in the physics model resulted in the source carrier having the lowest exclusion metric in 4 out of 10 trials, and again, the S-values were lower compared to accounting for attenuation in the model generation.

In Figure 4.22, the S-values for the source carrier are quite large in these encounters for both video and LiDAR. The large S-values are caused by an incomplete modeling of the attenuation present. The current implementation of the attenuation calculation only applies attenuation when an object is fully occluded. This is a limitation of this approach

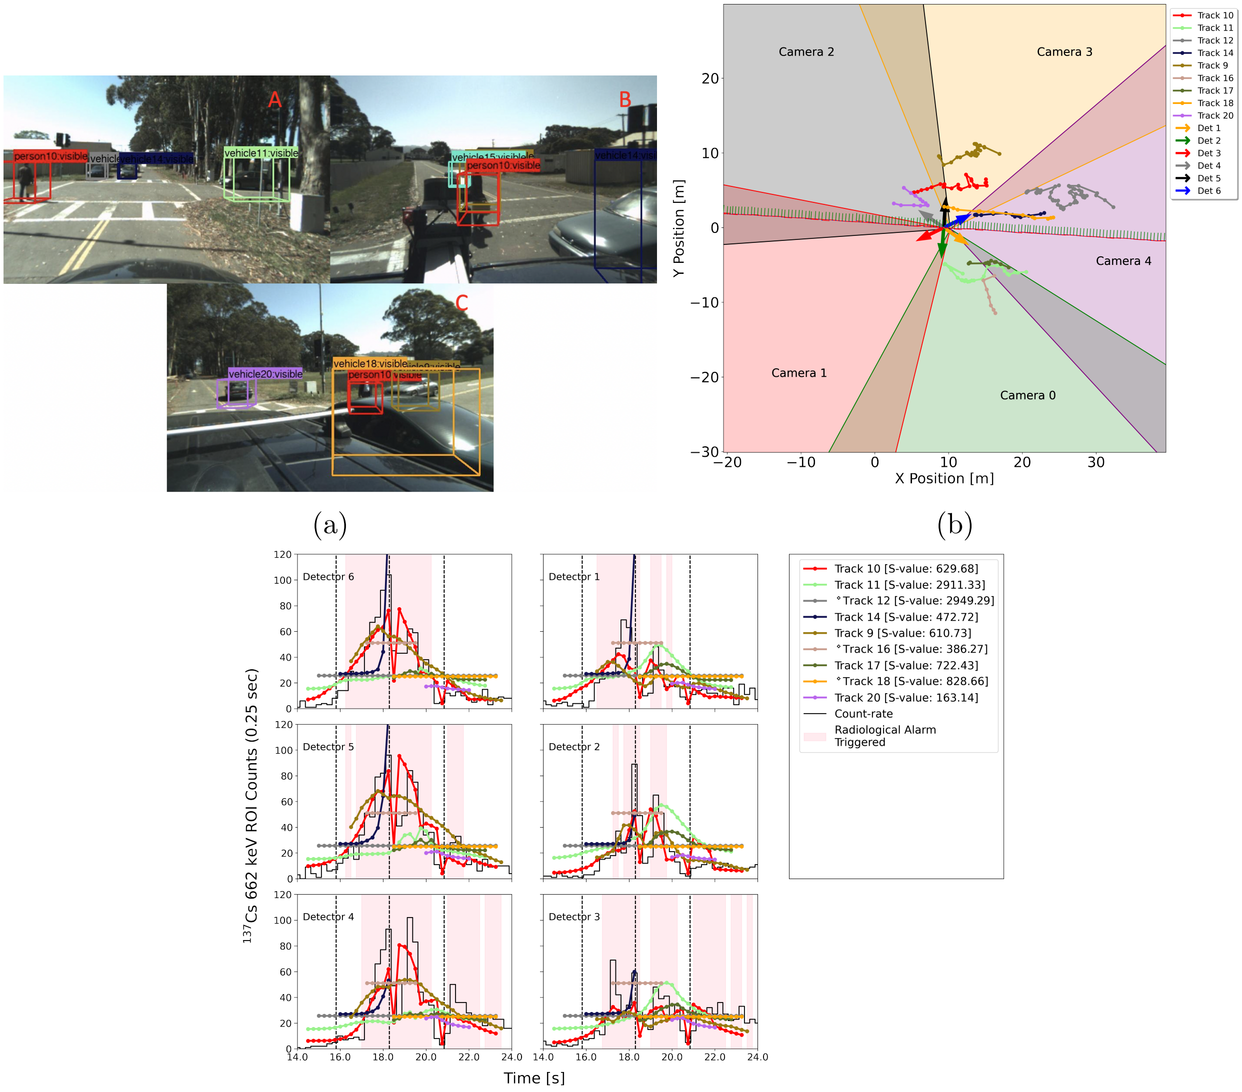

(a)

(b)

(c)

Figure 4.21: Performing source-object attribution in a mock urban environment using video. In this example, the LEMURS system drives straight past a mobile pedestrian (Person 10 - red bounding box) carrying a 1.8 mCi ^{137}Cs source inside of 2 cm Pb shielding, and when LEMURS passes Person 16, Vehicle 14 and Vehicle 18 occlude the source carrier from LEMURS at different points in time throughout the alarm encounter. (a) shows three different image frames as LEMURS drives past the source carrier In (b), a top-down view of the video trajectories overlaid on top if the FOV of each camera. The color-coding of the objects is consistent with Figure 4.19. In (c), the attribution analysis of this alarm encounter is shown. The count-rate data from each detector is labeled with the detector ID. The pink bar indicates the period of time the radiological alarm was triggered. The dashed vertical lines from left to right correspond with images ac in (a). The diamond indicates the best-fit model was better described by a background only model.

because attenuation does not occur instantaneously. Rather, the amount of attenuation present changes with time as the fraction of the solid angle covered by an occluding object between the detector and radiological source changes. Also, the attenuation from tracked objects is not known *a priori* and is assumed. This can cause either an overestimation or underestimation of the total attenuation imposed by an occluding object. Additionally, the findings suggest that the goodness-of-fit metric is not sensitive enough to attenuation effects in the count-rate data. In the alarm encounters, when attenuation was not accounted for in the physics model, the S-value was lower indicating a better association between the best-fit model and radiological data compared to the results when attenuation was accounted for in the model generation. With more complete modeling and a more sensitive goodness-of-fit metric, attenuation effects should further improve attenuation performance. Nonetheless, the results of this analysis demonstrate that situational awareness can be improved when the source is occluded. In a majority of the alarm encounters using video and LiDAR, a clear correlation between the best-fit model and the count-rate data existed even when the source was attenuated, and an operator monitoring in real-time would still be able to perform alarm adjudication.

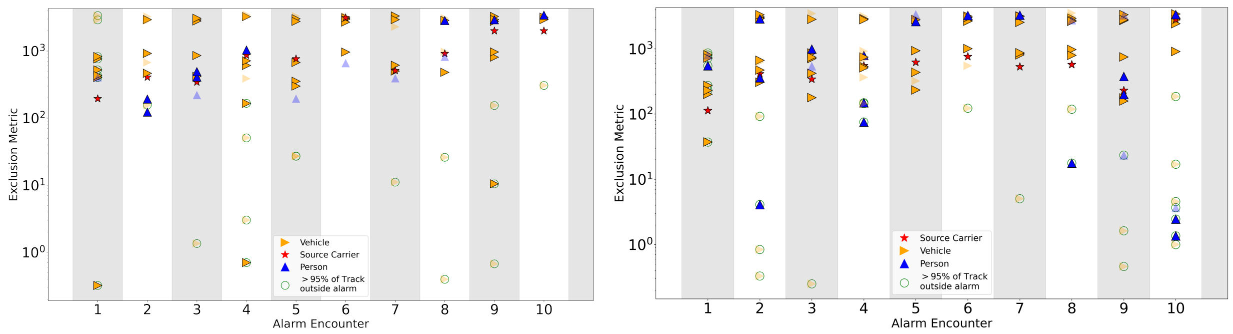

(a) (b)

Figure 4.22: Results of performing source-object attribution with source attenuation present in the scene for video (a) and LiDAR (b). In both (a) and (b), the faded (outlined) points indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

4.4 Improved Detection Sensitivity with Track-informed Optimized Integration Windows

With tracking information, optimal integration times to maximize SNR can be found using the formulation discussed in Section 2.10. Furthermore, for a given alarm encounter, certain detectors will be closer to the source carrier and will experience higher SNR than the detectors further from the source carrier. Thus, an optimal configuration of detectors should exist that should maximize SNR. The following two sections (Section 4.4.1 and Section 4.4.2) investigate optimizing integration windows to maximize SNR using a mobile detector system in the presence of a mobile source. Optimal integration windows are found either using the 6 NaI(Tl) detectors independently or summing the response of the 6 NaI(Tl) detectors. In both cases, position uncertainty of the object is accounted for with MCMC. Also, the analyses were only performed using the trajectories generated from tracking with SLAM.

4.4.1 10 mph Scenario

Figure 4.23 shows an example of applying the optimum integration window analysis for multiple detectors for the alarm encounter from Figure 4.9. The time-segments that should optimize SNR for the detector array for this alarm encounter are overlaid on the count-rate data. It is seen that the time-segments correspond with the time of closest approach between LEMURS and the source carrier. The spectral data from each detector's integration window is then combined to produce an optimal spectrum.

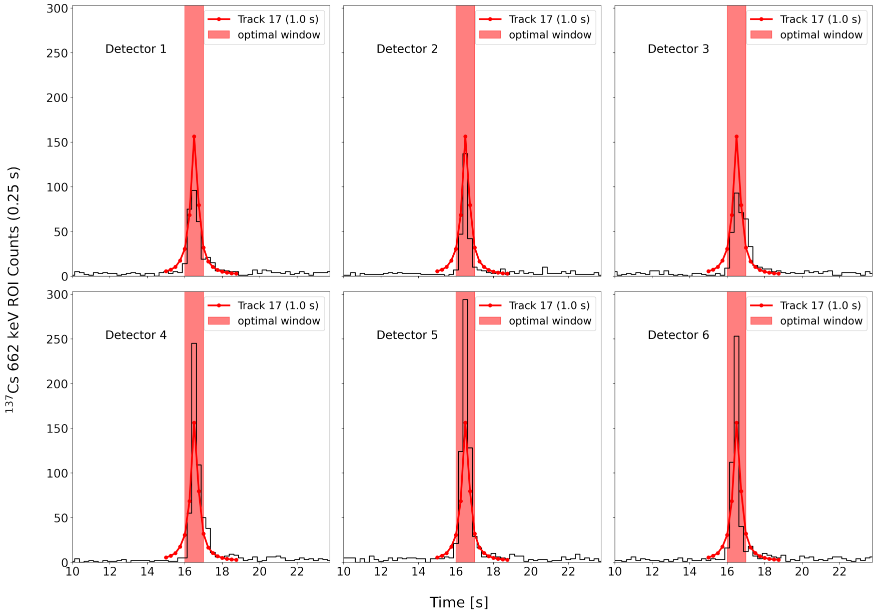

Figure 4.23: Optimized integration windows for all 6 NaI(Tl) detectors overlaid onto the count-rate data (black line) for the alarm encounter from Figure 4.9. The color coding for the source carrier (Track 17) is consistent with Figure 4.9.

The above approach can also be applied to the summed response of all 6 NaI(Tl) detectors. With the summed response, the detector array becomes a monolithic detector, and the optimal integration window is determined from the summed response.

Figure 4.24a and Figure 4.24b show 100 different samples drawn from the posterior distribution using MCMC along with the model that maximizes the likelihood for both video and LiDAR, respectively. The 100 samples demonstrate the variability of the best-fit model calculation when accounting for position uncertainty, and by constraining the priors in the MCMC calculation, the drawn samples correlate with the radiological data and do not deviate significantly from each other. For LiDAR, the MCMC sampling overestimates the radiological data before ($\sim$ 16 seconds) and after the radiological alarm ($\sim$ 17.5 seconds) demonstrating that the simple physics model is not perfect at modeling the expected count-rate. Nonetheless, the model that maximizes the likelihood for both video and LiDAR correlates well with the radiological data.

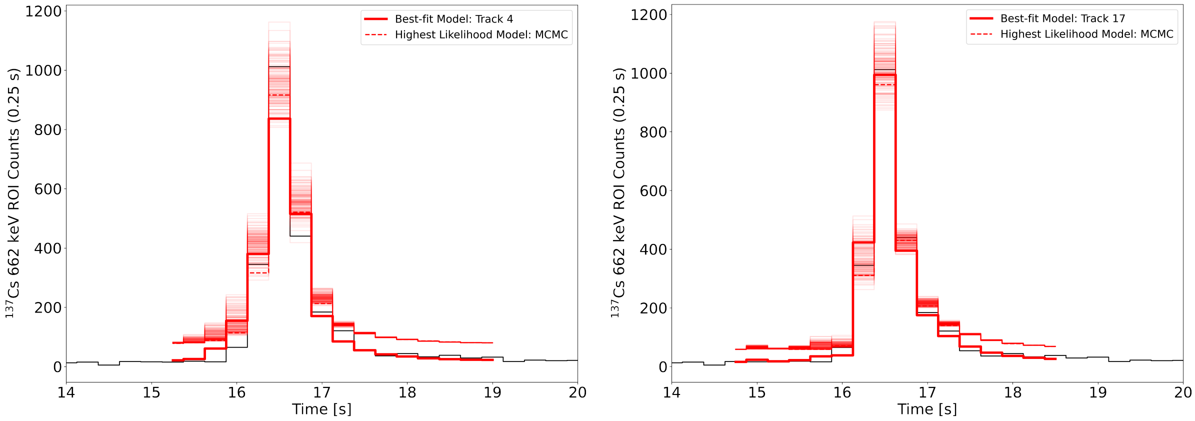

(a)

(b)

Figure 4.24: Applying MCMC to account for the position uncertainties in the source carrier track and in the optimal integration window formulation for video (a) and LiDAR (b). The trajectories are from the alarm encounter shown in Figure 4.8a and Figure 4.9a for video and LiDAR, respectively. The best-fit model for each track is shown, and the faded lines represent the different samples of the posterior distribution. The dashed-line is the MCMC sample that maximized the likelihood.

Figure 4.25a and Figure 4.25b show the calculated optimal integration windows overlaid on the count-rate data for both video and LiDAR, respectively. In both cases, the optimal integration window captures the time of closest approach. The optimal spectrum is produced by combining the spectral information within the time segments.

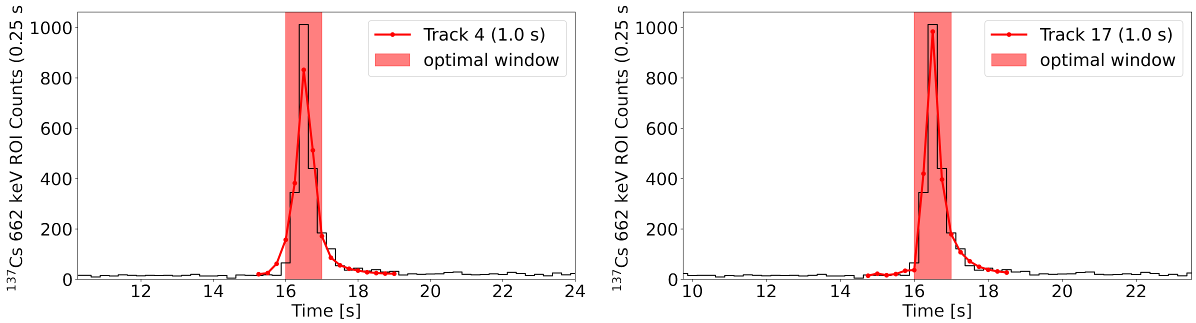

(a)

(b)

Figure 4.25: Optimal integration windows overlaid on the count-rate data for both video (a) and LiDAR (b). The trajectories are from the alarm encounter shown in Figure 4.8a and Figure 4.9a for video and LiDAR, respectively, and the color-coding in (a) and (b) is consistent with Figure 4.8a and Figure 4.9a, respectively.

It is seen in both Figure 4.23 and Figure 4.25 the optimal integration windows does not capture the anisotropic attenuation from the source carrier. For video, an anisotropic attenuation from an object is not accounted for in the simple physics model in Section 2.6.1 since heading information is not available. However, even for LiDAR where an attenuation correction is added to the data, there is not an improvement in the track-informed window capturing the anisotropic attenuation that is present.

An optimized spectroscopic search was applied to the optimal spectrum produced using the 6 NaI(Tl) detectors independently for all the trajectories and was compared to either the optimal spectrum created from the summed response of the source carrier or different fixed integration windows for video (Figure 4.26a) and LiDAR (Figure 4.26b), respectively. The color coding in Figure 4.26a and Figure 4.26b is consistent with Figure 4.8a and Figure 4.9a, respectively. The optimal integration window analysis using the 6 NaI(Tl) detector independently for the source carrier in video (Track 4) and LiDAR (Track 17) produce close to the largest anomaly value relative to the other trajectories in the scene. In both cases, it is seen that trajectories with similar trajectories to the source carrier have elevated anomaly values as well such as Track 2 and Track 5 for video and Track 5 in LiDAR. The optimal integration window for video improves detection sensitivity relative to a fixed integration window of 1.0, 3.0, 4.0, and 5.0 seconds. For LiDAR, the track-informed integration window is the optimal integration window compared to all fixed integration windows.

Using a summed detector response in both video and LiDAR, a track-informed integration window of 1 second is the optimal integration window compared to a fixed 1, 3, 4, or 5 second integration window but is less than a fixed integration window of 2 seconds. For both video and LiDAR, using all 6 NaI(Tl) detector independently to find an optimal configuration of detectors improves detection sensitivity relative to using the summed response of the detectors, which matches the hypothesis that for a given alarm encounter there exists an optimal configuration of detectors that will maximize SNR compared to summing the response of all the detectors together. The results of this analysis for video and LiDAR suggests that track-informed integration windows can improve the anomaly value, a proxy for detection sensitivity, relative to different fixed integration windows.

Figure 4.27 shows the results from all of the alarm encounters when LEMURS was traveling 10 mph for both video and LiDAR. With video trajectories (Figure 4.27a), using the 6 NaI(Tl) detectors individually and MCMC to produce the track-informed optimal integration window yielded an optimal window in 8 of the 19 trials, and for the other 11 trials the track-informed window is close to the optimal integration window in most of the encounters. When the response function is summed together and MCMC is applied to the data the track-informed optimal integration window is the optimal window in 2 of the 19 trials. The discrepancy between the two track-informed optimal integration windows is due to the higher amount of background that is present in the summed response. With the 6 NaI(Tl) detectors, the optimal window for each detector captures the time of closest approach when the signal is highest, which improves the anomaly value, a proxy for detection sensitivity. In 15 out of the 19 trials, using the 6 NaI(Tl) detectors individually produces a higher anomaly value compared to the summed response.

For LiDAR trajectories (Figure 4.27b), the results show that the track-informed integration window is the optimal window in 6 (1) of the 19 trials for optimal configuration of detectors (summed response). In both cases, the track-informed windows are close to the optimal integration window in the remaining 13 (18) trials. In this case, the optimal configuration of detectors produces a larger anomaly value compared to the summed response in

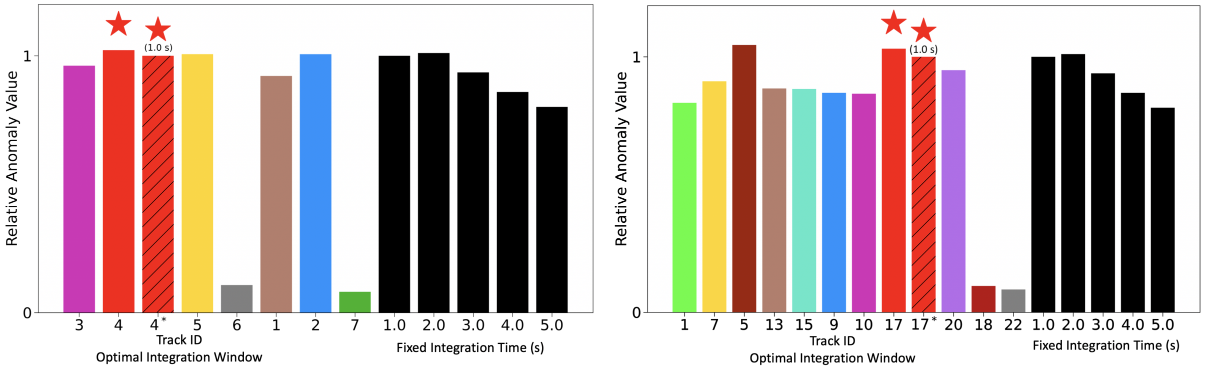

(a) (b)

Figure 4.26: Spectroscopic analysis for video (a) and LiDAR-based (b) trajectories using either the trajectory's optimized integration window or fixed integration windows. The trajectories for this analysis are from Figure 4.8a and Figure 4.9a for video and LiDAR, respectively. Also, the color coding in (a) and (b) is consistent with Figure 4.8a and Figure 4.9a, respectively. In addition, the $*$ indicates the spectroscopic analysis was performed using the summed response of the detector array, and the parentheses indicate the time duration of the optimal integration window for the summed response.

all of the alarm encounters. This result along with the video-based results aligns with the hypothesis that for a given alarm encounter there exists an optimal configuration of detectors that will maximize SNR compared to summing the response of all the detectors together.

Overall, from these transient alarm encounters, the track-informed integration approaches for both video and LiDAR produce optimal integration windows that are either the optimal integration window or are close to the optimal integration window. In most cases, a fixed integration window of 1 or 2 second produces the optimal integration window, but there is variability in which fixed integration window produces the best result. Even though the track-informed integration windows were not always the optimal window, the results from this analysis demonstrate that the track-informed integration approach can better inform integration times by adapting to the dynamics of the scene and relative motion of objects in a scene. In addition, the results suggest the track-informed integration windows can improve detection sensitivity on a mobile system relative to fixed integration windows.

To explore the track-informed optimal integration formulation in lower SNR environments, the spectra data in the above alarm encounters were binomial down-sampled. In order to properly compare the spectroscopic analysis using the individual detectors or summing the detectors into a monolithic detector, the binomial down-sampling was applied to the spectra data from each individual detector. Subsequently, the spectroscopic analysis was either applied to the individual detectors by summing the spectral data from each detector's respective optimal integration window or to the summed response from all 6 NaI(Tl) detectors. The spectra data was binomial down-sampled by 75% of the highest count-rate data within a given alarm window.

The results of this analysis for video-based and LiDAR-based trajectories is shown in Figure 4.28a and Figure 4.28b, respectively. In Figure 4.28a, the optimal configuration of

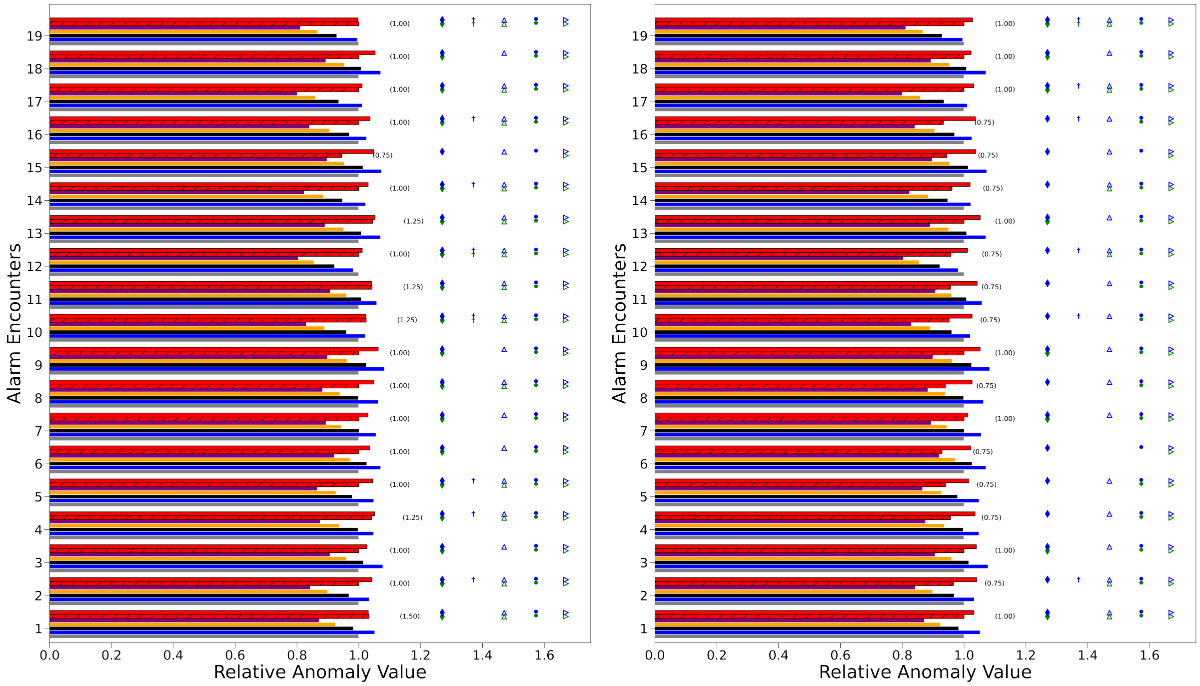

(a) (b)

Figure 4.27: Maximum relative anomaly values from each alarm encounter from Figure 4.11a and Figure 4.11b using either track-optimized time-windows or fixed integration windows for both video (a) and LiDAR (b), respectively. In (a) and (b), the solid red bars (red bars with hatch marks) indicate track-optimized windows using the 6 NaI(Tl) detectors independently (summing the response of the detectors). The grey, blue, black, orange, and purple bars indicate a fixed integration window of 1.0, 2.0, 3.0, 4.0, or 5 seconds, respectively. The blue (green) diamond, dagger, triangle, filled in circles, or right facing triangle indicate the encounters where the optimal integration window yielded a higher anomaly value than a 1, 2, 3, 4, or 5 second integration window, respectively, using the 6 NaI(Tl) detectors independently (summing the response of the detectors).

detectors improves the anomaly value relative to all fixed integration windows in 5 out of the 19 trials. In the remaining cases, the track-informed integration window produces an anomaly value that is close to the maximum anomaly value generated using a 1 second integration window. By performing the spectroscopic analysis on the summed response of the detectors, performance is degraded compared to the optimal configuration of detectors. None of the cases produce the maximum anomaly value. Additionally, the track-informed anomaly values are not close to the maximum anomaly value produced by the the fixed integration window of 1 second.

For LiDAR-based trajectories (Figure 4.28b), using the optimal configuration of detectors, the optimal integration window has the highest relative anomaly value in 6 of the 19 trials, and in the remaining cases, the track-informed integration window produces an anomaly value that is close to the maximum anomaly value produced by the 1 second fixed integration window. Similar to higher SNR environment presented earlier in this section, the optimal configuration window enables better detection sensitivity compared to the summed response.

The results from the lower SNR study support the hypothesis that using tracking information can better inform integration times. In Figure 4.27, the maximum anomaly value in a majority of the alarm encounters was generated using a 2 second fixed integration window; whereas in the lower SNR environment presented in Figure 4.28 a 2 second window only produced the maximum anomaly value in 1 alarm encounter. In this environment, a 1 second fixed integration window generated the maximum anomaly value in the majority of cases. Multiple fixed integration windows can be run in tandem to try and maximize detection sensitivity. However, it is not possible to account for the countless changes that can occur in a scene that could impact detection sensitivity with a fixed integration window. For example, if a source-carrying vehicle is stopped at a light near LEMURS for more than 10 seconds, a 1 or 2 second fixed integration window will give lower anomaly values compared to a longer integration window. This is the advantage of using the track-informed integration window which adapts to the scene and motion of objects. Overall, the findings from this analysis further suggests that detection sensitivity can be improved relative to fixed integration windows in low SNR environments.

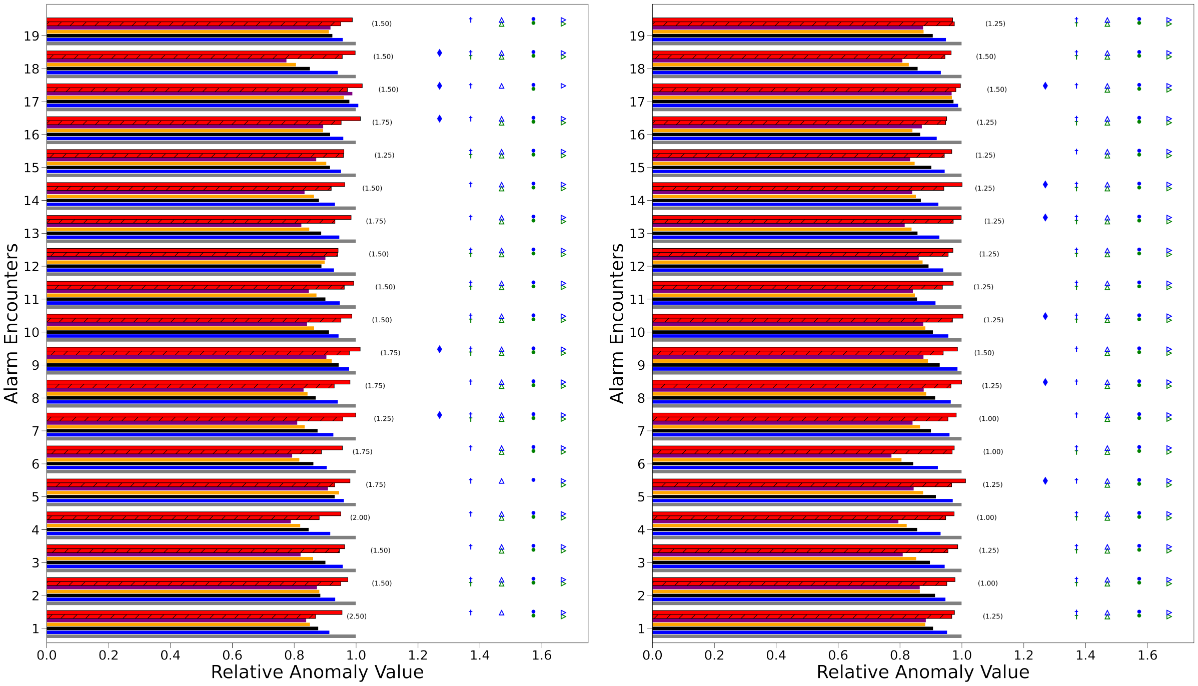

(a) (b)

Figure 4.28: Maximum relative anomaly values from each alarm encounter from Figure 4.11 using binomial down-sampled spectral data and applying either track-optimized time-windows or fixed integration windows for video (a) and LiDAR (b). In (a) and (b), the red bars with hatch marks (solid red bars) indicate track-optimized windows using summing the response of the detectors (6 NaI(Tl) detectors independently). The grey, blue, black, orange, and purple bars indicate a fixed integration window of 1.0, 2.0, 3.0, 4.0, or 5 seconds, respectively. The time duration for the summed response optimal window is provided in the parentheses. The green (blue) diamond, dagger, triangle, filled in circles, or right facing triangle indicate the encounters where the optimal integration window yielded a higher anomaly value than a 1, 2, 3, 4, or 5 second integration window, respectively, using the 6 NaI(Tl) detectors independently (summing the response of the detectors).

4.4.2 20 mph Scenario

The spectroscopic analysis from Section 4.4.1 was applied to the alarm encounters where both the source carrier and LEMURS were traveling at 20 mph relative to each other to further explore the optimal integration window analysis in more transient alarm encounters.

In Figure 4.29, the results of applying this analysis to all 23 alarm encounters for video and LiDAR are shown. For both video and LiDAR, the track-informed integration windows were calculated using either the summed response of the detector array or summing the optimal integration windows from each respective detector to produce an optimum spectrum. In Figure 4.29a, the track-informed integration window produces the maximum anomaly value in 10 (11) out of 23 trials compared to different fixed integration windows using the summed response of the detector array (summing the optimal integration windows from each respective detector to produce an optimum spectrum). There was inconsistent tracking with alarm encounter 1 and 3 causing degraded detection performance. The LiDAR-based track-informed integration window produces the largest anomaly value in 4 (5) out of 23 alarm encounters using the summed response of the detector array (summing the optimal integration windows from each respective detector to produce an optimum spectrum) shown in Figure 4.29b. However, in each case, the relative anomaly value for the track-informed integration window using either method is close to the largest anomaly value produced by a 1 second fixed integration window.

While it is expected that using an optimal configuration of detectors throughout an alarm encounter will produce higher anomaly values, in the 20 mph case, the time of closest approach between LEMURS and the source carrier is quicker resulting in the counts from the radiological source to be more concentrated around the time of closest approach, which reduces the contribution of background when a spectroscopic analysis is performed. As a result, the anomaly values using the summed response of the detector array were comparable to the summing the optimal integration windows from each respective detector to produce an optimum spectrum. Altogether, the results of this analysis show that the track-informed integration window can improve the anomaly value compared to different fixed integration windows using either approach for video or LiDAR.

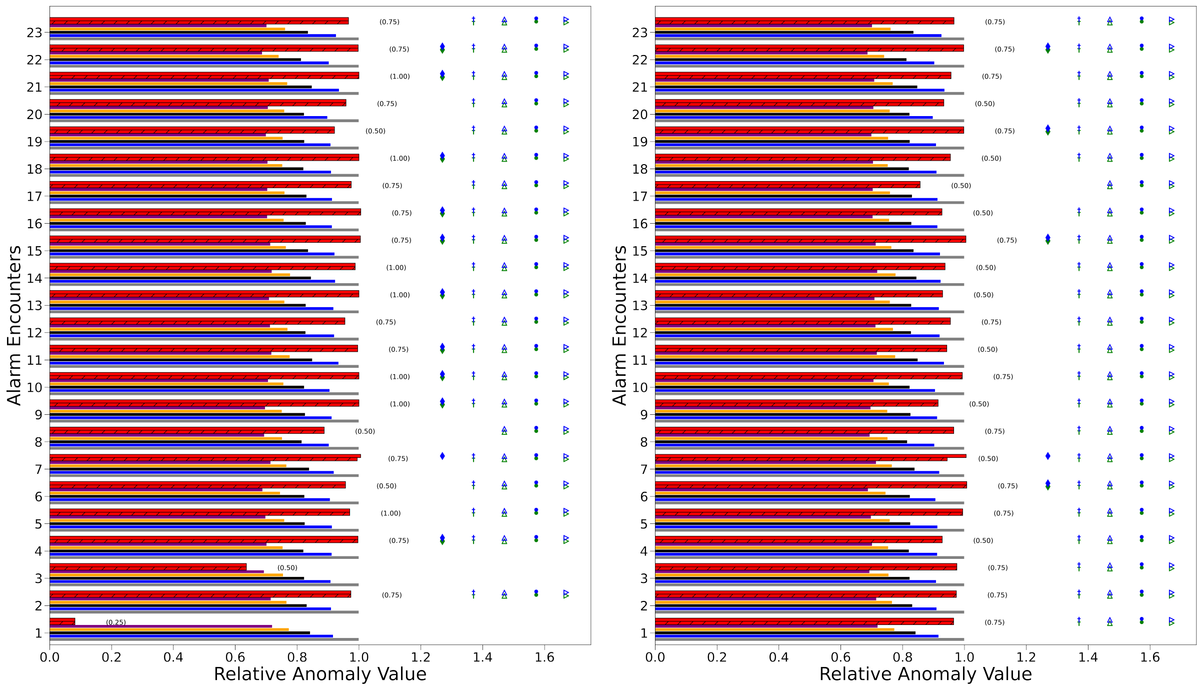

(a)

(b)

Figure 4.29: Maximum relative anomaly values from each alarm encounter from Figure 4.17a and Figure 4.17b using either track-optimized time-windows or fixed integration windows for both video (a) and LiDAR (b), respectively. In (a) and (b), the solid red bars (red bars with hatch marks) indicate track-optimized windows using the 6 NaI(Tl) detectors independently (summing the response of the detectors and applying MCMC). The grey, blue, black, orange, and purple bars indicate a fixed integration window of 1.0, 2.0, 3.0, 4.0, or 5 seconds, respectively. The time duration for the summed response optimal window is provided in the parentheses. The blue (green) diamond, dagger, triangle, filled in circles, or right facing triangle indicate the encounters where the optimal integration window yielded a higher anomaly value than a 1, 2, 3, 4, or 5 second integration window using the 6 NaI(Tl) detectors independently (summing the response of the detectors and applying MCMC).

4.5 Importance of Detector Response in Source-object Attribution

For mobile detector systems, the intrinsics of the vehicle (i.e. number of occupants, occluding material, etc.) can vary often, and all of these variations will impact the response function differently. However, it is not possible to properly account for the myriad different changes that can occur within the vehicle. In order to understand the importance of the detector response for source-object attribution, detector responses with varying levels of complexity were applied to the same alarm encounters considered in Section 4.3.1. In addition, the detector responses were applied in a degenerate scenario described in Section 4.2.1. It should be noted this analysis was only performed using SLAM to generate a consistent reference frame.

The responses compared in this analysis were an isotropic response (lowest fidelity), zero energy response, and the current LEMURS response function (highest fidelity) which is a Monte Carlo response described in Section 2.6.1.1.

The zero energy response functions were generated using OpenGL [49] (Figure 4.30a). As mentioned in Section 2.6.1.1, a zero-energy approach does not account for any gamma-ray physics (i.e. no scattering and infinite attenuation). The OpenGL response consisted of the detector array, vehicle frame, roof rack, and engine block - the major occluding structures in LEMURS (Figure 4.30a). Figure 4.30b shows a zero-energy response map for Detector 0 from Figure 4.30a. The coordinate system of the NaI(Tl) detector array in OpenGL was oriented such that the Z axis was up; however, in the URDF of LEMURS (Figure 2.3), the NaI(Tl) detector array is oriented with Z down. Thus, the directional response of OpenGL was rotated to correctly correspond with the URDF file.

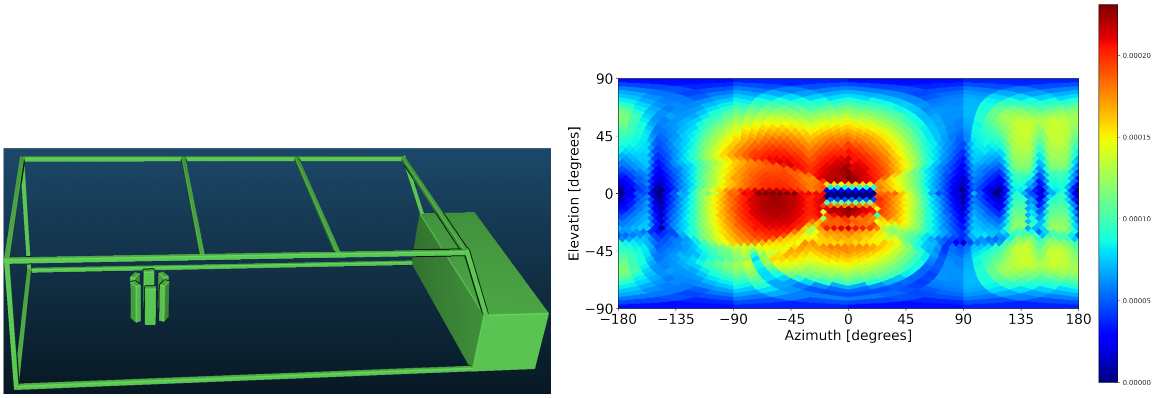

(a) (b)

Figure 4.30: LEMURS system modeled with OpenGL (a). Zero-energy directional response for Detector 0 from the LEMURS NaI(Tl) array is shown in (b). The response is in the reference frame of the NaI(Tl) array which is oriented with the z-axis down (see Figure 2.3).

The OpenGL and isotropic response functions were used to perform source-object attribution on the same alarm encounter from Figure 4.8 for video-based trajectories. The results from using the OpenGL response function are shown in Figure 4.31a. A clear correlation exists between the source carrier (Track 4) and the count-rate data and all of the other trajectories can be rejected. However, even though a correlation exists, the quality of the fit is diminished compared to the results in Figure 4.8c. The S-value for the source carrier is higher in this case. When the isotropic response function is used (lowest fidelity model), the quality of the fit is further reduced. Across all 6 detectors the best-fit model for Track 4 is the same with only the magnitude of the count rate varying between detectors. Since the response is the same in every direction, no modulation is present from the detector array, but the correct attribution is still made with this low fidelity response function while all other trajectories can be rejected. Applying these response functions to LiDAR-based trajectories (results not shown), a similar result is seen. The S-value for the source carrier increases as the fidelity of the response function is diminished; nonetheless, the correct attribution is still made for the OpenGL and isotropic response functions while all other trajectories can be rejected.

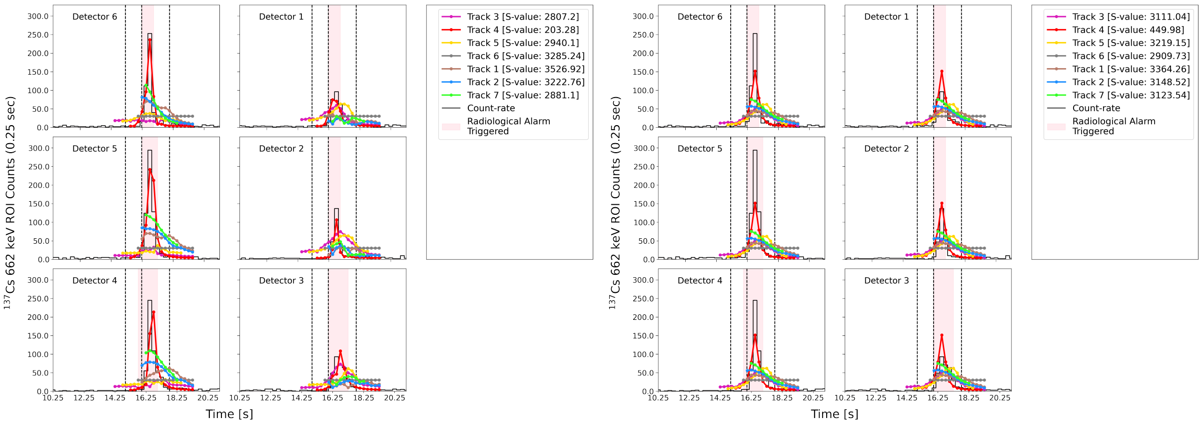

(a) (b)

Figure 4.31: Comparison of using different response functions with varying fidelity to perform source-object attribution. In (a), the response function was modeled with OpenGL and corresponds with Figure 4.30, and (b) is an isotropic response. Both (a) and (b) correspond with the alarm encounter from Figure 4.8.

The results show that using an isotropic response with LiDAR-based (video-based) trajectories the source carrier has the lowest exclusion metric in 17 (11) out of the 19 trials Figure 4.32b (Figure 4.32a). For LiDAR-based trajectories, this is comparable to the results using the current LEMURS response function (Figure 4.11b). However, the exclusion metric values are on average higher compared to the Monte Carlo approach for both video and LiDAR (Figure 4.11a and Figure 4.11b, respectively), but an operator monitoring the system would be able to correctly perform alarm adjudication in all 19 trials.

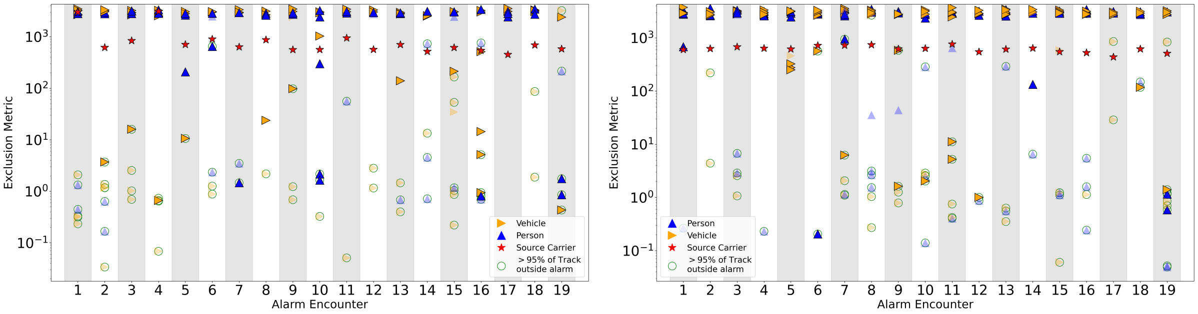

(a) (b)

Figure 4.32: Applying an isotropic response function to the alarm encounters presented in Figure 4.11 for video (a) and LiDAR (b), respectively. In (a), an isotropic response is used for video-based trajectories and (b) shows the results for LiDAR-based trajectories. In both cases, the trajectories were modeled and fit to the 6 NaI(Tl) detectors independently from LEMURS. The faded (outlined) points in (a) and (b) indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

The results for the zero energy response function show an improvement in attribution performance compared to the isotropic response for both video (Figure 4.33a) and LiDAR (Figure 4.33b). The source carrier has the lowest exclusion metric in 18 (16) out of 19 trials for LiDAR (video) trajectories, and a clear correlation between the best-fit model for the source carrier and radiological data exists in each trial, but the exclusion metrics are still higher than the current LEMURS response function.

Overall, this analysis demonstrates that the fidelity in the response function does contribute to the performance of the source-object attribution analysis. With a more detailed response function, the best-fit model for the object responsible for the radiological alarm correlates better with the radiological data. However, most of the localization and attribution performance is driven by proximity ($1/r^2$), and with any of the response functions, an operator monitoring in real time would be able to correctly perform alarm adjudication.

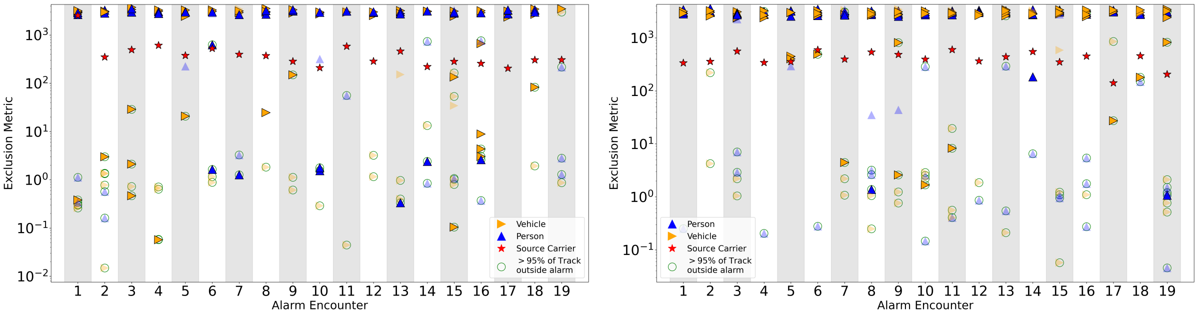

(a) (b)

Figure 4.33: Applying a zero-energy response function to the alarm encounters presented in Figure 4.11. In (a), a zero energy response function is used for video-based trajectories and (b) shows the results for LiDAR-based trajectories. In both cases, the trajectories were modeled and fit to the 6 NaI(Tl) detectors independently from LEMURS. The faded (outlined) points in (a) and (b) indicate best-fit models that were better described by a background only (source plus background) model. The green circles indicate tracks that are 95% or more outside of the radiological alarm.

In the above analyses, the proximity between LEMURS and the source carrier was small (less than 2 m) and the activity of the source present was high leading to large count rates in the detectors. To investigate weaker source encounters with lower counting statistics in the detectors, the count-rate data was binomial down-sampled. The count-rate data was down-sampled by a factor 25%, 75%, 90% of the highest count rate within the time of closest approach of the source carrier. The lowest Exclusion Metric was used as a metric for localization performance.

The results from this analysis are shown in Figure 4.34 for both video (Figure 4.34a) and LiDAR (Figure 4.34b). With both video and LiDAR, the localization performance for both the OpenGL and isotropic response functions decreases with weaker source encounters (i.e., higher percentage of binomial down-sampling); however, a clear correlation does exist in every alarm encounter between the source carrier and the radiological data. In the weaker source encounters, the duration of the radiological alarm window is about 0.5 seconds or less, which leads to encounters where a majority of trajectories have more than 95% of their trajectory outside of the alarm window including the source carrier. To limit rejecting too many trajectories in these situations, if the alarm window is not greater than 0.5 seconds trajectories that are outside of the alarm window are not rejected, which limits localization performance. Nonetheless, with the Monte Carlo response function, the localization performance is consistent across the different source strength encounters. The results from this analysis demonstrate that a higher fidelity response function enables more effective rejection of trajectories that are consistent with the radiological data.

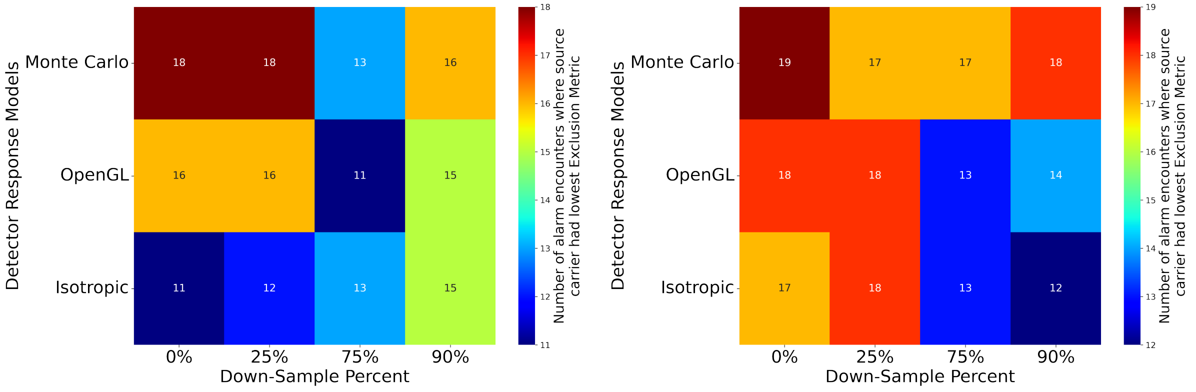

(a) (b)

Figure 4.34: Comparing the source-object attribution performance using response functions with varying levels of fidelity and binomial down-sampled data. The alarm encounters correspond with Figure 4.11 for both video and LiDAR. In total 19 alarm encounters were performed. The annotation in each square represents the number of trials the source carrier had the lowest exclusion metric. The x-axis shows the percentage of the highest count rate within a particular alarm encounter the spectra data was binomial down-sampled, and the y-axis represents the different response functions used to perform source-object attribution.

This topic was further explored using the degenerate scenario described in Section 4.2.1. Similar to the analysis above, localization performance was determined by the number of encounters where the source carrier had the lowest exclusion metric. This analysis was performed using LiDAR-based trajectories, and the analysis was performed with and without binomial down-sampled data.

Figure 4.35 shows the results of performing source-object attribution in the degenerate scenario with different response functions. The color coding is consistent with the objects shown in Figure 4.3. In both Figure 4.35a and Figure 4.35b, the lower fidelity detector response functions (OpenGL and Isotropic, respectively) have reduced localization performance compared to the higher fidelity model (Monte Carlo). With the isotropic response function, all directional information from the detector array is lost since the response is the same in all directions so the degeneracy between the source carrier and the vehicle directly across from the source carrier cannot be broken. The OpenGL response can break the degeneracy present between the source carrier and the vehicle directly across from the source carrier because there is modulation information from the detector array in the response function; however, there is large statistical variation in the exclusion metric for the source carrier in both Figure 4.35a and Figure 4.35b. This leads to the source being localized to both the source carrier and vehicle adjacent to the source carrier; whereas, the Monte Carlo response

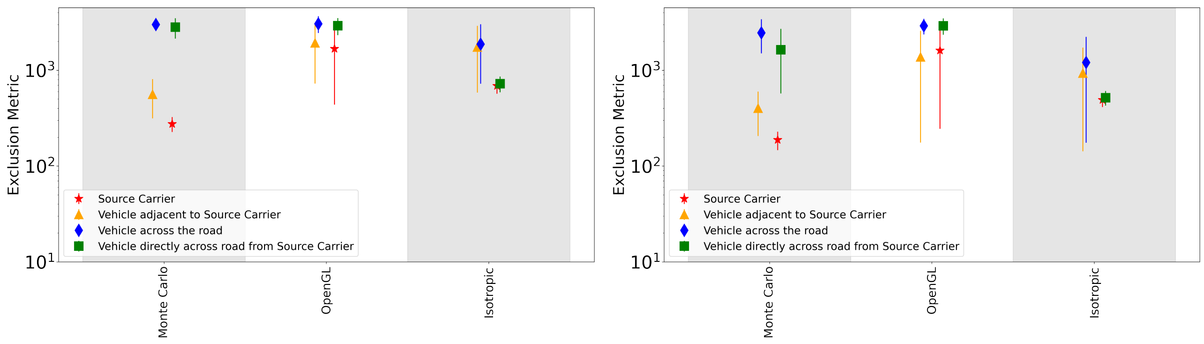

(a) (b)

Figure 4.35: Localization performance in a degenerate scenario (described in Section 4.2.1) using source-object attribution and detector response functions with varying levels of fidelity. In (a), the results of performing source-object attribution without binomial down-sampling the data is shown, and in (b) the count-rate data was down-sampled by 25% of the highest count rate within the time of closest approach by LEMURS. For both (a) and (b), the x-axis describes the fidelity of the detector response function, and the color coding is consistent with the objects shown in Figure 4.3.

function is able to correctly localize the source to the source carrier in both Figure 4.35a and Figure 4.35b.

The results from this analysis demonstrate that the additional information gained from a higher fidelity detector response function does enable better source localization information, especially in degenerate scenarios.

4.6 Importance of Angular Response in Source Localization

To better understand the impact of the angular response on source localization when tracking information is available, source-object attribution was performed using different detector array configurations. The different detector array configurations were generated from the 6 NaI(Tl) detector array of LEMURS. This was done by combining the radiation data from the individual detectors of the 6 NaI(Tl) detector array to create a left-right detector array, a 3 NaI(Tl) detector array, or a monolithic detector (Figure 4.36).

Localization performance of the detector array configurations are first compared to the localization performance of the full 6 NaI(Tl) detector array using source-object attribution. Subsequently, the source localization performance of the different detector array designs using source-object attribution are compared to performing source localization using a 3D image space that is discretized into cubic voxels and reconstructing the spatial distribution of the radiological source utilizing the ML-EM algorithm. In addition, source localization performance is compared to a gamma-ray reconstruction algorithm that assumes sparsity in the voxelized space [22]. All of the analyses were performed using SLAM.

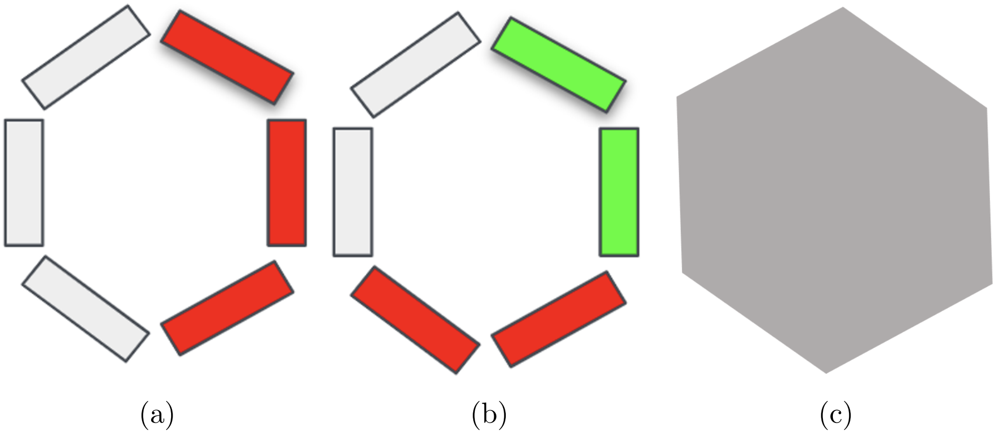
(a) (b) (c)

Figure 4.36: Different detector array configurations generated from the 6 NaI(Tl) detector array of LEMURS. In (a), a left-right detector array is shown and was generated by combining the radiation data from the red-highlighted detectors together to form the left detector, and the right detector was created by summing the response of the remaining three detectors. The 3 NaI(Tl) detector array (b) was generated by summing the detector response for each highlighted detector pair together producing three individual detectors. In (c), the monolithic detector was formed by summing the response from all 6 NaI(Tl) detectors together.

4.6.1 Source-object Attribution using Different Detector Array Configurations

4.6.1.1 Degenerate Scenario

Figure 4.37 shows the results of performing source-object attribution in the degenerate scenario described in Section 4.2.1 using the different detector array configurations. The spectra data in these scenarios has been down-sampled by a factor of 25% of the highest count rate within the alarm window. Similar to the results in Figure 4.4, the variation in the Kalman filter positions and the position-source offset cause the best-fit models to underestimate the magnitude of the source activity in Figure 4.37a – Figure 4.37d. The result of fitting all 6 NaI(Tl) detectors is shown in Figure 4.37a and is similar to Figure 4.4. In this case, the source is correctly attributed and localized to the source carrier (Track 39) while all the other trajectories can be rejected from the analysis. With the 3 detector array (Figure 4.37b) or left-right array (Figure 4.37c) the radiological source is correctly attributed to the source carrier while Track 44 and Track 45 can be rejected from the analysis. The result of performing source-object attributing with a monolithic detector is shown in Figure 4.37d. In this case, the degeneracy that exists between Track 39 and Track 45 cannot be broken. Track 45 is assigned a lower exclusion metric, which indicates that some modulation in the detector array is necessary to break degeneracies that exist in the LEMURS vehicle.

The above degenerate scenario was replicated 20 times. The results from all the source encounters are shown in Figure 4.38 using the different detector array configurations. In these alarm encounters, the count rates in the detectors were quite high due to the source activity and the proximity of LEMURS to the source. To better understand the importance

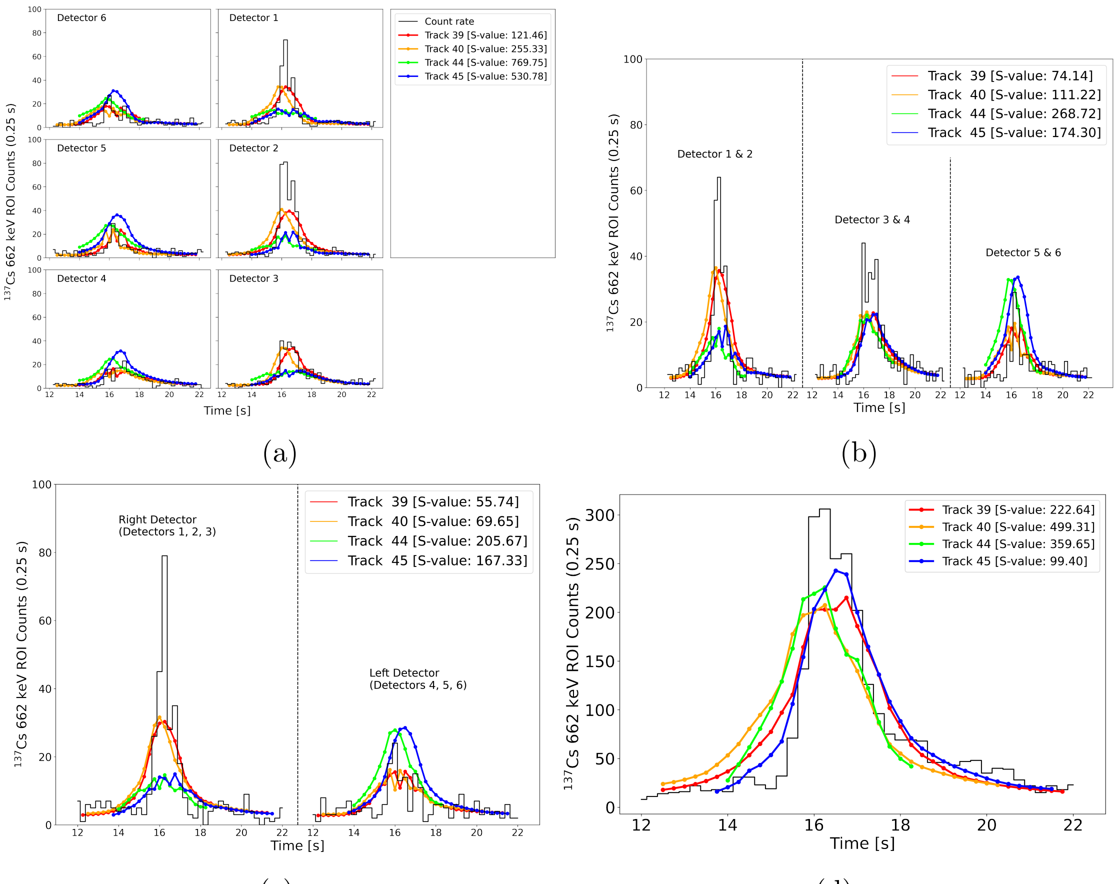

(a) (b)

(c) (d)

Figure 4.37: Performing source-object attribution using different detector array configurations generated from the 6 NaI(Tl) detector array of LEMURS. The trajectories are from Figure 4.3 and the color coding of the objects is consistent. The count-rate data (black line) in (a) – (d) has been binomial down-sampled by a factor of 1.3. In (a), the count-rate data from the full 6 NaI(Tl) detector array and the best-fit models to the count-rate data are displayed. The three detector array is shown in (b). The dashed lines are separating the count rates from the different detectors, and the annotation signifies the detectors summed together to create the response. In (c) and (d), the results of performing source-object attribution using the left-right detector array and monolithic detector are shown, respectively. Again, the annotation in (c) signifies the detectors that were summed together to generate the response.

of the angular response on source localization when tracking information is available, the spectra data was binomial down-sampled by a factor 25% or 75% of the highest count rate within the time of closest approach to investigate localization performance in the presence of weaker sources. Localization performance between the different array configurations was determined by using the lowest exclusion metric. The results show that using either a left-right or 3 NaI(Tl) detector array produces similar localization performance when compared to the full 6 NaI(Tl) detector array. Even in the presence of weaker source encounters, localization performance between the left-right, 3 NaI(Tl), or full 6 NaI(Tl) detector array is similar. For each configuration, the left-right degeneracy that exists between the source carrier and the vehicle across the street from the source carrier is broken. The incorrect attributions are between the source carrier and the vehicle adjacent to the source carrier. In these trials, the vehicle adjacent to the source carrier was assigned the lowest exclusion metric. Similar to the result shown in Figure 4.37d, the monolithic detector cannot break the degeneracies present, and only correctly identifies the source carrier in a few of the trials. In these cases, the vehicle directly across the street from the source carrier was assigned the lowest exclusion metric.

The findings from this degeneracy study suggest that when tracking information is available a simpler detector array design could provide similar radiological source localization compared to a detector array with a more complex angular response in low and high count-rate environments. However, with simpler array designs (i.e., less detectors), the count rate in the detector(s) will be less compared to the full 6 NaI(Tl) detector array. This study did not investigate how this reduction in detection efficiency will impact detection.

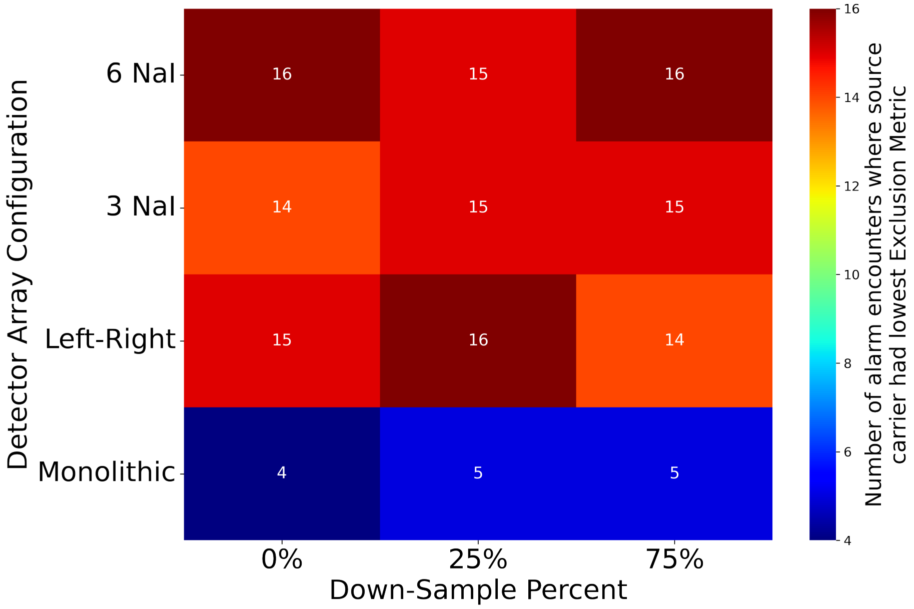

Figure 4.38: Comparing the source-object attribution performance using different detector array configurations and binomial down-sampled data. The alarm encounters correspond with Figure 4.3 for LiDAR. In total 20 alarm encounters were performed. The annotation in each square represents the number of trials the source carrier had the lowest exclusion metric. The x-axis shows the percentage of the highest count rate within a particular alarm encounter the spectra data was binomial down-sampled, and the y-axis represents the different detector array configurations generated by combining the radiation data from the individual detectors of the 6 NaI(Tl) detector array.

In the above analysis, only LiDAR-based trajectories were used to explore the importance of the angular response in source localization when tracking information is available. The inherent depth information from LiDAR enables more reliant tracking and consistent object position. With video, the source position varies greatly in these source encounters leading to incorrect associations between detections and tracks. This limits the ability to properly investigate the importance of the angular response in source localization when tracking information is available with video data. However, with a more reliable depth estimation, the findings should be similar to LiDAR-based trajectories.

4.6.1.2 Complex Scenario in a Mock Urban Environment

This topic was also extended to the complex scenario shown in Figure 4.8 and Figure 4.9. Source-object attribution was performed using either the full NaI(Tl) bar array, the left-right detector, the 3 detector NaI(Tl) array, or the monolithic detector and the attribution performance between these different configurations was compared. In addition, binomial down-sampling was performed to investigate attribution performance in weaker source interactions. The spectra data was down-sampled by a factor of 25%, 75%, or 90% of the highest count-rate within a given alarm encounter. The metric for attribution performance was the number of trials the source carrier was assigned the lowest exclusion.

The results shown in Figure 4.39a and Figure 4.39b compare the attribution performance of the different array configurations across different source strengths. The findings in Figure 4.39b suggest that a simpler array configuration could be used when tracking information is available to achieve similar localization performance to a more complex array design. The decreased localization performance at 75% and 90% down-sampled is driven by a reduced ability reject trajectories that are more than 95% outside of the alarm window. At 75% or 90% down-sampled, the alarm windows are about 0.75 seconds or less in duration. This results in trajectories that would otherwise be excluded from the analysis being included, which diminishes localization performance, but a clear correlation still existed and an operator would be able to make the correct attribution in real time.

In Figure 4.39b, the full 6 NaI(Tl) detector array offers the most consistent localization performance across the different down-sample percentages. However, the simpler array configurations – 3 NaI(Tl) detector array, left-right detector array, and monolithic detector – have similar localization capabilities. In this case, there is decreased localization performance at 75% down-sampled but only for the 3 NaI(Tl) and left-right detector array. Given the probabilistic nature of binomial down-sampling and the further reduction in counts at 90%, most of the trajectories that could not be excluded at 75% down-sampling were excluded because the trajectories were better described by a background-only model. This improved localization performance at 90% for both the 3 NaI(Tl) and left-right detector array. The alarm encounters considered here do not capture the degeneracies that are present in LEMURS, which explains the localization performance of the monolithic detector; however, the results for the monolithic detector demonstrate the ability of the source-object attribution analysis to improve the localization performance of a nondirectional detector.

Comparing Figure 4.39a and Figure 4.39b, LiDAR has improved localization performance in weaker source strength scenarios. This demonstrates the advantage of the inherent depth information with LiDAR. This allows for more consistent tracking, which enables more effective attribution performance.

Overall, the findings from this analysis for both video and LiDAR suggest that a simpler array configuration could be used when tracking information is available.

Another aspect in this analysis is the relative speed between LEMURS and other objects

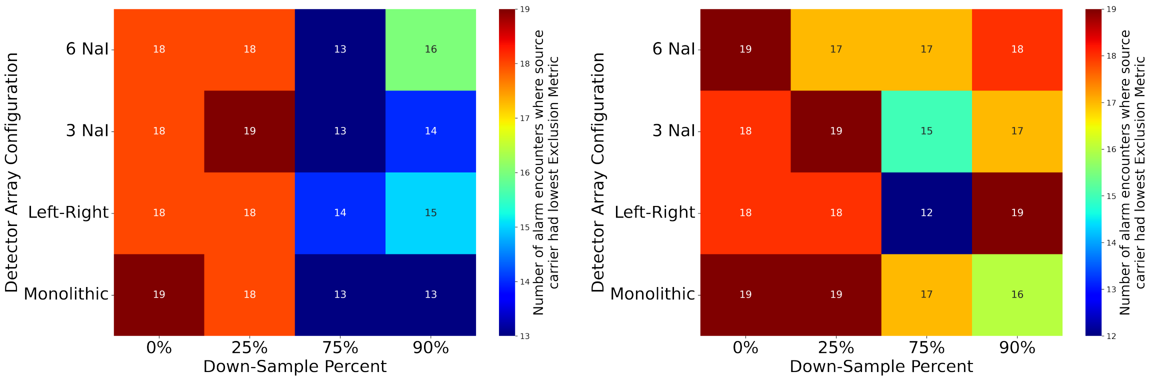

(a) (b)

Figure 4.39: Comparing the source-object attribution performance using different detector array configurations and binomial down-sampled data. The alarm encounters correspond with Figure 4.11 for both video and LiDAR. In total 19 alarm encounters were performed. The annotation in each square represents the number of trials the source carrier had the lowest exclusion metric. The x-axis shows the percentage of the highest count rate within a particular alarm encounter the spectra data was binomial down-sampled, and the y-axis represents the different detector array configurations generated by combining the radiation data from the individual detectors of the 6 NaI(Tl) detector array.

in the scene. As seen in Section 4.3.2, the track length of objects was reduced with a LEMURS speed of 20 mph compared to the results shown in Section 4.3.1. It is important to understand how a decreased track length could potentially impact localization performance across different array configurations and source strengths.

In Figure 4.40, the localization performance using different array configurations is shown for both video (Figure 4.40a) and LiDAR (Figure 4.40b). The results in Figure 4.40a show similar localization performance across the full NaI(Tl), 3 NaI(Tl), and left-right detector array configurations for the different source strength encounters, and a reduction in localization performance for the monolithic detector compared to the other three configurations. In Figure 4.40b, the results are similar; however, the monolithic localization performance is comparable to the other detector array configurations. This analysis further suggests that when tracking information is available a simpler detector array configuration (i.e., left-right or 3 NaI(Tl) detector array) offers similar localization performance when compared to a more complex array design (i.e. 6 NaI(Tl) bars arranged in a hexagonal arrangement).

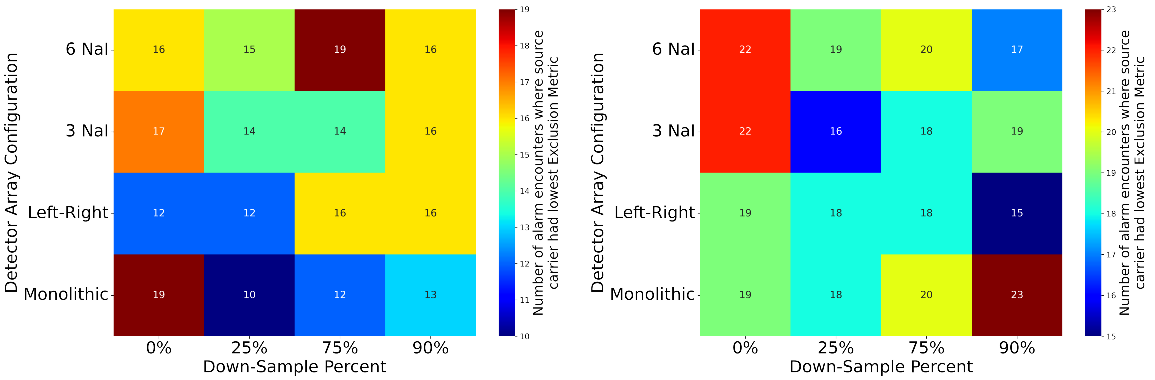

(a) (b)

Figure 4.40: Comparing the source-object attribution performance using different detector array configurations and binomial down-sampled data. The alarm encounters correspond with Figure 4.17 for both (a) and (b). In total 23 alarm encounters were performed. The annotation in each square represents the number of trials the source carrier had the lowest exclusion metric. The x-axis shows the percentage of the highest count rate within a particular alarm encounter the spectra data was binomial down-sampled, and the y-axis represents the different detector array configurations generated by combining the radiation data from the individual detectors of the 6 NaI(Tl) detector array.

4.7 Comparison of Source Localization Performance

With source-object attribution, the source is inherently localized to a particular object. The purpose of this section is to explore how this might improve localization performance compared to using different reconstruction algorithms – ML-EM or Gridded Point-Source Likelihood (GPSL). Unlike APSL described in Section 2.9, GPSL is a discrete Point-Source Likelihood (PSL) approach that only looks for a single radiological source. The discrete PSL approach is described in [23, 22]. Localization performance using different array configurations and binomial down-sampling will also be explored. Additionally, the reconstructions were performed using the photopeak ROI (617 keV - 707 keV) for ^{137}Cs.

For the analysis, the degenerate case described in Section 4.2.1 was used. The static scenario provided an exact true source location for every trial; whereas, with the mobile scenarios, the source location relative to LEMURS varied from trial-to-trial. The LiDAR point cloud was discretized into cubic 60 cm voxels, and the localization of the highest source weight voxel for the ML-EM algorithm or lowest z-score voxel for the GPSL algorithm was used as a metric to compare the localization performance of the different reconstruction algorithms to the source-object attribution analysis. The z-score in GPSL represents the confidence interval around the minimum loss voxel. If the localization of the voxel was located within an area defined around the source carrier, then the reconstruction algorithm correctly localized the source to the source carrier. In the cases where the voxel was located outside the defined area, a euclidean distance was used to determine the distance of the voxel from the ground truth voxel.

Figure 4.41a and Figure 4.41b show two different projections of the image reconstruction using GPSL. The defined area around the source carrier is marked in red. In this case, all 6 NaI(Tl) detectors were used in the reconstruction and the spectra data was not down-sampled. The image reconstruction correctly localizes the source intensity to the source carrier with the lowest z-score being both contained within the defined area of the source carrier and near the true source location. With ML-EM, the image reconstruction is more dispersed and the reconstructed source intensity is biased towards the vehicle (Figure 4.41c). Also, the reconstruction is more dispersed in the Z-axis (Figure 4.41d) compared to Figure 4.41b. ML-EM is unconstrained unlike GPSL and the algorithm assigns source intensity to voxels in order to explain the radiological data in the detectors. This causes the biasing towards the LEMURS vehicle. In this case, the highest source intensity voxel is located 1.95 m from the true source location.

The above analysis was extended to all alarm encounters in the degenerate scenario. The localization performance of the GPSL reconstructions for all the alarm encounters using varying levels of detector array information and binomial down-sampling amounts is shown in Figure 4.42. Results for the ML-EM reconstructions are not shown because the highest intensity voxel was not placed within the defined area for any of the alarm encounters. Comparing these results with the source-object attribution findings from Figure 4.38, GPSL provides similar localization performance. Across the different detector array configurations and source strengths, GPSL was able to correctly localize the source to the source carrier in a majority of the alarm encounters even with the monolithic detector. However, the localization performance of source-object attribution is significantly improved compared to ML-EM. With ML-EM, the source intensity was biased towards the vehicle. In the non-binomial down-sampled encounters, the highest source intensity voxel was on average about 2.8 m from the source carrier. In these cases, an operator would be able to determine the object responsible for the radiological alarm in real-time because the highest intensity voxel is near the source carrier. However, with the weaker source encounters, in particular 75% and 90% down-sampling, the average distance of the highest voxel from the true source location was 22.6 m, and no alarm adjudication could be performed in these cases. These scenarios have the most real world relevance since radiological search involves finding weakly-emitting sources.

The findings of this analysis demonstrate that source-object attribution does improve localization performance compared to using the ML-EM reconstruction algorithm and reconstructing to occupied voxels in the scene. In addition, if the level of localization offered by source-object attribution is sufficient (i.e., localization to an entire object), then source-object attribution has similar localization performance to GPSL. However, it should be noted explicitly, when analyzing source localization to particular voxels, GPSL enables superior localization compared to source-object attribution.

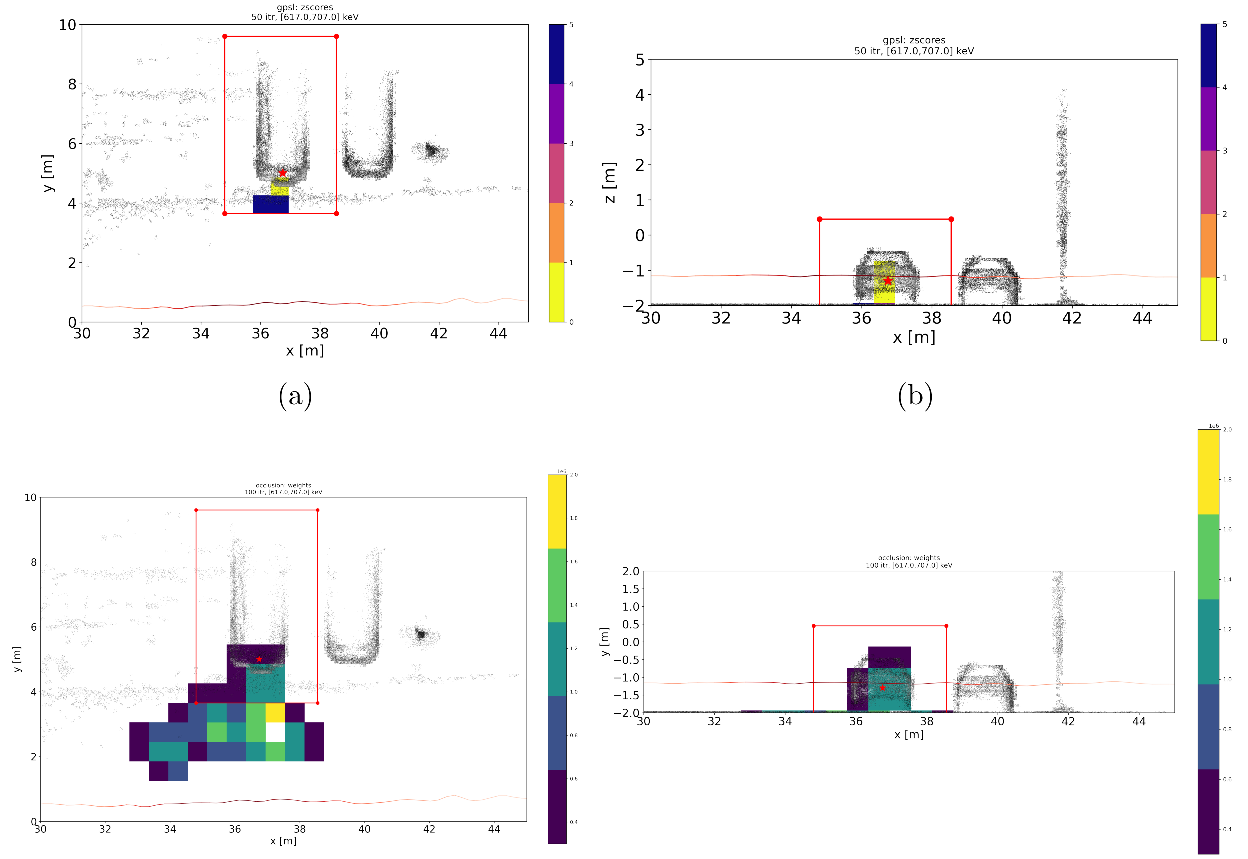

(a)

(b)

(c)

(d)

Figure 4.41: Image reconstruction using Gridded Point-Source Likelihood (a-b) or Maximum Likelihood Estimation Maximization (c-d) as LEMURS drives past a static ^{137}Cs source. The reconstructions were performed using all 6 NaI(Tl) detectors independently. In (a), a x-y projection of the image reconstruction is shown. The z-score represents the confidence interval around the minimum loss voxel. The $*$ indicates the true location of the source. The scene is discretized into 60 cm cubic voxels, and the red rectangle is the defined area around the source carrier. In (b), a x-z projection of the same scene as (a) is shown. In (c) [(d)], a x-y [x-z] projection using a 100 iterations of ML-EM is shown. The voxels are in units of source intensity.

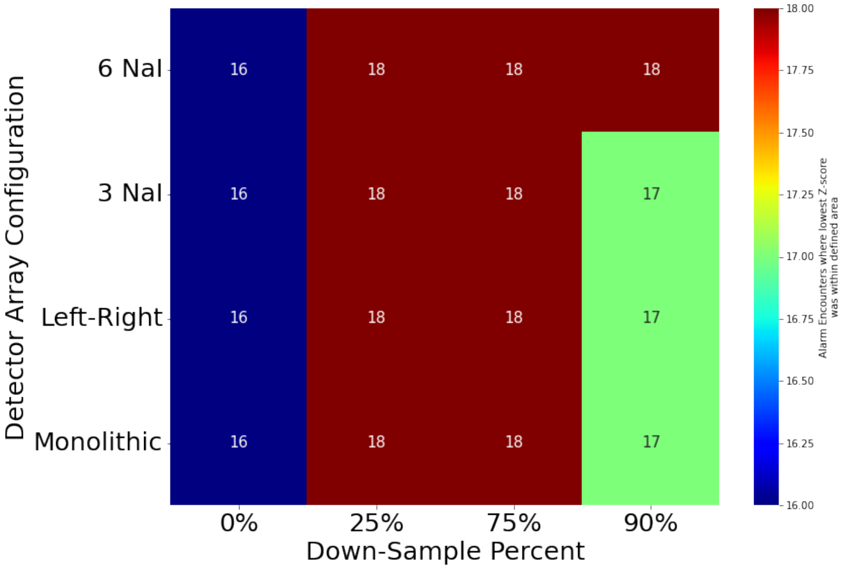

Figure 4.42: Results from all 20 alarm encounters using GPSL to perform reconstructions. The annotation in each square represents the number of trials the lowest Z-score was in the predefined area around the object responsible for the radiological alarm. The x-axis shows the percentage of the highest count rate within a particular alarm encounter the spectra data was binomial down-sampled, and the y-axis represents the different response functions used to perform the reconstructions.

4.8 Conclusions & Future Directions

On mobile platform systems, alarm encounters are transient and localization has to be performed quickly and efficiently. Source-object attribution enables new capabilities on a mobile platform by providing automatic associations between objects in the scene and the radiological data.

This work demonstrated that situational awareness could be improved in a mock urban environment in the presence of dynamic sources using SLAM to generate pose estimates in a world-fixed frame and video or LiDAR-based trajectories. The findings show that video or LiDAR offer similar tracking performance which enabled effective rejection of tracks that were inconsistent with the radiological data. This performance was seen for various LEMURS and source carrier vehicle speeds. In addition, situational awareness was improved using INS data along with video or LiDAR-based trajectories except for video in the case of increased LEMURS and vehicle speeds. The findings from the INS study demonstrated that if satellite interference is low using INS pose estimates to perform tracking and attribution offers similar capabilities to using SLAM with LiDAR-based trajectories. Furthermore, when source attenuation is present in a scene, situational awareness can be improved with video and LiDAR. This was demonstrated by tracking a dynamic pedestrian carrying a radiological source through temporary occlusions caused by mobile vehicles in a mock urban environment.

Additionally, the findings from this work suggests that detection sensitivity can be improved using track-informed integration windows compared to fixed integration windows. The track-informed integration windows from video and LiDAR trajectories were able to adapt to the dynamics of the scene and improved the anomaly value, proxy for detection sensitivity, in two different transient alarm encounter scenarios. Furthermore, an improvement in the anomaly value was seen in low SNR environments using track-informed integration windows.

Furthermore, the results showed that when tracking information is available a detector array with less angular resolution could be used to achieve similar localization performance as a detector array with more angular resolution. These findings have the potential to influence future detector system development where typically the detector array is configured without considering the contextual information. The findings also demonstrated that using a high fidelity detector response model enables more effective source-object attribution particularly in scenarios where degeneracies are present.

Finally, the results demonstrated that localizing the radiological source to objects in the scene improves localization performance compared to using a ML-EM reconstruction algorithm.

Future work is needed to fully explore using an INS in urban environments to produce pose estimates in a world-fixed frame and the effect this would have on tracking and attribution. While the pose estimates generated using the INS did not impact attribution performance using LiDAR, it is thought in a true urban environment with buildings there will be potentially degraded tracking and attribution performance. Urban environments are full of buildings, which will interfere with satellite signals more than the trees in the mock urban environment considered in this work. When surveying Broadway Street in Downtown Oakland between 8th Street and 14th Street – an area with multiple buildings that are higher than 5 floors – with LEMURS, VDOP and HDOP values that ranged from 5–50, respectively, were present, and during a 5.5 second interval within this survey, where LEMURS averaged 8 ± 1 mph, the INS had mean VDOP and HDOP values of 28 ± 17 and 41 ± 31, respectively. The prolonged satellite interference in this type of urban environment, where many building are present, should reduce the accuracy of the pose estimate, which could reduce tracking performance. Additionally, since the angular pose information will be degraded, performing attribution using all 6 detector independently will potentially have worse results than the findings in Section 4.3.1 and Section 4.3.2.

While situational awareness was improved by accounting for attenuation in the physics model, the S-values were quite large. In the current implementation, attenuation through an object in the scene is only calculated when an object is fully occluded. This work neglects how the fraction of the solid angle covered by an occluding object between the detector and radiological source changes with time, which influences the expected counts in the detector. While this is a reasonable assumption for the occlusion that occurs between two pedestrians, for objects with larger extents, such as vehicles or trucks, this assumption is not valid. To better account for how the expected time-dependent count-rate in the physics model changes as two objects pass by each other, how the fraction of the solid angle between the detector and radiological source changes as a function of time needs to be determined. This will require more complex experimental data and simulations to better understand both the attenuation effect of objects in the scene as well as the attenuation effect of the environment.

Also, attenuation from untracked objects is not accounted for in the model generation for a track. The object detection CNNs for video and LiDAR provide identification labels for a myriad of objects in a sensors FOV. This information could be utilized to determine source attenuation from untracked objects in the scene. This should further improve attribution performance but needs to be characterized in more detail.

Chapter 5

Conclusion

5.1 Summary

A key challenge in nuclear security is the detection and localization of radiological and/or nuclear material in urban environments, which is difficult due to the highly cluttered and dynamic nature of these environments. In urban environments, often weakly emitting gamma-ray sources need to be distinguished from a constantly fluctuating gamma-ray background. Once a radiological and/or nuclear alarm is triggered, operators monitoring systems in real-time need to be able to correctly attribute the detected source to the object responsible for the alarm quickly and efficiently. However, this becomes more difficult as both the number and speed of objects in a scene increases. By using contextual information, such as streaming video or LiDAR, object detection and tracking can augment the radiological data from detectors to improve the situational awareness of the scene and provide inherent localization. The aim of this work was to fuse contextual information (streaming video and LiDAR) with radiological data to provide automatic associations between physical objects in a scene and the radiological data to improve both situational awareness and detection sensitivity.

The results in Section 3 demonstrated a static system that is capable of performing object detection, tracking, and source-object attribution in real-time using either LiDAR or visual cameras. Using this system to perform source-object attribution, source localization could be performed quickly and efficiently in a challenging mock environment by creating associations between objects in a scene and the radiological data, and improved situational awareness was demonstrated using both video and LiDAR. Source-object attribution enabled a non-directional NaI(Tl) detector to provide localization information. Furthermore, by applying a track-informed integration window, the findings suggested that detection sensitivity can be improved compared to different fixed integration windows.

The work in Section 3 was extended in Section 4 to a mobile platform system. Source-object attribution capabilities were tested and demonstrated in the presence of static and mobile sources. In a mock urban environment, source-object attribution was successfully performed using video and LiDAR across different SNR scenarios. The findings from this work demonstrated improved situational awareness in transient alarm encounters using both video and LiDAR. Moreover, this work showed that contextual information enables a new paradigm for source localization and has the potential to influence detector array design when tracking information is available. Using the inherent localization provided by source-object attribution, detector array configurations with less angular resolution had similar localization performance compared to a hexagonal arrangement of NaI(Tl) detectors. Additionally, source-object attribution enabled better source localization across different detector array configurations with varying levels of complexity compared to a traditional ML-EM reconstruction algorithm. The work further suggested that track-informed optimal integration windows can improve detection sensitivity relative to fixed integration windows by better adapting to the dynamics of the scene and relative motion of objects in a scene. Finally, spectroscopic analysis was performed using a summed response of the detector array and an optimal configuration of detectors within the detector array. Improved detection sensitivity using the optimal configuration of detectors was shown compared to the summed response demonstrating an optimal configuration of detectors does exist that maximizes SNR.

5.2 Future Outlooks

This work demonstrated the capabilities of fusing contextual information with radiological data to perform real-time source-object attribution on static and mobile platform systems. Even though source-object attribution was correctly performed in this dissertation, the S-values were often quite large and did not reflect realistic probabilities, particularly for situations, where, when inspected by eye, a clear association between tracks and radiological data was obvious. Currently, the attribution analysis does not account for uncertainties related to trajectories. Tracks are considered to be fixed and assumed to be the source of the radiation signal resulting in large S-values. Furthermore, large S-values are often caused by an inadequate physics model. While some uncertainty, for example those originating from tracking, could be incorporated into an improved attribution metric, many other uncertainties are of systematic nature and are not known. In these cases, the S-value, providing a statistical assessment of the situation, might not be the optimal metric and another description might provide an answer closer to the system operator's needs.

While this research focused on source-object attribution on static or vehicle-based systems, the methodology presented here is platform-agnostic and could be applied to any type of contextual sensor system. Future work could explore applying source-object attribution using unmanned aerial vehicles. This would enable exploration of less constrained environments. One could imagine surveying a crowded parking lot to quickly identify any clandestine material present.

The work presented here focused on source localization of gamma-ray emitting radiological sources and did not explore neutron sources. Localization of spontaneous fission and (α, n) neutron sources through thermal capture with free-moving detector systems has been demonstrated [61]. Given that work, applying contextual-radiological data fusion with a neutron-sensitive detector for neutron source attribution has potential. However, it is highly dependent on the neutron source geometry, the scatter media, and sensitivity of the detector. Exploring these topics in greater detail should enable neutron localization using source-object attribution in future work.

Additionally, this work only explored point sources. In the case of a distributed source, if the gamma-rays are fairly localized to the object, then the span of the captured gamma-rays will be concentrated to the person/vehicle and nearby surface. It would be expected that the source-object attribution analysis would be able to localize the object responsible for the radiological alarm. Future research should consider how to handle distributed sources and the impact on source-object attribution performance.

The analyses presented in this dissertation suggest that object detection and tracking with LiDAR is currently more robust than a video-based approach. This is mostly driven by the direct availability of distance information. There are depth extractor CNNs that produce depth maps for monocular cameras. They should provide better distance information compared to the current implementation but are currently too GPU intensive to run in parallel with object detection on a low-power system. Given the rate of improvement of these depth-estimating algorithms, it is expected that a depth estimating CNN could be run along with an object detection, tracking, and source-object attribution analysis in the near future. Another approach to improving distance estimation for video data is to use a stereoscopic camera, which provides a depth map. An additional method to improving the distance estimate is fusing the information from LiDAR and video and using the inherent depth from LiDAR for video.

Furthermore, by using additional object appearance descriptions (features) derived from imagery, more reliable tracking could be achieved. Notably the issue of track switching could be resolved. An implementation of such a method that can already run on a low power system is described in [62]. This would also improve alarm adjudication for an operator because the features from imagery could be provided along with the track ID (i.e., a red truck could be more quickly identified compared to just track ID's).

Finally, SECOND/PointPillar's object detector doesn't work reliably for people standing near walls or large objects so the object may go undetected for LiDAR. It is believed this is because spurious points interfere with its anchoring algorithm. This topic was not investigated in this dissertation; however, a training set including more data with pedestrians nearby objects could potentially increase the networks sensitivity.

Chapter 6

Bibliography

[1] Martin J Berger and JH Hubbell. Xcom: Photon cross sections on a personal computer. Technical report, National Bureau of Standards, Washington, DC (USA). Center for Radiation , 1987.

[2] Roger Labbe. Kalman and bayesian filters in python. *Chap*, 7:246, 2014.

[3] Robert C Runkle, L Eric Smith, and Anthony J Peurrung. The photon haystack and emerging radiation detection technology. *Journal of Applied Physics*, 106(4):7, 2009.

[4] Charles G Bathke, Bartley B Ebbinghaus, Brian A Collins, Brad W Sleaford, Kevin R Hase, Martin Robel, Richard K Wallace, Keith S Bradley, John R Ireland, Gordon D Jarvinen, et al. The attractiveness of materials in advanced nuclear fuel cycles for various proliferation and theft scenarios. *Nuclear Technology*, 179(1):5–30, 2012.

[5] Daniel J Barnett, Cindy L Parker, David W Blodgett, Rachel K Wierzba, and Jonathan M Links. Understanding radiologic and nuclear terrorism as public health threats: preparedness and response perspectives. *Journal of Nuclear Medicine*, 47(10):1653, 2006.

[6] Charles G Bathke, Richard K Wallace, John R Ireland, MW Johnson, Kevin R Hase, Gordon D Jarvinen, Bartley B Ebbinghaus, Brad W Sleaford, Brian A Collins, Martin Robel, et al. An assessment of the attractiveness of material associated with a mox fuel cycle from a safeguards perspective. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2009.

[7] Roger C Byrd, Joel M Moss, William C Priedhorsky, Carolyn A Pura, Gary W Richter, Kevin J Saeger, W Robert Scarlett, Sara C Scott, and RL Wagner. Nuclear detection to prevent or defeat clandestine nuclear attack. *IEEE Sensors Journal*, 5(4):593–609, 2005.

[8] Robert J Nemzek, Jared S Dreicer, David C Torney, and Tony T Warnock. Distributed sensor networks for detection of mobile radioactive sources. *IEEE Transactions on Nuclear Science*, 51(4):1693–1700, 2004.

[9] Sean M Brennan, Angela M Mielke, David C Torney, and Arthur B Maccabe. Radiation detection with distributed sensor networks. *Computer*, 37(8):57–59, 2004.

[10] Angela M Mielke, Sean M Brennan, Mark C Smith, David C Torney, Arthur B Maccabe, and JF KarlinM. Independent sensor networks. *IEEE instrumentation & measurement magazine*, 8(2):33–37, 2005.

[11] RD Penny, WE Hood, RM Polichar, FH Cardone, LG Chavez, SG Grubbs, BP Huntley, RA Kuharski, RT Shyffer, Lorenzo Fabris, et al. A dual-sided coded-aperture radiation detection system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 652(1):578–581, 2011.

[12] Eric A Wulf, Bernard F Phlips, W Neil Johnson, Byron Leas, and Lee J Mitchell. Misti imaging and source localization. In *2008 IEEE Nuclear Science Symposium Conference Record*, pages 2413–2417. IEEE, 2008.

[13] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv*, 2018.

[14] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv e-prints*, page arXiv:2004.10934, April 2020.

[15] Y. Yan, Y. Mao, and B. Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10), 2018.

[16] Thomas P Karnowski, Mark F Cunningham, James S Goddard, Anil M Cheriyadat, Donald E Hornback, Lorenzo Fabris, Ryan A Kerekes, Klaus-Peter Ziock, and Timothy F Gee. Motion estimation accuracy for visible-light/gamma-ray imaging fusion for portable portal monitoring. In *Image Processing: Machine Vision Applications III*, volume 7538, page 75380F. International Society for Optics and Photonics, 2010.

[17] Klaus P Ziock, Jeff Collins, Mark Cunningham, Lorenzo Fabris, Timothy Gee, James Goddard, Frezghi Habte, and Thomas Karnowski. The use of gamma-ray imaging to improve portal monitor performance. *IEEE Transactions on Nuclear Science*, 55(6):3654–3664, 2008.

[18] Kristofer Henderson, Xiaomeng Liu, Kelsey Stadnikia, Allan Martin, Andreas Enqvist, and Sanjeev J Koppal. Proximity-based sensor fusion of depth cameras and isotropic rad-detectors. *IEEE Transactions on Nuclear Science*, 2020.

[19] Christopher G Wahl, Willy R Kaye, Weiyi Wang, Feng Zhang, Jason M Jaworski, Alexis King, Y Andy Boucher, and Zhong He. The polaris-h imaging spectrometer. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 784:377–381, 2015.

[20] Andrew Haefner, Ross Barnowski, Paul Luke, Mark Amman, and Kai Vetter. Handheld real-time volumetric 3-d gamma-ray imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 857:42–49, 2017.

[21] Daniel Hellfeld, Paul Barton, Donald Gunter, Andrew Haefner, Lucian Mihailescu, and Kai Vetter. Real-time free-moving active coded mask 3d gamma-ray imaging. *IEEE Transactions on Nuclear Science*, 66(10):2252–2260, 2019.

[22] Daniel Hellfeld, Tenzing HY Joshi, Mark S Bandstra, Reynold J Cooper, Brian J Quiter, and Kai Vetter. Gamma-ray point-source localization and sparse image reconstruction using poisson likelihood. *IEEE Transactions on Nuclear Science*, 66(9):2088–2099, 2019.

[23] Jayson R Vavrek, Daniel Hellfeld, Mark S Bandstra, Victor Negut, Kathryn Meehan, William Joe Vanderlip, Joshua W Cates, Ryan Pavlovsky, Brian J Quiter, Reynold J Cooper, et al. Reconstructing the position and intensity of multiple gamma-ray point sources with a sparse parametric algorithm. *IEEE Transactions on Nuclear Science*, 67(11):2421–2430, 2020.

[24] J.C Curtis, M Salathe, T Joshi, and K Bilton. Contextual sensing for mobile urban radiation search. *Abstract IEEE NSS-MIC*, Manchester, UK 2019.

[25] M. R. Marshall, D. Hellfeld, T. H. Y. Joshi, M. Salathe, M. S. Bandstra, K. J. Bilton, R. J. Cooper, J. C. Curtis, V. Negut, A. J. Shurley, and K. Vetter. 3-d object tracking in panoramic video and lidar for radiological source-object attribution and improved source detection. *IEEE Transactions on Nuclear Science*, 68(2):189–202, 2021.

[26] Glenn F Knoll. *Radiation detection and measurement.* John Wiley & Sons, 2010.

[27] Glenn F Knoll. Radiation detection and measurement. pages 779–786. John Wiley & Sons.

[28] Mark S Bandstra, Timothy J Aucott, Erik Brubaker, Daniel H Chivers, Reynold J Cooper, Joseph C Curtis, John R Davis, Tenzing H Joshi, John Kua, Ross Meyer, et al. Radmap: The radiological multi-sensor analysis platform. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 840:59–68, 2016.

[29] Daniel Hellfeld, Paul Barton, Donald Gunter, Lucian Mihailescu, and Kai Vetter. Optimization of a spherical active coded mask gamma-ray imager. In *2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop (NSS/MIC/RTSD)*, pages 1–3. IEEE, 2016.

[30] Joseph C Curtis, Reynold J Cooper, Tenzing H Joshi, Bogdan Cosofret, Thomas Schmit, John Wright, Jonathan Rameau, Daisei Konno, Daniel Brown, Forrest Otsuka, et al. Simulation and validation of the mobile urban radiation search (murs) gamma-ray detector response. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 2018.

[31] KJ Bilton, TH Joshi, MS Bandstra, JC Curtis, BJ Quiter, RJ Cooper, and K Vetter. Non-negative matrix factorization of gamma-ray spectra for background modeling, detection, and source identification. *IEEE Transactions on Nuclear Science*, 2019.

[32] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.

[33] Marko Bjelonic. YOLO ROS: Real-time object detection for ROS. https://github.com/leggedrobotics/darknet_ros, 2016–2018.

[34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

[35] Cheryl D Fryar, Qiuping Gu, Cynthia L Ogden, and Katherine M Flegal. Anthropometric reference data for children and adults; united states, 2011-2014. 2016.

[36] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom. Pointpillars: Fast encoders for object detection from point clouds. *CVPR*, 2019.

[37] H. Caesar et al. nuscenes: A multimodal dataset for autonomous driving. *arXiv:1903.11027v4*, 2019.

[38] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar slam. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278. IEEE, 2016.

[39] Advanced navigation. https://www.advancednavigation.com/products/spatial. (Accessed on 06/10/2021).

[40] Yaakov Bar-Shalom, X Rong Li, and Thiagalingam Kirubarajan. *Estimation with applications to tracking and navigation: theory algorithms and software*, chapter 5. John Wiley & Sons, 2004.

[41] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 3464–3468. IEEE, 2016.

[42] Leandro Pardo. *Statistical inference based on divergence measures.* Chapman and Hall/CRC, 2018.

[43] High-resolution os1 lidar sensor: robotics, trucking, mapping — ouster. https://ouster.com/products/os1-lidar-sensor/. (Accessed on 06/17/2020).

[44] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

[45] Linda D Meyers, Jennifer Pitzi Hellwig, Jennifer J Otten, et al. *Dietary reference intakes: the essential guide to nutrient requirements.* National Academies Press, 2006.

[46] Andreas Zoglauer, R Andritschke, and F Schopper. Megalib–the medium energy gamma-ray astronomy library. *New Astronomy Reviews*, 50(7-8):629–632, 2006.

[47] Sea Agostinelli, John Allison, K al Amako, John Apostolakis, H Araujo, P Arce, M Asai, D Axen, S Banerjee, G 2 Barrand, et al. Geant4a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[48] Krzysztof M Gorski, Eric Hivon, Anthony J Banday, Benjamin D Wandelt, Frode K Hansen, Mstvos Reinecke, and Matthia Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005.

[49] Dave Shreiner, Graham Sellers, John Kessenich, and Bill Licea-Kane. *OpenGL programming guide: The Official guide to learning OpenGL, version 4.3.* Addison-Wesley, 2013.

[50] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122, 1982.

[51] Steve Baker and Robert D. Cousins. Clarification of the use of chi-square and likelihood functions in fits to histograms. *Nuclear Instruments and Methods in Physics Research*, 221(2):437–442, April 1984.

[52] Sander Greenland. Valid p-values behave exactly as they should: Some misleading criticisms of p-values and their resolution with s-values. *The American Statistician*, 73(sup1):106–114, 2019.

[53] Gideon Schwarz et al. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.

[54] Jonathan Goodman and Jonathan Weare. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80, 2010.

[55] Daniel Foreman-Mackey, David W Hogg, Dustin Lang, and Jonathan Goodman. emcee: the mcmc hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306, 2013.

[56] e-con Systems. https://www.e-consystems.com/nvidia-cameras/jetson-agx-xavier-cameras/four-synchronized-4k-cameras.asp. (Accessed on 02/04/2020).

[57] Jetson agx xavier developer kit — nvidia developer. https://developer.nvidia.com/embedded/jetson-agx-xavier-developer-kit. (Accessed on 06/17/2020).

[58] digibase 14-pin pmt base with integrated bias supply, preamplifier, and mca with digital signal processing. https://www.ortec-online.com/products/electronics/photomultiplier-tube-bases/digibase. (Accessed on 08/12/2020).

[59] Stanford Artificial Intelligence Laboratory et al. Robotic operating system.

[60] David Mills, Jim Martin, Jack Burbank, and William Kasch. Network time protocol version 4: Protocol and algorithms specification. 2010.

[61] R Pavlovsky, JW Cates, WJ Vanderlip, THY Joshi, A Haefner, E Suzuki, R Barnowski, V Negut, A Moran, K Vetter, et al. 3d gamma-ray and neutron mapping in real-time with the localization and mapping platform from unmanned aerial systems and man-portable configurations. *arXiv preprint arXiv:1908.06114*, 2019.

[62] N. Wojke, A. Bewley, and D. Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649, 2017.