Tracking Radioactive Sources through Sensor Fusion of Omnidirectional LIDAR and Isotropic Rad-detectors

Kristofer Henderson*, Kelsey Stadnikia*, Allan Martin*, Andreas Enqvist*, Sanjeev Koppal*
*University of Florida
khenderson5@ufl.edu

Abstract

Tracking radioactive sources in large spaces has applications for homeland security, airport and port surveillance as well as military and security uses. Unfortunately, source localizing radiological detectors are extremely expensive, and those with low prices are isotropic - i.e. they integrate radiation from a sphere of directions centered at the sensor. In this paper, we show that omnidirectional depth sensors and isotropic radiological detectors have complementary strengths and can enable many applications. We model the source strength of radiological sources and integrate these with LIDAR measurements and a Kalman filter tracker. This enables applications such as tracking behind walls and detecting multiple radiological sources in the same scene.

1. Introduction

Nuclear material trafficking has the potential to cause serious harm to our national security. The volume of special nuclear material (SNM) needed to create a nuclear weapon is less than 1 liter, an amount small enough to easily fit in a suitcase or on a person. Nuclear trafficking threats are likely to occur at ports of entry. Radiological material is dangerous not only because of the risk of nuclear attack via radiation dispersal device (RDD) or nuclear weapon, but also because people near the trafficked material can be exposed to high doses of nuclear radiation [6]. The problem posed by nuclear material trafficking is not contained to a few isolated incidents, there have been 2500 reported incidents of nuclear material trafficking in the IAEA incident and trafficking database [1]. The current method of dealing with nuclear trafficking scenarios is to shut down large areas of the airport, potentially including all flight activity. This method is both impractical and expensive.

The primary tool to address trafficking are radiation detectors that measure nuclear radiation such as neutrons, gamma ray photons, etc. The number of radiation detector events increases with source strength, and decreases with the inverse square of the source-detector distance. Inexpensive versions of these detectors ($1,000-$5,000) are problematic, since they lack the ability to differentiate between multiple sources, and cannot measure the direction from the detector to the source.

Figure 1: The senor setup used in this paper, fusing a LIDAR with a radiation detector.

To overcome the limitations of the inexpensive detectors, radiological researchers have designed and built systems that fuse data across multiple sensors and include techniques such as satellite imagery and laser mapping [19, 24, 23, 33, 3, 32]. The main technique has been to increase angular resolution by using an array of isotropic radiation detectors with a coded aperture (made of lead slabs). However, this increases the cost of the system dramatically $\approx$ $100,000 and requires additional calibration. Both of these reasons prohibit these systems from being deployed to monitor large areas. On the other hand, reducing the number of sensor nodes in these systems increases the chances of occlusion, especially in crowded and visually cluttered environments such as airports and ports.

Recently, 3D sensor researchers [26] have proposed systems with a single, isotropic radiological sensor and commercial time-of-flight sensors, such as the Microsoft Kinect, which is on the order of \$100s. The approach is to have many of these devices around the airport or port, track all moving objects and computationally discern the location of the sources by fusing measurements from the radiological sensor. While some progress has been made, the proposed calibration in [26] has an error of $100cm$ (1 meter) in the relative position of the vision sensor and the radiological sensor, reducing its applicability to slow moving scenes.

In this paper, we propose a simple strategy to replace a large number of radiation detectors in the array with an *omnidirectional* depth sensor, as shown in Fig. 1. Our key idea is to combine the advantages of each sensor type. The radiological sensors, while isotropic, has the advantage of potentially tracking sources through visual obstacles such as walls. The vision sensor, while requiring line-of-sight, has a $360°$ field of view, which solves the calibration problem efficiently. This eliminates the error caused by calibration stated in [26] to be the major source for error in vision and radiation data fusion systems. This will allow the use of far fewer radiation detectors in the system, reducing the cost of the system while keeping the angular resolution of the array approach. We compromise slightly on the cost of the vision sensor (in the \$10000s, with similar sensors quickly dropping in cost), while noting that this is far lower than competing radiological techniques. Our contributions are the following demonstrations (please see the accompanying video for a full summary),

- We show that a LIDAR paired with a single radiological detector can easily identify a single moving source in a cluttered and fast moving environment, without knowing the source strength. We also demonstrate, for the first time, a simple and robust estimation that can detect and track multiple radiological sources, each of unknown strength.

- Building on our technique to detect, track and estimate source strength, we show occlusion-resistant applications by combining the vision and radiological measurements in the Kalman filter framework.

- The above first two contributions equal the functionality of arrays of radiological sensors, with only a single radiological sensor and a 3D vision sensor. By adding a second radiation detector to our system, we demonstrate a new capability of tracking a radiation source through occlusions.

1.1. Related Work

3D Vision and Radiological Sensing: Most efforts [19, 24, 23] at LIDAR and rad-detector fusion focus on rigidly constructed gantries in static scenes. Other efforts use coded apertures (where the random pattern is made of lead squares) to encode directionality [33] in the isotropic sensor, to enable stereo reconstruction [32] or to reduce noise [2]. Inferring material properties, and not just geometry, from the visual measurements allows for estimating background radiation and compensating for its effects [3]. We show how to replicate these types of capabilities, at a small fraction of the cost of the coded aperture systems, by inferring radioactive source strength and radial distance together, as byproducts of our fusion approach.

Sensor Fusion: Multi-modal sensor platforms allow the combination of thermal, acoustic, sonar, LIDAR, etc. This has a long history in 3D vision [8, 30] and recent efforts have had significant impact [13]. In the radiological sensor domain, using many rad-detectors have resulted in intelligent radiation sensor systems (IRSS) [7] which are based on distributed radiation sensors coupled with networked position data to detect and locate radiation sources, either using geometry [7, 4], statistical models [28, 17, 11, 12, 21] or combining rad-detection with electromagnetic induction data [18]. All of these efforts rely on a large number of radiation detectors and are prohibitively expensive, or suffer from low efficiency; our goal is to achieve similar results with just one or two rad-detectors. Closest to our work is [26] where sensor calibration is presented and its noise characteristics analyzed. In contrast, we show that an omnidirectional LIDAR enables simple user-driven calibration and allows for results such as tracking multiple sources and tracking through walls.

Near Lighting: Most computer vision techniques assume distant lighting, such as the sun or the sky. However, near lighting is present in indoor scenes, night scenes and underwater situations. Near lighting analysis has a rich history in vision by exploiting the movement of the sources [5] or switching source positions [31], and [22] provides a recent survey. In all these efforts, the light-source's inverse square fall-off was used to obtain depth cues. We exploit radiological source fall-off to track the source in 3D and estimate its source strength, by using an additional, omnidirectional LIDAR sensor.

1.2. Overview

Our hybrid sensor suite produces radiological measurements and 3D vision sensor-based tracking of moving objects in the scene. Both of these streams of information can be converted into *trajectories*, i.e. 1D signals that denote change in radial distance over time.

The core problem of sensor fusion in our application becomes analysis of these signals to find those tracked objects that carry or contain radioactive material or sources. When we tease apart the components of the radiological trajectory given the measured vision trajectories as priors, the problem is cast as non-blind signal separation. With multiple radiological sources, the problem specifically becomes a non-negative least squares analysis of the radiological trajectory, and we discuss heuristics for robust performance in the face of radiological intereflections, occlusions, noise and radiological background noise.

In practical application, hundreds of people and objects may be tracked in an airport or port, with only a few, if any, radiological sources present. This suggests a sparse signal analysis of the trajectories. In our simulations and real laboratory experiments, the number of sources were a large fraction of the total number of tracked objects, and we found that non-negative least squares with a sparsifying threshold worked well and did better than sparse signal reconstruction by using L1 optimization (shown in the supplementary material). In all our experiments, we processed either the raw trajectories or by using a similarity transform applied to appearance profiles described in [15]. Finally, a small number of false positives is acceptable when determining the identity of the trafficker as long as the correct trajectory is found, i.e. the right person is apprehended.

2. Experimental Setup and Kalman Filtering

We demonstrate our algorithms on a real experimental setup consisting of isotropic radiological sensors and an omnidirectional LIDAR, as in Fig. 1. We use the Velodyne HDL-32E LIDAR which generates 70k points per frame at 10Hz and at 2cm resolution. The main radiation source used in this paper is a Californium source. A Plutonium Beryllium source was used alongside the Californium source for the two-source experiments. Both sources are isotropic.

Our rad-detector uses organic liquid compound EJ-309 [9, 10] that has both a high flashpoint and low chemical toxicity when compared to other detector liquids. A cylindrical alumina cell holds the liquid, of dimensions 7.56-by-7.56 cm (diameter by height). Scintillation in the liquid due to radioactive energy is absorbed by a photomultiplier tube (ET-Enterprises 9821B), and then converted into an electronic signal with high gain. A 14 bit, 250 MHz, 16-channel digitizer data acquisition system (Struck SIS3316) was used where each input channel buffer was sufficient to acquire large amounts of data over a long experiment.

Tracking moving humans is a core problem of computer vision, and the number of efforts in this space are too numerous to list here; see [25, 27, 20] for good surveys of the area. We are focusing on situations where people are carrying radiological sources. To reduce the computational burden and decrease response times, we project the omnidirectional 3D LIDAR data to a 2D plane and apply Kalman filter tracking of the resultant blobs [14].

In all the results shown here, this preprocessing step occurs in simulation, on a PC. In Fig 2 we propose a simple addition to the imaging ASIC circuit for any LIDAR system,

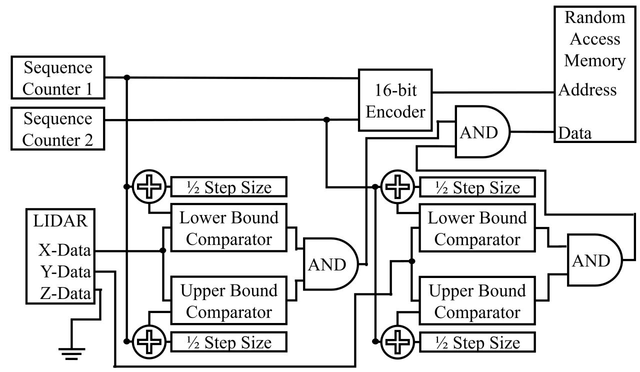

Figure 2: A high-level schematic diagram of the proposed circuit that produces top-down 2-D videos by integrating 3-D LIDAR data along the z-dimension.

that could perform the projection efficiently, in hardware. Most omnidirectional LIDAR systems recover 3D points of the scene, depicted in the diagram as the X, Y and Z, which are dimensions centered at the LIDAR as the origin and with an arbitrary rotation to the system axes. This circuit simply applies 2D projection on the $Z = 0$ plane followed by image discretization. The Z dimension is collapsed simply by ignoring it (grounding), and the two comparator components in the circuit discretize the X and Y space, allowing multiple LIDAR points to be represented in a single pixel in the new, grid-like 2D representation of the scene.

Before applying fusion, it is important to locate the radiological sensor in the LIDAR frame of reference. We envision a scenario where a space to be monitored (airport, docks, etc.) is instrumented with a small number of expensive radiological sensors and a larger number of cheaper vision sensors. These differences in number are because visual sensors are affected by opaque obstacles (e.g. walls) and require proper placement for full coverage.

Most previous efforts have focused on highly engineered setups [19, 24, 23] to estimate the relative pose of the LIDAR and radiological sensor. In the unstructured domain, [26] have proposed a calibration technique with a continuous wave time-of-flight (TOF) sensor. Since the method requires blind estimation of the radiological sensor location (i.e. self-calibration) errors in the optimization can result in location errors of the order of 1 meter. In contrast, an omnidirectional 360° vision system such as the LIDAR pictured in Fig. 1, can, in fact, directly view the radiological sensor. The only problem is disambiguating the radiological sensor in the field-of-view. To make the problem easier, we simply require that the user clicks on the location of the sensor, in the top-down 2D projected view that we render using the proposed ASIC. From this point on, we will assume that the location of the LIDAR is at the origin $(0, 0, 0)$ and the location of the radiological sensor is also known in the same frame, (S_x, S_y, S_z).

Exp	GT	Unaltered Signal			Transformed Signal		
		Det	Cor	y/n	Det	Cor	y/n
1	1	1	0.984	Yes	1	0.840	Yes
2	1	1	0.986	Yes	1	0.842	Yes
3	1	1	0.992	Yes	1	0.595	Yes
4	3	3	0.991	Yes	3	0.686	Yes
5	1	2	0.991	No	2	0.915	No
6	3	1	0.993	No	3	0.846	Yes
7	2	2	0.988	Yes	2	0.933	Yes
8	1	1	0.982	Yes	1	0.726	Yes
9	3	3	0.990	Yes	3	0.815	Yes
10	3	3	0.994	Yes	3	0.956	Yes
11	2	2	0.994	Yes	2	0.703	Yes
12	3	3	0.989	Yes	3	0.838	Yes
13	3	3	0.994	Yes	3	0.779	Yes
14	2	2	0.993	Yes	2	0.834	Yes
15	2	2	0.999	Yes	2	0.985	Yes

Table 1: Single source detection results both with and without the transformation from [15]. The trajectory indicated by GT (Ground Truth) is carrying the radiation source. The trajectory indicated by Det (Detected Signals) is *detected* to be carrying the radiation source. Cor indicated the correlation between the original radiological signal and the reconstructed signal. A correct detection is indicated by a "Yes" in the y/n column.

3. Single Source Tracking

In this section, we tackle the problem of finding and tracking a single *moving* radiological source. One of the main issues when tracking such a radiological source is that the measurements at the radiological sensor (usually called counts) depend on both the strength of the source (λ) and the source-sensor distance (R), i.e. the count rate is

$$C = \lambda(x, y, z) * e^{-\sigma R} * \frac{1}{R^2}. \tag{1}$$

λ is normally referred to as intrinsic efficiency, and is a function of the size of the radiological sensor and its internal efficiency, along with the room geometry, and the source. (x, y, z) is the 3D location of the source and $R = \sqrt{(x - S_x)^2 + (y - S_y)^2 + (z - S_z)^2}$ is the radial distance from source to radiological sensor. σ is the scattering and absorption of radiation in the medium; in this case, of air. In practice, the first two factors of Eq. 1 can be approximated as a slowly varying function of radial distance. For neutron radiation data the attenuation is only about 1% per meter of air distance travelled.

Note that this equation is similar to the near-lighting model [5, 16] in the case where the *camera directly images the source*, where the numerator would be the light-source intensity. The dependence of the "source strength" $\lambda(x, y, z) * e^{-\sigma R}$ on the scene is due to radiation's propensity to penetrate material and scatter about the scene. This produces different counts, especially if the source is near the floor or walls. Even in the absence of large objects or boundaries, there is a weak dependence on scattering through air, governed by the normally small parameter σ.

Of course, the other issue is that the low cost isotropic radiological sensors used in the experiments in this paper have no angular resolution and cannot tell where in the scene the detected material is located. Before we explain how to exploit the omnidirectional LIDAR to solve this problem, we make one more assumption about how the model for the "source strength", or $\lambda(x, y, z)$ can be made simpler, similar to previous models such as from [26]. We do approximate the numerator in Eq. 1 as a constant factor λ,

$$C = \lambda * \frac{1}{R^2}. \tag{2}$$

Consider a visually cluttered scenario where there is only a single source in the scene. This is usually the case in, say, an airport or port, where the chances of multiple security events happening on the same day in the same scene are highly unlikely. From the Kalman filter-based tracking, let us assume we have tracked all objects in the scene $O_1, O_2, ..O_n$ throughout the time of the experiment T. Each object would have a (x, y) trajectory in our 2D LIDAR space, for example O_1 : $((x_{11}, y_{11}), (x_{12}, y_{12}), ...(x_{1T}, y_{1T}))$, and this can easily be converted into radial distance *from the radiological sensor* (S_x, S_y, S_z), for example $O_1 : (R_{11}, R_{12}, ...R_{1T})$.

Now consider the count measurements $C_1, C_2, ..., C_n$ from the radiological sensor itself. These must, of course, be only dependent on the movement of the person carrying the single source and invariant to the other people in the scene. Inverting Eq. 2, we can convert these count measurements into radial measurements, modulated by the unknown source strength λ. Therefore the counts $C_1, C_2, ..., C_n$ from the rad-detector can be converted into radial distances $(R_{rad1}, R_{rad2}, ...R_{radT})$, given by $R_{radi} = \sqrt{\frac{\lambda}{C_i}}$.

We can normalize these values by taking the ratio of the radial measurements with, say, the minimal radial measurement corresponding the closest the source ever reaches the sensor. This would result in measurements that are independent of the source strength λ, as $R_{normi} = \frac{R_i}{R_{min}} = \sqrt{\frac{C_{max}}{C_i}}$. Similar normalization can also be applied in the objects tracked by the Kalman filter, for example O_{norm1} : $(\frac{R_{11}}{R_{1min}}, \frac{R_{12}}{R_{1min}}, ... \frac{R_{1T}}{R_{1min}})$. These trajectories are now independent of source strength and depend only on geometry. We can then use the cosine distance [29] between the radiological measurements and each of the n Kalman filter trajectories to give the best match.

Table 1 shows the results from 15 experiments. Each experiment has a group of two or three people moving in a

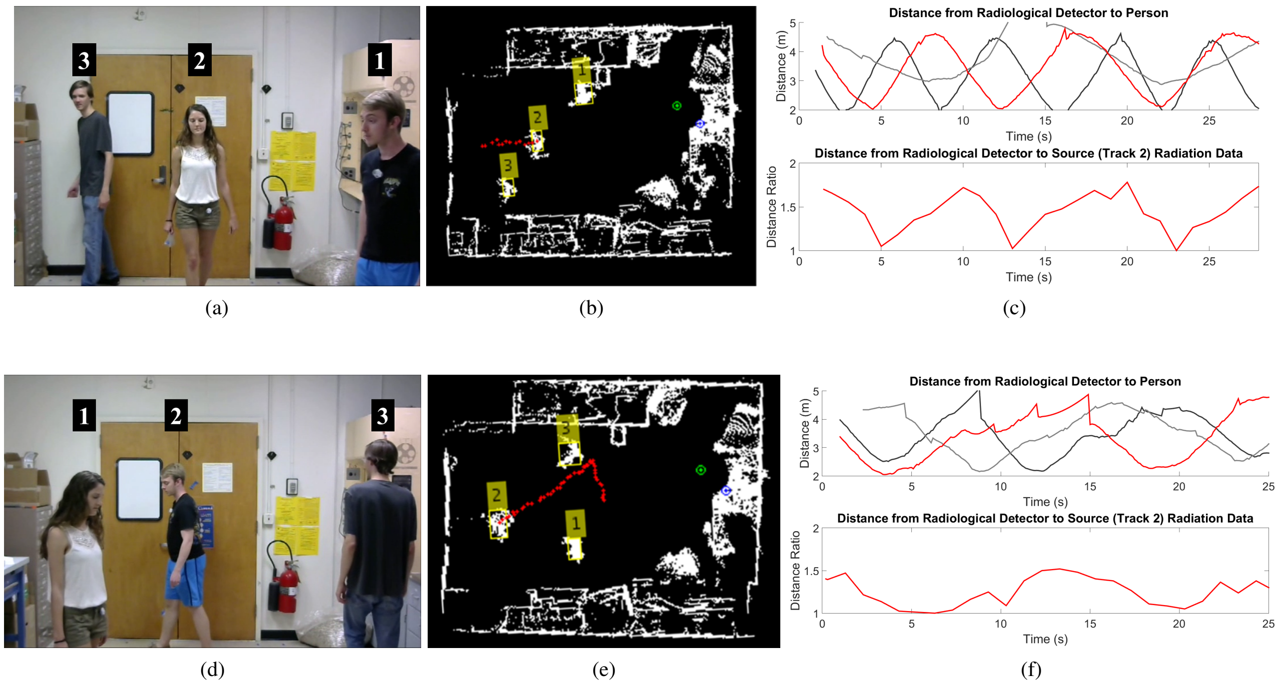

Figure 3: Results from two single-source tracking experiments. The person carrying the radiological source is Person 2. (top) The detection and tracking of a radiological source during periodic motion. (lower) The detection and tracking of a radiological source during random motion. (a,d) RGB image of the scene. (b,e) An Image of the LIDAR data's 2-D representation. (c,f) A plot of the distance to each person above a plot of the radiation data.

random motion in a laboratory environment. In 13 of those experiments, we correctly detected the radiological source holder, giving an accuracy of 86.67%. In Fig. 3 we show two of the experiments summarized in the table. In each of these two experiments Person 2 has the source. The final column shows plots of the trajectories (distance vs. time) for all three people in the scene, and under that a plot of the radiation data. There is only one trajectory (Person 2) that matches radiation data's waveform.

3.1. Multiple Source Tracking

In many environments, such as nuclear power plants or nuclear medicine clinics, there may be more than one radiological source being transported. To allow for better monitoring and awareness, it is useful to consider how to track multiple such sources. The measurements in the radiological sensor are linear in the number of counts C in Eq. 2, so, at some time t, given two sources i and j, the counts measured would be,

$$C(t) = C_i(t) + C_j(t) = \lambda_i * \frac{1}{R_i(t)^2} + \lambda_j * \frac{1}{R_j(t)^2}, \quad (3)$$

where R_i and R_j are the radial distances from each source to the radiological sensor. Now consider again, from the Kalman filter-based tracking, we have tracked all objects in the scene $O_1, O_2, ..O_n$ throughout the time of the experiment T. For each object, we can convert the radial distance from the radiological sensor (S_x, S_y, S_z), into a series of inverse-square fall-off terms, $O_1 : (\frac{1}{R_{11}^2}, \frac{1}{R_{12}^2}, ... \frac{1}{R_{1T}^2})$.

If we collect the radiological counts as a $T \times 1$ vector $\mathbf{C}$, and the radial inverse-square fall-off terms for each object as a $T \times n$ matrix $\mathbf{O}$, then

$$\mathbf{C} = \mathbf{O}\,\mathbf{I}, \quad (4)$$

where $\mathbf{I}$ is an $n \times 1$ vector such that

$$I(x) = \begin{cases} 0 & \text{if } x \neq i \text{ and } x \neq j \\ \lambda_x & \text{otherwise} \end{cases}$$

We used non-negative least-squares matrix factoring to solve Eq. 4, and sparsified $\mathbf{I}$ by setting all the values lower than a user defined threshold to zero, and declaring the trajectories corresponding to those remaining non-negative values as the ones with the radiological sources. In practice, we found that threshold setting was easy, and could perhaps be automated. We did not need to add additional constraints to enforce the sparsity of $\mathbf{I}$ during its estimation, although these can be used to increase robustness in the future. In Table 2 we show the first set of simulated experiments that we ran to test our approach. These were

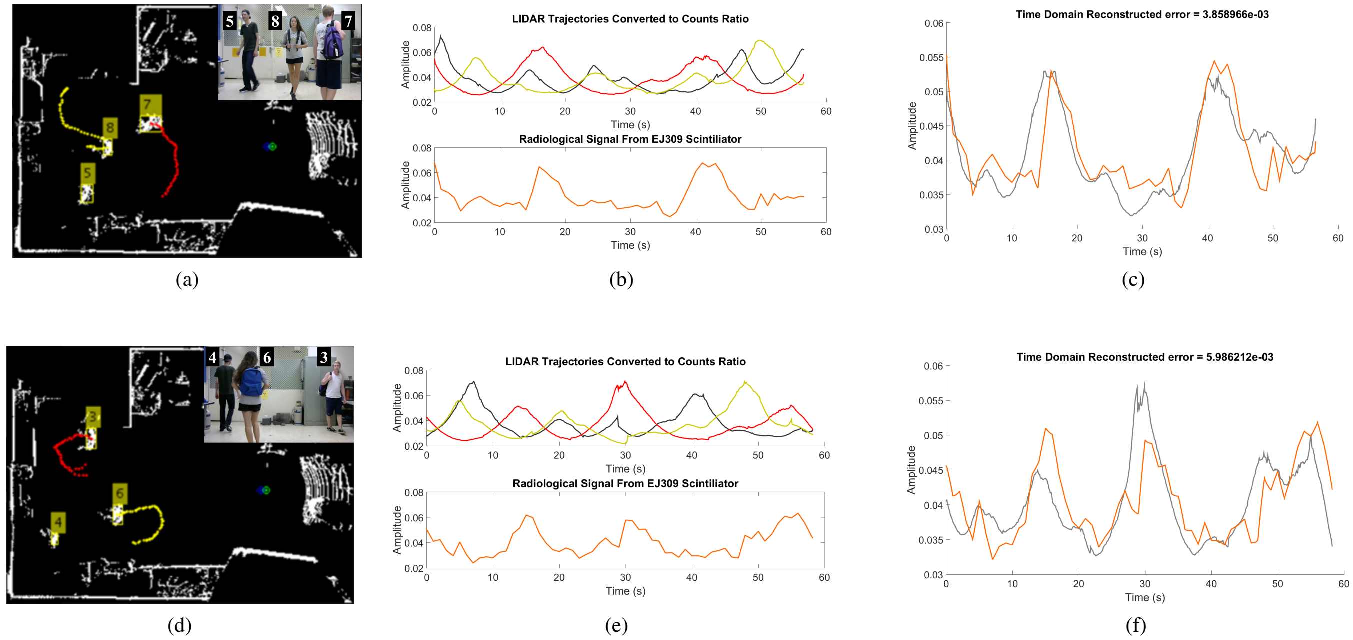

Figure 4: Results from two-source tracking experiments. Red and yellow indicate the people carrying radiation sources. (inset) RGB image of the scene. (a,d) An image of the LIDAR data's 2-D representation. (b,e) A plot of the distance to each person and a plot of the radiation data. (c,f) A plot of the radiation signal (orange) and the reconstructed signal (grey).

Exp	GT	Unaltered Signal		Transformed Signal	
		Det.	y/n	Det.	y/n
1	2, 6	2, 6	Yes	2, 6	Yes
2	2, 5	2, 5	Yes	2, 5	Yes
3	2, 6	2, 6	Yes	2, 6	Yes
4	2, 6	2, 6	Yes	2, 6	Yes
5	2, 5	2, 5	Yes	2, 5	Yes
6	3, 5	3, 5	Yes	3, 5	Yes
7	3, 6	3, 6	Yes	3, 6	Yes
8	3, 6	3, 6	Yes	3, 6	Yes
9	3, 5	3, 6	No	3, 5	Yes
10	2, 5	2, 5	Yes	2, 5	Yes
11	2, 5	2, 5	Yes	2, 3	No
12	2, 4	2, 4	Yes	2, 4	Yes
13	3, 6	3, 6	Yes	3, 6	Yes
14	3, 5	3, 5	Yes	3, 5	Yes
15	3, 5	3, 5	Yes	3, 5	Yes

Table 2: Two source detection results using simulated data for both unaltered and transformed data. Both sources have the same strength in these simulated data sets. The trajectories indicated by GT (Ground Truth) are carrying a radiation source. The trajectories indicated by Det (Detected Signals) are *detected* to be carrying a radiation source. A correct detection is indicated by a "Yes" in the y/n column.

Exp	GT	Unaltered Signal			Transformed Signal		
		Det.	Cor.	y/n	Det.	Cor.	y/n
1	1, 2	1, 2	0.996	Yes	1, 2	0.598	Yes
2	2, 1	2, 1	0.995	Yes	2, 1	0.579	Yes
3	1, 2	2, 1	0.993	Yes	1, 2	0.487	Yes
4	2, 3	3, 2	0.996	Yes	2, 3	0.607	Yes
5	1, 3	3, 1	0.991	Yes	1, 3	0.334	Yes
6	2, 3	3, 2	0.994	Yes	3, 2	0.102	Yes
7	1, 3	1, 3	0.996	Yes	1, 3	0.479	Yes
8	1, 3	1, 3	0.992	Yes	1, 2	0.613	No
9	1, 3	1, 3	0.996	Yes	1, 3	0.461	Yes
10	1, 3	1, 3	0.996	Yes	1, 3	0.462	Yes
11	2, 3	2, 3	0.994	Yes	2, 3	0.603	Yes
12	2, 3	2, 3	0.997	Yes	2, 3	0.617	Yes

Table 3: Two source detection results using real experimental data both with and without the transformation from [15]. The trajectories indicated by GT (Ground Truth) are carrying a radiation source. The trajectories indicated by Det (Detected Signals) are *detected* to be carrying a radiation source. A correct detection is indicated by a "Yes" in the y/n column. The first number in both the GT and Det columns corresponds to the stronger sources and the second number corresponds to the weaker source.

"real" simulations in the sense that we took real radiological counts from two actual experiments and added these using Eq. 3, and attempted to recover the ground-truth. The table has the ground-truth radiological pair, and the detected pair. Since this was a simulated setting, we could create groups of up to 7 "people" using vision trajectories from other experiments, to see how well our method can detect the two sources in a cluttered environment. A few false positives are acceptable if every trafficker is correctly identified. This means that a correct detection can be achieved despite a detected false positive as long as all ground truth trajectories are detected. In 15 experiments, we achieved only one incorrect detection, resulting in 93% separation and detection rate.

We then performed 12 actual experiments using two radiological sources, as shown in Table 3. The table shows, again the ground-truth radiological pair, and the detected pair, as well as the reconstructed distance measure from the cosine distance metric. Running our algorithm with the unaltered data yielded a detection rate of 100%. Fig. 4 shows two experiments from the set of 12. Each person wearing a backpack is carrying a radiation source. The center image shows three vision trajectories, one for each person in the scene, and the single radiological response (C in Eq. 3). The final column compares the original radiation signal shown in orange to the reconstructed signal using the weighted trajectories (C_i and C_j in Eq. 3) that have been inferred from the single radiological response C.

4. Blind Tracking

Radiological sources emit radiation that passes through materials that are opaque to light. Therefore, there is potential for tracking radiological sources through opaque obstacles, perhaps even far beyond visual range. However, the reason this is not done in the radiological community is because the constant λ assumption that we made in the earlier section does not hold. In other words, the source strength varies with radial distance, over time t, as

$$C(t) = \lambda(R(t)) * \frac{1}{R(t)^2} = F(R(t)). \tag{5}$$

Since counts are now a non-linear combination of source strength and distance, tracking single sources over long distances with a single radiological sensor is almost never done, since it would imply finding $F^{-1}(C(t))$ or inverting the above equation. In our case, however, we are using a combination of a radiological sensor and an omnidirectional LIDAR, which can allow us to break the dependency between source strength λ and radial distance R.

Consider a scene with a single moving radiological source, which we have identified using the methods described previously. For a given interval of time, we now know pairs of $(C(t), R(t))$, or $(F(R(t)), R(t))$. This is because the radial distances $R(t)$s are estimated directly by the LIDAR, and the counts come from the radiological sensor. We can therefore fit a parametric model (always linear in the results shown) to the $(F(R(t)), R(t))$ pairs. This data-driven approach makes it trivial to invert F, since we can extrapolate the model to predict the $R(t)$ associated with any $C(t)$ in any future time instance t.

The only problem remaining is tracking a person with a radial distance $R(t)$, since the person could be anywhere in a circle around the radiological sensor. To break this ambiguity, we have two strategies. In the case of short periods of occlusion, we use the previous velocity vector estimated by the Kalman filter, and intersect the circle with the ray defined by this velocity vector and the last detected LIDAR postion. In the top row of Fig. 5 we show an example of a person disappearing behind a large cubicle wall. The first image shows the color frame at the moment of occlusion. The second image is a 2-D representation of the LIDAR data before the person is occluded. In the third image, we intersect the circle (depicted in green) with the motion ray. This position is fed into the Kalman filter. The final image shows that the person reappears at the other end of the wall, with their Kalman filter label intact, i.e. we have tracked the same person through visual occlusion.

Obviously, the previous method fails if the person's motion deviates from the most recent Kalman filter estimated velocity vector. To go beyond this, we simply add a second radiological sensor and apply the same modeling algorithm as discussed earlier. Given the two radiological counts from each sensor, we can convert them into two estimated radial distances. These (in 2D) will intersect at two points, and we use the point closest to the last seen LIDAR position, followed by Kalman filtering to choose between the two intersection points. In the second row of Fig. 5 we show an example of a person going behind a wall, and then walking in a circle. The circular motion happens completely behind the wall and is not visible by the LIDAR. Yet, using the LIDAR to estimate the $(C(t), R(t))$ map for each sensor allow us to predict the location of the person. Notice the two green circles intersect behind the wall, and the red dots (that depict the recent trajectory) show a circular path.

5. Conclusion and Limitations

This is the first effort to fuse omnidirectional LIDARs with isotropic radiological sources. The impact of this work is that such a combination can be used in airports, ports, commercial areas and battlefields with net security benefits. However, we point out a few limitations that we hope to address in future work:

Source strength model: In single source tracking we assumed the source strength λ is constant, and in the blind tracking experiments we assumed it was modeled with a ra-

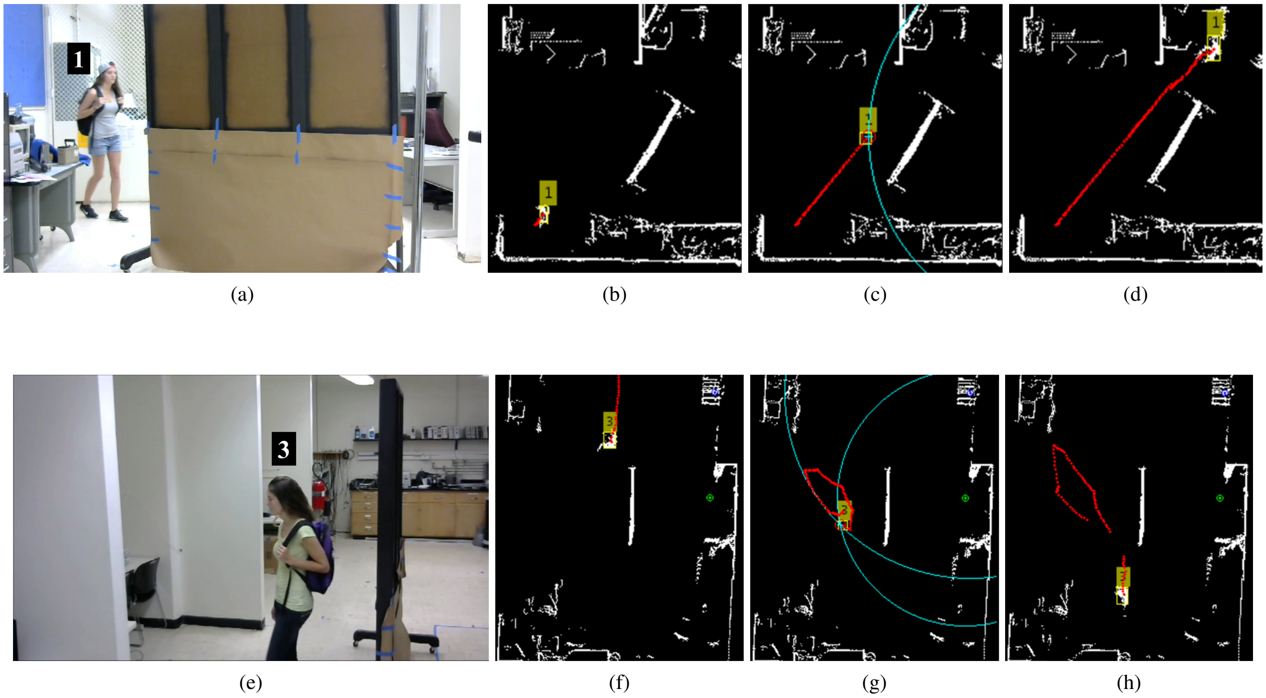

Figure 5: Blind tracking results using a single detector (top) and two detectors (lower). (a,e) RGB image of the scene. (b-d,f-h) A sequence of images from the 2-D LIDAR data's representation showing the person being tracked through the occlusion. The blue circles indicate the unidirectional distance from each detector to the person.

dial dependency $\lambda(R)$. Both of these approximations can break with scenes that have more interesting geometries than our laboratory settings, with complex material properties. A future goal is to learn the scene radiation background prior, from the LIDAR geometry itself.

Complex sources: Our algorithm assumes that the sources are solid. With different radiological materials, or even a volumetric source such as liquid or gas, these assumptions must change.

Static source and switching sources: Our algorithms require moving radiation sources. A static source becomes part of the background and is removed by background subtraction. This may cause problems if people carrying the sources "drop off" the package. Of course, we can detect when the radiological counts become constant, and can pick up the radiological trajectory again as the source starts moving. This is a fascinating direction for future work.

Occlusions of people: Our blind tracking results assume a static, homogeneous obstacle, such as a wall. However, there are obstacles with severe and dynamically changing radiation absorption, such as groups of people. In supplementary material we show an experiment tracking a person as they are occluded by a crowd. Even though our method allows tracking to function longer than with just the Kalman filter + LIDAR data, it still fails eventually. Modeling heterogeneous obstacles will be addressed in follow-up work.

Transformation Results: The transformation described in [15] does not appear to help or hinder the outcome of the detection algorithm in a significant way.

Failure Cases: Single-source and multi-source tracking failure cases are caused when two object trajectories are similar. This happens due to the different sampling rate of radiological (1HZ) and vision (10HZ) sensors. Combined with Poisson noise in the radiation counts, this explains visual discrepancies in Fig.3 and Fig.4. Notice that the correct match still shares trajectory extrema (maxima and minima), which allows signal correlation to find the answer.

6. Acknowledgements

This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number, 2014-DN-077-ARI083. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

References

[1] Application of nuclear forensics in combating illicit trafficking of nuclear and other radioactive material. In *IAEA TECDOC Series 1730*. IAEA, 2014. 1

[2] H. Barrow and J. Tenenbaum. Computer vision systems. *Computer vision systems*, page 2, 1978. 2

[3] B. Bernacki, J. Schweppe, S. Stave, D. Jordan, J. Kulisek, T. Stewart, and C. Seifert. Estimating radiological background using imaging spectroscopy. In *SPIE Defense+ Security*, pages 90880L–90880L. International Society for Optics and Photonics, 2014. 1, 2

[4] J.-C. Chin, D. K. Yau, N. S. Rao, Y. Yang, C. Y. Ma, and M. Shankar. Accurate localization of low-level radioactive source under noise and measurement errors. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*, SenSys '08, pages 183–196, New York, NY, USA, 2008. ACM. 2

[5] J. J. Clark. Active photometric stereo. In *Computer Vision and Pattern Recognition, 1992. Proceedings CVPR'92., 1992 IEEE Computer Society Conference on*, pages 29–34. IEEE, 1992. 2, 4

[6] T. B. Cochran and M. G. McKinzie. Detecting nuclear smuggling. *Scientific American*, 298(april):98 – 104, 2008. 1

[7] D. Cooper, R. Ledoux, K. Kamieniecki, S. Korbly, J. Thompson, J. Batcheler, S. Chowdhury, N. Roza, J. Costales, and V. Aiyawar. Intelligent radiation sensor system (irss) advanced technology demonstrator (atd). In *Homeland Security (HST), 2012 IEEE Conference on Technologies for*, pages 511–516, Nov 2012. 2

[8] J. L. Crowley. Principles and techniques for sensor data fusion. In *Multisensor Fusion for Computer Vision*, pages 15–36. Springer, 1993. 2

[9] S. T. . Eljen Technology, P. O. Box 870. `http://www.eljentechnology.com/images/stories/Data_Sheets/Liquid_Scintillators/EJ309datasheet.pdf`. 2012. 3

[10] A. Enqvist, C. C. Lawrence, B. M. Wieger, S. A. Pozzi, and T. N. Massey. Neutron light output response and resolution functions in ej-309 liquid scintillation detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 715(0):79 – 86, 2013. 3

[11] A. Gunatilaka, B. Ristic, and R. Gailis. On localisation of a radiological point source. In *Information, Decision and Control, 2007. IDC '07*, pages 236–241, Feb 2007. 2

[12] J. W. Howse, L. O. Ticknor, and K. R. Muske. Least squares estimation techniques for position tracking of radioactive sources. *Automatica*, 37(11):1727 – 1737, 2001. 2

[13] S. Izadi, D. Kim, O. Hilliges, D. Molyneaux, R. Newcombe, P. Kohli, J. Shotton, S. Hodges, D. Freeman, A. Davison, et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 559–568. ACM, 2011. 2

[14] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960. 3

[15] S. J. Koppal and S. G. Narasimhan. Clustering appearance for scene analysis. In *Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2006. 3, 4, 6, 8

[16] S. J. Koppal and S. G. Narasimhan. Novel depth cues from uncalibrated near-field lighting. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007. 4

[17] A. Liu, J. Bunn, and K. Chandy. An analysis of data fusion for radiation detection and localization. In *Information Fusion (FUSION), 2010 13th Conference on*, pages 1–8, July 2010. 2

[18] Z. Long, W. Wei, A. Turlapaty, Q. Du, and N. Younan. Fusion of radiation and electromagnetic induction data for buried radioactive target detection and characterization. *Nuclear Science, IEEE Transactions on*, 60(2):1126–1133, April 2013. 2

[19] L. Mihailescu, K. Vetter, and D. Chivers. Standoff 3d gamma-ray imaging. *Nuclear Science, IEEE Transactions on*, 56(2):479–486, 2009. 1, 2, 3

[20] T. B. Moeslund and E. Granum. A survey of computer vision-based human motion capture. *Computer vision and image understanding*, 81(3):231–268, 2001. 3

[21] M. Morelande and A. Skvortsov. Radiation field estimation using a gaussian mixture. In *Information Fusion, 2009. FUSION '09. 12th International Conference on*, pages 2247–2254, July 2009. 2

[22] T. Papadhimitri and P. Favaro. Uncalibrated near-light photometric stereo. In *BMVC*, volume 2, page 8, 2014. 2

[23] A. C. Raffo-Caiado, K.-P. Ziock, J. Hayward, S. Smith, J. Bogard, and C. B. Boehnen. Combining measurements with three-dimensional laser scanning system and coded-aperture gamma-ray imaging systems for international safeguards applications. In *2010 IAEA Symposium on International Safeguards*, pages 1–5, 2010. 1, 2, 3

[24] A. C. Raffo-Caiado, K.-P. Ziock, J. Hayward, S. Smith, A. Solodov, L. Mihailescu, K. Vetter, A. Dougan, M. Burks, J. G. Gonçalves, et al. Investigation of combined measurements with three-dimensional design information verification system and gamma-ray imaging systems for international safeguards applications. In *Proceedings 50th INMM Annual Meeting*, pages 12–16, 2009. 1, 2, 3

[25] C. Raghavachari, V. Aparna, S. Chithira, and V. Balasubramanian. A comparative study of vision based human detection techniques in people counting applications. *Procedia Computer Science*, 58:461–469, 2015. 3

[26] P. Riley, A. Enqvist, and S. J. Koppal. Low-cost depth and radiological sensor fusion to detect moving sources. In *3D Vision (3DV), 2015 International Conference on*, pages 198–205. IEEE, 2015. 2, 3, 4

[27] A. W. Smeulders, D. M. Chu, R. Cucchiara, S. Calderara, A. Dehghan, and M. Shah. Visual tracking: An experimental survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1442–1468, 2014. 3

[28] A. Sundaresan, P. Varshney, and N. Rao. Distributed detection of a nuclear radioactive source using fusion of correlated decisions. In *Information Fusion, 2007 10th International Conference on*, pages 1–7, July 2007. 2

[29] P.-N. Tan et al. *Introduction to data mining*. Pearson Education India, 2006. 4

[30] M. Ye, Y. Zhang, R. Yang, and D. Manocha. 3d reconstruction in the presence of glasses by acoustic and stereo fusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4885–4893, 2015. 2

[31] T. E. Zickler, P. N. Belhumeur, and D. J. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision*, 49(2-3):215–227, 2002. 2

[32] K.-P. Ziock, E. C. Bradley, A. Cheriyadat, M. Cunningham, L. Fabris, C. Fitzgerald, J. Goddard, D. E. Hornback, R. A. Kerekes, T. P. Karnowski, et al. Performance of the roadside tracker portal-less portal monitor. *Nuclear Science, IEEE Transactions on*, 60(3):2237–2246, 2013. 1, 2

[33] K. P. Ziock, W. W. Craig, L. Fabris, R. C. Lanza, S. Gallagher, B. K. Horn, and N. W. Madden. Large area imaging detector for long-range, passive detection of fissile material. *Nuclear Science, IEEE Transactions on*, 51(5):2238–2244, 2004. 1, 2