Check for updates

OPEN

Expectation-propagation for weak radionuclide identification at radiation portal monitors

Yoann Altmann[1], Angela Di Fulvio[2] ✉, Marc G. Paff[3], Shaun D. Clarke[4], Mike E. Davies[5], Stephen McLaughlin[1], Alfred O. Hero[6] & Sara A. Pozzi[4]

We propose a sparsity-promoting Bayesian algorithm capable of identifying radionuclide signatures from weak sources in the presence of a high radiation background. The proposed method is relevant to radiation identification for security applications. In such scenarios, the background typically consists of terrestrial, cosmic, and cosmogenic radiation that may cause false positive responses. We evaluate the new Bayesian approach using gamma-ray data and are able to identify weapons-grade plutonium, masked by naturally-occurring radioactive material (NORM), in a measurement time of a few seconds. We demonstrate this identification capability using organic scintillators (stilbene crystals and EJ-309 liquid scintillators), which do not provide direct, high-resolution, source spectroscopic information. Compared to the EJ-309 detector, the stilbene-based detector exhibits a lower identification error, on average, owing to its better energy resolution. Organic scintillators are used within radiation portal monitors to detect gamma rays emitted from conveyances crossing ports of entry. The described method is therefore applicable to radiation portal monitors deployed in the field and could improve their threat discrimination capability by minimizing "nuisance" alarms produced either by NORM-bearing materials found in shipped cargoes, such as ceramics and fertilizers, or radionuclides in recently treated nuclear medicine patients.

The growing terrorism threat based on the use of special nuclear materials (SNMs), i.e., highly enriched uranium (HEU), weapons-grade plutonium (WGPu), or high-activity radiological sources has reinforced the need for improved population protection mechanisms. Nuclear security aims to deter and detect the smuggling of these materials across state borders. One major defense mechanism involves the installation of radiation portal monitors (RPMs) at border crossings. These RPMs typically consist of ^{3}He proportional counters embedded in polyethylene for neutron detection, and slabs of polyvinyl-toluene (PVT) scintillators for gamma-ray detection. Only a tiny fraction of the millions of vehicles and cargo containers entering a country like the United States are likely to be carrying radiological contraband. The International Atomic Energy Agency's Incident and Trafficking Database (ITDB) merely counts a few dozen reported successful interdictions of nuclear and radiological materials globally per year[1,2]. The ITDB provides only a partial picture of the number of smuggling attempts. The reported figures should be considered a lower bound of the number of successful interdictions, because they include only successful interdictions, voluntarily reported by the member states.

Complicating matters, the radiological contraband might be well shielded. In 2017, the United Nations Conference on Trade and Development estimated the global container port throughput at over 750 million 20-foot equivalent units[3]. As a consequence, RPMs are limited in measurement time to minimize unnecessary impediments to the flow of traffic and commerce. RPMs need to function rapidly while collecting sufficient data to positively identify the presence of a radiation source, which may produce a signal just slightly above the natural background.

[1]School of Engineering and Physical Sciences, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS, United Kingdom. [2]Department of Nuclear, Plasma, and Radiological Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, United States. [3]Los Alamos National Laboratory, Los Alamos, NM, 87545, United States. [4]Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, MI, 48109, United States. [5]School of Engineering University of Edinburgh King's Buildings, Edinburgh, EH9 3JG, United Kingdom. [6]Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, 48109, United States. ✉e-mail: difulvio@illinois.edu

Border protection agents screen inbound vehicles and cargo containers for suspicious levels of radiation relative to the background, and flag these for a more thorough secondary inspection. Detecting smuggled nuclear and radiological material is analogous to finding a needle in a haystack, whereby SNMs can be difficult to detect, quantify, and locate. Nuisance alarms are radiation alarms caused by sources of radiation that pose no security threat. Many common goods shipped across border crossings contain sufficient naturally occurring radioactive material (NORM) to set off gamma alarms in RPMs[4]. Medical isotopes are another growing source of nuisance alarms. A patient may emit sufficient gamma radiation for days or even weeks after a procedure to set off an RPM gamma alarm[4–7], depending on the nuclear medicine isotope used and its administered activity.

NORM-bearing cargo and nuclear medicine patients are significantly more prevalent than nuclear smugglers in cross-border traffic. Hence, customs and border protection agents spend an exorbitant amount of time processing nuisance alarms in secondary inspections that can last tens of minutes per offending vehicle or cargo container[8]. Due to the low signal to background ratio, simply alarming on the presence of a radioactive source is a challenge in itself for primary inspection. In the interest of saving time for both customs and border protection agents, as well as people crossing borders, combining primary and secondary inspections appears as an attractive solution. In an ideal scenario, the primary inspection would simultaneously detect, identify and quantify any source of radiation of interest, so that, for example, nuclear medicine patients avoid the discomfort caused by a lengthy secondary inspection. Identifying radionuclides, however, is even more sensitive to signal-to-background ratio than simply detecting the presence of a radiation source.

One concerning and challenging scenario involves the contextual presence of multiple radionuclides, i.e., mixed sources. In this case, strong NORM sources can mask a weaker SNM source, and further jeopardize the identification process. Gamma-ray spectroscopy inspections performed using inorganic scintillators or semiconductor detectors, such as NaI(Tl) or HPGe, respectively, are typically able to resolve most of the photopeaks, which serve as fingerprints of the present radionuclides, and therefore facilitate the nuclide identification in a mixed source scenario[9]. The vast majority of deployed RPMs utilizes instead organic scintillators, i.e., PVT, because of the high intrinsic efficiency of these detectors, their relatively low cost, and suitability to be produced in large shapes. The response of organic scintillator-based RPMs is not characterized by sharp photopeaks, but rather by smooth edges and continuum regions that result from Compton scattering interactions. Therefore, the spectral response of an RPM organic scintillator to a mixed source will essentially be a smooth linear combination of the responses to individual sources. It is hence challenging to identify all the components of the mixed source and estimate the relative activities of the constituent sources.

The performance of a portal monitor in terms of sensitivity, i.e., maximization of the positive detection rate, is a function of the detection efficiency of the system and its form factor, which should be optimized for a specific application. Paff and colleagues[8] have already shown that the system sensitivity can be optimized by selecting large detector panels. In this work, we focus on the capability of identifying multiple sources in a mixture of nuclides, following an alarm event.

The proposed method is also relevant to a number of other radiation identification and localization applications, such as radionuclide search with unmanned vehicles in a given environment, where the statistics of the signal of interest is poor compared to the background, because of short measurement time, distance between the detector and the source, low detection intrinsic efficiency and/or weakness of the source.

Algorithms for RPM signal unmixing

Radiation detection and characterization in the nuclear security area is challenging due to the low intensity of the signal of interest, typically much lower than the background. Two main detrimental components are added to the SNM signal of interest: spectra of additional NORM sources, either located inside the cargo or part of the natural background surrounding the portal monitor, and intrinsic observation Poisson noise (shot noise), which is not negligible for short measurement times and therefore can lead to poor signal-to-noise ratios. Bayesian inference is particularly attractive in such challenging scenarios, and advances in approximate methods[10,11] allow complex models to be used with computational times compatible with real-time constraints.

Bayesian approaches to detect, classify, and estimate smuggled nuclear and radiological materials are not a new consideration[6,12], and were extensively studied for the development of the Statistical Radiation Detection System at Lawrence Livermore National Laboratory. This group has used Bayesian model-based sequential statistical processing techniques to overcome the low signal-to-background ratio that complicates traditional gamma spectroscopy techniques with high-resolution HPGe and inorganic scintillation detectors[13,14]. Bayesian approaches have also been applied to radionuclide identification for NaI(Tl) detectors using a wavelet-based peak identification algorithm with Bayesian classifiers[15], for LaBr$_3$(Ce) using a sequential approach[16], and to HPGe detectors using non-parametric Bayesian deconvolution to resolve overlapping peaks[17]. Bayesian approaches have been recently investigated for the detection of single and mixed gamma sources with short measurement times[12]. The use of related machine-learning-based methods was also recently demonstrated for source identification in spectra recorded using inorganic scintillators[18].

Results

In this study, we considered two types of organic scintillation detectors, based on liquid EJ-309 and stilbene crystal, respectively, as detailed in the "Methods" section. The functional difference between the two detectors most relevant to this work is their energy resolution as illustrated in Fig. 1, which depicts their integral-normalized response to a ^{201}Tl (left) and to a ^{99m}Tc (right) source. This figure shows that stilbene exhibits sharper Compton edges than EJ-309, thanks to its better energy resolution.

Table 1 lists the 11 nuclides that were measured using the two different detectors, and the relative fractions used to generate synthetic mixtures and assess the performance of the new algorithm. For each mixture, several
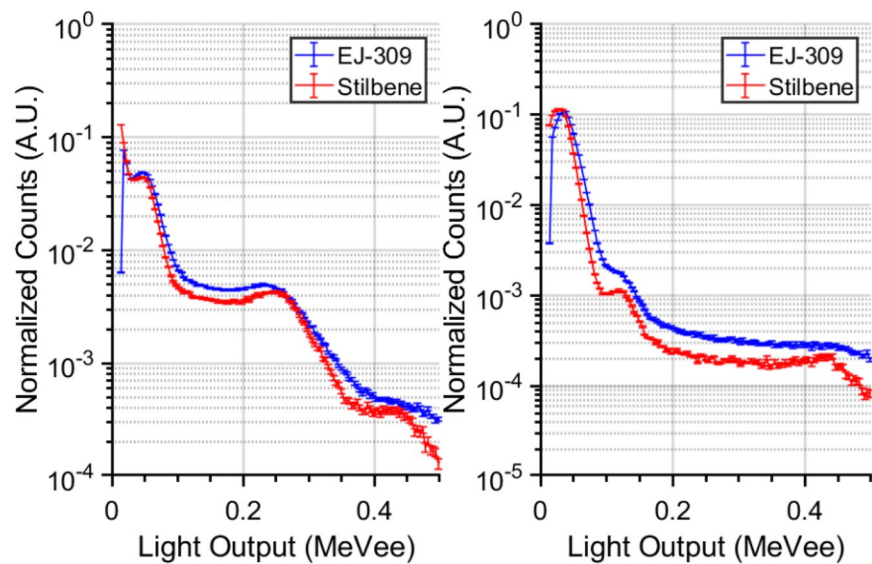
Figure 1. Comparison of light output spectra of ^{201}Tl (left) and ^{99m}Tc (right) sources measured using the EJ-309 (blue curves) and the stilbene (red curves) scintillators. For comparison purposes, the spectra have been normalised to integrate to one.

	^{241}Am	^{133}Ba	^{57}Co	^{137}Cs	^{123}I	^{201}Tl	^{67}Ga	I^{131}	^{99m}Tc	^{111}In	WGPu
Mixture 1	1/6	0	0	5/6	0	0	0	0	0	0	0
Mixture 2	0	1/3	0	1/3	0	0	1/3	0	0	0	0
Mixture 3	0	0	0	0	0	0	1/3	0	1/3	0	1/3
Mixture 4	0	1/4	0	1/4	1/4	0	0	0	1/4	0	0
Mixture 5	0	1/2	0	0	0	0	0	1/3	0	0	1/6
Mixture 6	0	0	0	0	0	0	0	0	5/6	0	1/6
Mixture 7	0	0	0	0	0	0	0	0	6/7	0	1/7
Mixture 8	0	0	0	0	0	0	0	0	7/8	0	1/8
Mixture 9	0	0	0	0	0	0	0	0	8/9	0	1/9

Table 1. Composition of the nine mixtures tested in this work. The fractions represent the ratio of the detected counts associated with each source.

data sets were created to obtain spectra with a total counts from 500 to 500 k, where the observation noise was modeled by Poisson noise.

We compared the unmixing performance of the new algorithm, referred to as MMSE$_{BTG}$, to that of two Bayesian strategies, namely the maximum a-posteriori (MAP) and the minimum mean squared error (MMSE$_{L1}$) approaches presented in[12]. These two approaches, denoted by MAP$_{L1}$ and MMSE$_{L1}$, respectively, are detailed in the "Methods" section.

As the metric for estimation accuracy, we used the root-mean-square error (RMSE)

$$RMSE = \sqrt{\frac{\|\mathbf{z} - \hat{\mathbf{z}}\|_2^2}{N}}$$

(1)

between the known nuclide fractions $\mathbf{z}$ and their estimated values $\hat{\mathbf{z}}$, where N is the number of nuclides in the spectral library. Figure 2 compares the RMSEs obtained by the three methods mentioned above for the nine mixtures of Table 1, as the total number of counts increases (from 500 to 1 M) and using the stilbene detector. The new MMSE$_{BTG}$ method generally provides more robust results, compared to the MAP$_{L1}$ and MMSE$_{L1}$ approaches, yielding consistently lower RMSEs. The MMSE$_{BTG}$ RMSE becomes comparable to the MAP$_{L1}$ RMSE when only 500 counts are measured, and when the mixture contains nuclides with spectral similarities, e.g., ^{123}I and ^{99m}Tc in the fourth mixture.

The RMSEs obtained using the MMSE$_{BTG}$ algorithm and simulated data show overall comparable performances using either detector (see Fig. 3). The results using the stilbene detector are slightly better, i.e., present lower RMSEs, especially for mixtures of three or more nuclides, e.g., WGPu, ^{99m}Tc, and ^{67}Ga (mixture 3). This result is expected because of the better energy resolution of stilbene, compared to EJ-309. Similar results have been obtained with the two other competing methods.

3
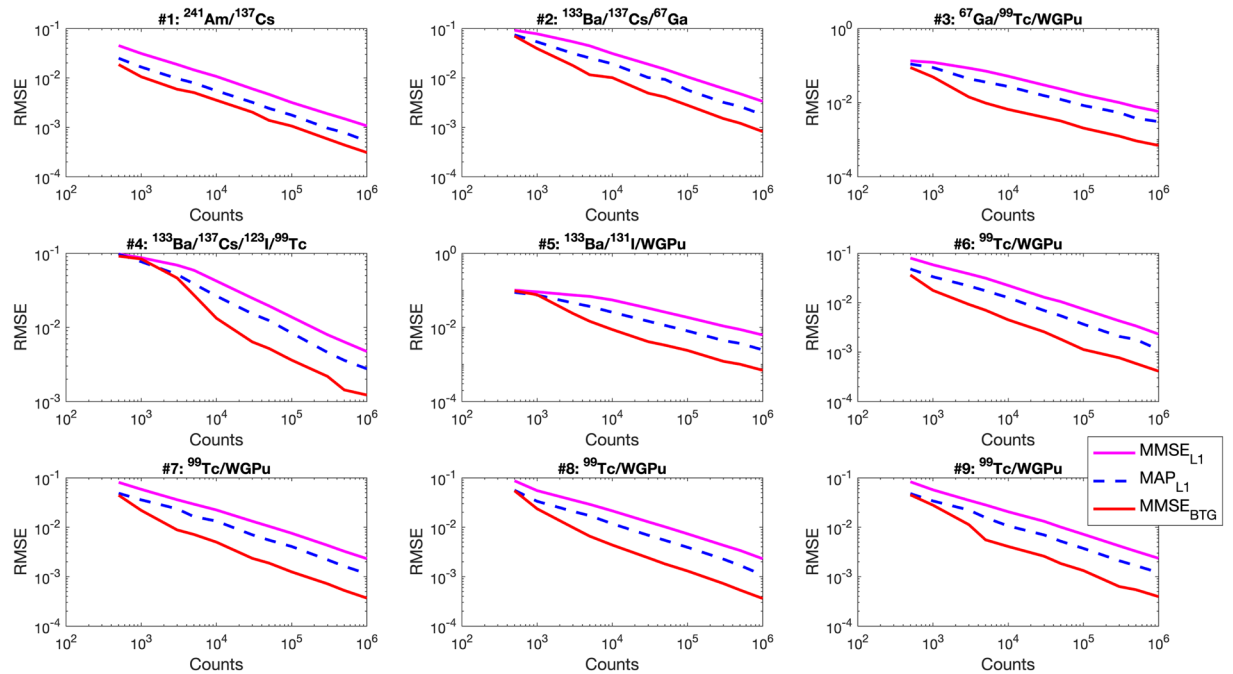
Figure 2. RMSEs obtained with the MMSE$_{BTG}$, MMSE$_{L1}$ and MAP$_{L1}$ algorithms for the mixtures described in Table 1 and measured with the stilbene detector, as a function of the detection counts.
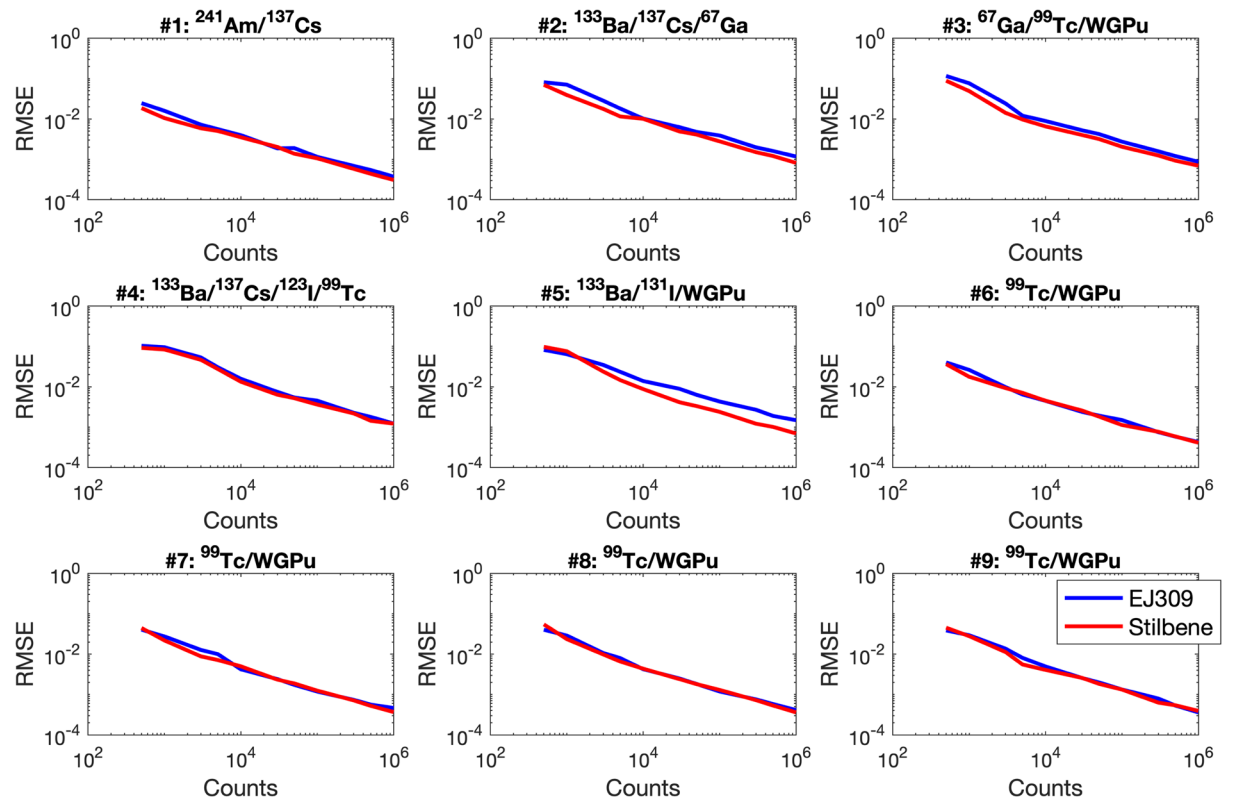
Figure 3. RMSE of the MMSE$_{BTG}$ algorithm calculated using the stilbene (red curves) and EJ-309 (blue curves) experimentally measured data.

A significant advantage of the proposed MMSE$_{BTG}$ algorithm is that it directly provides uncertainty quantification, i.e., the estimated probability of the presence of each source from the library. The MMSE$_{BTG}$ algorithm generates theses estimates from the posterior distribution, which are not directly available from the MAP$_{L1}$ and
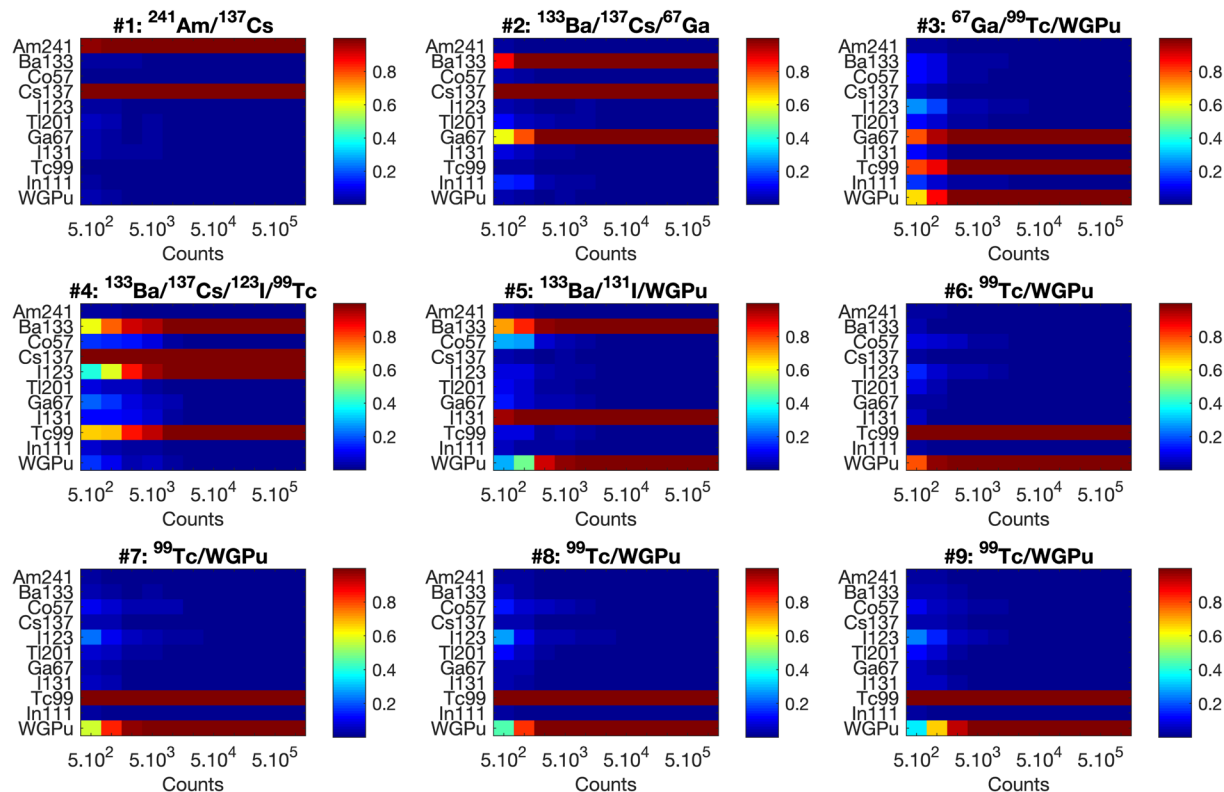
Figure 4. Estimated marginal posterior probabilities of the presence of each of the library sources in the mixture as a function of the total photon count. The probabilities depicted have been obtained using the stilbene detection spectra, have been computed individually across 100 noise realizations, and finally averaged over these 100 noise realizations.

$MMSE_{L1}$ algorithms. If the measured spectrum consists of more than 1000 counts, the algorithm correctly identifies with high probability the nuclides in the mixture and its performance slightly degrades as the number of sources that are present increases and the overall gamma counts per source decrease (see Fig. 4). In addition to providing estimated probabilities of source presence, the $MMSE_{BTG}$ yields superior performance, compared to the MAP_{L1} and $MMSE_{L1}$ algorithms used in Fig. 2. Furthermore, in contrast to the proposed $MMSE_{BTG}$ approach, the $MMSE_{L1}$ and MAP_{L1} algorithms require tuning of a threshold for source detection, whose optimal value (in terms of probabilities of false alarm and detection) is difficult to tune in practice, as it depends on the counts and the mixture composition. For this reason, we only report here the detection results obtained using the $MMSE_{BTG}$ approach.

In Fig. 4, an increasing number of isotopes not present in the actual mixture is identified as potentially present, when there are few gamma counts. For example, for sparse spectra (<1,000 counts) containing WGPu and ^{99m}Tc (mixtures 6–9), the algorithm suggests the potential presence of ^{123}I. This can be explained by the similarity of the spectra of ^{123}I and ^{99m}Tc (as shown in Fig. 5). The discrimination of these two nuclides becomes easier as the gamma counts increase.

Mixtures 6–9 simulate a specific scenario, where a WGPu source is detected together with an increasing amount of ^{99m}Tc, which is the most commonly used medical radioisotope and could, therefore, be used to mask (in terms of relative counts) the presence of WGPu. The results illustrate that the estimated probability of presence of WGPu decreases as its proportion decreases in the mixture (from mixture 6 to mixture 9), as could be expected.

Figure 6 shows the empirical WGPu alarm rate, i.e., the fraction of the measurements containing WGPu for which the estimated probability of WGPu presence is larger than 50%, as a function of the total photon counts (top) and WGPu counts (bottom) for the different WGPu -based mixtures, using the stilbene detector. With a target WGPu alarm rate of 80%, a few hundreds of counts from the WGPu source would set off the portal alarm, even in the presence of up to three other highly-radioactive masking sources. The highest number of approximately 3000 overall counts to trigger an alarm state is needed for mixture 5, which includes WPGu, ^{133}Ba, and ^{131}In. In similar irradiation conditions, in the presence of a mixed source, a detector similar to the one investigated would record approximately 130 counts during a 3-s vehicle scan time[8]. Assuming that the intrinsic efficiency scales with the volume of the detector and factoring an efficiency loss of 10% due to non-ideal light collection, a relatively small 2752 cm^3 single module used in portal monitors[19] would record approximately 3100 counts during a 3-s acquisition of a mixture of ^{133}Ba, ^{131}In, and WGPu. This acquisition time would be sufficient to set an alarm condition in the portal monitor. Regarding computational costs, the three competing methods ($MMSE_{BTG}$,
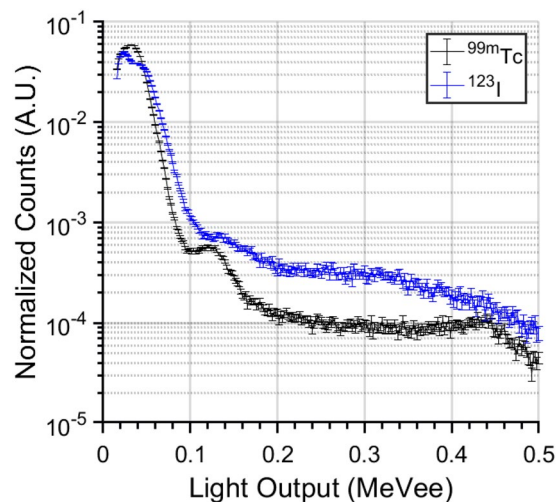
Figure 5. Comparison of integral-normalized light output spectra emitted by ^{123}I and ^{99m}Tc measured by the EJ-309 detector. The zero-lagged cross-correlation coefficient between the two spectra is 0.97.
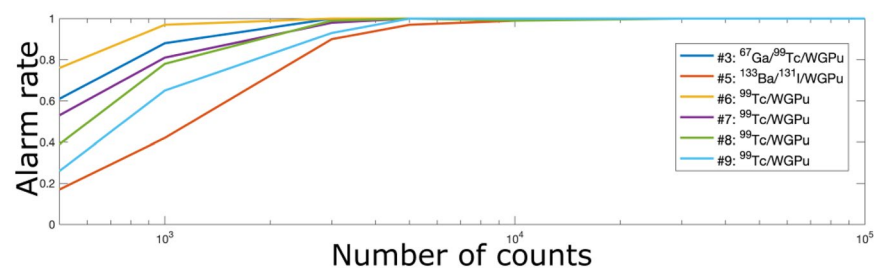
Figure 6. Comparison of the WPGu alarm rates with the stilbene detector in the presence of mixtures with WGPu.

$MMSE_{L1}$ and MAP_{L1}) have been implemented using Matlab 2017b running on a MacBook Pro with 16 GB of RAM and a 2.9 GHz Intel Core i7 processor. Since the $MMSE_{L1}$ is a simulation-based algorithm (see "Methods" section), its computational cost is significantly higher than the two other methods and it requires 66 s to analyze one spectrum (using 5000 iterations and assuming at most 11 sources in the mixture), on average. This prevents its use within portal monitors. Conversely, MAP_{L1} only takes 50–110 ms per spectrum and is the fastest method. Our new algorithm $MMSE_{BTG}$ is slower (approximately 1 s per spectrum) but still compatible with real-time monitoring. While slower than MAP_{L1}, $MMSE_{BTG}$ provides better estimates and allows automatic source detection and uncertainty quantification.

Discussion

RPMs must be able to detect weak SNM sources masked by a stronger NORM or nuisance radiation source. In this work, we overcame the limited energy resolution of organic scintillators by applying a new Bayesian algorithm to decompose and identify mixed gamma-ray sources. Bayesian algorithms proved to be useful tools to improve the source detection accuracy even with limited statistics (few counts) and poor signal-to-background ratios.

The proposed Bayesian $MMSE_{BTG}$ technique is designed to allow more accurate source identification and quantification in the presence of one or more masking nuclides, with cumulative count integrals as low as 500 counts. The automated identification obtained with the $MMSE_{BTG}$ method is more robust than using the MAP_{L1} and $MMSE_{L1}$ algorithms, which require unpractical parameter tuning. The main benefit of the proposed method is a more sensitive model that captures the sparsity of the mixing coefficients. The application of the $MMSE_{BTG}$ reduces, for instance, the average root-mean-square error between real and estimated nuclide fractions to 0.0177, compared to 0.0334 for MAP_{L1}, and 0.0584 for $MMSE_{L1}$ for the sixth mixture, containing ^{99m}Tc and WGPu, with only 1000 detection events. Our study also confirmed the importance of detector energy resolution. The stilbene crystal exhibits a better energy resolution than EJ-309 and, as a result, the stilbene data yielded a slightly better quantification accuracy, compared to EJ-309. Therefore, a slight improvement in the nuclide identification accuracy can be achieved by improving the energy resolution of the detector. Energy resolution improvement can be achieved either by using different materials, as we have shown in this study, and also by optimizing the detector's light collection geometry[20]. A relevant feature of organic scintillators is their sensitivity to both neutrons and gamma rays. Neutron and gamma-ray interactions in the organic scintillators are distinguishable through pulse

	EJ-309[39]	Stilbene[40,41]
Chemical Formula	n.a.	$C_{14}H_{12}$
Geometric isomerism	n.a.	Trans
H:C ratio	1.25	0.86
Density (g cm^{-3})	0.959	1.15
Light output (% anthracene)	80%	120%
Maximum wavelength (nm)	424	382
Scintillation efficiency (photons/1 MeVee)	12300	9760
Flash point (°C)	144	n.a.

Table 2. Physical properties of EJ-309 and stilbene detectors.

shape discrimination. The neutron signature was not used in this work but could be further exploited to aid the classification of fissile and other neutron emitting materials.

In this paper, we have applied new Bayesian algorithms for the identification of source mixtures that are not shielded. While this scenario applies to pedestrian portal monitors, it would be interesting to study the algorithm performance when sources are transported with other goods, or deliberately shielded. Effectively shielding SNMs and intense gamma-ray emitting radionuclides, such as ^{137}Cs and ^{60}Co, would require a combination of low- and high-atomic-number elements. The current algorithm could be enhanced by coupling it to spectra recon- struction methods that we have recently developed[21], to account for the spectral effect of shielding materials, given their known gamma-ray and neutron attenuation coefficients, as proposed by Lawrence and colleagues[22]. It should also be noted that containers carrying covert or overt amounts of metal are likely to prompt secondary inspections. For example, cargo containers carrying a large number of metal items typically undergo radiation inspection because orphan sources are often improperly disposed of as scrap metal and can be cast into metal parts[23]. Conversely, electro-magnetic inspection is performed on cargoes that are declared metal-free, and would promptly identify covert metal items.
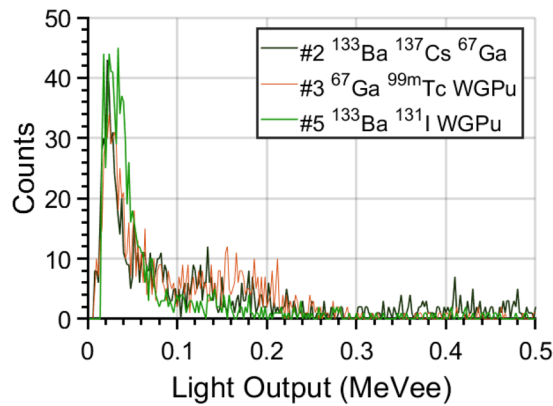
Methods

Over the past years, our group has developed several radionuclide identification algorithms for EJ-309-based portal monitors[6,8,12]. This work proposes a novel computational Bayesian method for source identification that we have applied to both liquid EJ-309 and solid-state trans-stilbene scintillators. In this section, we first detail how our measured data have been collected and then the principle of the new computational method.

Experimental methods.   We have used two detectors: an EJ-309 organic liquid scintillator (7.6-cm diameter by 7.6-cm height) by Eljen Technology, and a cylindrical trans-stilbene crystal (5.08-cm diameter by 5.08-cm height) produced using the solution-growth technique by Inrad Optics. The detection system used can be easily scaled up to be a pedestrian portal by using an array of detector cells. Despite the similar composition, EJ-309 and stilbene exhibit different properties (see Table 2). Noticeably, EJ-309 has a higher scintillation efficiency and higher density, compared to stilbene, which determines its higher intrinsic detection efficiency[24]. However, the stilbene crystal shows a favorable energy resolution, defined as the full width at half maximum (FWHM) of a spectrum peak in response to the energy deposited in the detector by monoenergetic charged recoils, divided by its centroid. This improved energy resolution can enhance isotope identification accuracy using stilbene over EJ-309. Note that the energy resolution of a scintillation detector is affected by both the scintillating material and the light collection and conversion process. The energy resolution at 478 keVee of stilbene and EJ-309 detectors of the same size as those used in this work is 9.64 ± 0.06[25] and 19.33 ± 0.18[26], respectively.

For completeness, Fig. 7 depicts the light output spectra of some of the mixtures analyzed, when approxi- mately 1000 counts were acquired. Despite the spectra consisting of different nuclides, their overall distribution as a function of light output is similar. This effect is due to the scatter-based detection of organic scintillators and the low counting statistics.

We measured a variety of sources, including ^{241}Am, ^{133}Ba, ^{57}Co and ^{137}Cs sources with activities of approxi- mately 500 kBq. The WGPu source (180 MBq) was measured at the Zero-Power Research Reactor of the Idaho National Laboratory[27]. In addition, 260 kBq liquid solution samples of the medical isotopes, i.e., ^{99m}Tc, ^{111}In, ^{67}Ga, ^{123}I, ^{131}I, and ^{201}Tl were measured at the University of Michigan C.S. Mott Children's Hospital.

The on-the-fly radionuclide identification algorithms used in this work rely on a library of nuclides that is assumed to include the species potentially present in the mixtures. The detection of unknown sources is out of the scope of this work and is left for future work. The isotope library used in this work consists of a collection of light output spectra acquired over one hour to reduce shot-noise effects. As the two detectors exhibit slightly different light responses, calibration was necessary to detect the same portion of the energy spectrum with both detectors. The detectors were gain-matched using a 3.3-MBq ^{137}Cs source, by aligning the ^{137}Cs Compton edge to 1.8 V in the pulse-height detector response. Lower and upper detection thresholds of 40 keV-electron-equivalent (keVee) and 480 keVee, respectively, were applied to both stilbene and EJ-309 detectors light output spectra. The electron-equivalent light output of a pulse in a scintillator, measured in electron-equivalent electron Volts, or eVee, refers to the energy required for an electron to produce a pulse with equivalent light output.

Figure 7. Comparison of light output spectra of mixtures 2, 3, and 5. The total number of counts in each distribution is approximately 1000.

Computational Method. *Bayesian estimation: competing methods.* Bayesian methods rely on exploiting the posterior distribution of variables of interest, by combining the observed data with additional prior information available about those variables. Here, we are interested in finding the coefficients associated with a set of nuclides. Numerous strategies have been proposed to solve this problem and, before introducing the proposed method, we first discuss the two methods used in[12], namely, the MAP_{L1} and the $MMSE_{L1}$ methods, to motivate the new $MMSE_{BTG}$ method.

Consider an observed spectral response $\mathbf{y} = [y_1, ..., y_M]^T$ observed in M non-overlapping energy bins ($M = 232$ for all the results presented here), which is associated with a mixture of up to N known sources whose individual spectral responses are denoted by $\{\mathbf{A}_{:,n}\}_{n=1,...,N}$ and gathered in the $M \times N$ matrix $\mathbf{A} = [\mathbf{A}_{:,1}, ..., \mathbf{A}_{:,N}] = [\mathbf{A}_{1,:}^T, ..., \mathbf{A}_{M,:}^T]^T$. Each $\mathbf{A}_{m,:}$ is a row vector gathering the spectral responses of the N known sources in the mth energy bin. Note that the spectral signatures are normalised such that they integrate to one and that this normalization has been performed using spectra measured with long integration times to reduce as much as possible shot-noise effects during the normalization. The amount/coefficient associated with the nth source is denoted by x_n and the N coefficients are gathered in the vector $\mathbf{x} = [x_1, ..., x_N]^T$. A classical approach to source separation is to assume as a first approximation, a linear mixing model which can be expressed in matrix/vector form as $\mathbf{y} \approx \mathbf{A}\mathbf{x}$. This model assumes that all the radiation sources present in the scene that are not included in the matrix $\mathbf{A}$ can be neglected. To avoid environment-dependent results, the background is neglected here. Our aim here is to study the nuclide identification and quantification in scenarios where the integration time is short and thus when the number of gamma detection events is low. In such cases, the observation noise corrupting each measurement can be accurately modeled by Poisson noise, leading to Poissonian form of the likelihood.

$$f(y_m|\mathbf{x}) = (\mathbf{A}_{m,:}\mathbf{x})^{y_m} \exp[-\mathbf{A}_{m,:}\mathbf{x}]/y_m!, \quad \forall m = 1, ..., M. \tag{2}$$

Since $\mathbf{A}$ is known, it is omitted in all the conditional distributions hereafter. Note that Eq. (2) implies that the sources present a fixed activity (or are static) during the integration time. In more complex scenarios, more complex models such as compound Poisson models might be used. Conditioned on the value of $\mathbf{x}$, the entries of $\mathbf{y}$ are independently distributed, i.e., $f(\mathbf{y}|\mathbf{x}) = \prod_{m=1}^{M} f(y_m|\mathbf{x}) = \prod_{m=1}^{M} f(y_m|\mathbf{A}_{m,:}\mathbf{x})$. Bayesian methods for spectral unmixing rely on additional prior information available about $\mathbf{x}$ to enhance its recovery from $\mathbf{y}$. Such methods formulate a priori information through a prior distribution $f(\mathbf{x})$ and the estimation of $\mathbf{x}$ can then be achieved using the posterior distribution $f(\mathbf{x}|\mathbf{y}) = f(\mathbf{y}|\mathbf{x})f(\mathbf{x})/f(\mathbf{y})$. The maximum a posteriori (MAP) estimate can be obtained by solving the following optimization problem

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmax}} f(\mathbf{x}|\mathbf{y}) \tag{3}$$

while the minimum mean squared error (MMSE) estimate, or posterior mean, can be obtained by computing the expectation $E_{f(\mathbf{x}|\mathbf{y})}[\mathbf{x}]$. Using a product of independent exponential prior distributions for $\mathbf{x}$, leads to a model that is based on an ℓ_1-norm penalty. This is the model used in our preliminary work[12]. In that work, we compared two approaches, namely, MAP estimation and MMSE estimation, leading to two algorithms, MAP_{L1} and $MMSE_{L1}$, respectively. It is important to mention that this choice of sparsity model is primarily motivated by the fact that the problem in (3) is convex and can be solved efficiently. While the $MMSE_{L1}$ algorithm is based on Markov chain Monte Carlo (MCMC) methods and allows the estimation of a posteriori confidence intervals (which are not directly available with the MAP_{L1} method), we showed[12] that the proportions estimated were generally worse than when using MAP_{L1}. This is primarily due to the fact that although exponential prior distributions promote sparse MAP estimates; this family of distributions is not sparsity promoting (it only tends to concentrate the mass of the distribution around the origin). Hence, the resulting probabilistic estimates, such as means or covariances are questionable[28]. This observation is also confirmed with the results in Fig. 2. Our previous study[12] also showed that

by constraining $K \leq N$ the maximum number of sources present in each mixture, it is possible to further improve the unmixing performance using MAP_{L1}. This improvement however comes at a high computational cost as it requires comparing all the possible partitions of K sources, out of N sources in the original spectral library. This becomes rapidly intractable as N increases. It also requires a level of supervision (to set K properly) which is incompatible with practical, real-time applications.

Alternative prior model for sparse mixtures. In this work, we first propose to use an alternative, more efficient, sparsity-promoting prior model for **x**. Precisely, we consider the following Bernoulli-truncated Gaussian (BTG) model

$$
\begin{aligned}
f(x_n|w_n) &= (1 - w_n)\delta(x_n) + w_n \mathcal{N}_{\mathbb{R}^+}(x_n; 0, \sigma_n^2), &&\forall n = 1, \ldots, N \\
f_n(w_n = 1) &= \pi_n, &&\forall n = 1, \ldots, N,
\end{aligned}
\tag{4}
$$

where $\delta(\cdot)$ denotes the Dirac delta function which is equal to 1 when $x_n = 0$ and 0 elsewhere and where $\mathcal{N}_{\mathbb{R}^+}(x_n; 0, \sigma^2)$ is a truncated Gaussian distribution, defined on $\mathbb{R}^+$ to enforce the non-negativity of the elements of **x**. Moreover, 0 and σ^2 are respectively the mean and variance of the Gaussian prior truncation. In Eq. (4), w_n is a binary variable which relates to the presence ($w_n = 1$) or absence ($w_n = 0$) of the nth source and the probability π_n is the prior probability of presence of the nth source. More precisely, the first line in Eq. (4) reduces to a mass at 0 enforcing $x_n = 0$ if $w_n = 0$ (source absent) and to a truncated Gaussian distribution if $w_n = 1$ (source present).

The joint prior model can then be expressed as $f(\mathbf{x}, \mathbf{w}) = \prod_{n=1}^{N} f(x_n|w_n)f_n(w_n)$ and the proposed unmixing algorithm aims at estimating jointly $(\mathbf{x}, \mathbf{w} = [w_1, \ldots, w_N]^T)$, i.e., at performing jointly the source identification (through **w**) and quantification (through **x**). Note that $\{\pi_n\}_n$ and $\{\sigma_n^2\}$ are assumed to be known here and can be used-defined. For the prior probabilities of presence, we set $\pi_n = 1/N$, $\forall n$ as we expect a limit number of sources to be simultaneously present in the mixture, while we do not wish to promote any specific source. While arbitrary large values could in principle be used for the variances $\{\sigma_n^2\}$, reflecting the lack of information about the activity of the sources to be detected, this strategy can lead to poor detection[29]. If the variances cannot be set from prior knowledge, an alternative approach, adopted here consists of adjusting it using the current observation, in an empirical Bayes fashion. Since the matrix **A** is normalised, the variances $\{\sigma_n^2\}$ should scale with the photon counts, provided that few sources are expected simultaneously in the mixture. In this work we set $\sigma_n^2 = 0.1 \sum_{m=1}^{M} y_m$ for each source and for all the results presented and did not observed unexpectedly poor detection results.

Using the Bayes' rule, the joint posterior distribution of $(\mathbf{x}, \mathbf{w})$ is given by $f(\mathbf{x}, \mathbf{w}|\mathbf{y}) = f(\mathbf{y}|\mathbf{x})f(\mathbf{x}, \mathbf{w})/f(\mathbf{y})$. Unfortunately, the posterior means $E_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}[\mathbf{x}]$ and $E_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}[\mathbf{w}]$ associated with this posterior distribution are intractable analytically and the traditional approach to exploit the posterior distribution consists of using a simulation method (as used in the MMSE_{L1} algorithm). In particular, constrained Hamiltonian Monte Carlo methods[30] have been investigated to solve regression problems in the presence of Poisson noise[12,31] (see also[32] for comparison of samplers). However, efficient sampling from $f(\mathbf{x}, \mathbf{w}|\mathbf{y})$ is very difficult due to the Poisson likelihood (2) coupled with the multimodality of the $f(\mathbf{x}, \mathbf{w}|\mathbf{y})$ induced by the joint model $f(\mathbf{x}, \mathbf{w})$. Indeed, adopting a Gibbs sampling strategy to sample iteratively from $f(x_n, w_n|\mathbf{y}, \mathbf{x}_{\backslash n}, \mathbf{w}_{\backslash n})$, where $\mathbf{w}_{\backslash n}$ contains all the elements of **w** but w_n, leads to poor mixing properties for the resulting Markov chain and thus prohibitively long chains. Similarly, block Gibbs samplers yield low acceptance rates and also poor mixing properties.

Proposed algorithm using variational inference. In this paper, we adopt an approximate Bayesian method and build an approximate distribution $Q(\mathbf{x}, \mathbf{w}) \approx f(\mathbf{x}, \mathbf{w}|\mathbf{y})$ whose moments are much simpler to evaluate than those of $f(\mathbf{x}, \mathbf{w}|\mathbf{y})$. In particular, for the identification of the nuclides present in a mixture, one important quantity is $E_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}[\mathbf{w}]$, the vector of marginal a posteriori probabilities of presence of each nuclide. For the quantification of the nuclides, interesting quantities are the posterior mean and covariance of **x**, i.e., $E_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}(\mathbf{x})$ and $\text{Cov}_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}(\mathbf{x})$. While the posterior mean is used as point estimate for the mixing coefficients, the posterior covariance matrix of **x** can be used to assess which sources are the most difficult to quantify. Here, we use the so-called expectation propagation (EP) method[33] to provide approximate point estimates, e.g., $E_{Q(\mathbf{x}, \mathbf{w})}(\mathbf{x}) \approx E_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}(\mathbf{x})$ and $E_{Q(\mathbf{x}, \mathbf{w})}[\mathbf{w}] \approx E_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}[\mathbf{w}]$, as well as approximations of the covariance of the posterior distribution of **x**, i.e., $\text{Cov}_{Q(\mathbf{x}, \mathbf{w})}(\mathbf{x}) \approx \text{Cov}_{f(\mathbf{x}, \mathbf{w}|\mathbf{y})}(\mathbf{x})$. While less well known than other Variational Bayes (VB) techniques, the EP has several recognized advantages[34]. It is particularly well suited to fast distributed Bayesian inference on partitioned data, giving it a high potential for real-time implementation.

The EP framework used for regression with Gaussian noise[35] and generalised linear models[36], approximates each exact factor $f(y_m|\mathbf{A}_{m,:}\mathbf{x}) = q_m(\mathbf{A}_{m,:}\mathbf{x})$ (resp. $f(x_n|w_n) = g_n(x_n, w_n)$) with a simpler factor $\tilde{q}_m(\mathbf{A}_{m,:}\mathbf{x})$ (resp. $\tilde{g}_n(x_n)\tilde{h}_n(w_n)$) so that

$$
\begin{aligned}
f(\mathbf{x}, \mathbf{w}|\mathbf{y}) &\propto \prod_{m=1}^{M} q_m(\mathbf{A}_{m,:}\mathbf{x}) \prod_{n=1}^{N} g_n(x_n, w_n)f(w_n) \\
&\approx \prod_{m=1}^{M} \tilde{q}_m(\mathbf{A}_{m,:}\mathbf{x}) \prod_{n=1}^{N} \tilde{g}_n(x_n)\tilde{h}_n(w_n)f(w_n) = Q(\mathbf{x}, \mathbf{w}),
\end{aligned}
\tag{5}
$$

where all the approximate factors belong to the same family of distributions. Here, in a similar fashion to the work by Hernandez-Lobato *et al.*[37], the approximate factors dependent on **x** are Gaussian and those associated with each w_n are discrete probabilities (see Fig. 8). This choice allows a more computationally attractive EP algorithm
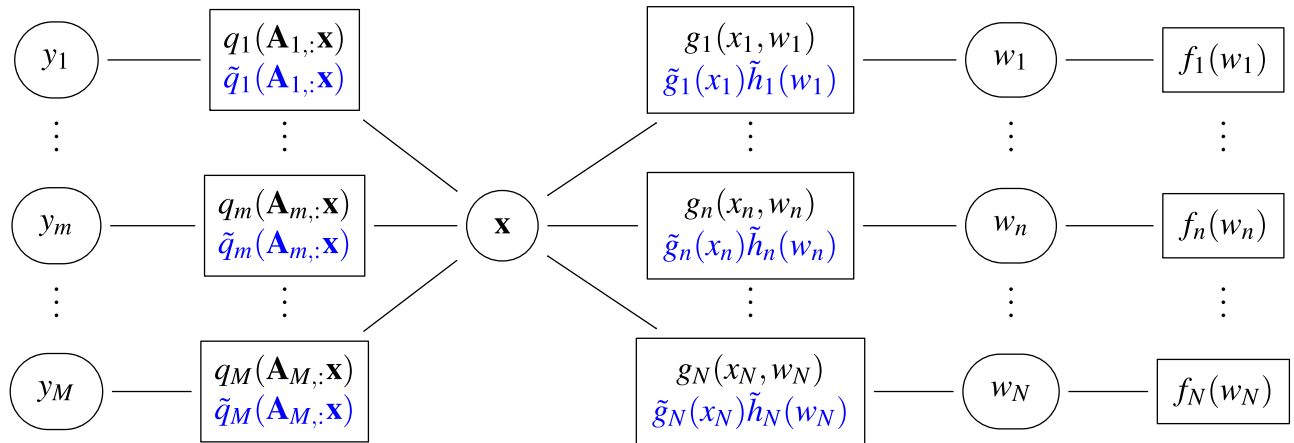
Figure 8. Factor graph used to perform EP-based sparse spectral unmixing. The circles (resp. rectangular boxes) represent the variable (resp. factor) nodes and the approximate factors are shown in blue.

and direct access to the moments of the posterior distribution. Moreover, it is important to note that using the splitting $f(x_n|w_n) \approx \tilde{g}_n(x_n)\tilde{h}_n(w_n)$, the approximate distribution $Q(\mathbf{x}, \mathbf{w})$ can be written $Q(\mathbf{x}, \mathbf{w}) = Q_x(\mathbf{x})Q_w(\mathbf{w})$, i.e., the approximation does not explicitly capture the correlation a posteriori between $\mathbf{x}$ and $\mathbf{w}$. Nonetheless, this type of separable approximation is classically used in variational inference and the parameters of $Q_x(\cdot)$ and $Q_w(\cdot)$ are in practice highly dependent. To optimize $Q(\mathbf{x}, \mathbf{w})$ so that $f(\mathbf{x}, \mathbf{w}|\mathbf{y}) \approx Q(\mathbf{x}, \mathbf{w})$, EP sequentially refines the factors $\{\tilde{q}_m(\mathbf{A}_{m,:}\mathbf{x})\}_m$ and $\{\tilde{g}_n(x_n), \tilde{h}_n(w_n)\}_n$ by minimizing the following Kullback-Leibler (KL) divergences

$$\begin{cases} \min_{\tilde{q}_m} KL(q_m(\mathbf{A}_{m,:}\mathbf{x})Q^{\backslash m}(\mathbf{x}, \mathbf{w}) \| \tilde{q}_m(\mathbf{A}_{m,:}\mathbf{x})Q^{\backslash m}(\mathbf{x}, \mathbf{w})), & \forall m = 1, \dots, M \\ \min_{\tilde{g}_n, \tilde{h}_n} KL(g_n(x_n, w_n)Q^{\backslash n}(\mathbf{x}, \mathbf{w}) \| \tilde{g}_n(x_n)\tilde{h}_n(w_n)Q^{\backslash n}(\mathbf{x}, \mathbf{w})), & \forall n = 1, \dots, N \end{cases} \tag{6}$$

where the so-called cavity distributions satisfy $Q^{\backslash m}(\mathbf{x}, \mathbf{w}) = Q(\mathbf{x}, \mathbf{w})/\tilde{q}_m(\mathbf{A}_{m,:}\mathbf{x})$ and $Q^{\backslash n}(\mathbf{x}, \mathbf{w}) = Q(\mathbf{x}, \mathbf{w})/(\tilde{g}_n(x_n)\tilde{h}_n(w_n))$. Solving the first row of Eq. (6) reduces to matching the mean and covariance of $Q_x(\mathbf{x})$ and of the so-called tilted distributions $\int q_m(\mathbf{A}_{m,:}\mathbf{x})Q^{\backslash m}(\mathbf{x}, \mathbf{w})d\mathbf{w}$, $\forall m$. In the work by Ko *et al.*[38], the authors showed that these problems can be solved analytically by computing sequentially one-dimensional integrals (see also[11] for additional details). The second row of Eq. (6) can be solved by using the method presented by Hernández-Lobato and colleagues[37]. Since the approximation of $g_n(x_n, w_n)$ is separable ($\tilde{g}_n(x_n)\tilde{h}_n(w_n)$), it is sufficient to compute the mean of w_n with respect to the tilted distribution $\int g_n(x_n, w_n)Q^{\backslash n}(\mathbf{x}, \mathbf{w})d\mathbf{w}$ as well as the mean and covariance of the tilted distribution $\int g_n(x_n, w_n)Q^{\backslash n}(\mathbf{x}, \mathbf{w})d\mathbf{w}$, which in turn reduces to computing the first and second-order moments of $\iint g_n(x_n, w_n)Q^{\backslash n}(\mathbf{x}, \mathbf{w})d\mathbf{w}d\mathbf{x}_{\backslash n}$, with respect to x_n. This last distribution can be shown to be a Bernoulli truncated Gaussian distribution whose moments can be computed analytically. Finally, in a similar fashion to the procedure proposed by Hernández-Lobato and colleagues[37], we used a damping strategy to reduce convergence issues. We fixed the damping factor to $\varepsilon = 0.7$ and did not observe convergence issues with this value. When the algorithm has converged, we obtain $Q_x(\mathbf{x})$ which is a multivariate Gaussian distribution, and $Q_w(\mathbf{w})$ which is a product of N independent Bernoulli distributions, whose parameters have been optimized via the EP algorithm such that $f(\mathbf{x}, \mathbf{w}|\mathbf{y}) \approx Q(\mathbf{x}, \mathbf{w})$ The approximate posterior mean and covariance matrix of x are given by the mean and covariance matrix of $Q_x(\mathbf{x})$, respectively. To compute the estimated mixture fractions in Eq. (1), from any estimated mixture coefficients $\hat{\mathbf{x}}$ (e.g., by MMSE_{BTG}, MAP_{L1} or MMSE_{L1}), we then consider $\hat{\mathbf{z}} = \hat{\mathbf{x}}/\|\hat{\mathbf{x}}\|_1$. The parameters of the Bernoulli distributions in $Q_w(\mathbf{w})$ provide the approximate marginal posterior probabilities of presence, for each source. Thus, the source identification can be preformed using $Q_w(\mathbf{w})$, without resorting to thresholding the estimated mixture coefficients. Choosing the most appropriate decision rule for the source identification based on the marginal posterior distribution ultimately reduces to choosing an acceptable threshold for the probability of presence. Here, we consider a detection when the probability of presence is larger than the probability of absence, effectively using a marginal MAP criterion. If costs associated with the probabilities of false alarm and misdetection are available for each source, similar decision rules can also be easily derived using the output of the proposed method, based on a minimum cost criterion instead of the marginal MAP criterion. However, the study of such decision rules is out of scope of this paper. The current version of the algorithm is available at the url: https://gitlab.com/yaltmann/sparse_unmixing_poisson_noise_ep.

References

1. IAEA. IAEA Incident and Trafficking Database (ITDB), http://www-ns.iaea.org/downloads/security/itdb-fact-sheet.pdf.
2. Kouzes, R. T. *et al.* Naturally occurring radioactive materials and medical isotopes at border crossings. In *IEEE Nuclear Science Symposium Conference Record* (2003).

3. United Nations Conference on Trade and Development. UNCTADSTAT, https://unctadstat.unctad.org/wds/TableViewer/tableView.aspx?ReportId=13321 (2017).
4. Kouzes, R. T. & Siciliano, E. R. The response of radiation portal monitors to medical radionuclides at border crossings. *Radiation Measurements* (2006).
5. PNNL. Radiation detectors at U.S. ports of entry now operate more effectively, efficiently. Tech. Rep. (2016).
6. Paff, M. G., Di Fulvio, A., Clarke, S. D. & Pozzi, S. A. Radionuclide identification algorithm for organic scintillator-based radiation portal monitor. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2017).
7. Geelhood, B. D. *et al*. Overview of portal monitoring at border crossings. In *IEEE Nuclear Science Symposium Conference Record* (2003).
8. Paff, M. G., Clarke, S. D. & Pozzi, S. A. Organic liquid scintillation detector shape and volume impact on radiation portal monitors. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2016).
9. Di Fulvio, A., Shin, T. H., Hamel, M. C. & Pozzi, S. A. Digital pulse processing for NaI(Tl) detectors. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2015).
10. Seeger, M. & Nickisch, H. Fast convergent algorithms for expectation propagation approximate bayesian inference. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 652–660 (2011).
11. Altmann, Y., Perelli, A. & Davies, M. E. Expectation-Propagation algorithms for linear regression with poisson noise: application to photon-limited spectral unmixing. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)* (Brighton, United Kingdom, 2019).
12. Paff, M. *et al*. Identification of mixed sources with an organic scintillator-based radiation portal monitor. *J. Nucl. Mater. Manag.* **46**, 48–57 (2018).
13. Tandon, P. *et al*. Detection of radioactive sources in urban scenes using Bayesian Aggregation of data from mobile spectrometers. *Information Systems* (2016).
14. Penny, R. D. *et al*. Improved radiological/nuclear source localization in variable NORM background: An MLEM approach with segmentation data. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2015).
15. Sokolova, M. & Lapalme, G. A systematic analysis of performance measures for classification tasks. *Information Processing and Management* (2009).
16. Kruse, F. A. *et al*. The spectral image processing system (SIPS)-interactive visualization and analysis of imaging spectrometer data. *Remote Sensing of Environment* (1993).
17. Yuhas, R., Goetz, A. & Boardman, J. Descrimination among semi-arid landscape endmembers using the spectral angle mapper (SAM) algorithm. In *Summaries of the Third Annual JPL Airborne Geoscience Workshop*, vol. 1, 147–149 (JPL, 1992).
18. Kamuda, M., Zhao, J. & Huff, K. A comparison of machine learning methods for automated gamma-ray spectroscopy. *Nucl. Instrum. Methods Phys. Research, Sect. A: Accelerators, Spectrometers, Detect. Associated Equip.* **954**, 161385 (2020).
19. Ludlum Measurements, Inc. *Model 375 P-33 6-1Monitoring System* (2019).
20. Sosa, C. *et al*. Energy resolution experiments of conical organic scintillators and a comparison with geant4 simulations. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accelerators, Spectrometers, Detect. Associated Equip.* **898**, 77–84 (2018).
21. Zhu, H. *et al*. A hierarchical bayesian approach to neutron spectrum unfolding with organic scintillators. *IEEE Trans. Nucl. Sci.* **66**, 2265–2274 (2019).
22. Lawrence, C., Febbraro, M., Flaska, M., Pozzi, S. & Becchetti, F. Warhead verification as inverse problem: Applications of neutron spectrum unfolding from organic-scintillator measurements. *Journal of Applied Physics* **120** (2016).
23. Rigoni Garola, A. Muon tomography effectiveness in detecting orphan sources in scrap metal. *Nuovo Cimento della Societa Ital. di Fis. C.* **37**, 155–163 (2014).
24. Pozzi, S. A., Clarke, S. D., Paff, M., Di Fulvio, A. & Kouzes, R. T. Comparative neutron detection efficiency in He-3 proportional counters and liquid scintillators. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2019).
25. Sosa, C. *et al*. Improved neutron–gammadiscrimination at low-light output events using conical trans-stilbene. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accelerators, Spectrometers, Detect. Associated Equip.* **916**, 42–46 (2019).
26. Enqvist, A., Lawrence, C. C., Wieger, B. M., Pozzi, S. A. & Massey, T. N. 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* (2013).
27. A., D.-F. *et al*. Passive assay of plutonium metal plates using a fast-neutron multiplicity counter. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accelerators, Spectrometers, Detect. Associated Equip.* **855**, 92–101 (2017).
28. Gribonval, R., Cevher, V. & Davies, M. E. Compressible distributions for high-dimensional statistic. *IEEE Trans. Inf. Theory* **58**, 5016–5034 (2012).
29. Klumpp, J. & Brandl, A. Simultaneous source detection and analysis using a zero-inflated count rate model. *Health Physics* **109** (2015).
30. Brooks, S. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods (Taylor & Francis, 2011).
31. Altmann, Y. *et al*. Robust spectral unmixing of sparse multispectral lidar waveforms using gamma Markov random fields. *IEEE Trans. Comput. Imaging* **3**, 658–670 (2017).
32. Tachella, J., Altmann, Y., Pereyra, M. & Tourneret, J.-Y. Bayesian restoration of high-dimensional photon-starved images. In *Proc. European Signal Processing Conf. (EUSIPCO)* (Rome, Italy, 2018).
33. Minka, T. P. Expectation propagation for approximate bayesian inference. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, 362–369 (Morgan Kaufmann Publishers Inc., 2001).
34. Vehtari, A. *et al*. Expectation propagation as a way of life: A framework for bayesian inference on partitioned data, https://arxiv.org/abs/1412.4869v4 (2014).
35. Seeger, M. W. Bayesian inference and optimal design for the sparse linear model. *J. Mach. Learn. Res.* **9**, 759–813 (2008).
36. Kim, A. & Wand, M. P. On expectation propagation for generalised, linear and mixed models. *Australian N. Zealand J. Stat.* **60**, 75–102 (2018).
37. Hernández-Lobato, J., Hernández-Lobato, D. & Suárez, A. Expectation propagation in linear regression models with spike-and-slab priors. *Mach. Learn.* **99**, 437–487 (2015).
38. Ko, Y.-J. & Seeger, M. W. Expectation propagation for rectified linear poisson regression. In *Asian Conference on Machine Learning*, vol. 45 of *Proceedings of Machine Learning Research*, 253–268 (Hong Kong, 2016).
39. Eljen Technology. NEUTRON/GAMMA PSD EJ-301, EJ-309, https://eljentechnology.com/products/liquid-scintillators/ej-301-ej-309.
40. Baker, J. H., Galunov, N. Z. & Tarasenko, O. A. Neutron scintillation detectors for environmental, security and geological studies. In *2007 IEEE Nuclear Science Symposium Conference Record*, vol. 2, 1358–1364 (2007).
41. Inrad Optics. Scintinel™ Stilbene, https://www.inradoptics.com/scintinel-stilbene.

Author contributions

Y.A. developed the Bayesian algorithm and analyzed the data, A.D.F. conceived the general methodology, executed the experiments and analyzed the data, M.P. conceived and executed the experiments, S.P. and S.C. conceived the experiments and oversaw the project, A.H., M.D. and S.M. gave guidance for the algorithm development. Y.A., A.D.F. and M.P. wrote the manuscript. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to A.D.F.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.