

KamNet: An integrated spatiotemporal deep neural network for rare event searches in KamLAND-Zen^{*}

A. Li ^{1,†}, Z. Fu ^{1,2}, C. Grant, ³ H. Ozaki ^{1,4}, I. Shimizu ^{1,4}, H. Song, ³ A. Takeuchi, ⁴ and L. A. Winslow ^{1,2}

¹*Department of Physics and Astronomy, University of North Carolina, Chapel Hill, North Carolina 27514, USA*
and Triangle Universities Nuclear Laboratory, Durham, North Carolina 27708, USA

²*Laboratory of Nuclear Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

³*Department of Physics, Boston University, Boston, Massachusetts 02215, USA*

⁴*Research Center for Neutrino Science, Tohoku University, Sendai 980-8578, Japan*



(Received 6 March 2022; accepted 14 September 2022; published 30 January 2023)

Rare event searches allow us to search for new physics at energy scales inaccessible with other means by leveraging specialized large-mass detectors. Machine learning provides a new tool to maximize the information provided by these detectors. The information is sparse, which forces these algorithms to start from the lowest level data and exploit all symmetries in the detector to produce results. In this work we present KamNet, which harnesses breakthroughs in geometric deep learning and spatiotemporal data analysis to maximize the physics reach of KamLAND-Zen, a kiloton scale spherical liquid scintillator detector searching for $0\nu\beta\beta$. Using a simplified background model for KamLAND, we show that KamNet outperforms a conventional convolutional neural network (CNN) on benchmarking Monte Carlo simulations with an increasing level of robustness. Using simulated data, we then demonstrate KamNet's ability to increase KamLAND-Zen's sensitivity to $0\nu\beta\beta$ and $2\nu\beta\beta$ decay to excited states. A key component of this work is the addition of an attention mechanism to elucidate the underlying physics KamNet is using for the background rejection.

DOI: 10.1103/PhysRevC.107.014323

I. INTRODUCTION

Rare event searches provide a unique window on processes happening at energy scales beyond those currently accessible with accelerators up to and including the GUT scale. These experiments do this with highly customized detectors to reduce background and large masses to maximize exposure. By their nature the data coming from these detectors is sparse and algorithms for analyzing this data must maximize the available information. This is a natural application for machine learning but a different optimization from the big data applications that have been the main focus of the field.

Monolithic kiloton-scale liquid scintillator (LS) detectors like KamLAND exemplify this approach and have been the work horse of neutrino physics for many decades [1–9]. In the first phase of KamLAND, 1 kiloton of LS is contained in a 13-m-diameter balloon and this LS-filled balloon surrounded by mineral oil (acting as a buffer volume) is viewed by 1879 photomultiplier tubes (PMTs). The KamLAND-Zen experiment inherits the infrastructure of the KamLAND detector and deploys 24 tons of Xe-loaded liquid scintillator (XeLS) in a 3.80-m-diameter spherical inner-balloon at the center of the KamLAND detector.

This inner-balloon currently contains XeLS doped with 745 ± 3 kg of Xe to search for neutrinoless double-beta

decay ($0\nu\beta\beta$) [10]. Figure 1 (center) shows the KamLAND-Zen detector.

The first observation of $0\nu\beta\beta$ would prove that the neutrino is its own antiparticle, also known as a Majorana particle. This is a key ingredient for Leptogenesis [11], which describes the observed matter-antimatter asymmetry in our universe. The challenge of searching for $0\nu\beta\beta$ decay with monolithic LS detectors is that they function primarily as calorimeters and lack the sophisticated tracking and topological information provided by other technologies like Time Projection Chambers, Cherenkov ring imaging detectors, or silicon strip trackers. Enhancing monolithic LS detectors with the capability to discriminate between different event types based on tracking and topology would be a revolutionary advancement.

In this work, our goal is to provide a template for developing ML that makes use of all detector symmetries and information encoded in the low-level data. A key component of this work is the development of tools to interrogate the algorithm to discover which data are key to its performance. The resulting algorithm is KamNet. It is built on our experience with a conventional convolutional neural network (CNN) [12]. However, the conventional CNN lacks the ability to harness certain symmetries embedded in KamLAND-Zen data which suppresses the performance of neural network.

This state-of-the-art machine learning algorithm harnesses the symmetry of a spherical detector to discriminate between topological differences in energy deposits. With the power to discriminate between different event topologies, we then apply KamNet to different studies from monitoring data

^{*}Code Repository: <https://github.com/aobol/KamNet.git>

[†]Corresponding author: liaobo77@ad.unc.edu

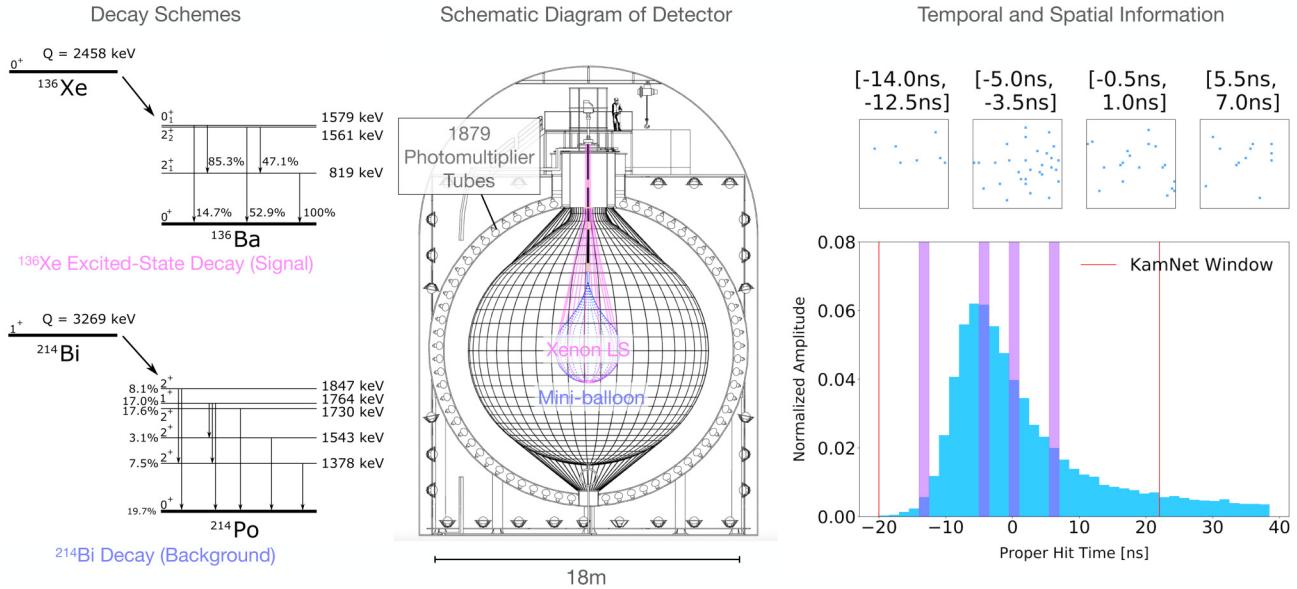


FIG. 1. Left: The decay schemes for ^{136}Xe and ^{214}Bi with branching ratios $<1\%$ are omitted for simplicity. Center: The schematic diagram of the KamLAND-Zen detector. Right: The distribution of PMT hit times for typical physics events. The spatial distribution of the PMT hit times highlighted in violet are shown above.

quality to enabling new analyses and, of course, reducing backgrounds in the $0\nu\beta\beta$ analysis. The paper is structured as follows. The first section explains the construction of the algorithm: Section II describes the low-level information encoded in the KamLAND-Zen data, Sec. III describes the different simulations used, and Sec. IV describes the network design. Section V outlines the application to the $0\nu\beta\beta$ analysis. We highlight the ability to interpret the attention mechanism's output in KamNet, thus helping to unveil what underlying physics is driving the discrimination power. Section VI introduces a future analysis that uses KamNet to extract the $2\nu\beta\beta$ decay to excited states signal for the first time. Finally, in Sec. VA we highlight KamNet's ability to self-interpret its decision via an attention mechanism, thus unveiling what underlying physics is driving the discrimination power. We highlight Sec. VA which presents the application of an attention mechanism to determine what low-level information is driving the discrimination power of KamNet.

II. KAMLAND-ZEN DATA

The level diagram for double- β decay and a common background the single β decay of ^{214}Bi are shown in Fig. 1 (left). These decays can proceed directly to the ground state where the β or β 's carry away all of the energy or via an excited state where one or more deexcitation γ 's are emitted. In LS, β 's deposit their energy in a very localized region, resulting in highly isotropic scintillation light. We define this type of event as a single vertex event. In comparison, γ 's can scatter multiple times with a mean free path in the scintillator on the order of ~ 10 cm. This multisite energy deposition results in a slightly less isotropic scintillation light emission and thus we define this type as closely spaced multivertex events.

Each event depositing energy in the XeLS (or LS) produces isotropic scintillation light accompanied by a relatively small

amount Cherenkov light. These photons propagate throughout the detector volume and eventually register hits on a subset of the 1879 PMTs. Among them, 1325 PMTs are 17-inch PMTs with fast timing resolution and contribute 22% of photo-coverage; the rest are 20-inch PMTs with relatively poorer timing resolution and contribute an additional 12% photo-coverage. In this way, the raw data for a single event is composed of PMT hit times and a spherical map of the positions of hit PMTs, see Fig. 1 (right). The raw data are sparse, but they encode key information about the underlying physics that can be utilized more efficiently by KamNet compared to traditional cut-based data analyses to disentangle strictly single-vertex events and closely spaced multivertex events in LS detectors.

III. SIMULATION PRODUCTION

Three different detector Monte Carlo (MC) simulations are used to study the performance of KamNet. Two of the simulations are written using the Reactor Analysis Toolkit (RAT) [13]. RAT is a simulation and analysis package that acts as an interface to GEANT4 [14,15]. The first is a simple KamLAND-Zen 400 simulation, referred to as sim-Fast, which is used for very fast benchmarking studies described in Sec. IVC. The second simulation is based on RAT, referred to as sim-RAT, which contains a more detailed model of the KamLAND-Zen 800 detector and is used to accurately characterize KamNet's classification performance for ^{136}Xe excited-state decays in KamLAND-Zen 800. Finally, the performance of KamNet on $0\nu\beta\beta$ is studied using the official KamLAND-Zen 800 detector MC, which is written purely in GEANT4 and has been carefully tuned to replicate the response of the real detector. The official simulation for KamLAND-Zen 800, referred to as sim-KLZ800, is quite resource intensive compared to sim-RAT, which is why it has

not yet been used for the excited-state study. However, we observe similar classification performance when KamNet is applied to sim-RAT and sim-KLZ800, so we have strong reason to believe the KamNet results for the ^{136}Xe excited-state decays in sim-RAT will carry over when it is applied to the real KamLAND-Zen 800 detector data.

Both sim-KLZ800 and sim-RAT incorporate the detailed geometry of KamLAND-Zen 800 from the innermost XeLS miniballoon out to the 1879 PMTs encased inside the 18-m-diameter stainless steel sphere as shown in the center of Fig. 1. The optical properties of all the inner detector materials, such as transparency, index of refraction, and reflectivity, are chosen to match those of the KamLAND-Zen detector. The scintillation and Cherenkov photon production is modeled using bench-top measurements of the index of refraction, light yield, absorption, re-emission, and quenching of XeLS and LS as input. After production, all optical photons are allowed to propagate throughout the inner detector volume until they are absorbed. The one difference between the simulations is that sim-KLZ800 has turned off Cherenkov photon production to increase performance without decreasing the MC-data agreement.

PMT photocathodes that absorb optical photons, accompanied by the production of one or more photoelectrons, are referred to as being “hit.” The time at which a PMT is hit by a photon, T_{raw} , is measured from the first PMT hit. In KamLAND-Zen, two corrections are applied to the raw hit time of each PMT to give what is called the proper hit time τ . The proper hit time of each PMT is calculated as follows:

$$\tau = T_{\text{raw}} - \text{TOF} - T_0, \quad (1)$$

where TOF is the photon time-of-flight from the event vertex to PMT position and T_0 is referred to as the proper start time. The event vertex, used to calculate the TOF, is reconstructed using the standard centroid fitter in RAT. By subtracting TOF from T_{raw} , we effectively move the vertex of each event to the center of the detector and correct for intraevent distortion of the scintillation time profile by the vertex position. The correction for intraevent distortion of the time profile due to varying energy deposits comes from subtracting T_0 , which is a fractional charge weighted sum of the differences between T_{raw} and TOF over all the PMTs. This is calculated as follows:

$$T_0 = \frac{\sum_i (T_{\text{raw}}^i - \text{TOF}^i) Q_i}{\sum_i Q_i}, \quad (2)$$

where Q_i is the charge on the i th PMT. The proper hit time is calculated inside an interval from -20 to 22 ns and is binned in increments of 1.5 ns in order to match the sampling time interval of the KamLAND-Zen readout electronics. Each proper hit time bin contains PMT hits that register photoelectron charge to a 38×38 spherical grid segmented in the azimuthal and polar angles, θ and ϕ . The resulting dimensions of the collection of spatiotemporal hit maps is $28 \times 38 \times 38$ in t , θ , and ϕ , respectively. A few slices of a typical spatiotemporal hit map, where the spherical hit maps are projected onto a two-dimensional (2D) plane, are shown in the rightmost plot of Fig. 1.

The energies, times, directions, and angular distributions of particles emitted during $2\nu\beta\beta$ ^{136}Xe decay to ground state, $2\nu\beta\beta$ ^{136}Xe decay to excited states ($0^+ \rightarrow 0_1^+$, $0^+ \rightarrow 2_1^+$, and $0^+ \rightarrow 2_2^+$), and other radioactive decay backgrounds, are simulated using the standard Decay0 event generator [16]. The simulated events are distributed uniformly within the roughly 3.8-m-diameter miniballoon spherical volume at the center of KamLAND-Zen detector.

IV. NETWORK DESIGN

In our first attempt to apply machine learning to a KamLAND-like detector, we used a conventional CNN to study its ability to reject muon spallation background. We achieved a rejection efficiency of 61% for ^{10}C while preserving 90% of ^{136}Xe [12]. However, the conventional CNN was originally designed for 2D planar images. Since KamLAND-Zen data are effectively a time series of spherical images, the conventional CNN lacks the ability to harness certain embedded symmetries which suppresses the neural network performance. In this paper, we fully redesigned our deep learning model to recognize the additional symmetries. The new model is referred to as KamNet since it is originally designed for KamLAND-Zen.

KamNet builds on our initial work and a complementary recurrent neural network (RNN) algorithm that was also designed for KamLAND-Zen [17]. It is inspired by two recent breakthroughs in geometric deep learning-based spherical analysis: S2CNN [18] and ConvLSTM [19]. S2CNN uses a group theory approach to handle spherical data input and ConvLSTM provides a mechanism for understanding time correlations among images.

A. Rotational symmetry

The conventional CNN scans rectangular filters throughout 2D images in a translation-invariant manner, and the translational invariance makes it highly efficient in analyzing planar images. KamLAND-Zen produces spherical PMT hit maps not rectangular images. Mapping the spherical surface of the detector image onto a 2D planar grid will necessarily introduce a space-varying distortion, which makes the translational weight sharing ineffective.

The spherical CNN [18] is introduced to address this issue. This model uses a group theory approach to approximate the $\text{SO}(3)$ rotational symmetry. In the spherical CNN, the input images undergo a spherical Fourier transform and the kernels undergo an $\text{SO}(3)$ Fourier transform. Both Fourier transforms outputs a 3D representation tensor in Euler angle space. The convolution is performed by multiplication within the Euler angle space to produce a feature map. Each cell in the feature map represents a global convolution between the entire image and filter viewed at a given Euler angle. By scanning the range of all possible rotation angles, the spherical CNN is able to harness the rotational invariance of KamLAND-Zen data.

B. Temporal symmetry

The other key symmetry is temporal symmetry. The input data contain a series of spherical hit maps segmented by time. In the conventional CNN, and spherical CNN, multiple input images are treated as channels. This treatment does not preserve the order of input, and it also does not properly handle the short- and long-range temporal correlations. A convolutional LSTM (ConvLSTM) model is introduced to resolve this issue [19].

Long short-term memory (LSTM) is a classical recurrent neural network model for time-series analysis [20]. The LSTM layer contains a hidden state and a cell state to store short-term and long-term information, respectively. The ConvLSTM [19] combines LSTM with a convolution to analyze the time-ordered images, or spatiotemporal data. The convolution operation allows the algorithm to understand the spatial correlation within each image, while the LSTM structure allows it to understand short-term and long-term correlations between images at different times.

Generally, the ConvLSTM is sufficient for handling temporal symmetry in a dataset, but the inclusion of the ConvLSTM into the spherical CNN model is not so straightforward. The problem originates from the timing profile of scintillation light shown in Fig. 1. Most of the PMT hits are registered within a few spherical hit maps near the scintillation peak, while all other hit maps contain few or zero PMT hits. Therefore, hit maps near the scintillation peak carry much more information than those far from the peak and more weight should be put on the hit maps containing most of the data. In order to accomplish this, an attention mechanism [21] is added to ConvLSTM layers and, as far as we know, the method described below is a unique development among the deep learning community.

When a series of spherical hit maps $I_{\text{in}}(t, \theta, \phi)$ are fed into the ConvLSTM layers, we obtain two tensors. The first is the collection of output images of the intermediate hidden state, denoted $I_{\text{hidden}}(c, t, \theta, \phi)$, and the second is the final output state of the ConvLSTM, denoted $I_{\text{output}}(c, t, \theta, \phi)$. The time indices are denoted by t and the channels of image are denoted by c . Next, the $I_{\text{output}}(c, t, \theta, \phi)$ tensor is expanded with a singleton in the time dimension and is used to calculate the attention score $S(t)$. The attention score is a weighting factor calculated for each time index as:

$$S(t) = \text{Softmax}[I_{\text{hidden}}(c, t, \theta, \phi) \circ W(c, t, \theta, \phi) \circ I_{\text{output}}(c, t, \theta, \phi)], \quad (3)$$

where $W(c, t, \theta, \phi)$ is the attention weight matrix learned during training. The multiplication of the three tensors $I_{\text{hidden}}(c, t, \theta, \phi)$, $W(c, t, \theta, \phi)$, and $I_{\text{output}}(c, t, \theta, \phi)$ in Eq. (3) is performed element-wise in a manner equivalent to the Hadamard product of matrices, and the Softmax function is performed along the time dimension. Finally, a context tensor $I_{\text{context}}(c, \theta, \phi)$ is obtained by:

$$I_{\text{context}}(c, \theta, \phi) = \sum_t S(t) I_{\text{hidden}}(c, t, \theta, \phi). \quad (4)$$

Thus, the input spatiotemporal hit maps $I_{\text{in}}(t, \theta, \phi)$ indexed by time are converted to context images $I_{\text{context}}(c, \theta, \phi)$ indexed

by channel. The content in each channel is controlled by the attention score $S(t)$, which allows the machine to zoom in to important time slices and leave out less important ones. In other words, each channel of the context image is a 2D representation of the spatiotemporal hit map, taking into consideration the temporal symmetry among different hit maps. The model that produces context images is called the AttentionConvLSTM layer, a schematic diagram of this layer is shown in Fig. 2(a).

After combining the AttentionConvLSTM layer and spherical CNN, we obtain the full KamNet model shown in Fig. 2(b). The input hit maps are first processed by the AttentionConvLSTM layer to produce 32 context images. The context images are then analyzed by the spherical CNN to produce feature vectors. Eventually, the feature vectors are classified by a fully connected network to produce a KamNet score for each event. The input event is more signal-like if KamNet assigns a high attention score, and vice versa.

C. Benchmarking dataset

In order to provide a direct comparison of the improved performance of KamNet over the conventional CNN, we use the dataset in our previous work [12] as a benchmark. As described earlier in Sec. III, sim-Fast is used to produce this benchmark dataset and is a simplified version of KamLAND-Zen 400. In order to speed up the simulation process, sim-Fast employs a gray-disk model for the PMTs. Both ^{136}Xe $0\nu\beta\beta$ and ^{10}C events are uniformly distributed within a 3-m-diameter miniballoon, which is surrounded by a 13-m-diameter LS balloon and a 2.5-m-thick mineral oil buffer volume. Photons are propagated throughout the XeLS miniballoon, LS balloon, and outer mineral oil buffer volume. Photons that reach the outer boundary of the mineral oil buffer and pass through the gray-disk PMTs are either accepted or rejected based on a uniform quantum efficiency (QE) applied across the gray disk.

The gray-disk PMT model allows us to vary the detector's QE and the photocoverage, affecting the number of scintillation photons collected. At high QE and the photocoverage, the increased number of scintillation photons provides more information to KamNet, leading to an easier (lower-pressure) classification task; conversely, low QE and photocoverage lead to a more challenging (higher-pressure) classification task. Therefore, we define QE and photocoverage as the two pressure parameters and vary them to generate 99 trials of training datasets. The performance of the deep learning models is demonstrated by training them simultaneously on the 99 trials and displaying the ^{10}C background rejection efficiencies at 90% ^{136}Xe $0\nu\beta\beta$ signal acceptance over pressure maps. The resulting pressure maps of the conventional CNN and KamNet are shown in Fig. 3.

As seen in Fig. 3, KamNet clearly outperforms the conventional CNN over all pressures. In terms of rejection efficiency, KamNet rejects 7.6% more ^{10}C events which is obtained by averaging the difference over the entire pressure map. At a photocoverage of 22% and quantum efficiency of 23%, which is a similar representation of the real KamLAND-Zen detector configuration, KamNet rejects 74.0% ^{10}C events compared

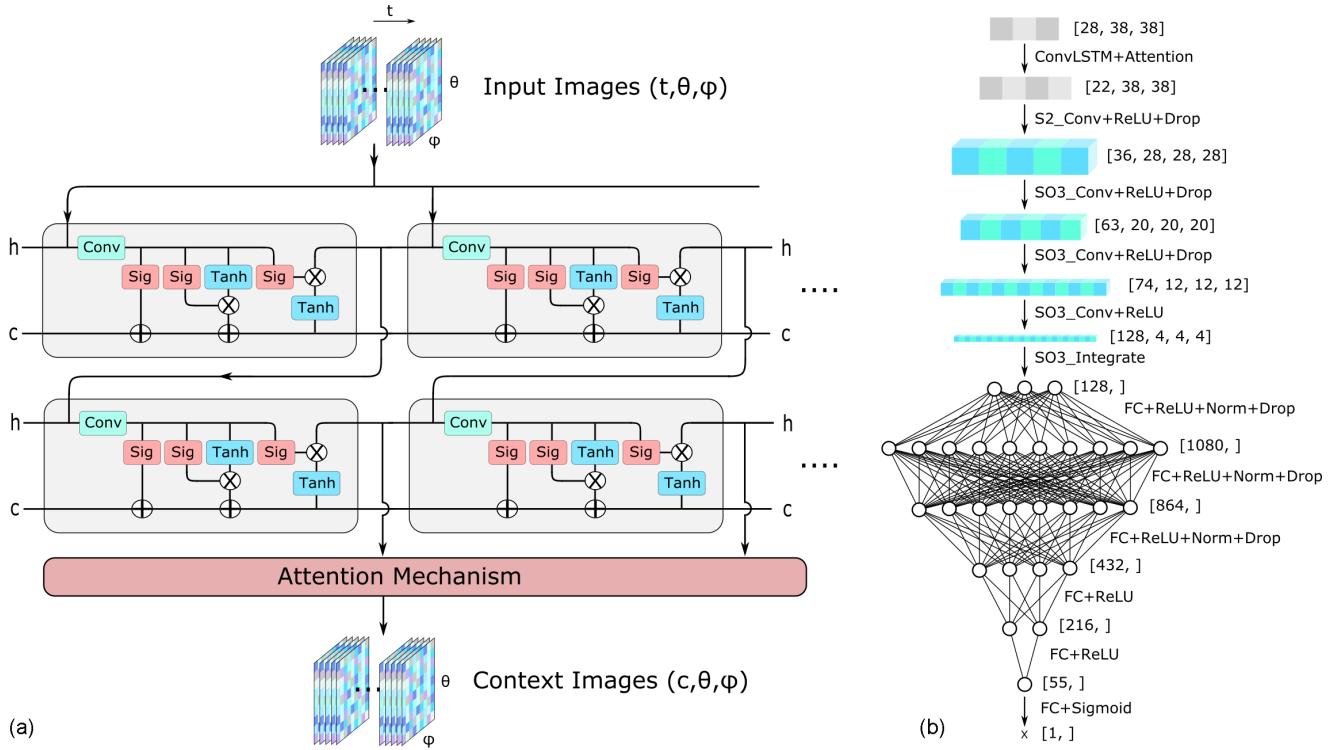


FIG. 2. (a) Schematic diagram of the AttentionConvLSTM layer, including two ConvLSTM layers and attention mechanism. (b) Diagram of KamNet.

to 61.5% for the conventional CNN. An ablation test was conducted to quantify the contributions of different parts of KamNet. With the AttentionConvLSTM module removed, KamNet's rejection efficiency drops from 74% to 71%, still 9.5% above conventional CNN. Therefore, the Spherical CNN provides the largest increase in performance. At the same time, the AttentionConvLSTM provides an additional performance boost while allowing us to interpret the attention score.

As we increase photocoverage and quantum efficiency, the increased number of detected photons will bring in more information that deep learning models can leverage to make classification decisions. A robust neural network model is expected to have monotonically increased performance correlated with the additional information, thereby producing a continuous improvement from the lower left to the upper right of the pressure map. As observed in Fig. 3, KamNet produces a much smoother transition across the pressure map, indicating that KamNet is much more robust to variation of the input data. We believe this stems from KamNet's ability to harness the spherical and spatiotemporal symmetries inherent in the data, while the conventional CNN appears to only be focusing on timing discrepancies.

D. Model comparison

The comparison between KamNet and the conventional CNN model was described in Sec. IV C. In addition to these two models, there are others that use scintillation time profiles to reject backgrounds in LS detector, namely the KamLAND-Zen RNN model [17] and a likelihood model applied on the SNO+ experiment [22]. These models use less information

than KamNet, since PMT hit location is ignored by collapsing 3D spatiotemporal data into 1D temporal data. The performance of the RNN model is compared to KamNet and conventional CNN on the benchmarking dataset. At a photo-coverage of 22% and quantum efficiency of 23%, the RNN model only rejects 52.3% of ^{10}C , underperforming both KamNet (74%) and conventional CNN (61.5%). Comparison to the likelihood method applied on SNO+ is not as straightforward since SNO+ and KamLAND-Zen have different hardware configurations. An indirect comparison can be made using Fig. 3(b), where KamNet rejects 87% ^{10}C at SNO+ photo-coverage of $\sim 50\%$ and quantum efficiency of 23%. Reference [22] claims a 75% rejection on ^{60}Co with negligible sacrifice, under the assumption that faster PMTs are used in a future SNO+ upgrade.

V. $0\nu\beta\beta$ DECAY

We now use the precisely tuned sim-KLZ800 to train KamNet for the $0\nu\beta\beta$ decay analysis. This study uses events in the energy window 2.0–3.0 MeV with radii $R < 157$ cm from the center of the XeLS miniballoon. The signal training set is $0\nu\beta\beta$ decay to ground-state events, and the background training set is ^{214}Bi events. We find that the training of KamNet on just one background produces nearly identical classification power compared to a training set formed from a mixture of all the simulated backgrounds. The choice of ^{214}Bi over other backgrounds is motivated by the fact that it is possible to extract a pure ^{214}Bi dataset from the detector using the prompt-delayed coincidence tagging of ^{214}Bi - ^{214}Po decays.

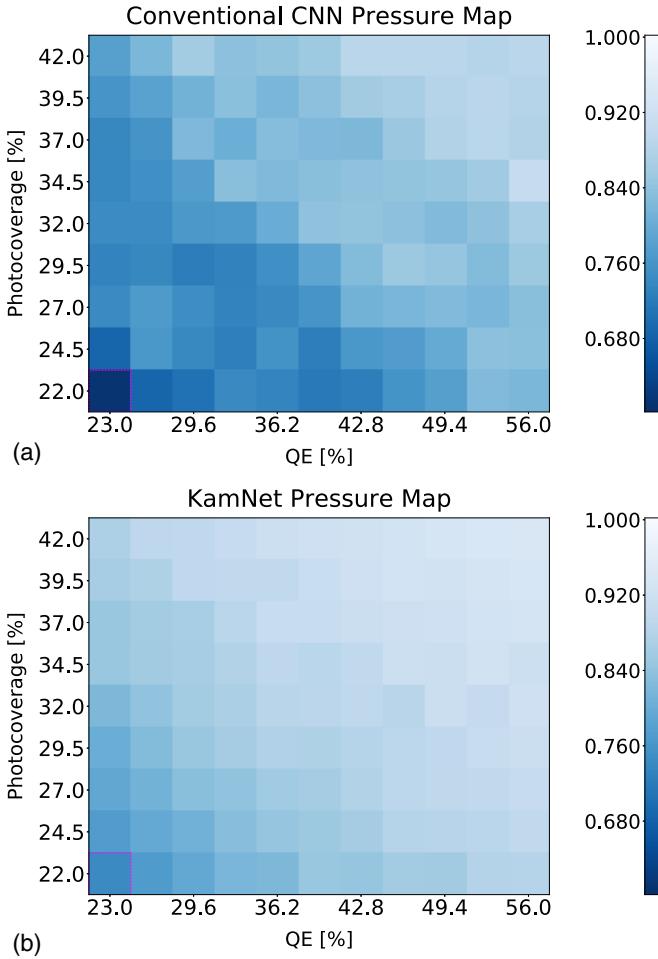


FIG. 3. (a) Pressure map produced by conventional CNN described in Ref. [12]. (b) Pressure map produced by KamNet on the same dataset. The values in each cell are the ^{10}C rejection efficiency at 90% ^{136}Xe acceptance. The lower left cell (circled in dashed red) represents the KamLAND-Zen detector configuration.

To check the MC-data agreement, we study signal-like $2\nu\beta\beta$ events and backgroundlike events in data and MC. For the signal-like $2\nu\beta\beta$ events, we select events from 1.7 to 2.2 MeV within the 157-cm radius, a region dominated by $2\nu\beta\beta$ events. For the background-like events, we use ^{214}Bi - ^{214}Po tagged events from 2.0 to 3.0 MeV within the 157 cm radius. The selection criteria are identically applied to both data and MC. KamNet is then applied to the MC and data events. The resulting KamNet scores are shown in Fig. 4. Excellent agreement is observed between the data and MC spectrum shape.

To quantify the MC-data agreement, we define a cut at the KamNet score corresponding to 90% $2\nu\beta\beta$ -MC acceptance. At this same KamNet score, the corresponding $2\nu\beta\beta$ data has a slightly smaller acceptance of 89.3% due backgrounds in the data. This is consistent with the results in this region presented in Ref. [1]. The acceptance of ^{214}Bi data events are 69.5% compared to 72.2% of ^{214}Bi MC events. ^{214}Bi data events are rejected more efficiently because 14.1% of the events originated on the balloon film. Film ^{214}Bi are rejected much

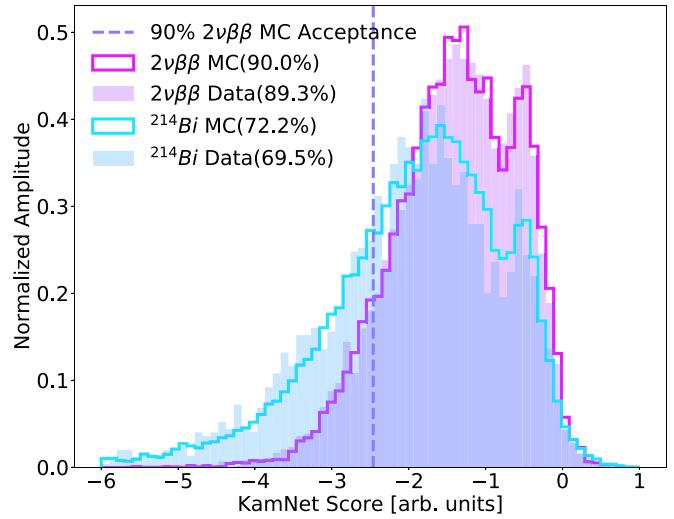


FIG. 4. KamNet score spectrum for MC and data of signal events (^{136}Xe $2\nu\beta\beta$) and background events (^{214}Bi). The number in parentheses shows the survival percentage if we make a cut at 90% ^{136}Xe MC acceptance (purple dashed line).

more efficiently by KamNet, as shown in Table I. Adding 14.1% of film ^{214}Bi into the ^{214}Bi MC will reduce the data-MC difference to only 1.7%.

A bootstrapping-based evaluation is conducted to quantify KamNet's statistical and systematic uncertainties. We generated 78 bootstrapped datasets and trained 78 bootstrapped KamNets to study the variation of output. Compared to 90% signal acceptance in base KamNet, the Gaussian-fitted acceptance of all bootstrapped KamNet reads $89.7\% \pm 0.58\%$, corresponding to a one-sigma agreement and 0.58% statistical uncertainty. We also conducted a series of systematic uncertainty studies involving energy, position, and data-MC comparisons. The combined statistical and systematic

TABLE I. Result of trained KamNet classifier on $0\nu\beta\beta$ analysis. The second column of the table indicates the type of decay the isotope undergoes. e^- indicates a strictly single vertex, β like events. $\beta^\pm + \gamma$ indicates a β decay with γ cascade, and LL stands for long-lived spallation backgrounds. The KamNet rejection percentage is evaluated at 90% signal acceptance (or equivalence speaking, 10% signal rejection). ROC AUC is the area under curve of receiver operating characteristic curve [23], higher AUC under ROC curve indicates the better separation between signal and background.

Isotopes	Type	ROC AUC	Rejection
$0\nu\beta\beta$	Signal	0.5	10%
Solar ν	e^-	0.49	9.5%
^{10}C	$\beta^+ + \gamma$	0.72	40.0%
^{214}Bi XeLS	$\beta^- + \gamma$	0.65	27.0%
^{214}Bi film	$\beta^- + \gamma$	0.83	58.8%
^{118}Sb	$\beta^+ + \gamma$, LL	0.59	18.3%
^{122}I	$\beta^+ + \gamma$, LL	0.61	22.2%
^{124}I	$\beta^+ + \gamma$, LL	0.67	30.6%
^{130}I	$\beta^- + \gamma$, LL	0.67	27.2%
^{132}I	$\beta^- + \gamma$, LL	0.66	28.5%

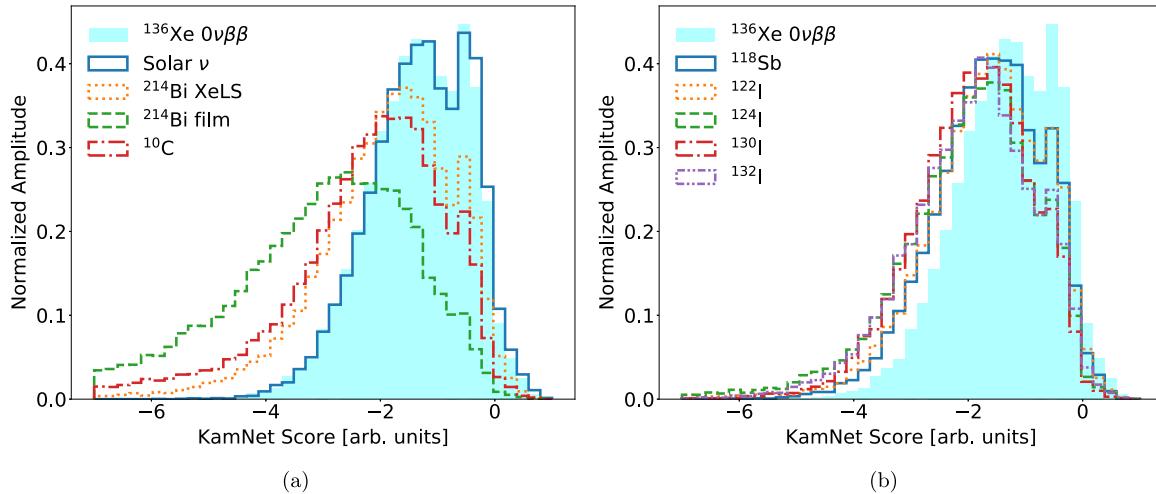


FIG. 5. (a) KamNet score spectrum for common backgrounds in KamLAND-Zen 800, including solar neutrino, ^{214}Bi , and ^{10}C backgrounds. (b) KamNet score spectrum for dominant long-lived spallation backgrounds in energy ROI. All histograms have been normalized to unity. Except for Solar ν , all backgrounds has lower KamNet score compared to ^{136}Xe , and thus they can be efficiently rejected by making cut on KamNet score.

uncertainty is roughly $\pm 3\%$ on the background rejection efficiency. This overall uncertainty is included in the upstream data analysis [1].

A. Network interpretability

The rejection power of KamNet comes from distinguishing strictly single-vertex events from closely spaced multivertex events such as β^\pm decay with γ cascade. This is demonstrated in Fig. 5(a) using sim-KLZ800 by comparing the KamNet score spectrum of the γ cascade decays like ^{214}Bi or ^{10}C with the β events produced by the elastic scattering of solar neutrinos. We again set a cut on the KamNet score corresponding to the acceptance of 90% of the $0\nu\beta\beta$ signal. Using this criterion, KamNet rejects 27.1% of ^{214}Bi events from within the XeLS and 59.6% of ^{214}Bi events from the miniballoon film. In comparison, the rejection of solar neutrino electron elastic scattering events is only 9.8%.

We can extend this study to long-lived spallation products, unstable light isotopes produced by high energy cosmic muons interacting in the liquid scintillator. Each long-lived spallation isotope undergoes a γ cascade with a unique relative intensity. It turns out that the isotope with highest relative intensity (^{90}Nb) has the highest rejection efficiency (35.1%), while the isotope with lowest relative intensity (^{137}Xe) has the lowest rejection efficiency (12.1%). Aggregating all long-lived isotopes, we obtain a Pearson correlation coefficient of 0.49 between rejection efficiency and the relative intensity of the γ cascade. This indicates that the rejection power of KamNet is moderately correlated with the γ cascade. In near future, we will conduct a systematic study to comprehensively understand how the γ cascade affect KamNet's performance.

The identification of closely spaced multivertex events can be directly visualized using the attention score in Eq. (3). The attention score is a probabilistic float point number assigned to each slice of the spatiotemporal hit map. KamNet relies heavily on time slices with high attention score to make

classification decisions, not as much on slices with low attention score.

Figure 6 shows the relative attention score for KamNet trained to reject different backgrounds. The shape of plot shows the scintillation time profile of KamLAND-Zen events. The bin width are 1.5 ns, identical to the input series of spherical hit maps. The first and last bins are populated since it contains overflow and underflow hits. The color of each bin indicates the relative attention score on this time slices as assigned by KamNet. In Fig. 6(a), KamNet is trained to reject ^{214}Bi backgrounds, thus most of the attention is placed between 5 and 10 ns when the γ cascade occurs. The second highest attention is placed near the scintillation peak. On the other hand, KamNet in Fig. 6(b) is trained to reject ^{10}C backgrounds. In this case, the highest attention is placed at the rising edge, since ^{10}C undergoes β^+ decay to produce a pair of γ . These γ s will Compton scatter and deposit their energy on a slightly longer timescale than β^- decay. In some case, the e^+ and e^- will form a positronium which further delay the energy deposition and adds on to the effect of rising edge. This effect was also observed in our previous work [12]. Besides the positronium effect, a 718-keV γ is also released in 0.7 ns after ^{10}C decays to ^{10}B excited states. KamNet also pay attention to the falling edge of scintillation time profile to capture this γ cascade.

The biggest difference between Figs. 6(a) and 6(b) is the last overflow bin. In Fig. 6(a), the last bin contains secondary effects such as absorption, reemission, scattering, and dark noise. These effects are identical in both signal and background events, and thus KamNet does not to pay attention to the last bin. On the contrary, KamNet places a significant amount of attention to the overflow bin in Fig. 6(b). KamNet in Fig. 6(b) is trained with ^{10}C with ortho-positronium half-life. Ortho-positronium possesses a lifetime of 3 ns in liquid scintillator [24]. Along with the delay caused by LS timing response, ortho-positronium allows the decay to delay and leak physical information into the last overflow bin, and

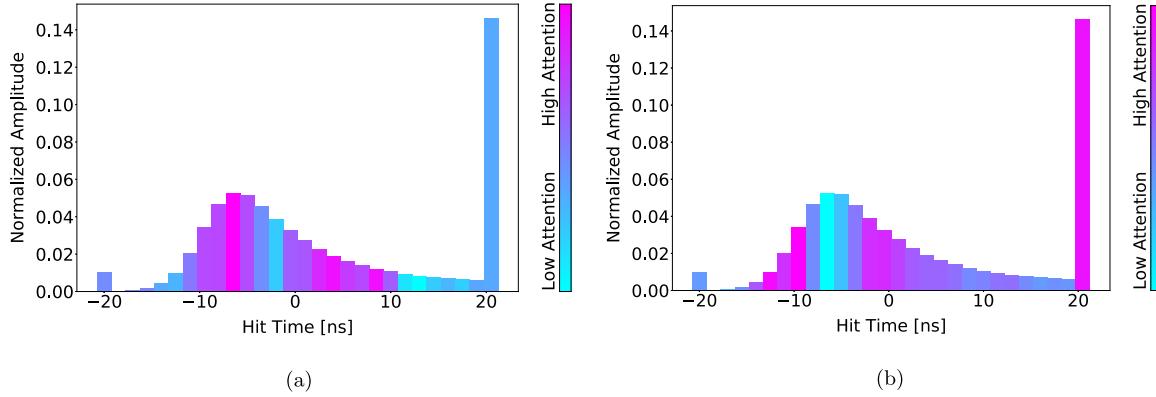


FIG. 6. The attention score plot of KamNet. The shape of plot shows the scintillation time profile of KamLAND-Zen events, and the color shows the relative attention score of each time slice. KamNet relies heavily on the high attention (magenta) region to make classification decisions, not as much on the low attention (cyan) region. In plot (a), KamNet is trained to reject ^{214}Bi backgrounds from $0\nu\beta\beta$ signal. In plot (b), KamNet is trained to reject ^{10}C backgrounds from $0\nu\beta\beta$ signal.

KamNet is able to extract that information to make classification decision.

B. Data cleaning

With the ability to interpret the network attention weights, we can now use it to identify periods of high background rates. In KamLAND, we observe that detector operating conditions, such as temperature fluctuations in the buffer liquids surrounding the PMTs or vibrations from construction on the deck above the detector, can lead to periods of increased convection inside the inner detector volumes. Convection in the miniballoon tends to pull contaminants off of the balloon surface and into the main $0\nu\beta\beta$ analysis volume. These contaminants are primarily closely spaced multivertex events KamNet can efficiently identify. Therefore, likelihood profiles of KamNet score are constructed based on the MC simulation of several representative isotopes, including both β -like signal and $\beta^\pm + \gamma$ backgrounds. The likelihood profiles are then fitted to data to extrapolate periods with high background concentration. This is a powerful tool that is used in conjunction with other monitoring tools and logbooks to veto these periods of data instability as described in Ref. [25].

C. Background rejection

On average, KamNet rejects 27% of internal backgrounds and 59% of film backgrounds. Furthermore, the background rejection of KamNet does not rely on high-level analysis cuts, like prompt-delayed coincidence tagging, or hardware upgrades. Thus, the background rejection factor has a multiplicative effect when applied to any standard physics analysis. Since long-lived spallation products are the major sources of background in the KamLAND-Zen region of interest, and an efficient coincidence tag is challenging because of their long half-life, KamNet plays a key role in pushing the KamLAND-Zen $0\nu\beta\beta$ limit forward. To evaluate the performance of KamNet, we selected five abundant long-lived spallation backgrounds within the KamLAND-Zen ROI.

The rejection efficiencies of those backgrounds are listed in Table I.

We now use KamNet's rejection power against XeLS backgrounds to estimate the expected sensitivity boost for the $0\nu\beta\beta$ search in KamLAND-Zen. To keep the calculation simple, a counting experiment model is used to estimate the sensitivity. This estimate is conservative compared to the fitting of the energy spectrum performed in a full-scale KamLAND-Zen analysis. In this model, KamLAND-Zen's sensitivity is proportional to S/\sqrt{B} , where S is the number of signal events and B is the number of background events in the ROI. If a classifier rejects 30% of the background while preserving 90% of the signal, then the sensitivity will be boosted by:

$$T_{0\nu}^{1/2} \propto \frac{90\%S}{\sqrt{(100\%-30\%)B}} = 1.076 \frac{S}{\sqrt{B}}. \quad (5)$$

Here 90% is the true positive rate (TPR) and 30% is the true negative rate. One minus the true negative rate is the false positive rate (FPR), indicating the percentage of backgrounds remaining after cut. The number 1.076 is the sensitivity factor S for this classifier, corresponding to a 7.6% increase in sensitivity. In KamLAND-Zen, we have more than one type of background in the ROI, so we have to take all of them into account. From preliminary fitting results [26], 58.2% of backgrounds are long-lived spallation backgrounds which can be efficiently rejected by KamNet, and 41.8% of backgrounds are irreducible $2\nu\beta\beta$ -like backgrounds. Therefore, we can estimate the sensitivity factor S with the following equation:

$$S = \frac{\text{TPR}}{\sqrt{58.2\% \text{FPR}_{LL} + 41.8\% \text{FPR}_{2\nu}}} \quad (6)$$

Since KamNet does not have any rejection power against $2\nu\beta\beta$ -like backgrounds, $\text{FPR}_{2\nu}$ is equivalent to TPR. TPR and FPR_{LL} are obtained by making cut at a given KamNet cutting threshold and evaluating on sim-KLZ800. Based on this calculation, we are able to evaluate the sensitivity factor on all possible cutting thresholds. While accepting 90% signal events, KamNet boost the $0\nu\beta\beta$ search sensitivity by 2.2%.

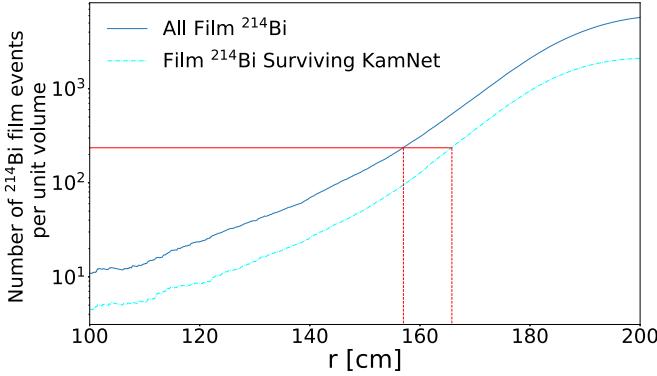


FIG. 7. Result of KamNet fiducial volume study based on sim-KLZ800. The dark blue curve shows the film background rate before KamNet cut, and the cyan dashed curve shows film background rate after KamNet cut.

This boost can be further enhanced by optimizing the energy selection to reduce the amount of $2\nu\beta\beta$ backgrounds.

The 2.2% boost is extremely conservative. The KamLAND-Zen 800 analysis is a multidimensional fit in energy and position. KamNet's ability to reject backgrounds with peaklike features in the ROI, amplifies its power in the region of interest providing much larger gains in sensitivity. For this reason, our future work focuses on a native propagation of the output of KamNet to the Bayesian extraction of the $0\nu\beta\beta$ result. This includes a modification of KamNet to produce joint Bayesian priors for each fitting spectrum and then integrate them into the Bayesian analysis.

D. Fiducial volume expansion

KamNet's rejection power against film backgrounds can also lead to an independent sensitivity boost. According to Table I, KamNet efficiently rejects 59% of film background while maintaining 90% signal acceptance. Since film background is the major limitation on KamLAND-Zen fiducial volume, KamNet enable us to expand fiducial volume to gain more exposure. Figure 7 shows the film background rate before and after KamNet cut is applied. Based on this figure, KamNet allows us to expand fiducial volume from $R < 157$ cm to $R < 165.8$ cm without increasing the film background level. This corresponds to a 17.7% increase in fiducial volume and exposure. In monolithic detectors, the sensitivity is proportional to the exposure in the following way:

$$T_{0\nu}^{1/2} \propto \sqrt{\frac{\alpha \epsilon M}{B \delta E}}, \quad (7)$$

where α is the isotopic abundance, ϵ is the detection efficiency, M is the exposure, B is the number of backgrounds, and δE is the energy resolution. Using this equation, we quote a 8.5% sensitivity boost from 17.7% increase in exposure. Once again, this is a conservative estimate.

The two sensitivity boosts we discussed above are mutually independent. Even within the expanded fiducial volume, KamNet can still efficiently reject XeLS background events. Therefore, the final sensitivity boost is quoted by multiplying

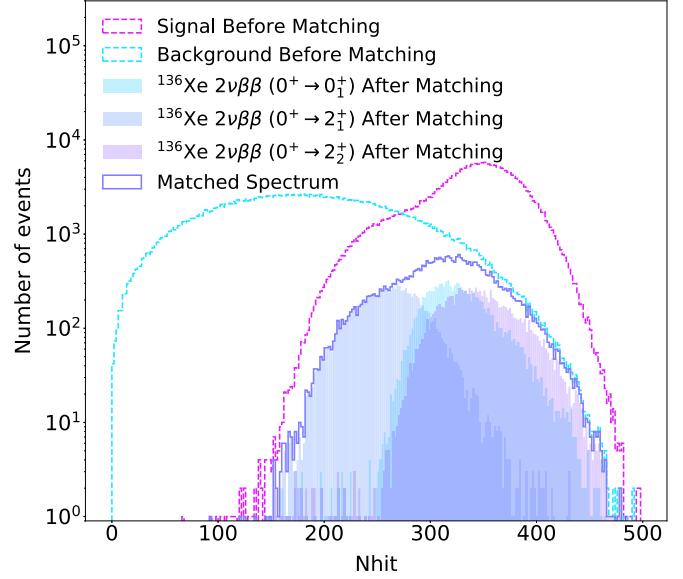


FIG. 8. Effect of Nhit matching on the signal and background Nhit distributions. The signal distribution is shown being equally divided into the three excited states.

2.2% XeLS sensitivity boost to 8.5% film sensitivity boost, resulting in a 10.8% overall sensitivity boost. With the help of KamNet, KamLAND-Zen unleashes its full detection power toward $0\nu\beta\beta$ decay.

VI. $2\nu\beta\beta$ DECAY TO EXCITED STATES DATASET

In order to quantify KamNet's performance on $2\nu\beta\beta$ decay of ^{136}Xe to excited states, we use the sim-RAT to generate training and validation datasets. The three main excited-state decays are simulated and aggregated together into one signal, and the decay to the ground state is simulated as the background. The dataset containing both signal and background events is split into training and validation datasets with a 7:3 ratio. KamNet is trained on the training dataset to produce a binary classifier and its performance is evaluated separately for each type of excited-state decay against the decay to the ground state.

Throughout this study we refer to the number of nonzero cells in the spatiotemporal hit map as "Nhit" where the total number of cells is equal to 40 432. Energy deposits in XeLS of about 2.5 MeV produce roughly 350 Nhits on average. Early studies showed that KamNet would attempt to classify events only by Nhit and ignore all other information when Nhit is very different between the signal and background events. Therefore, we perform something called "Nhit matching" prior to training. This means that both signal and background events are sampled from an Nhit distribution which is formed from the overlap between the signal and background Nhit distributions. During Nhit matching the signal Nhit distribution is divided equally among the three different excited-state decays. The matching loops through all possible Nhits and, at each step, randomly samples the same number of events from signal and background datasets without replacement. The signal and background Nhit distributions before, and

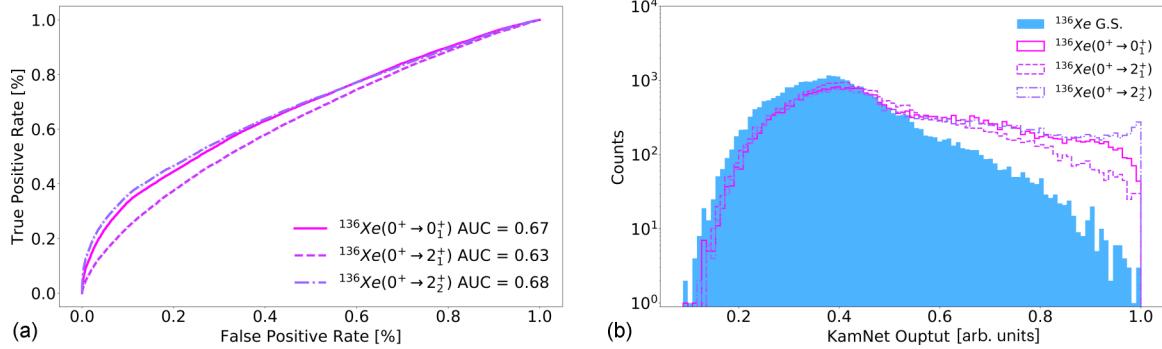


FIG. 9. (a) The ROC curve of KamNet output on each ^{136}Xe decay to excited state signal vs. ^{136}Xe ground-state backgrounds. (b) KamNet score spectrum of 3 ^{136}Xe decay to excited states and ^{136}Xe ground state.

after, Nhit matching are shown in Fig. 8. A total of 300 000 background events and 100 000 signal events forms the raw training dataset. After matching, the background and signal distributions each contain 33 498 events, and the events in the signal distribution are divided equally among the three excited states.

After preprocessing, the dataset is fed into KamNet for training in PyTorch [27]. KamNet is trained over 30 epochs and the binary cross-entropy loss is minimized with the ADAM optimizer [28]. After training, the result is evaluated in the form of receiver operating characteristic (ROC) curves. The ROC curve of each excited state signal is plotted in Fig. 9(a), where a higher area under curve (AUC) indicates better distinguishability between signal and background events. If the threshold is set to reject 70% of the background, then the signal acceptance efficiency ϵ for $(0^+ \rightarrow 0_1^+)$, $(0^+ \rightarrow$

$2_1^+)$, and $(0^+ \rightarrow 2_2^+)$ events are 54%, 49%, and 56%, respectively. Based on the decay half-lives provided in Ref. [29], we estimated the branching ratio r of $(0^+ \rightarrow 0_1^+)$, $(0^+ \rightarrow 2_1^+)$, and $(0^+ \rightarrow 2_2^+)$ to be 0.2970, 0.002366, and 3.246×10^{-8} at g_A of 0.60. Thus the overall signal efficiency is calculated using:

$$\epsilon_{\text{DES}} = \frac{\sum_i r_i \epsilon_i}{\sum_i r_i}. \quad (8)$$

The summation is conducted over three excited-state decay and the final ϵ_{DES} is 54.0%.

To evaluate the impact of KamNet's background rejection power on the excited-state analysis, we use a recent estimate of the backgrounds in KamLAND-Zen 800 [30]. The possible improvement on the excited-state analysis after applying KamNet is illustrated in Fig. 10. The spectrum amplitudes of the excited states are fixed using the lower half-life limits (at

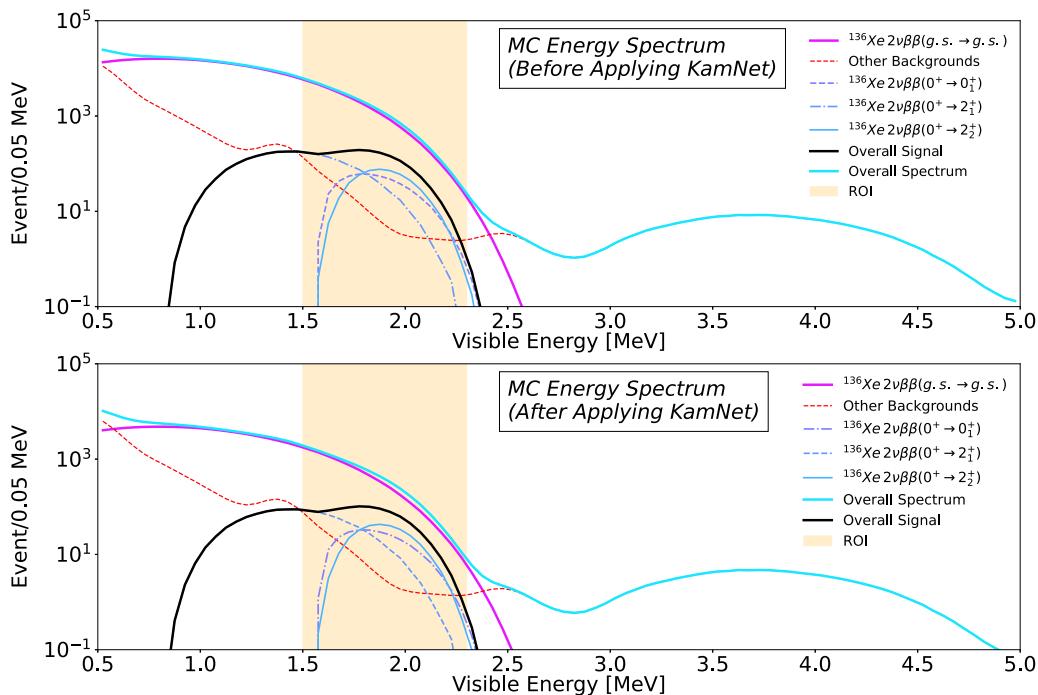


FIG. 10. The KamLAND-Zen 800 MC energy spectrum before applying KamNet (top) and after applying KamNet (bottom). The ratio between $2\nu\beta\beta$ DES and $2\nu\beta\beta$ DGS is obtained from Ref. [31]. The DES distribution corresponds to 90% confidence level upper limit.

90% C.L.) from an earlier analysis with KamLAND-Zen 400 [31]. The impact of KamNet on the excited-state analysis is estimated from the improvement in the signal-to-noise ratio (S/N). It is calculated using the sum of all three signals inside an energy window of 1.5–2.3 MeV. In this region, the excited-state decay signals are closely spaced multivertex events with γ cascades. The dominant background in this region is strictly single-vertex, $2\nu\beta\beta$ decay to the ground state, which constitutes roughly 99% of the background. In addition, backgrounds such as ^{214}Bi , ^{11}C , ^{122}I , ^{124}I , ^{130}I , and ^{118}Sb could also leak into this energy window. Unlike the situation in Sec. V, these additional backgrounds contain γ cascades and look more similar to the signal than to the dominant background. Due to limited time and computing resources, we chose to only use KamNet to reject the $2\nu\beta\beta$ decay to ground state and not run it over all of the additional backgrounds. Instead, we conservatively assume that KamNet will treat these backgrounds the same way it treats the excited-state signals. Therefore, we apply an acceptance factor of 0.56 (equal to the highest acceptance of the excited-state signals) over the entire “other backgrounds” spectrum to mimic this effect. Under these assumptions, the S/N before applying KamNet is 0.0691, and the S/N after applying KamNet is 0.1187, which corresponds to a 72% improvement.

VII. CONCLUSION AND OUTLOOK

Liquid scintillator detectors have been at the heart of many of the great discoveries in neutrino physics and have been a leading technology in the search for $0\nu\beta\beta$ decay. Their data is effectively a time series of images projected onto a sphere. In our previous work [12], the power of conventional CNN models had been demonstrated. In this work, we invented a novel deep learning model called KamNet for better performance. Leveraging recent breakthroughs in geometric deep learning and spatiotemporal data modeling, KamNet outperforms the conventional CNN on both rejection efficiency and robustness. With a standard detector configuration similar to the current KamLAND-Zen detector, we find KamNet can reject 74.0% of the ^{10}C background with 90% acceptance of the $0\nu\beta\beta$ -decay signal, surpassing 61.5% rejection from conventional CNN. Furthermore, by applying KamNet to ^{136}Xe

ground state and various excited states, we find we can boost the S/N ratio of ^{136}Xe excited-state decay search by 72%. Finally, with precisely tuned MC in KamLAND-Zen 800, we find KamNet can reject 27% XeLS backgrounds and 59% film backgrounds without any coincidence tagging or hardware upgrade. We conservatively estimated the $0\nu\beta\beta$ sensitivity boost from these background rejection to be 10.8%.

This work has focused on optimizing an algorithm for data from a spherical LS detector; however, the data-driven nature of KamNet allows a easy generalization to different detectors and different tasks. Furthermore, the network interpretation study we performed unravels the black-box nature of KamNet to reveal underlying physics. Our future work moves in two directions. We intend to perform a systematic interpretation study to rigorously unveil the origin of KamNet’s classification power and perhaps further improve the performance of the algorithm. We then plan to extend the reach of KamNet beyond event classification and background rejection. These include but not limited to KamNet-GAN for event generation, Self-supervised KamNet to provide Bayesian posterior distributions for spectrum fitting and Regressive KamNet for event reconstruction. These studies are benefiting from an abundance of work being done for other applications both inside and outside of particle and nuclear physics and this is just the beginning.

ACKNOWLEDGMENTS

This material is based on work supported by the National Science Foundation under Grants No. 2110720 and No. 2012964, and the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Award No. DE-SC0022339. This work is done in support of the KamLAND-Zen experiment and we thank our collaborators for their input. The KamLAND-Zen experiment is supported by JSPS KAKENHI Grants No. 21000001, No. 26104002, and No. 19H05803; the Dutch Research Council (NWO); and under the U.S. Department of Energy (DOE) Grant No. DE-AC02-05CH11231, as well as other DOE and NSF grants to individual institutions. We thank Prof. Julieta Gruszko for the support of Aobo Li’s work on KamLAND-Zen. This research was performed, in part, using the Boston University Shared Computing Cluster.

[1] A. Gando, Y. Gando, T. Hachiya, A. Hayashi, S. Hayashida, and H. E. A. Ikeda (KamLAND-Zen Collaboration), Search for Majorana Neutrinos Near the Inverted Mass Hierarchy Region with KamLAND-Zen, *Phys. Rev. Lett.* **117**, 082503 (2016).

[2] K. Eguchi *et al.* (KamLAND), First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance, *Phys. Rev. Lett.* **90**, 021802 (2003).

[3] Borexino Collaboration (BOREXINO), Experimental evidence of neutrinos produced in the CNO fusion cycle in the Sun, *Nature (Lond.)* **587**, 577 (2020).

[4] C. Arpesella *et al.* (Borexino), Direct Measurement of the Be-7 Solar Neutrino Flux with 192 Days of Borexino Data, *Phys. Rev. Lett.* **101**, 091302 (2008).

[5] G. Bellini *et al.*, Precision Measurement of the 7Be Solar Neutrino Interaction Rate in Borexino, *Phys. Rev. Lett.* **107**, 141302 (2011).

[6] M. Anderson *et al.* (SNO+), Search for invisible modes of nucleon decay in water with the SNO+ detector, *Phys. Rev. D* **99**, 032008 (2019).

[7] M. Anderson *et al.* (SNO+ Collaboration), Measurement of the ^8B solar neutrino flux in SNO+ with very low backgrounds, *Phys. Rev. D* **99**, 012012 (2019).

[8] F. P. An *et al.* (Daya Bay), New Measurement of Antineutrino Oscillation with the Full Detector Configuration at Daya Bay, *Phys. Rev. Lett.* **115**, 111802 (2015).

[9] F. An *et al.* (JUNO), Neutrino physics with JUNO, *J. Phys. G: Nucl. Part. Phys.* **43**, 030401 (2016).

[10] Y. Gando *et al.*, First results of KamLAND-Zen 800, in *Journal of Physics: Conference Series* (IOP Publishing, Bristol, UK, 2020), Vol. 1468, p. 012142.

[11] M. Fukugita and T. Yanagida, Baryogenesis without grand unification, *Phys. Lett. B* **174**, 45 (1986).

[12] A. Li, A. Elagin, S. Fraker, C. Grant, and L. Winslow, Suppression of cosmic muon spallation backgrounds in liquid scintillator detectors using convolutional neural networks, *Nucl. Instrum. Methods A* **947**, 162604 (2019).

[13] M. Askins, A. Bacon, Z. Bagdasarian, N. Barros, M. Bergevin, E. Callaghan, C. Grant, J. Gruszko, T. Kaptanoglu, J. Klein *et al.*, Letter of interest the rat (-pac) simulation and analysis code base, Snowmass2021 (2020).

[14] S. Agostinelli *et al.* (GEANT4), GEANT4: A Simulation toolkit, *Nucl. Instrum. Methods A* **506**, 250 (2003).

[15] J. Allison *et al.* (GEANT4), Geant4 developments and applications, *IEEE Trans. Nucl. Sci.* **53**, 270 (2006).

[16] O. A. Ponkratenko, V. I. Tretyak, and Y. G. Zdesenko, The Event generator DECAY4 for simulation of double beta processes and decay of radioactive nuclei, *Phys. At. Nucl.* **63**, 1282 (2000).

[17] S. Hayashida, Application of the rnn in the fundamental physics with kamland experiment, *Am. J. Comput. Sci. Inf. Technol.* **11** (2019).

[18] T. S. Cohen, M. Geiger, J. Köhler, and M. Welling, Spherical CNNs, International Conference on Learning Representations, [arXiv:1801.10130](https://arxiv.org/abs/1801.10130).

[19] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-C. Woo, Convolutional LSTM network: A machine learning approach for precipitation nowcasting, *Adv. Neural Inf. Process. Syst.* **28**, 802 (2015).

[20] S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**, 1735 (1997).

[21] D. Bahdanau, K. H. Cho, and Y. Bengio, Neural machine translation by jointly learning to align and translate, 3rd International Conference on Learning Representations, ICLR 2015, [arXiv:1409.0473](https://arxiv.org/abs/1409.0473).

[22] J. Dunger and S. D. Biller, Multi-site event discrimination in large liquid scintillation detectors, *Nucl. Instrum. Methods A* **943**, 162420 (2019).

[23] A. P. Bradley, The use of the area under the roc curve in the evaluation of machine learning algorithms, *Pattern Recogn.* **30**, 1145 (1997).

[24] T. Hachiya, e^+/e^- discrimination in KamLAND, *JPS Conf. Proc.* **12**, 010046 (2016).

[25] S. Abe *et al.* (KamLAND-Zen Collaboration), Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen, *Phys. Rev. Lett.* **130**, 051801 (2023).

[26] O. Hideyoshi, High sensitivity search for neutrinoless double-beta decay in KamLAND-Zen with double amount of ^{136}Xe , Tohoku University, 2020.

[27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai *et al.*, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2019), pp. 8024–8035.

[28] D. P. Kingma and J. A. Ba, Adam a method for stochastic optimization, ICLR (Poster), [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).

[29] P. Pirinen and J. Suhonen, Systematic approach to β and $2\nu\beta\beta$ decays of mass $a = 100\text{--}136$ nuclei, *Phys. Rev. C* **91**, 054309 (2015).

[30] A. Li, The Tao and Zen of Neutrinos: Neutrinoless Double Beta Decay in KamLAND-Zen 800, Ph.D. thesis, Boston University (2020).

[31] K. Asakura *et al.* (KamLAND-Zen), Search for double-beta decay of ^{136}Xe to excited states of ^{136}Ba with the KamLAND-Zen experiment, *Nucl. Phys. A* **946**, 171 (2016).