Overlap-aware segmentation for topological reconstruction of obscured objects

J. Schueler[1,*], H. M. Araújo[2,3], S. N. Balashov[3], J. E. Borg[4], C. Brew[3], F. M. Brunbauer[5], C. Cazzaniga[6], A. Cottle[7], D. Edgeman[1], C. D. Frost[6], F. Garcia[9], D. Hunt[8], M. Kastriotou[6], P. Knights[10], H. Kraus[8], A. Lindote[11], M. Lisowska[5], D. Loomba[1], E. Lopez Asamar[12], P. A. Majewski[3,2], T. Marley[2], C. McCabe[13], L. Millins[10,3], R. Nandakumar[3], T. Neep[10], F. Neves[11], K. Nikolopoulos[10,14], E. Oliveri[5], A. Roy[2], T. J. Sumner[2], E. Tilly[1], W. Thompson[1], M.A. Vogiatzi[2,3]

[1]Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM, 87131, USA
[2]Department of Physics, Blackett Laboratory, Imperial College London, London, SW7 2AZ, UK
[3]Particle Physics Department, STFC Rutherford Appleton Laboratory, Didcot, OX11 0QX, UK
[4]Luleå University of Technology, 97187 Luleå, Sweden
[5]CERN, 1211 Geneva 23, Switzerland
[6]ISIS Neutron and Muon Source, STFC Rutherford Appleton Laboratory, Didcot, OX11 0QX, UK
[7]University College London (UCL), Department of Physics and Astronomy, London WC1E 6BT, UK
[8]Department of Physics, Keble Road, University of Oxford, Oxford, OX1 3RH, UK
[9]Helsinki Institute of Physics, University of Helsinki, FI-00014 Helsinki, Finland
[10]School of Physics and Astronomy, University of Birmingham, Birmingham, B15 2TT, UK
[11]LIP – Laboratório de Instrumentação e Física Experimental de Partículas, University of Coimbra, P-3004-516 Coimbra, Portugal
[12]Departamento de Fisica Teorica, Universidad Autonoma de Madrid, 28049 Madrid, Spain
[13]Department of Physics, King's College London, London, WC2R 2LS, UK
[14]University of Hamburg, 22767, Hamburg, Germany
[*]Author to whom any correspondence should be addressed.

E-mail: jschueler1@unm.edu

Keywords: Segmentation-Regression, Deep learning, Scientific Imaging, Dark Matter, MIGDAL

Abstract

The separation of overlapping objects presents a significant challenge in scientific imaging. While deep learning segmentation-regression algorithms can predict pixel-wise intensities, they typically treat all regions equally rather than prioritizing overlap regions where attribution is most ambiguous. Recent advances in instance segmentation show that weighting regions of pixel overlap in training can improve segmentation boundary predictions in regions of overlap, but this idea has not yet been extended to segmentation regression. We address this with Overlap-Aware Segmentation of ImageS (OASIS): a new segmentation-regression framework with a weighted loss function designed to prioritize regions of object-overlap during training, enabling extraction of pixel intensities and topological features from heavily obscured objects. We demonstrate OASIS in the context of the MIGDAL experiment, which aims to directly image the Migdal effect–a rare process where electron emission is induced by nuclear scattering–in a low-pressure optical time projection chamber. This setting poses an extreme test case, as the target for reconstruction is a faint electron recoil track which is often heavily-buried within the orders-of-magnitude brighter nuclear recoil track. Compared to unweighted training, OASIS improves median intensity reconstruction errors from −32% to −14% for low-energy electron tracks (4-5 keV) and improves topological intersection-over-union scores from 0.828 to 0.855. These performance gains demonstrate OASIS's ability to recover obscured signals in overlap-dominated regions. The framework provides a generalizable methodology for scientific imaging where pixels represent physical quantities and overlap obscures features of interest. All code is openly available to facilitate cross-domain adoption.

1 Introduction

A pervasive challenge in scientific imaging is the decoupling of overlapping objects. Some examples include the blending of galaxies in crowded astronomical fields where overlapping light profiles can bias photometry and redshift estimations [1]; fluorescence microscopy, where multiple fluorophores can overlap within a diffraction-limited PSF [2]; diverse applications in medical imaging [3, 4, 5]; and particle detectors, where overlapping particle tracks can obscure features of interest such as decay vertices or faint secondaries [6, 7]. In these and other scientific imaging applications where pixels (or voxels) represent physically meaningful quantities, we aim to demonstrate that prioritizing regions of overlap in the training of pixel-intensity segmentation-regression models provides a path toward disentangling overlapping objects and improves the attribution of object-specific intensities even in regions where intensities of overlapping pixels span orders of magnitude.

Deep learning-based segmentation networks provide pixel-level reconstructions that have become essential tools across scientific imaging. Architectures such as U–Net [8] and Mask R-CNN [9] have enabled object-level separation in diverse domains, from biomedical imaging to astronomy, by mapping each pixel to a categorical label or probabilistic mask. Extending these architectures to segmentation-regression allows networks to predict continuous, physically meaningful quantities such as voxel-wise parameter maps in quantitative medical imaging [10] and flux maps in galaxy deblending [11, 12]. These models jointly reconstruct object topology and pixel-wise intensities, however they treat all regions equally rather than prioritizing overlap regions where attribution is most ambiguous. Occlusion-aware instance-segmentation frameworks such as MultiStar [13] and BCNet [14], on the other hand, explicitly emphasize regions of pixel-overlap in their training, but neither regress physical intensities. Here we introduce a novel framework called Overlap-Aware Segmentation of ImageS (OASIS) that extends the overlap-aware training paradigm to segmentation regression. By employing programmable weights in its loss function that are designed to specifically target regions of pixel overlap, OASIS is able to prioritize intensity attribution in the most ambiguous regions of object overlap during its training.

High resolution imaging-particle detectors like gas time projection chambers (TPCs) are widely used in nuclear and particle physics. These detectors are capable of reconstructing energy depositions along particle trajectories at a fine granularity, making their data suitable for testing OASIS. Some gaseous TPC applications with overlapping track signals where component-wise segmentation-regression could be applied are neutrinoless double beta decay search experiments like NEXT [15, 16, 17], optical TPC experiments measuring exotic nuclear decay processes [18, 19, 20], liquid argon TPC experiments like MicroBooNE [21], which already employs segmentation in its reconstruction chain [22, 23, 24, 25, 26], and the MIGDAL experiment's [27] rare event search for the Migdal effect. From a reconstruction standpoint, the latter is a particularly challenging case because the MIGDAL experiment's target signal is an image containing a bright nuclear recoil (NR) track that heavily overlaps with an orders-of-magnitude fainter electron recoil (ER) track. In order to measure the Migdal effect and meaningfully compare observations with theoretical predictions, reconstruction of the energy of the faint and heavily obscured ER track is essential, while reconstruction of its emission angle provides an additional test of theory. The extreme challenge of reconstructing the ER track makes the MIGDAL experiment an excellent testbed to demonstrate the performance and potential of OASIS.

We structure the rest of this work as follows: In Section 2 we detail the OASIS framework in a general context, including the backbone network we use (Section 2.1) and its loss function (Section 2.2). We then introduce the MIGDAL experiment (Section 3) and describe the details relevant for applying OASIS to images recorded by the MIGDAL experiment. In Section 4 we quantify OASIS's topological and intensity reconstruction performance for the faint ER component of Migdal-effect signals and compare it to the baseline of using an unweighted loss function. Finally in Section 5, we contextualize our results in terms of their impact to the MIGDAL experiment and beyond.

2 The OASIS Framework

The general objective of OASIS is to separate an input image consisting of k object classes into k object-specific intensity maps. OASIS's training is guided by a custom loss function that applies object and region-specific weights that are predefined for the application at hand. In the case of separating overlapping objects, different weights can be prescribed to differing regions of overlap. For example, if a task were to consist of three unique object-classes, labeled 1, 2, and 3, the four unique classes of object overlap ($1 \cap 2$, $1 \cap 3$, $2 \cap 3$, and $1 \cap 2 \cap 3$) could each be assigned their own overlap region weights.

Figure 1 shows a schematic of the OASIS framework. An input $n_x \times n_y$ intensity-map image containing k distinct object-classes is processed through a standard U–Net backbone network, which we describe in Section 2.1. The hierarchical features extracted by the backbone network are then passed into a segmentation-regression head that produces k output intensity maps of the same

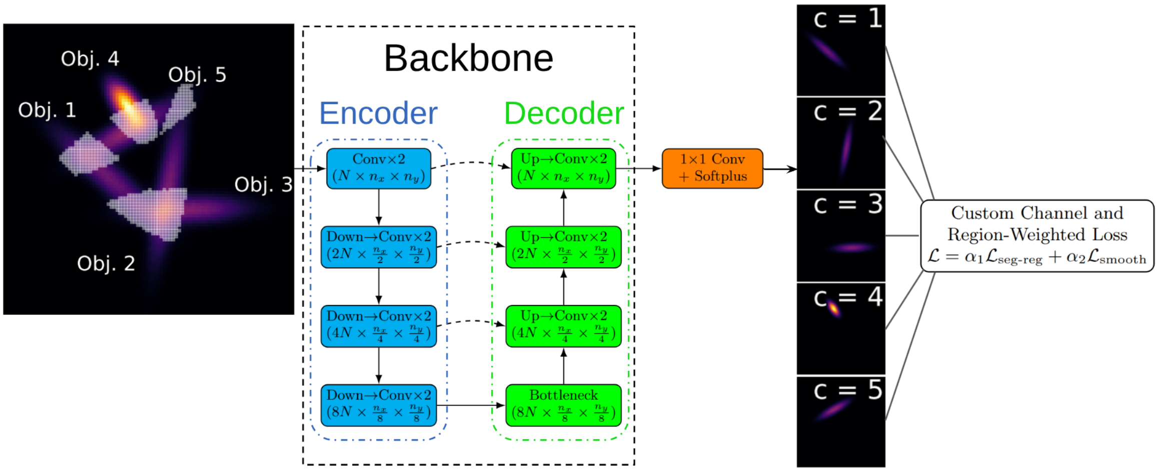

Figure 1. Schematic of OASIS. An input intensity-map image of dimension $n_x \times n_y$ containing k distinct object classes ($k = 5$ in this example) is passed as input into the network. The backbone network, where here we use U–Net, processes the images which are then passed into a segmentation-regression head, shown as the orange block. This maps the backbone network's output to k intensity maps of dimension $n_x \times n_y$, each corresponding to a distinct object class. During training, OASIS is optimized using a custom loss function that compares the predicted output intensity maps with truth. The loss function incorporates weights that can be tuned to assign higher penalties to reconstruction errors in overlap regions (white patches in the input image) and to channels with fainter objects.

dimensions, where each output map represents the predicted intensity contribution from one object class.

We train the network using a custom weighted loss function $\mathcal{L}$, described in Section 2.2, which can be tuned to assign higher penalties to reconstruction errors in overlap regions (white patches in Figure 1) and to channels corresponding to fainter objects. Such a weighting scheme prioritizes accurate reconstruction in regions where intensity attribution is most ambiguous due to overlapping sources with large intensity differences.

2.1 Network architecture

We employ a U–Net as our backbone architecture, though any backbone network suitable for segmentation regression could be substituted. Our U–Net consists of a four stage encoder path, a bottleneck layer, and a decoder path that mirrors the encoder path. Skip connections link the corresponding encoder and decoder stages as diagrammed in Figure 1 and explained below.

Given a single-channel[1] intensity-map input image, $\mathcal{I}(x, y) \in \mathbb{R}^{n_x \times n_y}$, the first encoder block applies two consecutive N-channel 3×3 convolutions, each followed by a group normalization over 8 groups and a SiLU [28] activation. We refer to this combination of operations as a *standard two-layer block* with an output depth of N-channels. The output of this first encoder block is an N-channel feature map at full spatial resolution, $\hat{\mathcal{I}}_{\text{enc1}} \in \mathbb{R}^{N \times n_x \times n_y}$. Each subsequent encoder block begins with a depthwise 3×3 convolution with stride 2 that halves the spatial resolution while maintaining the channel depth, followed by a standard two-layer block that doubles the output depth. After the final encoder block, the output feature map $\hat{\mathcal{I}}_{\text{enc4}} \in \mathbb{R}^{8N \times \frac{n_x}{8} \times \frac{n_y}{8}}$ is passed into the bottleneck block, which refines this output through another standard two-layer block that retains the same channel depth ($8N$). This output is then passed into the decoder path which mirrors the encoder path. Each decoder block performs the following steps

1. A bilinear interpolation that doubles the spatial resolution.
2. A 1×1 convolution that halves the channel depth.
3. Apply the skip connection: concatenate this output with the output feature map from the mirrored encoder block. This step doubles the number of channels.
4. Standard two-layer block that halves the output channel depth.

All together, each subsequent decoder block doubles the spatial resolution of the feature maps while halving their channel depth. The decoder path output $\hat{\mathcal{I}}_{\text{dec4}} \in \mathbb{R}^{N \times n_x \times n_y}$ is then passed

[1] In general the number of input channels need not be one, so multi-band images in astronomy, for example, could be used as input to OASIS.

through the segmentation head which is a 1×1 convolution with Softplus activation that projects the channel depth to the number of object-classes, k. The final output of our network is a k-channel intensity map, $\hat{\mathcal{I}} = [\hat{\mathcal{I}}_1, \ldots, \hat{\mathcal{I}}_k]^\top \in \mathbb{R}^{k \times n_x \times n_y}$, consisting of predictions of each individual object class's intensity contributions to the input image. When applying OASIS to the MIGDAL experiment (Section 3) we use 512×288 images, a base channel depth of $N = 32$, and $k = 2$ object-classes.

2.2 Loss Function

OASIS's loss function consists of a custom-weighted pixel-intensity regression loss term and a smoothness regularization term

$$\mathcal{L} = \alpha_1 \mathcal{L}_{\text{seg-reg}} + \alpha_2 \mathcal{L}_{\text{smooth}}, \tag{1}$$

where α_1 and α_2 are weights assigned to the two loss terms, which in this work we set as 1 and 0.01, respectively. To enhance training performance, input images, $\mathcal{I}_{\text{in}}(x, y)$, can optionally be masked by a pre-defined bounding box region $B \subseteq \mathcal{I}_{\text{in}}(x, y)$ that isolates regions of interest in the input data. Without loss of generality, we write our loss terms in terms of this bounding box region, B, and note that not applying a bounding box mask is equivalent to setting $B = \mathcal{I}_{\text{in}}$. Descriptions of the two loss function terms follow:

Channel-specific pixel and region-weighted regression loss, $\mathcal{L}_{\text{seg-reg}}$. This term is responsible for OASIS's segmentation-regression predictions and represents a channel and region-weighted mean absolute error minimization between channel predictions $\hat{\mathcal{I}}_{\text{c}}$ and ground truth $\mathcal{I}_{\text{c}}$, with extra weight given to predefined channels of interest and overlap regions of interest. Given a set of k unique object-classes, this loss aims to simultaneously optimize the channel-wise decomposition of predicted pixel positions and intensities for each object-class:

$$\mathcal{L}_{\text{seg-reg}} = \frac{\sum_{c=1}^{k} w_c \sum_{(x,y) \in B} W_{\text{region}}(x, y) \left| \hat{\mathcal{I}}_c(x, y) - \mathcal{I}_c(x, y) \right|}{\sum_{c=1}^{k} w_c \sum_{(x,y) \in B} W_{\text{region}}(x, y)} \tag{2}$$

Here w_c are channel (object)-specific weights, and $W_{\text{region}}(x, y)$ are region weights that can, for example, be applied to regions where truth input maps from specific channels have nonzero pixel intersection.

Smoothness loss, $\mathcal{L}_{\text{smooth}}$. This is a standard regularization term [29] that penalizes the model if neighboring pixel predictions vary by too much. Since this term is designed to be small but nonzero, its weighting factor α_2 should be small so that this term doesn't dominate the overall loss:

$$\mathcal{L}_{\text{smooth}} = \sum_{c=1}^{k} \sum_{(x,y) \in B} \Big(|\hat{\mathcal{I}}_c(x+1, y) - \hat{\mathcal{I}}_c(x, y)| + |\hat{\mathcal{I}}_c(x, y+1) - \hat{\mathcal{I}}_c(x, y)| \Big). \tag{3}$$

The OASIS framework is general and can be applied to any imaging problem requiring decomposition of overlapping sources into constituent intensity maps. The specific choice of loss function weights, preprocessing strategies, training procedures, and evaluation metrics depends on the characteristics of the application domain.

3 Application of OASIS to the MIGDAL experiment

Here we describe our application of OASIS to the MIGDAL experiment. We begin by introducing the Migdal effect and the MIGDAL experiment, and then provide parameters relevant to OASIS that are specific to its application to the MIGDAL experiment.

The Migdal effect occurs when the scattering-induced sudden displacement of a nucleus induces the emission of an atomic electron. First theorized by A. Migdal in 1939 [30], this rare effect in quantum mechanics gained renewed attention when Ibe et. al. [31] connected it to dark matter searches by tabulating probabilities of electron emission for several atomic species. Since then, numerous experiments have used the Migdal effect to extend their sensitivity to light dark matter [32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45], though, to date, the effect remains unobserved [46].

The MIGDAL experiment [27] aims to measure and characterize the Migdal effect using fast neutrons from a deuterium-deuterium generator incident on an optical time projection chamber filled with 50 Torr CF_4 gas. Images are recorded with a Hamamatsu ORCA Quest qCMOS

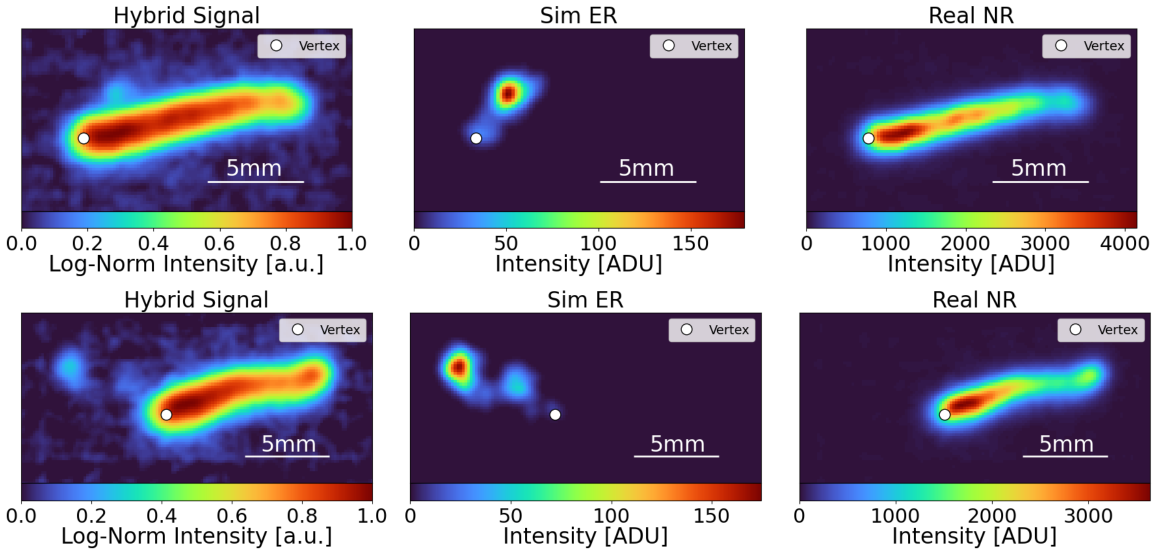

Figure 2. Illustration showing the hybrid signal simulation construction of two events: one with a 5.8 keV ER (top row) and the other with a 9.8 keV ER (bottom row). For each event, the hybrid signal (left column) is constructed by stitching a simulated ER track (middle column) with a real, measured NR track (right column). The stitching point is the truth vertex position of the simulated ER and the estimated vertex of the real NR (both shown as white dots).

camera [47]. In this work, we apply OASIS to both real nuclear recoil (NR) tracks and hybrid simulated images where simulated electron recoil (ER) tracks are stitched onto real NRs, mimicking the Migdal effect signal. Further details of the MIGDAL experiment and detector are given in Refs. [27, 48, 49].

In the region of interest for the MIGDAL experiment's search (ER energies between ∼4 keV and ∼15 keV), the Migdal effect cross section increases exponentially with decreasing ER energy [50]. This poses a challenge for searches in low pressure electron drift gas mixtures for two reasons. First, the length of ER tracks decreases with energy, and second, diffusion is relatively high in such gas mixtures. It is therefore common for ER tracks to be almost entirely obscured by the diffuse boundary of the accompanying NR track. In these cases it is extremely challenging to detect the ER signal, let alone extract sufficient information to reliably reconstruct its energy and angle. Measurements of the ER energy, as well as its angle with respect to the NR, are necessary to validate theoretical models, so overcoming these limitations is critical.

Throughout the rest of this work, we evaluate OASIS's performance at reconstructing Migdal-effect ER signals that overlap with orders-of-magnitude brighter NRs. To do this, we simulate Migdal-effect signal events using the hybrid approach adopted in Ref. [49], where events are formed by stitching the truth vertex of a simulated ER track to the predicted vertex of a randomly-selected real, measured NR track.[2] Figure 2 illustrates this construction for two example events with ER energies of 5.8 keV (top row) and 9.8 keV (bottom row). In this figure, the left column shows the hybrid signal (input to OASIS) while the middle and right columns show the ground-truth ER-only and NR-only intensity maps that serve as training targets. The simulation pipeline for generating ER tracks and the YOLO-based [53] pipeline used to define bounding box regions surrounding real NRs are described in Ref. [49]. OASIS's task is to decompose each hybrid input image, $\mathcal{I}_{\mathrm{in}}(x, y)$ into $k = 2$ output intensity maps: $\hat{\mathcal{I}}_{\mathrm{ER}}(x, y)$ representing the ER contribution, and $\hat{\mathcal{I}}_{\mathrm{NR}}(x, y)$, representing the NR contribution, to each pixel's intensity. For simplicity, the hybrid-simulated signal images formed in this work consist of exactly one ER and one NR. Further details about our image simulations, preprocessing, and labeling follow.

3.1 Image Simulation, Preprocessing, and Labeling

The hybrid signal simulation approach offers several advantages for training and evaluation. First, using real NRs allows the network to learn authentic detector context while avoiding Sim2Real gaps [54]; previous MIGDAL simulation studies [48, 49] have shown excellent pixel-level intensity and topological agreement between simulated and measured ER tracks, but comparable fidelity has not been achieved for NRs [55]. Second, the hybrid approach produces clean ground truth labels:

[2] In Ref. [49], we use the pixel of highest intensity in the NR track as a proxy for the its vertex. We now assign NR vertices more accurately using an algorithm based on curvilinear ridge detection [51, 52] that will be presented in future work.

since detector noise is already present in the real NR images and the simulated ERs are noiseless, we obtain truth ER labels, $\mathcal{I}_{\text{truth,ER}}$, consisting purely of signal without double-counting noise. Finally, using simulated ERs eliminates ER vertex reconstruction uncertainties which are often larger than those of NRs. All together, the hybrid approach provides realistic signal simulation with truth labels that allow us to directly evaluate OASIS's segmentation and intensity reconstruction performance.

To make OASIS compatible out-of-the-box with MIGDAL's camera search software [56], before training and evaluation, we preprocess each hybrid-simulated image using the same binning and smoothing procedure used for their search pipeline. Specifically, each image is 4×4 binned and Gaussian smoothed with a 9×9 kernel and $\sigma = 4/3$, and all negative pixels are set to 0. The resulting images of dimension $n_x = 512$ and $n_y = 288$ are then normalized using a log-scaling procedure designed to compress the large dynamic range of pixel intensities while preserving faint ER features. For each input image $\mathcal{I}(x, y) \in \mathbb{R}^{512\times288}$, we determine a scale factor s given by the 99th percentile of the pixel intensities of $\mathcal{I}$. Then, the following transformation is applied, yielding a scaled image with pixel values restricted to the unit interval:

$$\mathcal{I}_{\text{in}}(x, y) = \frac{\log_{10}\left(1 + \frac{\mathcal{I}}{s}\right)}{\max \log_{10}\left(1 + \frac{\mathcal{I}}{s}\right)}. \tag{4}$$

This scaled image, $\mathcal{I}_{\text{in}}(x, y)$, is what we input into our U–Net and the transformation from $\mathcal{I}$ to $\mathcal{I}_{\text{in}}$ preserves contrast in the low-intensity regime, where ER tracks appear, and also prevents high-intensity NR pixels from completely dominating the dynamic range.

Ground truth labels are constructed from the truth ER and NR portions of hybrid-simulated signal images. All together, event-data labels are stored with the following five quantities:

1. $\mathcal{I}_{\text{in}} \in \mathbb{R}^{512\times288}$
2. $(\mathcal{I}_{\text{truth,ER}}, \mathcal{I}_{\text{truth,NR}}) \in \mathbb{R}^{2\times512\times288}$
3. s
4. $\max_{xy}(\mathcal{I}_{\text{in}}(x, y))$
5. A binary mask $B(x, y)$ defined to be 1 for the contents within the minimal enclosing bounding box of the event content of $\mathcal{I}_{\text{in}}(x, y)$, and 0 elsewhere.

For each hybrid event, we define the bounding box region B as the minimal enclosing bounding box of both the NR bounding box from our YOLO search pipeline [49] and the truth boundary of the simulated ER. Quantities 1., 2., and 5. are those relevant for optimizing our loss function and therefore the only relevant quantities for training, while 3. and 4. are used to map predictions back to physically meaningful units at the inference stage. At inference time, log-scaled channel predictions, $\hat{\mathcal{I}}_{c,\log}$, where $c \in \{\text{ER}, \text{NR}\}$, can be mapped back to raw intensity units by inverting the transformation:

$$\hat{\mathcal{I}}_c = s \cdot \left(10^{\hat{\mathcal{I}}_{c,\log} \cdot \max_{xy}(\mathcal{I}_{\text{in}}(x,y))} - 1\right), \tag{5}$$

ensuring that intensities can be reconstructed in physically meaningful units.

3.2 Model training

OASIS's main task is to extract the contents of an ER embedded within an NR in order to reconstruct characteristics of electron emission via the Migdal effect in order to compare with theory. It is therefore important for our model to be trained both on samples including ERs embedded within NRs (Migdal-effect signal events) and samples involving just NRs, in order for OASIS to learn the proper context necessary to avoid the hallucination of a non-existent ER signal. With this in mind, we utilize a mixed-class training sample consisting of 20,000 hybrid signal images and 20,000 NR-only images. Like the NRs used for the hybrid signal sample, the NR images used in the NR-only sample are also real NR tracks, but they are distinct from those used in the hybrid signal sample. Ground truth labels for the signal sample are 2-channel log-scaled (Equation (4)) intensity maps of the ER-only and NR-only channel of the input event, while for the NR-only sample, the ground truth ER channel is set to zero.

To train our U–Net, we split both classes of our training sample into a 90%–10% training–validation split. We train with a batch size of 16 and model weights are updated with the AdamW algorithm [57] using a learning rate of 2×10^{-4} and weight decay of 1×10^{-4}. Training is terminated using early stopping [58] with a patience of 10 epochs, meaning that training is

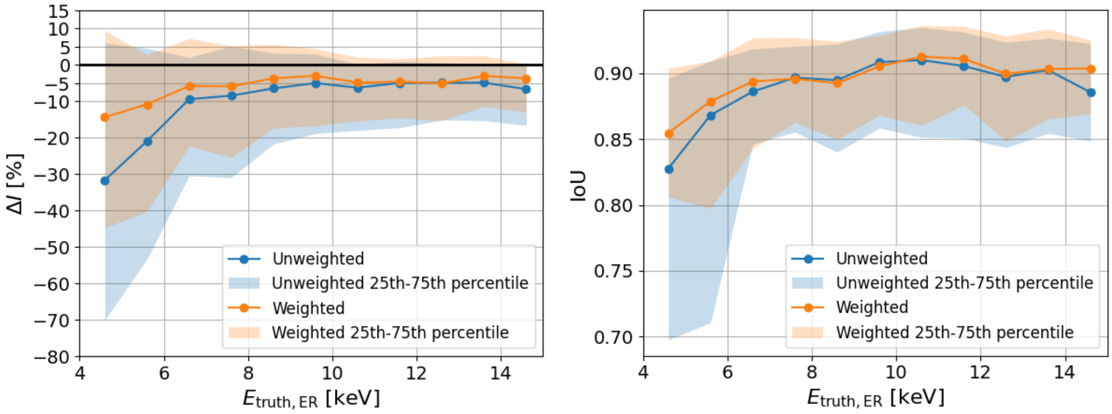

Figure 3. Comparison of OASIS's track reconstruction performance between the unweighted and weighted training campaigns. Left: Median (points) and 25th-75th percentile ranges (error bands) of the percent error of predicted ER track intensity compared to truth, ΔI, versus truth ER energy. Right: Median (points) and 25th-75th percentile ranges (error bands) of pixel IoU between OASIS's predicted ER and truth versus truth ER energy. In both cases, training OASIS with weights prioritizing regions of overlap lead to significant improvements at low ER energies, where pixel overlap is more significant in the ER channel.

terminated when there are 10 successive epochs over which the validation loss does not improve. Our model training and evaluation are accelerated with an Nvidia RTX 3090 consumer graphics card and are implemented in PyTorch [59].

4 Results applied to the MIGDAL experiment

We evaluate OASIS's performance on a test set consisting of 2,086 signal events (ER+NR) and 7,159 NR-only events. The test set was separated out from the training sample before training, so none of the test set events were used in training or validation.

4.1 Loss function weight reconstruction performance study

We first compare OASIS's ER intensity and topological reconstruction performance when being trained with and without region and channel-specific weights in its loss function. In both cases we evaluate OASIS's performance on the signal event test set using the network-weights corresponding to the best-achieved validation loss of that training campaign. We refer to these training campaigns as the "Weighted" and "Unweighted" trainings, respectively.

The weighted training campaign uses channel weights $w_{\mathrm{ER}} = 5$, $w_{\mathrm{NR}} = 1$, and region weights $W_{\mathrm{region}} = 3$ for ER-only pixels, $W_{\mathrm{region}} = 1$ for NR-only pixels, and $W_{\mathrm{region}} = 6$ for overlap pixels with nonzero values in both the ER and NR channels. These weights prioritize reconstruction of the fainter ER signal, with the highest weight assigned to overlap regions where signal attribution is most challenging. The ER channel receives a weight of 5 to counterbalance the orders-of-magnitude brightness difference between NR and ER tracks, while the overlap region weight of 6 ensures the network learns to disentangle signals where both are present. These weight values were selected based on preliminary performance comparisons across several weighting schemes. By contrast, the unweighted training campaign sets all channel and region weights to 1.

Figure 3 shows the results of this study both in terms of intensity attribution and topological reconstruction performance. The left panel shows each training campaign's median (points) and 25th-75th percentile ranges (error bands) of OASIS's percent error in *track* intensity reconstruction, $\Delta I \equiv \frac{\hat{I}_{\mathrm{ER}} - I_{\mathrm{truth,ER}}}{I_{\mathrm{truth,ER}}} \times 100\%$, where $I \equiv \sum_{xy} \mathcal{I}(x, y)$ represents track intensity, versus the ground truth energy of the ER track. The benefits of OASIS's overlap region weights become immediately apparent when comparing the weighted and unweighted training's intensity reconstruction performance in the lowest 4-5 keV energy bin. Compared to higher ER energies, events in this bin have a much larger proportion of the ER track toplogy buried beneath the footprint of the NR track, meaning these events, on average, have proportionately more ER-NR overlap. The fact that in this bin, the weighted training campaign significantly outperforms the unweighted training campaign in intensity attribution, ($\Delta I_{\mathrm{median,weighted}} = -14\%$ versus $\Delta I_{\mathrm{median,unweighted}} = -32\%$) shows that prioritizing the faint ER channel and regions of overlap in training achieves the desired effect of improving ER reconstruction in the overlap region where the overall signal is dominated by the NR.

The same conclusion can also be drawn for topological reconstruction. To quantify topological reconstruction, we use the pixel intersection-over-union (IoU) metric which we define as

$$\mathrm{IoU} \equiv \frac{|\hat{\mathcal{I}}_{\mathrm{ER}} \cap \mathcal{I}_{\mathrm{truth,ER}}|}{|\hat{\mathcal{I}}_{\mathrm{ER}} \cup \mathcal{I}_{\mathrm{truth,ER}}|}, \tag{6}$$

where the numerator represents the number of shared pixels between the predicted and truth ER intensity maps, while the denominator represents their union. Similar to intensity attribution, in IoU score, Figure 3 shows that the weighted training campaign significantly outperforms the unweighted training in the low energy ER regime where ER-NR overlap in the ER channel dominates. This is especially evident when comparing the 25th percentile IoU scores between the two training campaigns.

From these results it is clear that OASIS's use of region and channel weights significantly improves both pixel-level intensity attribution and topological reconstruction of low energy ER tracks. This improvement in the low-energy regime is particularly important because the Migdal effect cross section increases exponentially with decreasing ER energy [50], making accurate reconstruction in the 4-6 keV range essential for achieving sufficient statistics to test theoretical predictions.

Before making further conclusions, we next test OASIS's ER construction performance on the NR-only test set for both training campaigns. This test is important because it sheds light on false positive cases where OASIS hallucinates non-existent ER signal. To test this, we evaluate the predicted ER track intensity, $\hat{I}_{\mathrm{ER}}$, for true NR-only events, so $I_{\mathrm{truth,ER}} = 0$ for all of these events. It is inevitable that OASIS will always predict at least a small amount of ER signal, so we set a modest false positive detection threshold of 3 keV of visible energy, which is well below the 4 keV lower bound of the energy range analyzed in this study (see Figure 3). This means that if OASIS detects an ER channel intensity corresponding to at least 3 keV of visible energy from an NR-only input image, then we mark that event as a false positive ER detection. Running this test on both training campaigns, we observe false positive rates of 0.5% and 1.5% for the unweighted and weighted training sets, respectively. This tells us that while the weighted training significantly improves low energy ER reconstruction, it comes at the cost of it more likely reconstructing non-existent signal. In general, it is important to balance the tradeoffs between signal extraction performance and false positives, however for the purposes of the MIGDAL experiment, OASIS was designed to extract the most accurate representation of ERs that were *already* detected by MIGDAL's search pipeline. The higher false positive ER signal rate of the weighted training campaign is therefore not a concern because OASIS is not being used to search for Migdal effect candidates. Given this, we spend the remainder of this section presenting results from the weighted training campaign because it outperforms the unweighted campaign's ER reconstruction performance, especially at low energies where the ER channel tends to be dominated by regions of pixel-overlap. We first examine OASIS's reconstruction performance on individual test set events to illustrate typical reconstruction quality and failure modes.

4.2 Event-level reconstruction

Figure 4 illustrates OASIS's performance on four examples from the test set. For the shown input image, panels (a)-(d) of this figure show comparisons between OASIS's predicted ER and the ground truth ER associated with the input image. In panels (a)-(c), we compare (1) OASIS's *track* intensity prediction with truth, ΔI, and (2) assess the consistency of axial angular reconstruction by computing the angle between axial directions – determined by applying the same angle-finding algorithm – computed on the predicted and truth ER, $\Delta\phi \in [0^\circ, 90^\circ]$. Details of the angle finding algorithm are discussed in Appendix A.

Panel (a) shows a typical example of an event where the ER track is long enough and directed away from the NR so that a significant portion of its topology lies outside of the NR. In this case, OASIS does an excellent job reconstructing both the intensity ($\Delta I = -1\%$) and capturing directional information near the ER's vertex ($\Delta\phi = 1.8^\circ$). The example in panel (b) is noteworthy because the ER is hardly visible by eye in the input image, yet OASIS still successfully extracts the ER ($\Delta I = +6\%$, $\Delta\phi = 2.8^\circ$). Panel (c) illustrates a typical example where our model does not perform well. In this event, the truth ER travels inward toward the center of the NR. Despite weighting the ER-NR overlap region most heavily in our network training, this is most difficult region for OASIS to perform well, so it incorrectly predicts an ER poking slightly away from the NR. For this event, we find $\Delta I = -84\%$, and $\Delta\phi = 33^\circ$, hinting at a correlation between energy and angular reconstruction performance, which will be discussed later. Finally, panel (d) highlights typical reconstruction performance for an NR-only input event. In this case OASIS predicts about

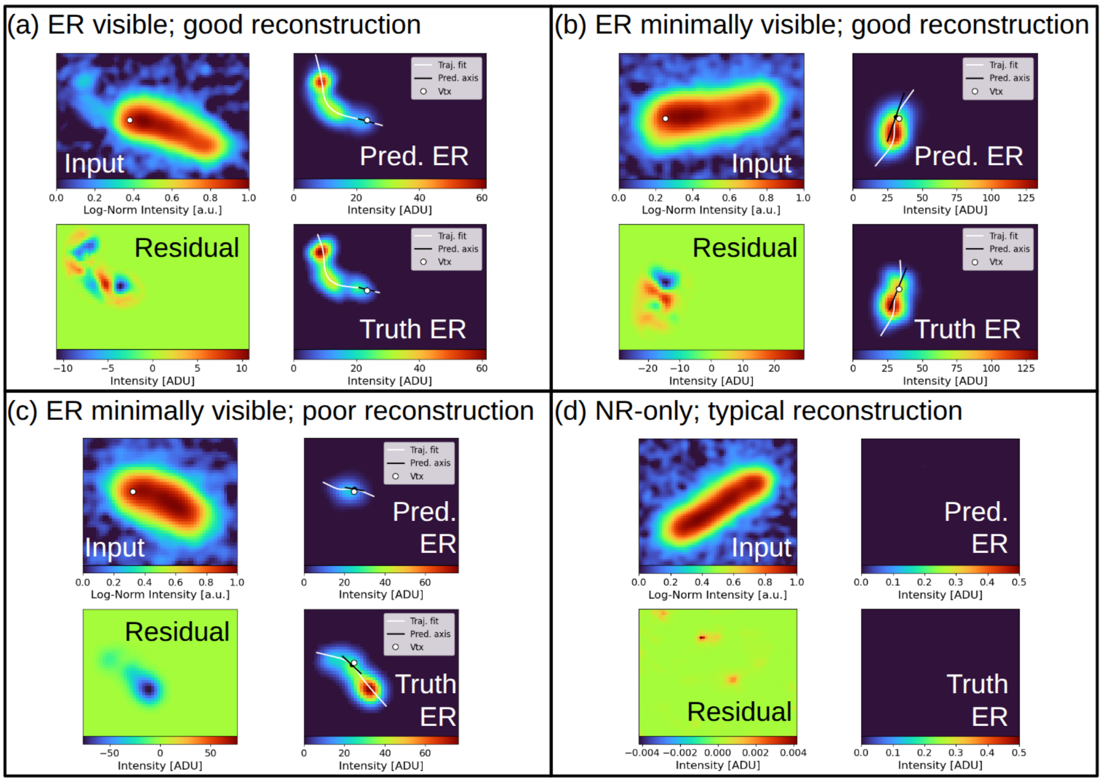

Figure 4. Four test set examples comparing OASIS's ER reconstruction performance to truth. In panels (a)-(c) both the truth and predicted ERs have principal curves shown in white with estimated directional axes shown in black. Panel (d) is an NR-only input event. Predicted ER and truth ER energies for each panel are: (a) $\hat{E}_{\mathrm{ER}} = 5.2\,\mathrm{keV}$, $E_{\mathrm{truth,ER}} = 5.2\,\mathrm{keV}$. (b) $\hat{E}_{\mathrm{ER}} = 5.5\,\mathrm{keV}$, $E_{\mathrm{truth,ER}} = 5.2\,\mathrm{keV}$. (c) $\hat{E}_{\mathrm{ER}} = 0.89\,\mathrm{keV}$, $E_{\mathrm{truth,ER}} = 5.7\,\mathrm{keV}$. (d) $\hat{E}_{\mathrm{ER}} = 0.6\,\mathrm{eV}$, $E_{\mathrm{truth,ER}} = 0\,\mathrm{eV}$, where $\hat{E}_{\mathrm{ER}} \equiv \frac{\hat{I}_{\mathrm{ER}}}{I_{\mathrm{truth,ER}}} E_{\mathrm{truth,ER}}$.

0.6 eV of visible energy which is effectively zero compared to the keV-scale ERs of interest in this study, so we treat this example as OASIS successfully determining that there was no ER in this input image. Having illustrated event-level performance, we now assess OASIS's intensity reconstruction across the full test set to quantify systematic trends and overall accuracy.

4.3 Intensity reconstruction performance

The green histogram in the left panel of Figure 5 shows the overall distribution of reconstructed ER track intensities for the NR-only sample. Noting the logarithmic scale and strong peaking near 0, we observe that OASIS correctly predicts virtually no ER signal in the vast majority of NR-only events. By contrast, in the hybrid signal sample, some amount of ER signal is nearly always detected as it should be.

The right panel of Figure 5 shows the correlation between OASIS's ER track intensity predictions with truth for events satisfying MIGDAL's search region of interest of $4 \leq E_{\mathrm{truth,ER}} \leq 15\,\mathrm{keV}$. We make the following observations from this distribution. First, the red line indicates perfect performance ($\hat{I}_{\mathrm{ER}} = I_{\mathrm{truth,ER}}$) and much of the data is tightly centered around this line. Second, while the mode of the $\hat{I}_{\mathrm{ER}}$ distribution at most slices of $I_{\mathrm{truth,ER}}$ aligns with perfect performance, at lower intensities, the model is more likely to under-predict the ER's intensity, as was already shown in Figure 3. Beyond intensity reconstruction, accurate angular information is essential for comparing observations with theoretical Migdal emission models. We therefore assess OASIS's angular reconstruction consistency using principal curve analysis.

4.4 Angular consistency performance

To assess OASIS's angular reconstruction performance, we fit principal curves (details in Appendix A) to the predicted and truth ER intensity maps for events in the signal sample of the test set satisfying $4 \leq E_{\mathrm{truth,ER}} \leq 15\,\mathrm{keV}$. Of these 1,750 events, principal curves were successfully fit on 1,639 of them (94%). Our interest in this study is assessing the topological similarity between OASIS's predictions and truth in a way that captures relevant detector performance metrics. We therefore use the axial angle between the reconstructed axis of the truth and predicted track

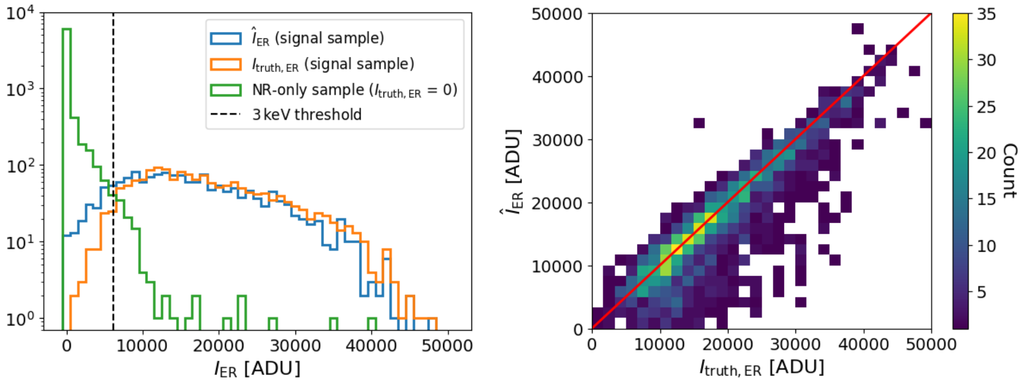

Figure 5. Left: Histograms of reconstructed and truth ER intensities for the test set signal sample (blue and orange, respectively), and reconstructed ER intensities for the NR-only sample (green). The black dashed line represents the 3 keV false positive threshold for ER detection in the NR-only sample. Right: Model's predicted ER track intensity ($\hat{I}_{\mathrm{ER}}$) versus truth ER track intensity ($I_{\mathrm{truth,ER}}$) for all test set ERs satisfying $4 \leq E_{\mathrm{truth,ER}} \leq 15\,\mathrm{keV}$. The red line shows equality between prediction and truth illustrating that OASIS, on average, nearly perfectly reconstructs energy.

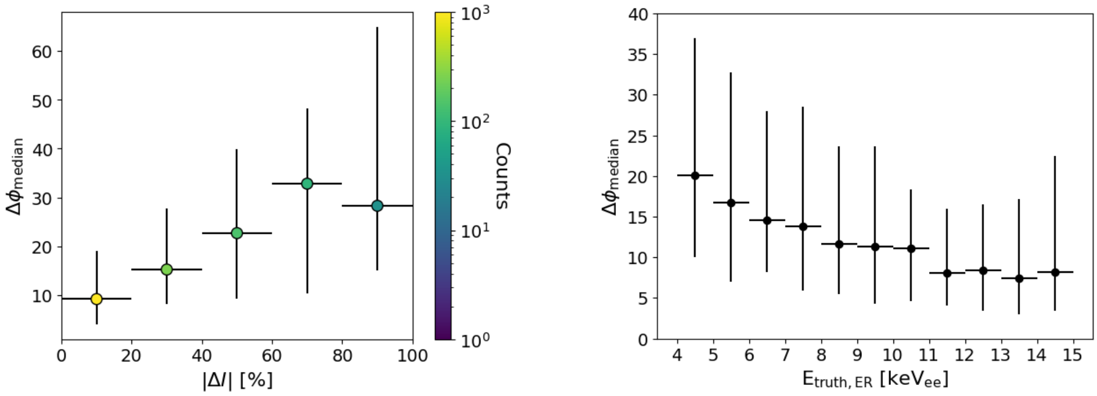

Figure 6. Left: Median (points) and inter-quartile range (error bars) of angular consistency $\Delta\phi$ versus $|\Delta I|$. Right: Median (points) and inter-quartile range (error bars) or $\Delta\phi$ versus truth ER energy.

topologies, $\Delta\phi \in [0, 90°]$, as our heuristic for topological similarity. The left panel of Figure 6 shows the median (points) and 25th-75th percentile values (errorbars) of $\Delta\phi$ versus the absolute percentage difference in reconstructed intensity versus truth, $|\Delta I|$. It is immediately clear that angular reconstruction performance is strongly correlated with intensity reconstruction performance. This tells us that in general, OASIS's ER intensity reconstructions are spatially correlated and tend toward the "correct answer" both in terms of energy and angle. While this is expected because we designed our loss function to simultaneously optimize event segmentation and energy reconstruction, this result confirms this optimization in a physically meaningful way.

The right panel of Figure 6 shows the median and 25th-75th percentile ranges of angular consistency versus truth ER energy. While we do not compute angular resolutions with respect to true recoil directions, the fact that median angular reconstruction consistencies are within about $20°$ over all energies $\geq 4\,\mathrm{keV}$ is a promising first step toward demonstrating sensitivity in measuring angular distributions of Migdal ER emissions. Indeed, if the angular resolution of a detector for isolated ERs is determined by some algorithm, then the angular resolution of ERs emitted via the Migdal effect can be estimated by adding $\Delta\phi$ – determined by the same algorithm – in quadrature to the detector's angular resolution.

5 Discussion

The regions where objects overlap are precisely where attribution is most ambiguous and informative. Yet these critical regions are typically given no special consideration in segmentation-regression training. OASIS addresses this by introducing explicit overlap-region weighting into the segmentation-regression loss function. By assigning both channel and region-specific weights to areas where multiple objects intersect, the framework directs training attention toward the regions where accurate attribution is most critical. The weighted loss function is computationally efficient, requiring no architectural changes to standard segmentation-regression networks, and is agnostic to the specific backbone architecture used.

When tasked with reconstructing the MIGDAL experiment's faint ER signal – one that is often buried within an orders-of-magnitude brighter NR signal – we showed that weighting OASIS's loss function to prioritize both regions of pixel overlap and ER-channel reconstruction significantly improves intensity attribution and event topology reconstruction. In particular, in the 4-5 keV ER energy regime where ER-NR overlap tends to dominate the ER signal, assigning an ER channel weight of 5, an ER-only region weight of 3, and an overlap region weight of 6, while leaving all NR-weights set to 1 in OASIS's loss function, improves both median predicted ER track intensity attribution errors from -32% to -14% and median IoU scores from 0.828 to 0.855 compared to an unweighted loss function. Linking this topological reconstruction performance to more physically meaningful quantities, we observe angular consistencies derived from principal curve fits to within 20° in this low energy regime. Given the exponentially rising Migdal emission probability with decreasing ER energy, achieving sensitivity to reconstruct ER energies and angles in this $\leq$5 keV regime is essential for achieving sufficient statistics to meaningfully test theoretical predictions. OASIS's angular sensitivity in this regime shows promise toward the current generation MIGDAL detector's capability of testing theoretical models of Migdal effect angular production, which we will report on in future work.

Beyond MIGDAL, the principles underlying OASIS extend to other scientific imaging domains where overlapping objects must be separated. In astronomy there are publicly accessible datasets like the Galaxy Zoo DECaLS database [60], which contains hundreds of thousands of images that are pre-annotated for the purpose of training models for tasks like "deblending" overlapping galaxy images in crowded fields. The fact that such datasets are publicly available and that intensity attribution in overlapping regions is an essential component of galaxy deblending makes it an attractive avenue to apply OASIS. The OASIS framework can be readily extended to support this, as well as other diverse applications: n-channel inputs enable multi-band astronomical imaging, while full 3D voxel support (3 spatial dimensions + intensity) would enable application to volumetric medical imaging and upgraded particle detectors with complete 3D reconstruction capabilities. All together, OASIS's ability to recover information from challenging regions of object-overlap gives it the potential to improve scientific image analyses across disciplines.

Data Availability Statement

All code supporting this study is openly available at https://github.com/jschuel/OASIS.

Acknowledgments

This work has been supported by the UKRI's Science & Technology Facilities Council through the Xenon Futures R & D programme (awards ST/T005823/1, ST/T005882/1, ST/V001833/1, ST/V001876/1), Consolidated Grants (ST/W000636/1, ST/X006042/1, ST/T000759/1, ST/W000652/1, ST/S000860/1, ST/X005976/1), and TM's and LM's PhD scholarships (ST/T505894/1, ST/X508913/1); by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award Number DE-SC0022357; by the U.S. National Science Foundation under Award number 2209307; ET acknowledges the Graduate Instrumentation Research Award funded by the U.S. Department of Energy, Office of Science, Office of High Energy Physics; by the Portuguese Foundation for Science and Technology (FCT) under award number PTDC/FIS-PAR/2831/2020; and by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 841261 (DarkSphere) and No. 101026519 (GaGARin). ELA acknowledges the support from Spanish grant CA3/RSUE/2021-00827, funded by Ministerio de Universidades, Plan de Recuperacion, Transformacion y Resiliencia, and Universidad Autonoma de Madrid. KN acknowledges support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2121 "Quantum Universe" – 390833306. We are grateful to the Particle Physics Department at RAL for significant additional support which made this project possible. Thanks are also due to the CERN RD51 collaboration for their support through Common Project funds, hardware tests and training, and useful discussions. We would also like to thank the ISIS facility for technical assistance and for hosting this experiment. We thank Master's students M. Handley (Cambridge), R. Hafeji (Surrey), and G. Buzzard (Liverpool).

A Appendix A: Directional reconstruction algorithm

When assessing the directional reconstruction performance of OASIS, our aim is to compare angles reconstructed in a consistent way between our model's predicted ER image and the truth ER image. To this aim, any angular reconstruction algorithm can be used to estimate $\Delta\phi$. Because electron recoil trajectories are varied, and often non-linear, we opt to evaluate $\Delta\phi$ using an algorithm that captures track curvature. The principal curve fitting technique [61] does just this and has a readily-adaptable implementation [62] in the R programming language [63]. Using this implementation, our angular reconstruction procedure for a given ER intensity map is as follows:

1. Sample 20,000 coordinates drawn from the non-zero bin-centers of the input intensity map, weighted by the log-scale intensity from Equation (4). We do this because the principal curve fitting implementation does not support weights as inputs, but does properly weight duplicate coordinates.

2. Generate a principal curve from the 20,000 input coordinates. We apply a convergence threshold of 0.25 [62] and interpolate the curve with a smoothed spline.

3. Construct a principal axis as the tangent line to the location along the principal curve that has the closest Euclidean distance to the truth ER vertex.

After performing steps 1.–3. on both OASIS's predicted ER intensity map and on the truth ER intensity map, we compute $\Delta\phi$ as the angle corresponding to the dot product between the two normalized axes, folded to the domain $\Delta\phi \in [0°, 90°]$.

Our principal axis construction is idealized in the sense that we utilize the truth ER's vertex position from simulation. ER vertices are challenging to accurately pinpoint, however since the ER and NR share a vertex in the Migdal effect, we can construct principal axes for Migdal ERs at our estimate of the NR vertex position. Our collaboration recently developed an algorithm based on our previous work on curvilinear ridge detection [51, 52] that correctly reconstructs NR vertices to sub-pixel-level precision (publication forthcoming). Assuming this vertex placement algorithm behaves similarly on OASIS's predicted *NR* outputs to how it does on raw NRs, then we expect ER directions evaluated at the point on a principal curve nearest to the predicted NR vertex to be similar to those evaluated nearest to the truth ER vertex.

References

[1] P. Melchior et al. SCARLET: Source separation in multi-band images by Constrained Matrix Factorization. *Astronomy and Computing*, 24:129, July 2018.

[2] Artur Speiser et al. Deep learning enables fast and dense single-molecule localization with high accuracy. *Nature Methods*, 18(9):1082–1090, 2021.

[3] Qikui Zhu, Bo Du, and Pingkun Yan. Boundary-weighted Domain Adaptive Neural Network for Prostate MR Image Segmentation. *arXiv e-prints*, page arXiv:1902.08128, February 2019.

[4] Hao Jiang, Rushan Zhang, Yanning Zhou, Yumeng Wang, and Hao Chen. DoNet: Deep De-overlapping Network for Cytology Instance Segmentation. *arXiv e-prints*, page arXiv:2303.14373, March 2023.

[5] Huiyu Li et al. Deep Distance Map Regression Network with Shape-aware Loss for Imbalanced Medical Image Segmentation. *arXiv e-prints*, page arXiv:2501.09116, January 2025.

[6] Saúl Alonso-Monsalve et al. Deep-learning-based decomposition of overlapping-sparse images: application at the vertex of simulated neutrino interactions. *Commun. Phys.*, 7(1):173, 2024.

[7] A. Abed Abud et al. Neutrino interaction vertex reconstruction in DUNE with Pandora deep learning. *Eur. Phys. J. C*, 85(697):697, 2025.

[8] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241. Springer, 2015.

[9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. *arXiv e-prints*, page arXiv:1703.06870, March 2017.

[10] Angel Torrado-Carvajal et al. Dixon-VIBE Deep Learning (DIVIDE) Pseudo-CT Synthesis for Pelvis PET/MR Attenuation Correction. *Journal of Nuclear Medicine*, 2018.

[11] Colin J. Burke et al. Deblending and classifying astronomical sources with Mask R-CNN deep learning. *MNRAS*, 490(3):3952–3965, December 2019.

[12] Ran Zhang et al. Deblending overlapping galaxies in DECaLS using transformer-based algorithm: A method combining multiple bands and data types. *PASA*, 41:e035, May 2024.

[13] Florin C. Walter, Sebastian Damrich, and Fred A. Hamprecht. MultiStar: Instance Segmentation of Overlapping Objects with Star-Convex Polygons. *arXiv e-prints*, page arXiv:2011.13228, November 2020.

[14] Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Deep Occlusion-Aware Instance Segmentation with Overlapping BiLayers. *arXiv e-prints*, page arXiv:2103.12340, March 2021.

[15] V. Alvarez et al. NEXT-100 Technical Design Report (TDR): Executive Summary. *JINST*, 7:T06001, 2012.

[16] M. Kekic et al. Demonstration of background rejection using deep convolutional neural networks in the NEXT experiment. *JHEP*, 01:189, 2021.

[17] A. Simón et al. Boosting background suppression in the NEXT experiment through Richardson-Lucy deconvolution. *JHEP*, 07:146, 2021.

[18] K. Miernik et al. Optical Time Projection Chamber for imaging nuclear decays. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 581(1-2):194–197, oct 2007.

[19] Miernik, K. et al. Two-proton radioactivity of ^{45}Fe. *Eur. Phys. J. A*, 42(3):431–439, 2009.

[20] M. Pomorski et al. Proton spectroscopy of ^{48}Ni,^{46}Fe, and ^{44}Cr. *Phys. Rev. C*, 90:014311, Jul 2014.

[21] R. Acciarri et al. Design and Construction of the MicroBooNE Detector. *JINST*, 12(02):P02017, 2017.

[22] P. Abratenko et al. Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE. *Phys. Rev. D*, 103(5):052012, 2021.

[23] P. Abratenko et al. Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber. *Phys. Rev. D*, 103(9):092003, 2021.

[24] P. Abratenko et al. Search for an anomalous excess of charged-current quasielastic νe interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction. *Phys. Rev. D*, 105(11):112003, 2022.

[25] Vitor F. Grizzi, Margaret Voetberg, Giuseppe Cerati, Hadi Meidani, and V. Hewes. NuGraph2 with Context-Aware Inputs: Physics-Inspired Improvements in Semantic Segmentation. 9 2025.

[26] A. Aurisano, V. Hewes, G. Cerati, J. Kowalkowski, C. S. Lee, W. Liao, D. Grzenda, K. Gumpula, and X. Zhang. Graph neural network for neutrino physics event reconstruction. *Phys. Rev. D*, 110(3):032008, 2024.

[27] H. M. Araújo et al. The MIGDAL experiment: Measuring a rare atomic process to aid the search for dark matter. *Astropart. Phys.*, 151:102853, 2023.

[28] Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning. *arXiv e-prints*, page arXiv:1702.03118, February 2017.

[29] Vania V. Estrela, Hermes Aguiar Magalhaes, and Osamu Saotome. Total Variation Applications in Computer Vision. *arXiv e-prints*, page arXiv:1603.09599, March 2016.

[30] A. Migdal. Ionizatsiya atomov pri yadernykh reaktsiyakh (Ionization of atoms in nuclear reactions). *ZhETF*, 9:1163–1165, 1939.

[31] Masahiro Ibe, Wakutaka Nakano, Yutaro Shoji, and Kazumine Suzuki. Migdal Effect in Dark Matter Direct Detection Experiments. *JHEP*, 03:194, 2018.

[32] D. S. Akerib et al. Results of a Search for Sub-GeV Dark Matter Using 2013 LUX Data. *Phys. Rev. Lett.*, 122(13):131301, 2019.

[33] E. Armengaud et al. Searching for low-mass dark matter particles with a massive Ge bolometer operated above-ground. *Phys. Rev. D*, 99(8):082003, 2019.

[34] Z. Z. Liu et al. Constraints on Spin-Independent Nucleus Scattering with sub-GeV Weakly Interacting Massive Particle Dark Matter from the CDEX-1B Experiment at the China Jinping Underground Laboratory. *Phys. Rev. Lett.*, 123(16):161301, 2019.

[35] E. Aprile et al. Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T. *Phys. Rev. Lett.*, 123(24):241803, 2019.

[36] Liron Barak et al. SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper-CCD. *Phys. Rev. Lett.*, 125(17):171802, 2020.

[37] G. Adhikari et al. Searching for low-mass dark matter via the Migdal effect in COSINE-100. *Phys. Rev. D*, 105(4):042006, 2022.

[38] E. Armengaud et al. Search for sub-GeV dark matter via the Migdal effect with an EDELWEISS germanium detector with NbSi transition-edge sensors. *Phys. Rev. D*, 106(6):062004, 2022.

[39] P. Agnes et al. Search for Dark-Matter–Nucleon Interactions via Migdal Effect with DarkSide-50. *Phys. Rev. Lett.*, 130(10):101001, 2023.

[40] K. Abe et al. Direct dark matter searches with the full data set of XMASS-I. *Phys. Rev. D*, 108(8):083022, 2023.

[41] J. Aalbers et al. Search for new physics in low-energy electron recoils from the first LZ exposure. *Phys. Rev. D*, 108(7):072006, 2023.

[42] Di Huang et al. Search for Dark-Matter–Nucleon Interactions with a Dark Mediator in PandaX-4T. *Phys. Rev. Lett.*, 131(19):191002, 2023.

[43] M. F. Albakry et al. Search for low-mass dark matter via bremsstrahlung radiation and the Migdal effect in SuperCDMS. *Phys. Rev. D*, 107(11):112013, 2023.

[44] Prakruth Adari et al. First Direct-Detection Results on Sub-GeV Dark Matter Using the SENSEI Detector at SNOLAB. *Phys. Rev. Lett.*, 134(1):011804, 2025.

[45] K. Aggarwal et al. Probing Benchmark Models of Hidden-Sector Dark Matter with DAMIC-M. *Phys. Rev. Lett.*, 135(7):071002, 2025.

[46] Jingke Xu et al. Search for the Migdal effect in liquid xenon with keV-level nuclear recoils. *Phys. Rev. D*, 109(5):L051101, 2024.

[47] Orca-quest qcmos camera c15550-20up technical note. Technical report, Hamamatsu Photonics, 05 2022.

[48] MIGDAL Collaboration. Performance of the Optical Time Projection Chamber of the MIGDAL Experiment. *[In Progress]*, 2025.

[49] J. Schueler et al. Transforming a rare event search into a not-so-rare event search in real-time with deep learning-based object detection. *Phys. Rev. D*, 111(7):072004, 2025.

[50] Peter Cox, Matthew J. Dolan, Christopher McCabe, and Harry M. Quiney. Precise predictions and new insights for atomic ionization from the Migdal effect. *Phys. Rev. D*, 107(3):035032, 2023.

[51] Carsten Steger. An Unbiased Detector of Curvilinear Structures. *IEEE Trans. Pattern Anal. Mach. Intell.*, 20:113–125, 1998.

[52] Elizabeth Tilly and Magnus Handley. 3D track reconstruction of low-energy electrons in the MIGDAL low pressure optical time projection chamber. *JINST*, 18(07):C07013, 2023.

[53] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics YOLOv8, 2023.

[54] Sebastian Höfer et al. Sim2real in robotics and automation: Applications and challenges. *IEEE Transactions on Automation Science and Engineering*, 18(2):398–400, 2021.

[55] D. Edgeman et al. Optical effects in Gaseous Electron Multipliers. *[In Progress]*, 2025.

[56] Jeffrey Schueler. `migYOLO`, July 2024. Available at https://doi.org/10.5281/zenodo.12628437 and https://github.com/jschuel/migYOLO.

[57] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. *arXiv e-prints*, page arXiv:1711.05101, November 2017.

[58] Lutz Prechelt. *Early Stopping — But When?*, pages 53–67. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[59] Adam Paszke et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

[60] Mike Walmsley et al. Galaxy Zoo DECaLS: Detailed visual morphology measurements from volunteers and deep learning for 314 000 galaxies. *MNRAS*, 509(3):3966–3988, January 2022.

[61] Trevor Hastie and Werner Stuetzle. Principal curves. *Journal of the American Statistical Association*, 84(406):502–516, June 1989. https://web.stanford.edu/~hastie/Papers/principalcurves.pdf.

[62] Robrecht Cannoodt. princurve 2.0: Fit a principal curve in arbitrary dimension, June 2018.

[63] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2023.