



PAPER

OPEN ACCESS

RECEIVED
10 December 2025REVISED
28 April 2026ACCEPTED FOR PUBLICATION
14 May 2026PUBLISHED
8 June 2026

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Overlap-aware segmentation for topological reconstruction of obscured objects

J Schueler^{1,*}, H M Araújo^{2,3}, S N Balashov³, J E Borg⁴, C Brew³, F M Brunbauer⁵, C Cazzaniga⁶, A Cottle⁷, D Edgeman¹, C D Frost⁶, F Garcia⁹, D Hunt⁸, M Kastriotou⁶, P Knights¹⁰, H Kraus⁸, A Lindote¹¹, M Lisowska⁵, D Loomba¹, E Lopez Asamar¹², P A Majewski^{2,3}, T Marley², C McCabe¹³, L Millins^{3,10}, R Nandakumar³, T Neep¹⁰, F Neves¹¹, K Nikolopoulos^{10,14}, E Oliveri⁵, A Roy², Timothy J Sumner², E Tilly¹, W Thompson¹ and M A Vogiatzi^{2,3}

¹ Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87131, United States of America

² Department of Physics, Blackett Laboratory, Imperial College London, London SW7 2AZ, United Kingdom

³ Particle Physics Department, STFC Rutherford Appleton Laboratory, Didcot OX11 0QX, United Kingdom

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, United Kingdom

⁷ University College London (UCL), Department of Physics and Astronomy, London WC1E 6BT, United Kingdom

⁸ Department of Physics, Keble Road, University of Oxford, Oxford OX1 3RH, United Kingdom

⁹ Helsinki Institute of Physics, University of Helsinki, FI-00014 Helsinki, Finland

¹⁰ School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, United Kingdom

¹¹ LIP—Laboratório de Instrumentação e Física Experimental de Partículas, University of Coimbra, P-3004-516 Coimbra, Portugal

¹² Departamento de Física Teórica, Universidad Autónoma de Madrid, 28049 Madrid, Spain

¹³ Department of Physics, King's College London, London WC2R 2LS, United Kingdom

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 order(s)-of-magnitude brighter nuclear recoil track. Compared to unweighted segmentation regression, we demonstrate OASIS's novel overlap region-targeted loss function weight to be the single most important training weight for improving intensity and topological reconstructions of the low-energy electron tracks that tend to be most dominated by pixel overlap. Averaging over eight training campaigns, we further show the addition of overlap-targeted weights to improve median intensity reconstruction errors from -41.1% to -13.3% for these low-energy electrons. These performance gains demonstrate OASIS as a generalizable methodology for recovering obscured signals in overlap-dominated regions. All code is openly available to facilitate cross-domain adoption.

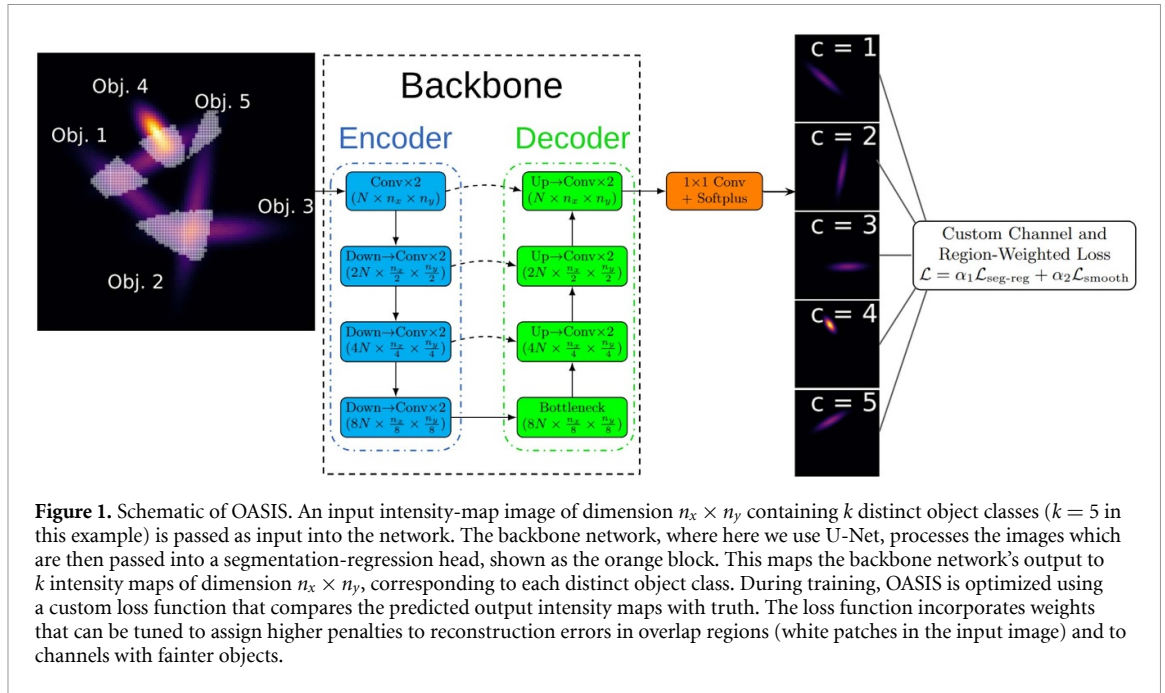
1. Introduction

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 [3–5]; 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 objective 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 overlap-aware training to segmentation regression. Specifically, OASIS employs programmable weights in its loss function that target both objects and regions. Object-targeted weights prioritize pixel-wise intensity attribution for specific object classes and are relatively standard in computer vision. Region-targeted weights, however, are novel and can be tuned to specifically target regions of pixel overlap, allowing training to prioritize intensity attribution in the most ambiguous pixels. The combination of overlap-aware training and segmentation-regression is well suited for high resolution particle detectors like gas time projection chambers (TPCs), where detector images can contain overlapping ionization signals with physically meaningful intensities.

Gas TPCs are widely used in nuclear and particle physics, where their high granularity imaging capability enables reconstruction of energy deposition along particle trajectories. Example applications in which overlapping track signals arise, and where component-wise segmentation-regression could therefore be valuable include neutrinoless double beta decay searches like NEXT [15–17], optical TPC studies of exotic nuclear decay processes [18–20], liquid argon TPC experiments such as MicroBooNE [21], which already employs segmentation in its reconstruction chain [22–26], and searches for the Migdal effect in nuclear scattering by the MIGDAL and MARVEL experiments [27, 28]. From a reconstruction standpoint, searches for the Migdal effect are a particularly challenging case because the target signal is an image containing a bright nuclear recoil (NR) track that heavily overlaps with an order(s)-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. In this work, we focus on the MIGDAL experiment, whose high-resolution imaging readout and strongly overlapping NR-ER topology provide a direct and challenging 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 present the results of systematic ablation studies comparing intensity and topological reconstruction of overlap-dominated ER tracks across 36 trained OASIS models with distinct training-weight configurations. We then use the best performing model from this study to further quantify OASIS's intensity and angular reconstruction performance. 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 cases 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 backbone network¹⁵, 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 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), as well as regions and channels corresponding to specific objects. Loss function weights that disproportionately penalize misattribution of pixels within regions of object-overlap effectively prioritize reconstructing objects in regions where intensity attribution is most challenging.

2.1. Network architecture

The U-Net we employ 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-channel¹⁶ 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 [29] 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, $\hat{\mathcal{I}}_{\text{enc}1} \in \mathbb{R}^{N \times n_x \times n_y}$. 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 $\hat{\mathcal{I}}_{\text{enc}4} \in \mathbb{R}^{8N \times \frac{n_x}{8} \times \frac{n_y}{8}}$ 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

¹⁵ We employ a standard U-Net backbone in this work, however other suitable backbone networks could be applied.

¹⁶ 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.

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 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 $\hat{\mathcal{I}} = [\hat{\mathcal{I}}_1, \dots, \hat{\mathcal{I}}_k]^T \in \mathbb{R}^{k \times n_x \times n_y}$. 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 novel loss function consists of a custom-weighted pixel-intensity regression loss term and a smoothness regularization term

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

where α_1 and α_2 are weights assigned to the two loss terms. 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 $\hat{\mathcal{I}}_c$ and ground truth \mathcal{I}_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:

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

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 [30] 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 does not dominate the overall loss:

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

The specific choice of loss function weights, preprocessing strategies, training procedures, and evaluation metrics depends on the characteristics of the application domain. While the OASIS framework is general and can be applied to any imaging problem requiring decomposition of any number of overlapping sources into constituent intensity maps, the number of tunable loss function weights in OASIS scales as $n + 2^n - 1$, where n is the number of target object classes to reconstruct. For problems involving the separation of many object classes, the large number of tunable weights presents a practical limitation in our ability to systematically study the effects different combinations of loss function weights have on performance. This partly motivates our choice of Migdal effect searches as the application used to demonstrate OASIS. In addition to presenting a genuinely difficult reconstruction problem, where a faint ER track must be separated from an order(s)-of-magnitude brighter and heavily overlapping NR track,

the problem contains only two objects of interest. As a result, OASIS has only five adjustable loss function weights in this application.

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 Migdal in 1939 [31], this rare effect in quantum mechanics gained renewed attention when Ibe *et al* [32] 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 [33–46], and the effect was observed for the first time in 2026 [28].

The MIGDAL experiment [27] aims to measure and characterize the Migdal effect using fast neutrons from a deuterium-deuterium generator incident on an optical TPC filled with 50 Torr CF_4 gas. Images are recorded with a Hamamatsu ORCA Quest qCMOS camera [47], which captures 2048×1152 images of the scintillation light produced by neutron scattering events that are amplified by two glass gas electron multipliers (GEMs) [48]. Images are acquired in continuous rolling shutter mode at a rate of 120 frames per second and view an $8.0 \times 4.5 \text{cm}^2$ region through an EHD-25 085-C F0.85 lens [49]. At nominal operating GEM gains, the 16-bit pixels of the CMOS sensor provide the dynamic range necessary to simultaneously image and reconstruct both sub-4 keV ERs and up to $\mathcal{O}(\text{MeV})$ -scale fast-neutron induced NRs without saturating. In this work, we apply OASIS to both real NR tracks and hybrid simulated images where simulated ER tracks are stitched onto real NRs, mimicking the Migdal effect signal. Further details of the MIGDAL experiment, detector, and optical camera readout system are given in [27, 50, 51].

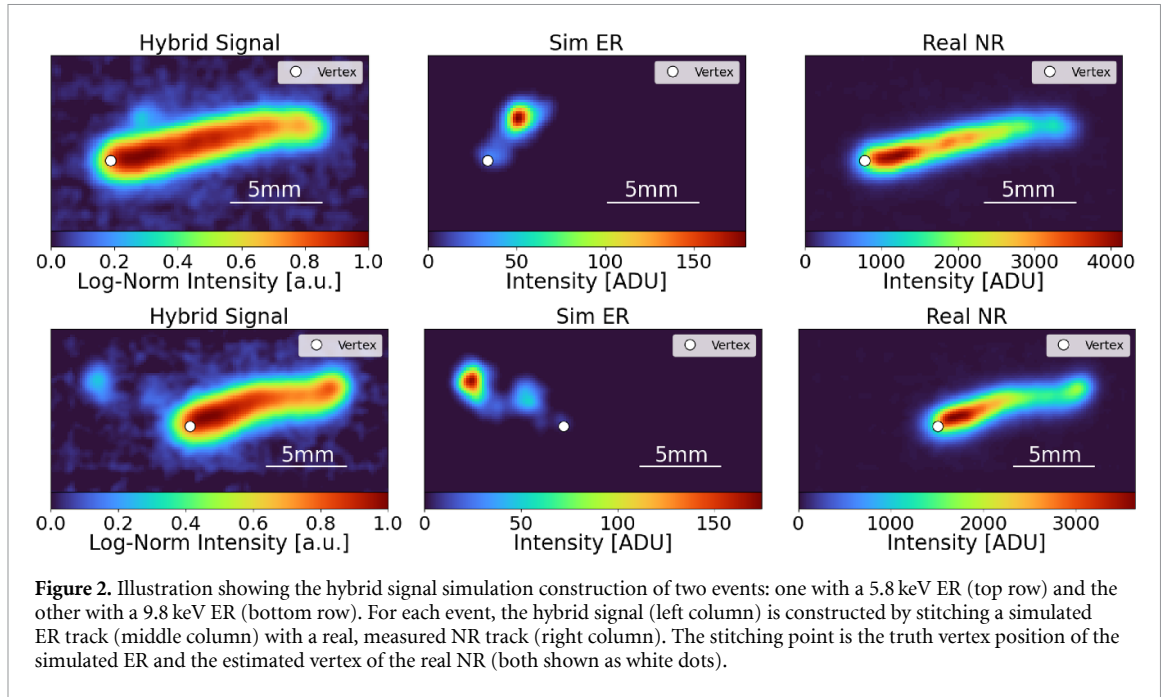
In the region of interest for the MIGDAL experiment’s search (ER energies between $\sim 4 \text{keV}$ and $\sim 15 \text{keV}$), the Migdal effect cross section increases exponentially with decreasing ER energy [52]. This poses a challenge for searches in low pressure electron drift gas mixtures for 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 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 order(s)-of-magnitude brighter NRs. To do this, we simulate Migdal-effect signal events using the hybrid approach adopted in [51], 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 track¹⁷. 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 [55] pipeline used to define bounding box regions surrounding real NRs are described in [51]. OASIS’s task is to decompose each hybrid input image, $\mathcal{I}_{\text{in}}(x, y)$ into $k = 2$ output intensity maps: $\hat{\mathcal{I}}_{\text{ER}}(x, y)$ representing the ER contribution, and $\hat{\mathcal{I}}_{\text{NR}}(x, y)$, 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 [56]; previous MIGDAL simulation studies [50, 51] have shown excellent pixel-level intensity and topological agreement between simulated and measured ER tracks, but comparable fidelity has not been achieved for NRs [57]. Second, the hybrid approach produces clean ground truth labels: since detector noise is already present in the real NR images and the simulated ERs are noiseless, we obtain truth ER labels,

¹⁷ In [51], we use the pixel of highest intensity in the NR track as a proxy for its vertex. We now assign NR vertices more accurately using an algorithm based on curvilinear ridge detection [53, 54] that will be presented in future work.



$\mathcal{I}_{\text{truth,ER}}$, 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 [58], 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 $\mathcal{I}(x, y) \in \mathbb{R}^{512 \times 288}$, 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:

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

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}_{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. $\mathcal{I}_{\text{in}} \in \mathbb{R}^{512 \times 288}$
2. $(\mathcal{I}_{\text{truth,ER}}, \mathcal{I}_{\text{truth,NR}}) \in \mathbb{R}^{2 \times 512 \times 288}$
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 [51] and the truth boundary of the

simulated ER¹⁸. 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:

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

ensuring that intensities can be reconstructed in physically meaningful units.

3.2. Model training

OASIS's main task in this application 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 and is therefore an empty image.

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 [59] using a learning rate of 2×10^{-4} and weight decay of 1×10^{-4} . Training is terminated using early stopping [60] with a patience of 10 epochs, meaning that training is terminated when there are 10 successive epochs over which the validation loss does not improve. Ultimately, we perform 36 distinct training campaigns, each with different combinations of loss function weights, w_c and W_{region} , to systematically study the effects these weights have on OASIS's ability to reconstruct the faint ER channel. Each campaign was trained and evaluated across pairs of either Nvidia TITAN RTX or Nvidia A100 80GB graphics cards using PyTorch's [61] Distributed Data Parallel platform [62]. We do not perform any additional architecture or hyperparameter tuning, as our primary interest is quantifying the effects OASIS's novel loss function weights have on ER track reconstruction.

4. Results applied to the MIGDAL experiment

We evaluate OASIS's performance on a test set consisting of 2086 signal events (ER+NR) and 7159 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 weighting ablation study

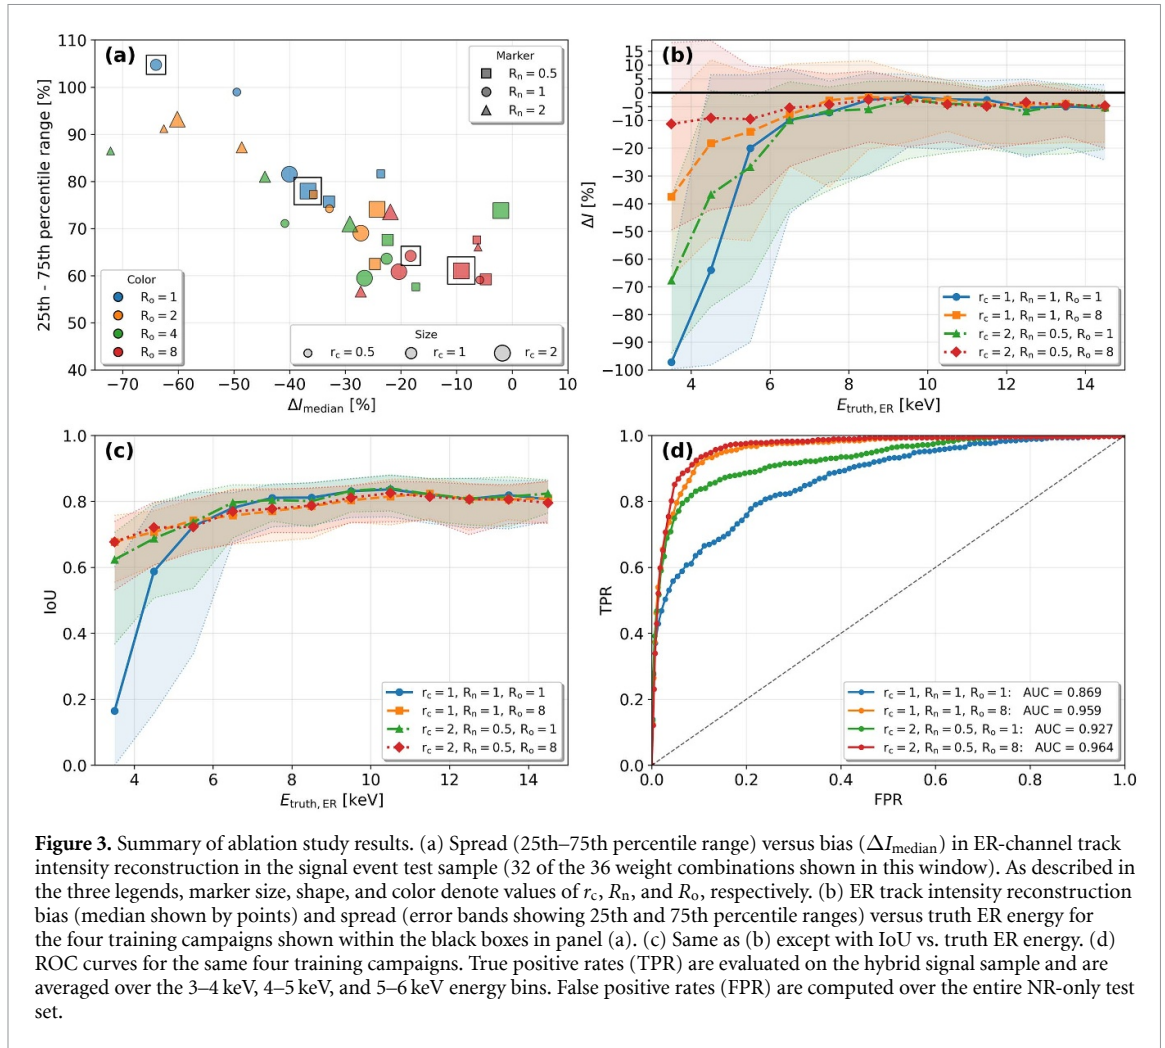
To evaluate the effects OASIS's region and channel-specific loss function weights have on ER intensity and topological reconstruction, we perform a systematic ablation study across 36 training campaigns. The weights we vary are the ER and NR channel-specific weights w_{ER} and w_{NR} , respectively, and the three region-specific weights W_{ER} , W_{NR} , and W_{overlap} , consisting of ER-only, NR-only, and overlap region pixels, respectively. To reduce the dimension of weight combinations, we set baselines of $w_{\text{NR}} = 1$ and $W_{\text{ER}} = 1$ and then define the following three ratios to vary in our ablation studies

$$r_c \equiv \frac{w_{\text{ER}}}{w_{\text{NR}}}, \quad R_n \equiv \frac{W_{\text{NR}}}{W_{\text{ER}}}, \quad R_o \equiv \frac{W_{\text{overlap}}}{W_{\text{ER}}}. \quad (6)$$

The results from evaluating OASIS models trained with all 36 unique combinations of $r_c \in [0.5, 1, 2]$, $R_n \in [0.5, 1, 2]$, and $R_o \in [1, 2, 4, 8]$, are summarized in figure 3 and described in detail below.

We are most interested in testing OASIS's ability to pixel-wise reconstruct the intensity map of the faint ERs buried within the significantly brighter NR. To this end, the metrics we reconstruct are (1) the

¹⁸ In real data where there is no knowledge of the truth ER boundary, B could be defined as the NR bounding box determined by YOLO with an appropriate amount of additional padding to contain any expected Migdal ER in a given search region of interest. For a significant majority of hybrids with sub-6 keV ERs evaluated in this work (those that tend to be most dominated by ER-NR overlap), it turns out the NR bounding box found by YOLO is equivalent to the hybrid bounding box, so no additional padding is needed. See the discussion in section 4.1 for more details.



bias and spread of the percent error of *track* intensity reconstruction compared to truth, (2) the pixel intersection-over-union (IoU) of the reconstructed ER track compared to truth, and (3) the false positive rate (FPRs) of ERs reconstructed in the NR-only test sample. More explicitly, we define the track intensity reconstruction percent error as

$$\Delta I \equiv \frac{\hat{I}_{\text{ER}} - I_{\text{truth,ER}}}{I_{\text{truth,ER}}} \times 100\%, \quad (7)$$

where $I \equiv \sum_{xy} \mathcal{I}(x,y)$ represents track intensity, and IoU as

$$\text{IoU} \equiv \frac{|\hat{\mathcal{I}}_{\text{ER}} \cap \mathcal{I}_{\text{truth,ER}}|}{|\hat{\mathcal{I}}_{\text{ER}} \cup \mathcal{I}_{\text{truth,ER}}|}, \quad (8)$$

where the numerator represents the number of shared pixels between the predicted and truth ER intensity maps, while the denominator represents their union. When computing IoU, we only include pixels with intensity of at least 1 ADU¹⁹. We report the bias and spread of both ΔI and IoU as their median and 25th–75th percentile range, respectively. Taken together, these two quantities characterize both ER intensity attribution and topological reconstruction performance.

Figure 3(a) summarizes intensity attribution bias and spread across all ablations for ER tracks in the 4–5 keV energy bin. We intentionally investigate a low energy ER bin because the region of pixel overlap between the ER and NR, on average, encompasses a larger proportion of low energy ER tracks compared to high energy tracks. The effects of overlap-weighted training should therefore be most pronounced for low energy ER tracks, which is confirmed in panels (b) and (c) of figure 3. Independent of R_n and r_c ,

¹⁹ Because we train with real-data NRs and simulated ERs, noise from the ORCA quest readout is only present in the NR channel. As a result, OASIS reconstructs ERs without noise, hence the low 1 ADU threshold for IoU masking.

Table 1. Bias (ΔI_{median}) and spread (25th-Curie grant agreement-75th percentile range) of track intensity reconstruction in the 4–5 keV ER-energy bin of the hybrid signal test set, averaged across all ablations of r_c and R_n for each value of overlap-weight ratio R_o . Increasing overlap penalty weights in the OASIS’s loss function significantly improves the intensity reconstruction performance of low energy ER tracks that are proportionately dominated by regions of pixel overlap, compared to higher energy tracks.

R_o	Bias (%)	Spread (%)
1	−41.1	86.8
2	−39.5	78.6
4	−30.8	70.2
8	−13.3	63.1

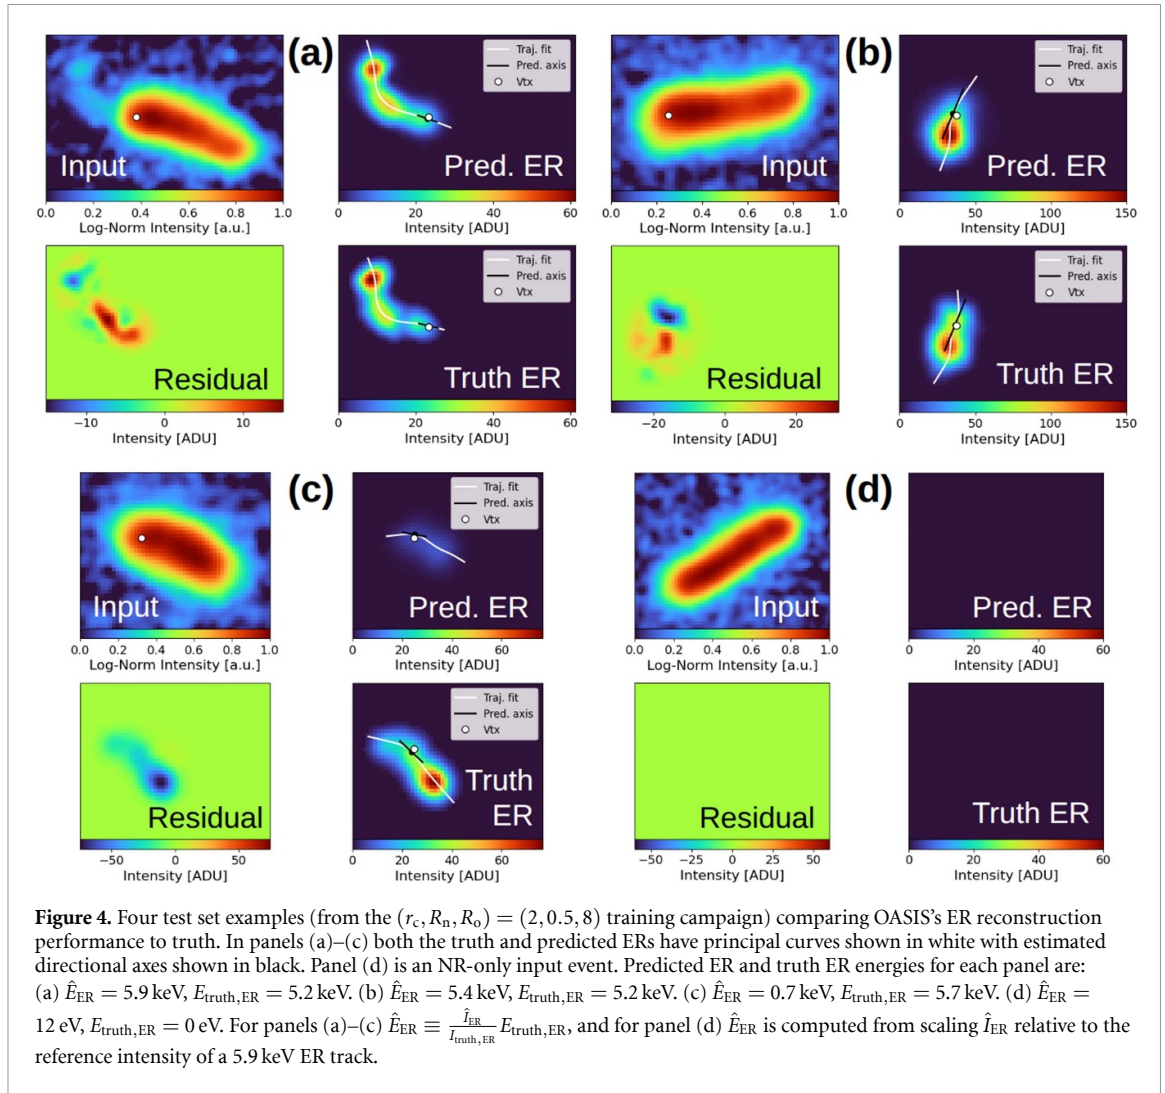
it is immediately apparent from figure 3(a), and quantified in table 1, that increasing R_o improves both bias and spread OASIS’s track intensity reconstruction error.

Narrowing in on the four ablations shown in figures 3(b) and (c) allows for a holistic comparison across all ER energies between the case of unweighted segmentation regression, the usage of weights prioritizing ER reconstruction, and the usage of weights prioritizing overlap-region reconstruction. Comparing the traces with blue circles and green triangles, we observe an improvement in intensity reconstruction and a significant improvement in topological IoU score in the 3–4 keV and 4–5 keV ER energy bins. This result shows that applying channel and region-specific weights that explicitly prioritize ER reconstruction over NR reconstruction in OASIS’s training improves performance over unweighted training in the ERs that overlap most with NRs. OASIS’s overlap region weights, however, yield more significant improvements in overlap-dominated ER reconstruction than weights prioritizing ER reconstruction, as is shown by the trace with orange squares. Finally, the trace with red diamonds includes weights that prioritize both ER reconstruction as well as overlap-targeted weights, yielding the best intensity reconstruction performance for low energy ERs where attribution is most challenging. For example, in the 4–5 keV bin, where the ER-NR overlap region tends to comprise a large proportion of the ER track, OASIS trained with $(r_c, R_n, R_o) = (2, 0.5, 8)$ achieves $\Delta I_{\text{median}} = -9.15\%$ and $\text{IoU}_{\text{median}} = 0.72$, compared to $\Delta I_{\text{median}} = -64.0\%$ and $\text{IoU}_{\text{median}} = 0.59$ for unweighted segmentation regression²⁰.

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 [52], making accurate reconstruction in the 3–6 keV range essential for achieving sufficient statistics to test theoretical predictions. We additionally note that around 80% of events in the hybrid signal sample with sub-6 keV ERs are evaluated by OASIS with bounding box mask regions that are identical to YOLO’s NR bounding box prediction. Therefore, the inputs corresponding to a significant majority of low ER-energy events do not utilize any information at inference that would not be available in a real-data pipeline. This gives confidence in the potential for OASIS’s overlap region reconstruction performance gains to generalize to real data.

The final test of this ablation study compares OASIS’s ER construction performance on both the hybrid signal and NR-only test sets to evaluate true and false positive detection rates, respectively. By design, OASIS should always output a nonzero amount of ER signal, even if the ER is not present, so modest track intensity thresholds must be applied to distinguish positive versus negative ER detections. Figure 3(d) shows ROC curves derived from varying the reconstructed intensity thresholds of the ER-channel from which a positive or negative ER detection is decided. These ROC curves are generated for the four ablations we have narrowed our focus to, with true positive rates (TPR) averaged over the 3–4 keV, 4–5 keV, and 5–6 keV truth ER energy bins of the hybrid signal samples, FPRs evaluated on the NR-only sample. Using the area under the ROC curve as our performance metric, we observe the same relative ranking in performance as figures 3(b) and (c). In other words, the overlap-region weights introduced in OASIS are the single most important loss function weight for improving TPR at fixed FPR, and the training campaign that weights both ER reconstruction and overlap region reconstruction performs best.

²⁰ For clarity, unweighted segmentation regression has $(r_c, R_n, R_o) = (1, 1, 1)$ which is equivalent to the segmentation-regression loss term optimizing the mean absolute error for each object.

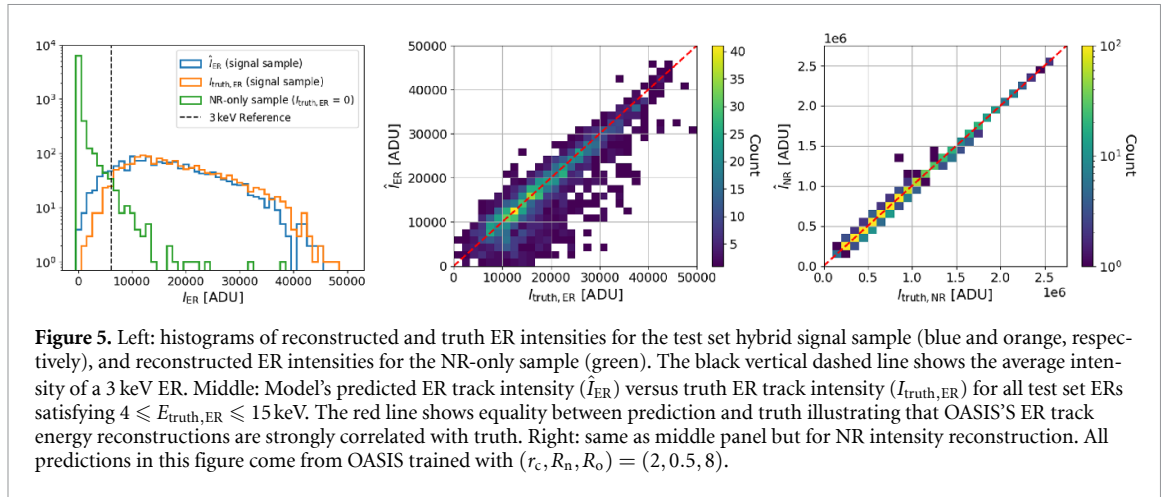


In summary, our systematic ablation study shows that ER-targeted- and especially overlap-region targeted- loss function weights improve ER intensity and topological reconstruction, and also improve true positive detection rates at fixed FPRs. Given that the $(r_c, R_n, R_o) = (2, 0.5, 8)$ training campaign is the best performing for all metrics evaluated in detail in this ablation study, we spend the remainder of this section presenting more detailed results from this campaign. Before further quantitative studies, we first examine this training campaign'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 intensity maps, $\Delta\phi \in [0^\circ, 90^\circ]$. We call the $\Delta\phi$ performance metric *angular consistency*, as it is a measure of similarity between reconstructed angles on both the ground truth ER image and the ER image reconstructed by OASIS. We note the distinction between this and angular resolution, which compares a reconstructed track angle with the true initial angle of the ER. OASIS's goal is to faithfully separate an ER from an overlapping NR, not to reconstruct the true ER direction, so angular consistency is a more practical performance measure for OASIS's ER topology reconstruction. Details of the angular reconstruction algorithm from which we compute angular consistencies are discussed in [appendix](#).

Panel (a) of figure 4 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 ($\Delta I = +13\%$) and capturing directional information near the ER's vertex ($\Delta\phi = 5.7^\circ$). 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 ($\Delta I = +4\%$, $\Delta\phi = 0.9^\circ$). Panel (c) illustrates a typical example where our model does not perform well. In this event, the truth ER travels inward toward the brightest regions of the NR. Despite weighting the ER-NR overlap region most heavily in our network training, OASIS predicts too little overall intensity. For this event, we find $\Delta I = -88\%$, 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 12 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, strong peaking near 0, and that the lowest truth ER energy in the hybrid signal sample is 3 keV, 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, non-negligible amounts of ER signal are nearly always detected.

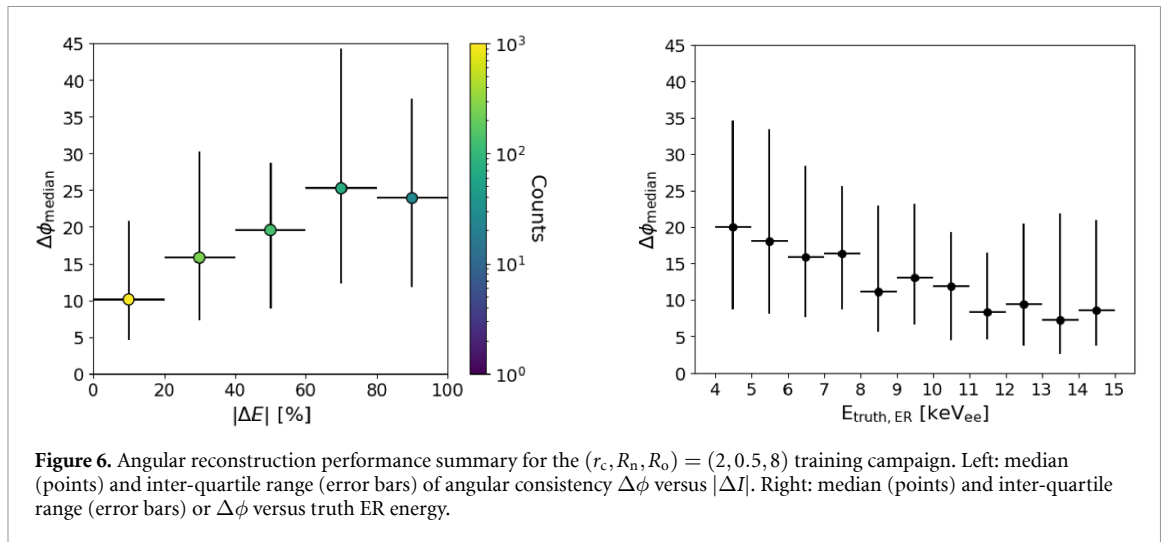
The middle panel of figure 5 shows the correlation between OASIS's ER track intensity predictions with truth for events satisfying MIGDAL's search region of interest of $4 \leq E_{truth,ER} \leq 15$ keV. We make the following observations from this distribution. First, the red line indicates perfect performance ($\hat{I}_{ER} = I_{truth,ER}$) and much of the data is tightly centered around this line. Second, while the mode of the \hat{I}_{ER} distribution at most slices of $I_{truth,ER}$ aligns with perfect performance, at lower intensities, the model is more likely to under-predict the ER's intensity, which was also shown in figure 3(b).

Finally, the right panel shows the correlation between OASIS's NR track intensity predictions with truth in the hybrid signal sample. Given that NR signal intensities overwhelmingly dominate over ER signals in all images, we expect near perfect NR intensity reconstruction, which is what we observe.

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) to the predicted and truth ER intensity maps for events in the signal sample of the test set satisfying $4 \leq E_{truth,ER} \leq 15$ keV. Of these 1750 events, principal curves were successfully fit on 1571 of them (90%). 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 topologies, $\Delta\phi \in [0, 90^\circ]$, 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 20° over all energies ≥ 4 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 its 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.

When tasked with reconstructing the MIGDAL experiment's faint ER signal—one that is often buried within an order(s)-of-magnitude brighter NR signal—we demonstrated that weighting OASIS's loss function to prioritize regions of pixel overlap between the ER and NR yields substantial improvement in intensity attribution, event topology reconstruction, and true positive detections of low energy ERs at fixed FPRs. In particular, in the 4–5 keV ER energy regime where ER-NR overlap tends to dominate the ER signal, averaging over all combinations of weights targeting ER/NR reconstruction, r_c and R_n , at fixed overlap weight ratios, R_o (equation (6)), we find ΔI_{median} to improve from -41.1% at $R_o = 1$ to -13.3% at $R_o = 8$ (table 1). Furthermore, compared to unweighted training, the addition of OASIS's novel overlap region-targeted weight in its loss function is the single most important adjustment in improving the reconstruction of low energy ER tracks across all metrics considered, as evidenced by comparing the blue and orange traces in figures 3(b)–(d). Linking OASIS's topological reconstruction performance to more physically meaningful quantities, with the choice of $(r_c, R_n, R_o) = (2, 0.5, 8)$, we observe angular consistencies derived from principal curve fits to within 20° , down to ER energies of 4 keV. Given the exponentially rising Migdal emission probability with decreasing ER energy, achieving sensitivity to reconstruct ER energies and angles in this 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 [63], 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. We are additionally exploring the possibility of extending OASIS’s overlap-aware loss weights to 1D waveform data, where the task is to sample-wise reconstruct a physical signal channel and a background channel. Applying OASIS to such a task could potentially enable weak signal extraction in regimes where traditional thresholding fails. 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.

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).

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://github.com/jschuel/OASIS> [67].

Appendix. 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 electron recoil (ER) image and the truth ER image. To this aim, any angular reconstruction algorithm can be used to estimate $\Delta\phi$. Because ER 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 [64] does just this and has a readily-adaptable implementation [65] in the R programming language [66]. 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 [65] 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 $\Delta\phi \in [0^\circ, 90^\circ]$.

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 [53, 54] 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.

ORCID iDs

J Schueler  0000-0002-2722-6953
H M Araújo  0000-0002-5972-2783
F M Brunbauer  0000-0003-0514-8189
A Cottle  0000-0003-4274-3622
D Edgeman  0009-0002-8836-3801
F Garcia  0000-0002-4023-7964
M Kastriotou  0000-0003-1010-2396
P Knights  0000-0003-2981-705X
A Lindote  0000-0002-7965-807X
P A Majewski  0000-0002-1371-8228
T Marley  0000-0003-3985-5367
C McCabe  0000-0002-4767-821X
L Millins  0009-0002-5784-2199
K Nikolopoulos  0000-0002-3048-489X
Timothy J Sumner  0000-0002-3572-600X
E Tilly  0000-0002-8694-8219
W Thompson  0009-0000-5661-5379

References

- [1] Melchior P, Moolekamp F, Jerdee M, Armstrong R, Sun A-L, Bosch J and Lupton R 2018 SCARLET: source separation in multi-band images by constrained matrix factorization *Astron. Comput.* **24** 129
- [2] Speiser A *et al* 2021 Deep learning enables fast and dense single-molecule localization with high accuracy *Nat. Methods* **18** 1082–90
- [3] Zhu Q, Du B, and Yan P 2019 Boundary-weighted domain adaptive neural network for prostate MR image segmentation (arXiv:1902.08128)
- [4] Jiang H, Zhang R, Zhou Y, Wang Y, and Chen H 2023 DoNet: deep de-overlapping network for cytology instance segmentation (arXiv:2303.14373)
- [5] Huiyu Li *et al* 2025 Deep distance map regression network with shape-aware loss for imbalanced medical image segmentation (arXiv:2501.09116)
- [6] Alonso-Monsalve S, Sgalaberna D, Zhao X, Molines A, McGrew C and Rubbia A 2024 Deep-learning-based decomposition of overlapping-sparse images: application at the vertex of simulated neutrino interactions *Commun. Phys.* **7** 173
- [7] Abud A A *et al* 2025 Neutrino interaction vertex reconstruction in DUNE with Pandora deep learning *Eur. Phys. J. C* **85** 697
- [8] Ronneberger O, Fischer P and Brox T 2015 U-Net: convolutional networks for biomedical image segmentation *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (Springer) pp 234–41
- [9] Kaiming H, Gkioxari G, Dollár P, and Girshick R 2017 Mask R-CNN (arXiv:1703.06870)

- [10] Torrado-Carvajal A et al 2018 Dixon-VIBE deep learning (DIVIDE) pseudo-CT synthesis for pelvis PET/MR attenuation correction *J. Nucl. Med.* **60** 429–35
- [11] Burke C J, Aleo P D, Chen Y-C, Liu X, Peterson J R, Sembroski G H and Lin J Y-Y 2019 Deblending and classifying astronomical sources with Mask R-CNN deep learning *Mon. Not. R. Astron. Soc.* **490** 3952–65
- [12] Zhang R et al 2024 Deblending overlapping galaxies in DECaLS using transformer-based algorithm: a method combining multiple bands and data types *Publ. Astron. Soc. Aust.* **41** e035
- [13] Walter F C, Damrich S, and Hamprecht F A 2020 MultiStar: instance segmentation of overlapping objects with star-convex polygons (arXiv:2011.13228)
- [14] Lei K, Tai Y-W, and Tang C-K 2021 Deep occlusion-aware instance segmentation with overlapping bilayers (arXiv:2103.12340)
- [15] Alvarez V et al 2012 NEXT-100 technical design report (TDR): executive summary *J. Instrum.* **7** T06001
- [16] Kekic M et al 2021 Demonstration of background rejection using deep convolutional neural networks in the NEXT experiment *J. High Energy Phys.* **JHEP01(2021)189**
- [17] Simón A et al 2021 Boosting background suppression in the NEXT experiment through Richardson-Lucy deconvolution *J. High Energy Phys.* **2021** 146
- [18] Miernik K et al 2007 Optical time projection chamber for imaging nuclear decays *Nucl. Instrum. Methods Phys. Res. A* **581** 194–7
- [19] Miernik K et al 2009 Two-proton radioactivity of ^{45}Fe *Eur. Phys. J. A* **42** 431–9
- [20] Pomorski M et al 2014 Proton spectroscopy of ^{48}Ni , ^{46}Fe and ^{44}Cr *Phys. Rev. C* **90** 014311
- [21] Acciarri R et al 2017 Design and construction of the MicroBooNE detector *J. Instrum.* **12** 02017
- [22] Abratenko P et al 2021 Semantic segmentation with a sparse convolutional neural network for event reconstruction in MicroBooNE *Phys. Rev. D* **103** 052012
- [23] Abratenko P et al 2021 Convolutional neural network for multiple particle identification in the MicroBooNE liquid argon time projection chamber *Phys. Rev. D* **103** 092003
- [24] Abratenko P et al 2022 Search for an anomalous excess of charged-current quasielastic νe interactions with the MicroBooNE experiment using Deep-Learning-based reconstruction *Phys. Rev. D* **105** 112003
- [25] Grizzi V F, Voetberg M, Cerati G, Meidani H, and Hewes V 2025 NuGraph2 with context-aware inputs: physics-inspired improvements in semantic segmentation
- [26] Aurisano A, Hewes V, Cerati G, Kowalkowski J, Lee C S, Liao W, Grzenda D, Gumpula K and Zhang X 2024 Graph neural network for neutrino physics event reconstruction *Phys. Rev. D* **110** 032008
- [27] Araújo H M et al 2023 The MIGDAL experiment: measuring a rare atomic process to aid the search for dark matter *Astropart. Phys.* **151** 102853
- [28] Difan Y et al 2026 Direct observation of the Migdal effect induced by neutron bombardment *Nature* **649** 580–3
- [29] Elfving S, Uchibe E, and Doya K 2017 Sigmoid-weighted linear units for neural network function approximation in reinforcement learning (arXiv:1702.03118)
- [30] Estrela V V, Magalhaes H A, and Saotome O 2016 Total variation applications in computer vision (arXiv:1603.09599)
- [31] Migdal A 1939 Ionizatsiya atomov pri yadernykh reaktsiyakh (Ionization of atoms in nuclear reactions) *Zh. Eksp. Teor. Fiz.* **9** 1163–5 (available at: <https://journals.aps.org/prd/pdf/10.1103/PhysRevD.107.035032>)
- [32] Ibe M, Nakano W, Shoji Y and Suzuki K 2018 Migdal effect in dark matter direct detection experiments *J. High Energy Phys.* **JHEP03(2018)194**
- [33] Akerib D S et al 2019 Results of a search for sub-GeV dark matter using 2013 LUX data *Phys. Rev. Lett.* **122** 131301
- [34] Armengaud E et al 2019 Searching for low-mass dark matter particles with a massive Ge bolometer operated above-ground *Phys. Rev. D* **99** 082003
- [35] Liu Z Z et al 2019 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** 161301
- [36] Aprile E et al 2019 Search for light dark matter interactions enhanced by the Migdal effect or Bremsstrahlung in XENON1T *Phys. Rev. Lett.* **123** 241803
- [37] Barak L et al 2020 SENSEI: direct-detection results on sub-GeV Dark matter from a new skipper-CCD *Phys. Rev. Lett.* **125** 171802
- [38] Adhikari G et al 2022 Searching for low-mass dark matter via the Migdal effect in COSINE-100 *Phys. Rev. D* **105** 042006
- [39] Armengaud E et al 2022 Search for sub-GeV dark matter via the Migdal effect with an EDELWEISS germanium detector with NbSi transition-edge sensors *Phys. Rev. D* **106** 062004
- [40] Agnes P et al 2023 Search for dark-matter–nucleon interactions via Migdal effect with darkside-50 *Phys. Rev. Lett.* **130** 101001
- [41] Abe K et al 2023 Direct dark matter searches with the full data set of XMASS-I *Phys. Rev. D* **108** 083022
- [42] Aalbers J et al 2023 Search for new physics in low-energy electron recoils from the first LZ exposure *Phys. Rev. D* **108** 072006
- [43] Huang Di et al 2023 Search for dark-matter–nucleon interactions with a dark mediator in PandaX-4T *Phys. Rev. Lett.* **131** 191002
- [44] Albakry M F et al 2023 Search for low-mass dark matter via bremsstrahlung radiation and the Migdal effect in SuperCDMS *Phys. Rev. D* **107** 112013
- [45] Adari P et al 2025 First direct-detection results on sub-GeV dark matter using the SENSEI detector at SNOLAB *Phys. Rev. Lett.* **134** 011804
- [46] Aggarwal K et al 2025 Probing benchmark models of hidden-sector dark matter with DAMIC-M *Phys. Rev. Lett.* **135** 071002
- [47] ORCA-Quest qCMOS camera C15550-20UP technical note 2022 *Technical Report* Hamamatsu Photonics
- [48] Takahashi H, Mitsuya Y, Fujiwara T and Fushie T 2013 Development of a glass GEM *Nucl. Instrum. Meth. A* **724** 1–4
- [49] EHD Imaging GmbH, Zum Rennplatz 15, D-49401 Damme, Germany 2020 EHD F0.85 Lenses (available at: www.ehd.de/wp-content/uploads/2024/05/EHD-F085-Lenses.pdf)
- [50] MIGDAL Collaboration 2025 Performance of the optical time projection chamber of the MIGDAL experiment *[In Progress]*
- [51] Schueler J et al 2025 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** 072004
- [52] Cox P, Dolan M J, McCabe C and Quiney H M 2023 Precise predictions and new insights for atomic ionization from the Migdal effect *Phys. Rev. D* **107** 035032
- [53] Steger C 1998 An unbiased detector of curvilinear structures *IEEE Trans. Pattern Anal. Mach. Intell.* **20** 113–25
- [54] Tilly E and Handley M 2023 3D track reconstruction of low-energy electrons in the MIGDAL low pressure optical time projection chamber *J. Instrum.* **18** C07013
- [55] Jocher G, Chaurasia A and Qiu J 2023 Ultralytics YOLOv8

- [56] Höfer S et al 2021 Sim2real in robotics and automation: applications and challenges *IEEE Trans. Autom. Sci. Eng.* **18** 398–400
- [57] Edgeman D et al 2025 Optical effects in gaseous electron multipliers (GEMs) (arXiv:2604.27399)
- [58] Schueler J 2024 migYOLO (available at: <https://github.com/jschuel/migYOLO>)
- [59] Loshchilov I and Hutter F 2017 Decoupled Weight Decay Regularization (arXiv:1711.05101)
- [60] Prechelt L 2012 *Early Stopping—But When?* (Springer) pp 53–67
- [61] Paszke A et al 2019 Pytorch: an imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems* p 32
- [62] Shen Li, et al 2020 Pytorch distributed: experiences on accelerating data parallel training (arXiv:2006.15704)
- [63] Walmsley M et al 2022 Galaxy Zoo DECaLS: detailed visual morphology measurements from volunteers and deep learning for 314 000 galaxies *Mon. Not. R. Astron. Soc.* **509** 3966–88
- [64] Hastie T and Stuetzle W 1989 Principal curves *J. Am. Stat. Assoc.* **84** 502–16
- [65] Robrecht Cannoodt. princurve 2.0: Fit a principal curve in arbitrary dimension 2018
- [66] R Core Team 2023 *R: A Language and Environment for Statistical Computing* R Foundation for Statistical Computing,
- [67] Schueler J (MIGDAL Collaboration) OASIS Sample Data4 *OASIS_sample_data5* OASIS (available at: https://lobogit.unm.edu/jschueler1/oasis_sample_data/)