AI-directed voxel extraction and volume EM identify intrusions as sites of mitochondrial contact

Benjamin S. Padman,[2,5,*] Runa Lindblom,[1,3] and Michael Lazarou[1,2,3*]

[1]Walter and Eliza Hall Institute of Medical Research, Parkville, Victoria, Australia
[2]Department of Biochemistry and Molecular Biology, Biomedicine Discovery Institute, Monash
[3]Department of Medical Biology, University of Melbourne, Melbourne, Victoria, Australia
University, Melbourne, Australia
[5]Present address: Telethon Kids Institute, Perth Children's Hospital, Nedlands, Western Australia,
Australia

Correspondence: Benjamin S. Padman (benjamin.padman@uwa.edu.au), Michael Lazarou
(lazarou.m@wehi.edu.au)

ABSTRACT

Membrane contact sites (MCS) establish organelle interactomes in cells to enable communication and exchange of materials. Volume electron microscopy (vEM) is ideally suited for MCS analyses, but semantic segmentation of large vEM datasets remains challenging. Recent adoption of artificial intelligence (AI) for segmentation has greatly enhanced our analysis capabilities. However, we show that organelle boundaries, which are important for defining MCS, are the least confident predictions made by AI. We outline a segmentation strategy termed AI-directed Voxel Extraction (AIVE), that combines AI predictions with image electron signals to confidently segment membrane boundaries irrespective of the AI model used. We demonstrate the precision conferred by AIVE by applying it to the quantitative analysis of organelle interactomes from multiple FIB-SEM datasets. Through AIVE, we discover a previously unknown category of mitochondrial contact that we term the mitochondrial intrusion. We hypothesise that intrusions serve as anchors that stabilize MCS and promote organelle communication.

Introduction

Organelles allow eukaryotic cells to compartmentalize cellular pathways that would normally be incompatible in a shared environment. However, compartmentalization also isolates key processes that depend on one another. To link isolated pathways, cells dynamically regulate interactions between organelles via physical points of contact called Membrane Contact Sites (MCSs) [1], that mediate inter-organelle communication, material exchange, tethering and structural support [2]. For example, cardiolipin synthesis at the inner mitochondrial membrane requires the precursor lipid phosphatidic acid, which is synthesised by the endoplasmic reticulum (ER) [3]; without adequate MCSs between mitochondria and the ER, cardiolipin synthesis would halt. MCSs are now widely accepted to be essential for cellular homeostasis [4], and MCSs have been reported between almost every conceivable combination of organelles [5-8]. Collectively, the interaction network between these various organelles is referred to as the "organelle interactome" [9].

Meaningful analysis of the organelle interactome poses a major technical challenge because MCSs represent small points of contact between comparatively large structures. Optical microscopy is suitable for imaging large interacting structures, but the separation distances of an MCS at 10 nm to 35 nm [10] are smaller than the resolvable limits of light. Biochemical methods also have limitations, since most sample preparation methods typically remove structural and spatial context from the MCS prior to analysis [11]. This is why volume electron microscopy (vEM) techniques are frequently referred to as the *"gold standard"* for studying MCS structure [10].

Volume EM (vEM) methods based on scanning electron microscopy (SEM), such as Focussed Ion Beam SEM (FIB-SEM) or Serial Block Face SEM (SBF-SEM), enable the ultrastructural visualisation of membrane structures within cellular scaled volumes. However, analysis of large vEM datasets poses a major challenge since the structures detected require semantic segmentation (hereafter segmentation) and object classification before they can be quantitatively analysed, and this process has historically relied on the manual labelling of each organelle membrane by hand to define their boundaries. This reliance on human input represents a major bottleneck in the throughput of vEM methods, and the traditional approaches to image segmentation can no longer keep pace with the rate of data acquisition [12], which increases every year [13]. Given these challenges, it can be argued that the rapid adoption of Machine Learning and Deep Learning based approaches for vEM was inevitable [14, 15]. AI-based image segmentation approaches are now widely accepted in the field of vEM, but they are not without their own potential flaws when it

comes organelle boundaries. Segmentation aims to define where an object is, but it also aims to

define where an object isn't. The transition between the presence or absence of an object defines

its bounding surface in a vEM dataset. The challenge arises when organelle surfaces segmented via

AI-based approaches represent locations where a trained model was equally certain about the

presence and absence of that object, and this can become problematic for MCS analyses.

Here, we demonstrate that the least certain model predictions made via AI-based image

segmentation occur predominantly at the bounding surface of an object, which is the most

important location for MCS analyses. This boundary uncertainty causes variable results between AI

models, and directly impacts the reliability and accuracy of ultrastructural measurements. We

outline and benchmark a strategy called AI-directed Voxel Extraction (AIVE) designed to negate

this limitation. Through a series of benchmarking experiments, we demonstrate AIVE's capacity to

reduce the variability between different AI models and confer greater consistency in MCS analyses

in cultured cells and tissues alike. We conduct a comprehensive analysis of FIB-SEM datasets

acquired from HeLa cells revealing the existence of a previously undescribed morphological

category of mitochondrial MCSs that we term mitochondrial intrusions. The benefits of

mitochondrial intrusions include increased surface area of membrane contacts between

organelles that may serve as an anchor to enhance crosstalk efficiency.

Results

AI based approaches for the segmentation of volume EM data

AI based image segmentation uses information learned from training data to create statistical or deductive frameworks (termed a model), so that predictions can be made about new unseen data. Given the large datasets that can be generated using FIB-SEM imaging of cells, there are clear advantages in using AI based approaches to automate membrane segmentation [16]. To demonstrate, we used a test dataset of approximately 3.5 µm³ in volume (Dimensions; X, 2.4 µm (740 px); Y, 1.2 µm (370 px); Z, 1.2 µm (121 slices)), that originates from a larger 655 µm³ dataset that will be discussed later in greater detail (Fig.4-8). The test dataset was used to train a random forest model to identify cellular membranes detected via FIB-SEM imaging (Fig.1A-B). This trained model generated probability (P) values indicating the likelihood that a membrane is present at each voxel in the dataset (Fig.1C). To complete the segmentation of these membranes (Fig.1D-E), the scalar P values were binarized by assigning a probability threshold (Fig.1D) or an isosurface value (Fig.1E) to generate a 3D result. The decision threshold is a critical parameter in AI-based image segmentation, because it defines the boundary of each segmented object, thereby defining an object's volume, surface area and contact with other objects. When selecting a threshold value, a mid-point P value of 50% appears to be a logical choice (Fig D-E), but 50% probability values represent where the model was equally certain about the presence and absence of an object, making it the least certain prediction. Deep learning methods like convolutional neural networks (CNNs) do not use a pre-chosen decision threshold, but the output layer of a CNN learns parameters akin to a threshold in order to generate a binary result [17, 18]. Regardless of the methods used to predict and binarize results, the decision thresholds control the ultimate conclusions drawn by a model, and by extension, the boundary of each object.

When presented with the same data, different AI algorithms make different predictions about the boundaries of an object, and these differences can affect interpretations of organelle ultrastructure and membrane contacts [19]. To demonstrate this, we trained six machine learning models based on different algorithms of varying complexity to detect membranes in the same test dataset from Fig.1B. All six models were trained in the detection of membranes using the exact same training samples and feature data before application to the same test dataset (Fig.1F), and therefore any variability between results would be due to their algorithmic differences. Each model was capable of membrane detection (Fig.1F), but key differences existed between their

results. By visualizing the differences in predicted values (ΔV) from each model (Fig.1G; ΔV), it can

be demonstrated that most of the predictions agree between models, but the disagreements that

do occur are almost exclusively localized to the membrane boundaries (Fig.1G). This uncertainty

about the position of a membrane boundary would negatively impact the reliability of any

organelle contact analysis, since it confounds the measurement of membrane separation

distances.

AIVE: a tool for accurate automated segmentation of cellular membranes

To improve the accuracy of membrane boundaries obtained from AI segmentation methods, and

to generate greater consensus across different AI algorithms, we developed an approach termed

AI-directed voxel extraction (AIVE)[16, 20]. AIVE combines the automated segmentation outputs

of an AI model with the ground truth of the volume EM data. Applying AIVE involves two major

steps: 1) Segmentation of cellular membranes with an AI (as in Fig 1), and then 2) multiplication of

the AI membrane prediction with the original EM data to extract only those voxels belonging to a

membrane. The benefit of AIVE for membrane boundary detection is that the regions of uncertain

AI prediction are accounted for by the intensity of signals in the EM data. For example, at a

boundary where the AI is uncertain, if the EM voxels have little to no intensity then the AI

prediction of the uncertain region will not contribute to the output (Fig S1A).

To demonstrate the overall AIVE approach, the sample EM data and AI predictions from Fig 1 were

applied to the AIVE workflow. First, the AI prediction was smoothened by applying a 3D Gaussian

blur to reduce variability between predictions (Fig.S1C). The EM data was processed using

Contrast Limited Adaptive Histogram Equalisation (CLAHE) (Fig.2C) to normalise differences in

image contrast that can occur between different data acquisitions, and median blurred to reduce

signal shot noise. In the final step, the voxel values from the AI predictions and the raw EM data

(Fig.2B-C) were multiplied to generate the AIVE output (Fig.2D). By multiplying voxel values from

two independent dynamic ranges (Fig.S1A), any input voxels with a value of zero were

automatically excluded from the output, whereas electron signals predicted by the model to

represent membranes were retained for segmentation (Fig 2D and Fig S1A). Importantly, this

strategy is not limited to membranes, since any cellular content that generates an electron signal

can be isolated (Fig.S1B). Like any other segmentation technique, a threshold value is required to

generate 3D surfaces from the AIVE dataset (Fig.S2A-C). it is important to note that values near

the centre of the output dynamic range (50%; 8-bit value of 128) are not equivalent to the 50%

probability threshold value discussed earlier (Fig.1D). This is because AIVE results are the product

of multiplying two independent dynamic ranges (Fig.S1A), so the centre of both dynamic ranges is

in fact 25% (50% × 50% = 25%; 8-bit value of 64). The premise of AIVE is therefore akin to a signal

masking strategy, in which electron signals detected during EM are used for the segmentation of

object boundaries instead of using AI predictions alone (Fig.1G). Doing so improves the accuracy,

precision, and consistency of membrane segmentation, as we will demonstrate through a series of

benchmarking experiments.

Benchmarking AIVE

AIVE enables outputs to be compared between different AI algorithms

We had observed large discrepancies between membrane boundary positions segmented by

different AI models in the same EM dataset (Fig.1F). To benchmark AIVE and demonstrate its

benefits, the same AI model predictions shown in Fig.1F were applied to the AIVE workflow

(Fig.2E). Upon application of AIVE, the difference in values (ΔV) between each model greatly

reduced (compare Fig.1F-G to Fig.2E-F). Quantitation of the ΔV per image slice also demonstrated

that the variability between results can be more than halved by AIVE, though we note that some

of this improvement was due to the initial blurring of the predictions (Fig.S1C). Next, we assessed

how the differences between boundaries from each model would affect 3D spatial measurements.

This was done by selecting 30 random points near membranes (distance probes; Fig.2H; see also

Video 1), then measuring the distance between each point and their nearest membrane (Fig.S2D).

By measuring the range (Fig.2I) and relative absolute deviation (Fig.2J) of distance values recorded

between models for each probe, the reliability of measurements under each condition can be

quantified. By both measures (Fig.2I-J), AIVE processing greatly improved the consistency of

measurements between models, outperforming the AI predictions alone with or without the

application of a Gaussian blur. Together, these analyses show that AIVE enables greater consensus

between AI models at regions of uncertainty (membrane boundaries), and that AIVE can be

applied with any AI model chosen by a user.

AIVE improves the accuracy and precision of organelle classification

The accuracy of organelle membrane segmentation is fundamentally important for the analysis of

organelle interactions. The AIVE framework can easily integrate organelle classification prior to the

voxel extraction step. This simply requires the preparation of binary labels that map the identity of

organelles. For example, if a mitochondrion of interest were present in the dataset, its identity can

be defined by drawing a label around that approximate region (Fig.3A). While this process may

seem identical to conventional classification, AIVE differs in that it does not use these classification

labels to define the boundaries of any organelles. Instead, the class labels are used to mask

regions of the AI predictions and merely identify them as belonging to a specific organelle (Fig.3B-

C). The resulting data is then processed via AIVE resulting in segmentation of membranes

belonging to a specific organelle, which in the sample case is a mitochondrion (Fig.3D-E).

Class labels can be assigned manually by a human or automated using AI. To demonstrate this, we

assessed the classification labels generated by a human analyst and two U-Net CNNs in the

detection of one mitochondrion (Fig.3F) from the test dataset in Fig.1A. Classification by a human

analyst was conducted manually on two consecutive days, while the two U-nets differed in their

randomization seed during training. Mitochondrial labelling was generally similar between the

classification approaches (Fig.3F), but the human classifications tended to extend further than

necessary beyond the perimeter of mitochondria, whereas the U-Nets occasionally appeared to

miss portions of the mitochondrion (Fig.3F; arrowheads).

To benchmark the benefits of AIVE, the binary class labels from Fig 3F were used for organelle

segmentation using four different approaches: (1) Traditional instance segmentation, in which

binary class labels are used to define object boundaries (Fig.3G); (2) AI Extraction (AIE), where the

class labels are used to mask the AI predictions without voxel extraction (Fig.3H); (3) direct Voxel

Extraction (VE), which uses the class labels to mask the normalized EM data without using AI

predictions (Fig.3I); and, (4) AIVE (Fig.3J).

To understand the reconstruction results, we must first address the characteristics and limitations

of conventional 3D segmentation. Specifically, 3D surfaces generated using binary segmentation

alone (Fig.3G) demonstrate a noteworthy artifact called "terracing" (Fig.3G; see 3D insets), which

presents as artificial plateaus [21]. The terracing is caused by abrupt (binary) changes in voxel

value, which affect the local shape since contours can only be represented by a limited pool of

surface topologies (2^8; 256) [22], of which only 15 are unique [23]. In contrast, the additional

information present within a 8-bit scalar datasets, such as those generated by AIE, VE and AIVE

(Fig. 3H-J) can theoretically adopt >$1.8×10^{19}$ (256^8) different polygon configurations at each voxel

and interpolate fractional angles with spatial anisotropy relative to the voxel grid [23]. Therefore,

by moving from binary labels to scalar values we can generate a more natural contour from our

voxel grid, as shown by the surface details shown in Fig. 3H-J compared to Fig. 3G.

Next, we benchmarked AIVE against voxel extraction (VE) and AI extraction (AIE). VE prevents

terracing (Fig. 3H), but the scalar data also includes irrelevant electron signals. Similarly, AIE

reduces terracing (AIE; Fig.4I), but the disadvantages of relying on AI predictions alone for

membrane boundaries have already been outlined (Fig.1-2). To directly show how the different

approaches influence the reliability of contact site analyses, we quantified the variability of

distance measurements between classification labels for each of the scalar 3D datasets (Fig.3H-J).

We conducted 3D measurements (Fig.3K) using a new set of 3D distance probes nearer to a

mitochondrion (Fig.S1D; see also Video 2), since the original 3D probe locations had no relation to

a specific organelle (Fig.2H). As demonstrated by the range (Fig.3L) and relative absolute deviation

(Fig.3M) of distance values recorded by each 3D probe, the distance measurements recorded for

AIVE datasets were very consistent between each of the differing class labels as opposed to VE

and AIE (Fig.3F). Therefore, AIVE substantially improves the precision and consistency of 3D spatial

measurements in volume EM datasets.

We next conducted a benchmarking experiment to compare the accuracy of measurements

generated via AIVE to a known standard point of reference. Nuclear Pore Complexes (NPCs) are

large protein channels in the nuclear envelope which mediate cargo transport between the

interior and exterior of the nucleus. NPC structure is highly conserved among eukaryotes, which is

why they are often used as a reference standard for benchmarking microscopy methods [24]. The

test dataset used in figures 1-3 originated from a larger cellular volume that also contained a

portion of nuclear envelope (Fig.4B). Using the manual classification strategy described in Fig. 3A-J,

we processed a portion of the nuclear envelope containing 18 individual nuclear pores (Fig.S1E),

via VE (Fig.3N), AIE (Fig.3O), or AIVE (Fig.3P). The same 18 pores were isolated from each

processed dataset and aligned using the exact same transforms (see methods for further detail;

Fig.S1F) to compare each processing method (Fig.3Q-T) against recent NPC structures [25, 26].

Data processed by VE alone (Fig.3N) rarely yielded a detectable pore because they were often

occluded by irrelevant electron signals (Fig.3Q). In contrast, central transport channels were

consistently detected in the data processed AIE alone (Fig.3O) or AIVE (Fig.3P), but the average

diameter of those pores differed between processing methods (Fig.3Q); 42.9 nm for AIE alone

compared to 52.4 nm for AIVE. We related these pore measurements to reported values from two

independent studies from 2021 (arrowheads in Fig.3Q), which used cryo-electron tomography to

resolve NPC structures *in cellulo* [25, 26]. Schuller *et al.* reported that human NPCs have a central

channel diameter of 57 nm (Fig.3Q; black arrowhead) [25]. The AIE processed data differs from

this value by 14.1 nm, whereas the AIVE results differed by only 4.6 nm. The study by Zimmerli *et*

al. was not conducted in human cells [26], but we include their findings because they provide

valuable context for the distribution of our measurements (Fig.3Q; value range indicated by white

arrowheads). Unlike data processed using VE or AIE alone, the majority of the AIVE processed

nuclear pores had a diameter within the published reported range. Renderings of the median

averaged pores from each processing method are also provided (Fig.3R-T; shown in purple), for

visual comparison with the human NPC structure published by Schuller *et al.* (Fig.3R-T; shown in

orange [25]). Collectively, our benchmark experiments show that AIVE processing substantially

improves both the precision and accuracy of spatial measurements in volume EM datasets.

Mapping organelle interactomes using AIVE

Compositional analysis of cellular membranes via AIVE

Having outlined the technical basis and advantages of AIVE, we next demonstrated some of the

biological insights AIVE can yield when applied to large volume datasets. We aimed to characterise

mitochondrial interactions with other cellular organelles in three FIB-SEM datasets from HeLa cells

that were processed using AIVE (Set #1, Fig.4B; Set #2, Fig.4C; Set #3, Fig.4D; see Video 3). As

mentioned previously, the test dataset used in our benchmarking experiments (Fig.1-3) originated

from a small region (Fig.4A) within one of these datasets. AIVE data can be subjected to a wide

range of analyses (Fig.4-8), the simplest being direct quantification of membrane volume (Fig.4F-

H). Set #1 (654.0 μm^3 total volume of dataset) contained the highest proportion of organellar

membranes (84.7 μm^3; 12.9% of total volume). Due to the prominence of the cellular nucleus in

sets #2 & #3, organellar membranes represented a substantially lower proportion of their total

volume (Set #2, 858.8 μm^3 total, 59.4 μm^3 organellar membrane; Set #3, 640.9 μm^3 total, 39.1 μm^3

organellar membrane). The nucleus in Set #2 is noteworthy (Fig.4C), as it also contained multiple

examples of nucleoplasmic reticulum tubes spanning the interior of the nucleus [27]. Excluding the

nuclear envelope from subsequent analyses reveals striking similarities in the distribution of

membrane between organelles across the datasets (Fig.4E-G), with the ER being the most

abundant category of organellar membrane, followed closely by mitochondrial membranes.

Intracellular vesicles and Golgi apparatus were the next most abundant categories of organellar

membrane, but Golgi apparatus was not detected Set #2 (Fig.4F); coincidentally, Set #2 also

contained a lower proportion of early and late endosomal membranes than in the other two

datasets. To explore these datasets in greater detail, we next turned our attention to individual

mitochondria.

Quantitative analysis of mitochondrial morphology via AIVE.

The three datasets (Fig.4B-D) collectively contained 186 individual mitochondria. Morphometric

analysis of these mitochondria based on their membranes alone would have limited value, since it

wouldn't include the mitochondrial interiors or bulk morphology, which is why our machine

learning models were also trained to classify materials other than membrane. For example, the

prediction results for "Matter" (Class 3), are specifically designed to represent any electron dense

non-membranous materials such as the mitochondrial matrix (Fig.S1B). This additional class

allowed us to "fill" the mitochondria by multiplying the matter class with the mitochondrial class

labels, then adding the result to the AIVE membrane outputs for the corresponding mitochondria.

In addition, the filling process enables imputation of the matrix of each mitochondrion.

Morphometric analyses, including measurements of mitochondrial volume, membrane surface

area, volume of matrix, shape, and length were applied to all 186 mitochondria. Based on the

morphometric analyses, six outlier mitochondrial morphologies were identified (Fig.5A-F). This

included the smallest mitochondrion, whose total volume was only 0.035 μm^3 (Fig.5A). As a point

of comparison, the smallest mitochondrion was approximately one-hundredth the volume of the

largest mitochondrion (Compare Fig.5A with Fig.5B), demonstrating a wide range of mitochondrial

size. Unsurprisingly, the largest mitochondrion also possessed the greatest membrane surface

area (25.079 μm^2), largest volume of membrane (1.863 μm^3), and largest volume of matrix (1.605

μm^3) of any detected mitochondrion (Video 4). Our analyses of mitochondrial morphologies also

included two different metrics of mitochondrial "length" (see Video 7), in which the most

elongated and longest mitochondria were identified. The most elongated mitochondrion (Fig.5D;

11.89 aspect ratio, 9.36 μm longest-optimal path) was identified by measuring mitochondrial

ellipsoid aspect ratios (major axis/minor axis; Fig.5J), whereas the longest mitochondrion (Fig.5E;

2.91 aspect ratio; 11.229 μm long) was determined by skeletonizing the bulk mitochondrial

structure, then measuring the longest-optimal path through the skeleton (Fig.5K). Mitochondrial

sphericity was also quantified (Fig.5L; Video 5, left), leading to the identification of the most

spherical mitochondrion (Fig.5F; sphericity = 0.860). The wealth of morphological data that was

obtained also allowed us to identify the most average mitochondrion (Fig.5G, and Video 5, right),

that had morphological measurements closest to the average value across every metric quantified.

Interestingly, we also noted that mitochondrial membrane volume correlated closely (R > 0.99)

with matrix volume in all datasets (Fig.5I), demonstrating that internal mitochondrial architecture

is tightly controlled despite large variations in morphology.

The morphometric analyses also led to the identification of mitochondrial nanotunnels based on

the separation of mitochondrial matrix volumes within a mitochondrion (Fig.5C; Video 5, center).

The presence of mitochondrial nanotunnels in HeLa cells was unexpected since nanotunnels are

typically observed in cellular/tissue environments in which mitochondrial movement is physically

constrained [28]. To show an example of nanotunnels in a physically constrained tissue, and

demonstrate AIVE's utility in tissue samples, we conducted FIB-SEM and AIVE of murine skeletal

muscle tissue (Fig.S3). Numerous examples of mitochondrial nanotunnels were observed (Fig.S3A-

C; Video 6), in which mitochondria remained interconnected over distance by reaching around

myofibrils via nanotunnels (Fig.S3, Video 6). In addition, we observed extensive sarcoplasmic

reticulum networks that contacted mitochondrial nanotunnels at sites around myofibrils.

Collectively, these analyses show the extensive mitochondrial morphology analyses that can be

conducted using AIVE, across cultured cells and tissue.

Quantitative analysis of the mitochondrial organelle interactome via AIVE.

Next, we aimed to investigate the interactions between mitochondria and other organelles,

including MCS, by quantifying the separation distances of each mitochondrion from the ER, early

endosomes (EE), late endosomes (LE), lipid droplets (LDs) and general cytosolic vesicles (VE) (Fig.6).

By individually indexing each mitochondrion, we were also able to quantify the separation

between each mitochondrion and the remainder of the mitochondrial network.

First, we determined the closest approach made between organelles by measuring the minimum

separation distance for each mitochondrion relative to each category of organelle (Fig.6A-F).

Closest approach values for the EE, LE, and LD organelle categories varied greatly between

mitochondria (Fig.6C-E), which were often several microns away from these categories of

organelle (Fig.6C-E). Next, the number of mitochondria in contact with each category of organelle

was determined by applying the standard criteria often described for membrane contact sites of <

35 nm separation [10]. On average, we found that 12.13% of mitochondria were in contact with an

early endosome, 24.96% were in contact with a late endosome, and 10.24% were in contact with a

LD (Fig.6G). In contrast, every mitochondrion was found to be in contact with the ER (Fig.6G), and

the mitochondria-ER closest approach distances never exceeded 5nm (Fig.6B). These results are in

close agreement with the findings of Friedman et al., [29] who reported that all cellular

mitochondria maintain continuous contact with the ER. Additional categories of mitochondrial

membrane contact were also observed, and they included interactions with other mitochondria

and cytosolic vesicles (Fig.6G). These results demonstrate the utility of AIVE in measuring organelle

interactomes, while also revealing the extent of organelle contacts, and how they vary across

different organelles.

Mitochondrial intrusions are morphological platforms for organelle interaction.

Using distance-based criteria to define MCS is a valuable tool. However, we also wanted to

account for the magnitude of the contact area relative to the size of an organelle because it gives

context to the breadth of a membrane contact site. For example, small point-like interactions with

randomly distributed vesicles (Fig.6G), cannot be distinguished from a membrane contact that

covers a substantially greater surface area. We therefore sought to refine our quantitative

analyses of MCS by accounting for the size of interaction sites between mitochondria with the ER,

vesicles (VE), and other mitochondria (Fig.6H-M). By grouping distance measurements into 5nm

brackets ranging between 0 to 250 nm, the percentage of a mitochondrion's surface from each

separation distance was determined (Fig.6K-M). The analysis revealed a prominent peak in the 0

to 5 nm distance bracket for mitochondrial (Fig.6K) and ER distance measurements (Fig.6L;

Arrowheads), but not for cytosolic vesicles (Fig.6M), that were typically further away from

mitochondria and occupied little of the mitochondrial surface.

The value peak detected in the 0-5nm range caught our attention because it is a large surface of

contact that cannot occur between round convex spheroids (Fig.S4A & C), unless one of those

objects partially envelopes the other (Fig.S4B & D). Indeed, our analysis of the 0-5nm membrane

contacts revealed striking mitochondrial morphologies that enabled membrane contact via

organelle intrusion (Fig.7). The intrusions were characterised by an invagination of the outer

mitochondrial membrane to form a narrow cavity that protruded into the interior of a

mitochondrion (Fig.7; see Videos 8-9). Mitochondrial intrusions frequently contained ER

membranes (Fig.7A-C; Video 8), that typically occupied the entire intrusion cavity. Interestingly,

the ER membranes often displayed higher levels of osmium staining (EM data; Fig.7A-C), indicating

that the membranes are protein rich or have a unique lipid profile. Mitochondrion-to-

mitochondrion intrusions were also observed, in which the exterior membranes of one

mitochondrion intruded into an adjacent mitochondrion (Fig.7D-E; Video 9). Unlike the ER-to-

mitochondrion intrusions, membranes in this category of intrusion did not display increased

osmium staining (Fig.7D-E). We also identified a mitochondrion that was intruding into another

mitochondrion while simultaneously receiving an intrusion from the ER (Fig.7D; Fig.8A, See box

indicating Set #1, mitochondria 61 & 89). This demonstrates that mitochondria can simultaneously

give and receive intrusions, potentially enabling inter-organelle communication. Mitochondrial

intrusions lacking an intruding membrane within the cavity were also observed (Fig.7F), indicating

that intrusion cavities may form independently from an intruding membrane.

We asked why mitochondrial intrusions have not previously been reported using established EM

approaches. Additional morphological analyses revealed key features that can provide an

explanation why intrusions have not been observed until now. Of the 186 mitochondria surveyed,

we detected a total of 50 mitochondrial intrusions among 36 individual mitochondria (Fig.8A),

equating to 0.269 intrusions per mitochondrion on average. No mitochondrial morphological traits

appeared to correlate with the presence of an intrusion (Fig.8B), implying that the formation of a

mitochondrial intrusion is stochastic. Indeed, the average number of intrusions per mitochondrion

(0.269) was used to model a Poisson distribution that closely approximates the observed values

(Fig.8C). Compared to the mitochondrion, the internal volume of the average intrusion cavity was

miniscule (Fig.8D), and the largest sphere capable of fitting within these cavities would have a

radius no larger than 50.6 nm (Fig.8E), yet they often extended several hundred nanometers

inward (Fig.8F).

A mitochondrial intrusion can only be confidently identified when viewed in its entirety. This is

because without evidence of the entrance, an intrusion would be indistinguishable from a

membrane inclusion (see Fig.7D), and without evidence of the deeper intrusion, the cavity

entrance would appear to be a slight membrane indentation (see Fig.7A). We sought to

demonstrate this by calculating the probability of detecting a fully intact mitochondrial intrusion

via TEM imaging by using Buffon's needle problem [30]. Buffon's needle problem is a geometrical

probability and is formulated as follows: If a needle of known length (L) is randomly dropped onto

an array of regularly spaced lines separated by a known distance (d), what is the probability that

the needle will overlap a line after landing (Fig.S4E). In the present context, the "needles"

represent mitochondrial intrusions with an average length (L) of 367.5 nm (Fig.8F), the space

between lines represent TEM sections of known thickness (d), and the needles that overlap a line

represent incomplete intrusions that extended into adjacent sections. The probability of detecting

a fully intact intrusion can therefore be considered the inverse of Buffon's needle problem (Fig.8G;

Fig.S4E). For a hypothetical 100nm TEM section (d = 100 nm) containing a mitochondrial intrusion,

there is less than a 1 in 10 chance (P = 0.087) that the intrusion will remain fully intact within that

section (solid line in Fig.8G). Under these conditions, an observer would be twice as likely to

observe only half of the intrusion (P = 0.177964; see dashed line, Fig.8G), which would be

indistinguishable from a membrane inclusion or indentation of the outer membrane. Even if the

section thickness was equal to the length of an intrusion (L = d), the probability of the structure

being intact remains low (P = 0.363; solid line in Fig.8G). It is important to note that the real

probability of observing an intrusion is likely even lower, since they were only present in a fraction

of the mitochondria (Fig.8C). In contrast, there would be a 95% chance (P = 0.953) of an intrusion

remaining intact if the Z-depth of the sampled volume spanned 5 μm (d = 5000 nm), which would

not be possible using TEM, but easily achieved via FIB-SEM. These calculations demonstrate why

mitochondrial intrusions have not previously been identified using historical TEM.

Taken together, by using AIVE we were enabled to discover that HeLa cell mitochondria can

contact the ER (and other mitochondria) via intrusion sites (Fig.7 & 8), which provide a larger

interface for surface interactions not possible between round objects (Fig.7; Fig.S4). We have also

shown that these intrusion sites are more likely to be indistinguishable from a mitochondrial

membrane inclusion when viewed by TEM (Fig.8G).

Discussion

In this study, we have outlined and benchmarked a technique termed AIVE for volumetric EM.

AIVE enables electron signals detected via FIB-SEM to be processed for rapid 3D reconstruction

and analysis with high fidelity. AIVE's fidelity is achieved by using the ground truth of EM data in

the final output. In AI research, ground truth is used to describe the preferred result from a

trained model [31], but the original definitions of ground truth specify that it originates from

objective empirical measurements [32]. In the context of this manuscript, the back-scattered

electron signals acquired via FIB-SEM imaging represented the most objective and empirical

measurements that could be used for ground truth. AIVE therefore uses an objective definition of

ground truth by linking AI predictions to real empirical measurements. Since its inception, AIVE has

been applied to volumetric EM analyses of cultured cells [16, 20] and animal tissues alike (Fig.S6).

However, in the absence of benchmarking, it was unclear why a cell biologist might choose to

apply AIVE versus AI predictions alone. Here, we demonstrate that membrane boundaries are the

least certain areas of AI predictions. Through extensive benchmarking, we make a case for the

benefits of AIVE to define membrane boundaries with high fidelity that provides benefits for the

study of organelle ultrastructure and MCSs.

Mitochondrial membrane ultrastructure is tightly linked to the biochemical functions of a

mitochondrion [33]. For example, mitochondrial damage and dysfunction are linked to a range of

abnormal mitochondrial morphologies including cristae disturbances and network fragmentation

[34-36]. In mitochondrial diseases, abnormal morphologies can manifest as osmiophilic inclusions

of concentric membranes and "onion-shaped" cristae within the mitochondria [34, 35], while in

other cases cristae structures are diminished [37]. Mitochondrial ultrastructure can also change in

response to the metabolic activity of mitochondria [38], in which mitochondria undertaking

oxidative reactions retain defined cristae whereas those enriched in reductive reactions lack

cristae. The intimate association between mitochondrial structure and function lends itself to

exploration using volumetric EM. AIVE can therefore be used in various contexts, ranging from

understanding mitochondrial ultrastructure in different tissues and disease states, to

understanding the relationship between mitochondrial morphology and physiological and

biochemical status. Cristae are intricate structures of the mitochondrial inner membrane that

come in all manner of unusual shapes and sizes depending on cell type and metabolic status [37].

For example, cristae in astrocytes appear to take the shape of triangles, whereas cristae in adrenal

cortex cells are highly circular [33]. Cristae analyses such as these have typically been conducted

using 2D TEM. By applying volume EM and AIVE, we anticipate that additional fascinating features

of cristae structure and biology will be revealed.

Mitochondrial structures can dynamically adapt according to their cytosolic environment. For

example, mitochondria in muscle tissue form nanotunnels that enable connectivity of the

mitochondrial network around the physical obstruction of myofibrils [28, 39](Fig. S6). Nanotunnel connectivity can increase in murine cardiac myocytes with defective calcium release from the sarcoplasmic reticulum [40], indicating an adaptive role for nanotunnels. Through AIVE, we were able to expand on previous studies of nanotunnels and show the extensive ultrastructure of mitochondrial nanotunnel networks in murine skeletal muscle (Fig. S6B). We also revealed extensive MCSs between nanotunnels and the sarcoplasmic reticulum (Figure S6C), that may play a role in calcium buffering during contraction [40]. Our findings demonstrate the benefits of collecting large EM volumes combined with AIVE in identifying long distance organelle connections, including mitochondria via nanotunnels, and the nucleus via nucleoplasmic tubes (Fig 4C).

Another area of cellular biology illuminated by AIVE and volume EM was the visualization and quantitation of MCSs. Through AIVE, we were enabled to undertake high resolution analyses of organelle interactomes via MCSs (Figs. 5-7). The value of AIVE for MCS analyses was best demonstrated by the discovery of a hitherto unknown form of mitochondrial contact that we term the mitochondrial intrusion (Fig. 7). Through intrusions, mitochondria were observed to contact other mitochondria or the ER. Intrusion relays in which a mitochondrion simultaneously contacted the ER and other mitochondria were also observed. These interactions go beyond the binary MCSs that are typically described and highlights the complexity of organelle contacts that can be revealed by volume EM and AIVE. Indeed, in a previous application of AIVE with volume EM, wholesale organelle-ome changes were observed in response to perturbation of a single organelle family [20]. We note that mitochondrial intrusions form basally, and they are distinct from previously reported mitochondrial cavities [41]. Mitochondrial cavities, induced by fragmenting mitochondria through depolarization, differ greatly from intrusions in size and morphology, and in contact characteristics including the lack of intramitochondrial contacts [41].

What might be the benefits of mitochondrial intrusions as a form of MCS? One answer is increased surface area for interaction. The exterior of a mitochondrion is typically round and convex [42, 43], but this is the least efficient geometry for mediating surface-to-surface interactions [44] (Fig.S3A). Mitochondrial intrusions can maximize the surface area of an MCS (Fig.S3B), and therefore promote the efficiency of inter-organelle communication [45]. Another potentially important benefit of intrusions is that the increased contact area can serve as an anchor to stabilise an otherwise fleeting MCS. Once formed, the very close proximity of membranes within intrusions (0-5nm) can feasibly support tethers that support lipid transfer including the mitochondrial proteins

PTPIP51 and MIGA2 [46-48], that could be facilitated by Mfn1/2 for intramitochondrial intrusions,

and Mfn2 and VAPB for ER-mitochondria intrusions [49-51]. In contrast, MCSs involving IP3Rs

(inositol 1,4,5-triphosphate receptors) that facilitate calcium transfer between ER and

mitochondria might be too bulky for the very close contacts made by intrusions [1, 52].

Overall, we demonstrate the utility of AIVE combined with volumetric EM for understanding the

ultrastructure and interactome of cellular organelles, while also demonstrating its capacity to

reveal fascinating and unexpected features of cell biology, including the identification of

mitochondrial intrusions as a form of MCS.

Acknowledgements:

495 We acknowledge Gediminas Gervinskas and the Ramaciotti Centre for Cryo-Electron Microscopy,

496 Monash University for technical support with electron microscopy. And we thank Matthew Eramo

497 and Mike Ryan for the mouse tissue FIB-SEM dataset, and acknowledge Monash Animal Research

498 Platform, Monash University.

Methods

Cell culture

HeLa cells were cultured Dulbecco's modified eagle medium (DMEM) supplemented with 10 %

(v/v) FBS (Cell Sera Australia), 1 % Penicillin-Streptomycin, 25 mM HEPES, 1x GlutaMAX (Life

Technologies) and 1x non-essential amino acids (Life Technologies). Samples were prepared by

culturing HeLa cell monolayers on a polymer film substrate (Aclar film; 203.2 µm thick; ProSciTech),

which was heat-welded to the plastic surface of a 10 cm tissue culture dish with a soldering iron.

Three dishes were prepared, and UV sterilized before seeding approximately 6×10^{6} HeLa cells

onto the Aclar film in each dish. The cells were allowed to attach under normal culture conditions

for 48 hrs, with replacement of the culture medium 1 hr prior to fixation.

Sample preparation of HeLa cells for FIB-SEM

The samples were chemically fixed with 4 % paraformaldehyde in 0.1 M phosphate buffer (pH 7.2) at 37 °C for 1 hour, with overnight postfixation with 2.5 % glutaraldehyde in 0.1 M sodium cacodylate buffer at 4 °C. Each polymer film was detached from their tissue culture dish before being transferred to a polypropylene tray (lid of a pipette tip box) where they were immobilized by heat-welding with a soldering iron. A BioWave Pro microwave system (Pelco) was used for microwave assisted sample processing in all subsequent stages. Osmication was conducted using a modified OTO method [53] with three microwave duty-cycles (120 s on, 120 s off) at 100 W under vacuum where indicated in the following stages; 2 % (w/v) OsO_4, 1.5 % (w/v) $K_3Fe(CN)_6$ in 0.1 M cacodylate buffer (pH 7.4) at 4° for 2 hours, followed by microwave processing; three MilliQ water rinses; 1 % (w/v) Thiocarbohydrazide in water with microwave processing; three MilliQ water rinses; then, 2 % OsO_4 in water with microwave processing. After further MilliQ water rinses, the samples were *en bloc* stained with 2 % (w/v) aqueous uranyl acetate [54], followed by Walton's Lead Aspartate [55], with each stain requiring three microwave duty-cycles (120 s on, 120 s off) at 100 W under vacuum. Microwave assisted dehydration (150 W for 40 s per stage) was performed by graduated ethanol series (80 %, 90 %, 95 %, 100 %, 100 % (w/v)) followed by propylene oxide (100 %, 100 % (w/v)). The samples were infiltrated with a modified Araldite 502/Embed 812 resin of the following composition (described for 10 ml resin); 2.20 g Araldite 502, 3.10 g Embed 812, 3.05 g NMA (nadic methyl anhydride), 3.05g DDSA (dodecenylsuccinic anhydride), and 400 μL BDMA (benzyldimethylamine). Infiltration was conducted using a graduated concentration series in propylene oxide (25 %, 50 %, 75 %, 100 %, 100 % (v/v); 180 s at 250 W under vacuum). Final embedding was conducted with minimal resin, which was polymerized at 60 °C for 48 hrs.

The Aclar film was peeled from each of the resin embedded cell monolayers, before using a hammer to sub-divide the resin into a random assortment of shards that were pooled into a single collection. Sample shards were blindly taken from the sample pool, deposited on a small droplet of the Araldite 502/Embed 812 resin on the tip of a 3.2 mm diameter aluminium rod; another droplet of resin was then placed on top of that shard so that another randomly selected shard could be placed on top, and the process was repeated to embed sandwiches of randomly oriented cells at least three monolayers thick, before final polymerization at 60 °C for 24 hr. The stacked cell monolayers were hand trimmed with a razor blade before trimming with a glass knife, to expose multiple monolayers for 3D FIB-SEM imaging.

Data acquisition and pre-processing

FIB-SEM imaging was conducted three separate times on randomly selected cells, one cell was chosen from each cell monolayer. Imaging was conducted using a cryo-Helios G4 UX FIB-SEM (FEI) at room temperature using the TLD detector in immersion lensing mode for back-scatter electron imaging at 3.255 nm per pixel (2 kV, 100 pA, 3 μs dwell time). Ion milling was conducted at 10 nm per slice (Gallium, 30 kV, 9 nA) using the Auto Slice and View (v4.1; FEI) software. All datasets were spatially registered (rigid transform) using the virtual stack registration plugin in FIJI (v1.53t) [56, 57]. The signal normalized inputs for AIVE were generated using a pseudo-3D implementation of contrast limited adaptive histogram equalisation (CLAHE), in which three different 2D CLAHE calculations were made from the front (XY), top (XZ), and side (YZ) of the registered dataset. Each 2D calculation was made using a modified version of the CLAHE plugin released by Stephan Saalfeld in 2009; this modified plugin accounts for the anisotropy of voxels in the XZ and YZ directions by allowing the use of rectangular sliding window during CLAHE. The modified plugin ("CLAHE_Anisotropic.class") and script for batch automation ("Batch_3DCLAHE_AnisotropicXYZ.ijm") are both provided on GitHub (https://github.com/BenPadman/AIVE.git).

Sample preparation & AIVE of mouse muscle tissue

Wildtype C57BL/6J mouse was anaesthetised at 3 weeks of age prior to dissection of the tibialis anterior muscle, which was immersion fixed in 4% PFA in 0.1M phosphate buffer at room temperature (1 hour). Tissue was further dissected under fixative to extract 1mm cubed samples, which were immersed in secondary fixative (4% PFA, 2.5% GA in 0.1M phosphate buffer) for 24hrs. All subsequent sample processing stages and imaging conditions were conducted as described above for cultured cell samples. Two separate FIB-SEM datasets of skeletal muscle were acquired. A machine learning model was trained on 3 classes of voxel designated "Cytosol", "Fibrils" and "Membrane", using 120,000 samples per class (360,000 total) to train the final model with a 10-fold test cross-validation result of over 97% self-accuracy. Mitochondria were manually classified in MIB as described above, and all remaining membranes were designated as sarcoplasmic reticulum. The final AIVE results were calculated as described above, and reconstructed in 3D by marching cubes with an isovalue of 64. All animal work was approved by the Monash Animal Research Platform Ethics Committee (MARP #17628), Monash University, Australia, and conducted in compliance with the specified ethics regulations.

Computational hardware

All presented analyses were conducted on one computer equipped with consumer-grade hardware, which included an Intel® Core™ i9-9900KF CPU clocked at 3.60GHz (16 cores), 64Gb of RAM at 2133MHz, one 2 Tb SSD (Samsung 970 EVO Plus) for short term storage during machine learning feature calculation, two 8Tb HDD's in RAID1 for long-term storage of machine learning features and results, and an Nvidia GeForce RTX 3080 Ti for deep learning and 3D rendering.

Machine learning

Machine learning was conducted within the Waikato Environment for Knowledge Analysis (WEKA v3.9)[58]. Training features were extracted from the registered image data using the 3D Trainable Weka Segmentation (TWS) plugin for ImageJ [59, 60]. As proposed in the original manuscript [59], we wrote a series of custom scripts for the TWS plugin to expand and automate its capabilities; these scripts have been provided on GitHub (https://github.com/BenPadman/AIVE.git). The scripts were designed to minimize RAM storage requirements for general machine learning, by ensuring that all image features required for training and evaluation are only calculated once prior to local storage on a hard disk. Instead of calculating all features for the entire stack simultaneously, the scripts only inspect a portion of the stack at a time while including the minimum number of flanking slices required by the radius of a 3D image filter. This enables features to be calculated for image stacks of arbitrary length at the cost of wasted computation on flanking slices. Furthermore, the script generates features incrementally by grouping image filter calculations into subsets, temporarily storing the results on a local hard disk, then merging them to generate final feature stacks for each slice (see script: "PART1_3DFeatSplitter_Sigma8.bsh"). Our approach also includes the calculation of additional custom 2D features (see script: "PART2_2DFeatSplitter.bsh"), which are merged with the original 3D feature stacks (see script: "PART3_Combine3Dand2DFeatures.bsh") allowing us to calculate 175 features per slice. The scripts required to execute this process are provided. In a departure from the TWS plugin all training annotations were generated externally, by using the Microscopy Image Browser (MIB) to generate label images where each training class is defined by the numerical value of a voxel (value of 1 for class 1, value of 2 for class 2, etc). A script was then used to import each training label image into TWS with its corresponding feature stack (Fig. 1G), to extract feature measurements from a predetermined number of randomly selected voxels per class (see script: "CorePartA_ExtractTrainingData_ArbitraryClassNum.bsh"). The feature data extracted via this

script can then be pooled using the "append" function in WEKA, then imported into WEKA Explore

for analysis and machine learning to generate various trained models, as was shown in Figures 1-3.

These trained models are loaded by a final script

("CorePartB_ApplyClassifierToFeatureStacks.bsh"), which is designed to evaluate pre-existing

feature stacks.

These scripts were used for machine learning throughout the manuscript as follows. The analyses

in figures 1-3 used models trained to detect three classes of material; "Sol" (Class 1), representing

granular protein content of the cytoplasm and nucleus; "Matter" (Class 2), representing electron-

dense homogenous materials that do not belong to a membrane (i.e. mitochondrial matrix);

"Memb" (Class 3), representing genuine cellular membranes. The division of non-membranous

materials into two classes was aimed to account for the class imbalance problem; Training

annotations were drawn by hand in MIB. For the analyses shown in figures 1-2, feature stacks for

each of the three altered test datasets were generated separately before using a shared pool of

training labels to extract feature measurements for 1500 voxels per-class-per-slice from eight

slices (12000 samples per class, 36000 samples per stack, 108000 feature measurements total).

For the training of the six models in Figures 1-2; RF (Random Forest), J48 (the java compatible

extension of Ross Quinlan's C4.5 classifier), MLP (Multi-Layer Perceptron), DT (Decision Table),

JRip (java compatible propositional rule-based RIPPER), and PART (Projective Adaptive Resonance

Theory neural network), algorithm specific parameters are enclosed within the training logs we

have provided for each model (See GitHub: https://github.com/BenPadman/AIVE.git). Differences

between model outputs were calculated using ImageJ, by calculating the absolute difference in

predicted values (ΔV) at each voxel between each model and every other model; results were

visualized by averaging all ΔV comparisons for a given model when compared to all others.

Data shown in Figures 4-8 used pooled feature data from all three datasets to train a random

forest classifier as described above, except that training involved five unique classes of material.

The first three of these classes were used to detect non-membrane materials: "Void" (Class 1),

representing the homogenous extracellular regions that lack an electron signal; "Sol" (Class 2),

representing granular protein content of the cell and nucleus; "Matter" (Class 3), representing

electron-dense homogenous materials that do not belong to a membrane (i.e. mitochondrial

matrix). The remaining two classes (4 & 5) were sub-categories of membrane: "Memb" (Class 4),

representing the majority of cellular membranes, and "Vesc" (Class 5) to represent vesicular

membranes; the membrane sub-classes needed to be merged together to generate a final AIVE

result, but separation into two classes allowed the classification of cellular vesicles by machine

learning alone.

Organelle classification for AIVE

Some of the raw classifier outputs described above were exploited to assist the classification of

organelles. For example, dilation of the "Void" class was used to assist classification of the plasma

membrane, whereas the Sol class labels in the nucleus were used to assist annotation of the

nuclear envelope. Manual labelling in MIB was then used to classify the Golgi apparatus, early

endosomes, late endosomes, lipid droplets. The mitochondria were also classified in this way, but

multiple sub-classes were used to ensure that adjacent un-fused mitochondria were considered

independently from one another, thus allowing each individual mitochondrion to be indexed for

independent analyses. The ER was classified through a process of elimination; it was the only

remaining unclassified membrane after classification of all other organelles. Final class labels were

exported from MIB as a label image. The label images containing all class labels were separated

into individual binary stacks, with each image representing one class of organelle. These binary

stacks were blurred with a 10nm gaussian filter. The 32-bit predictions generated for "Memb" and

"Vesc" (Classes 4 &5; described above) were combined, then multiplied with the binary stacks and

blurred with a 10nm gaussian filter, before merging the results with 3D CLAHE results calculated

earlier. This generates the final AIVE results for each individual organelle class.

U-Net training

Two U-Nets with 3D anisotropic architecture (same padding) were trained via stochastic gradient

descent with momentum (SGDM; 0.001 weight decay; 0.9 momentum), using 80×80×80 voxel

input patches (1040 x 1040 x 1600nm in XYZ), 5×5×5 filters, leaky ReLu activation layers (0.001

scale), and a dice pixel classification layer. Both U-Nets were structured with 16 input filters and 4

layers, and data augmentations (both 2D & 3D) were applied to 80% of the input patches during

training with a piecewise learning schedule (initial learn rate of 0.01; dropping by 80% every 5

epochs) for 35 epochs (512 iterations per epoch). The only difference between these two U-Nets was the random seed value used during network initialization; The first U-Net used a random seed of 1337, while the second used a value of 1338 (Fig.3F). To minimize segmentation discontinuities between adjacent prediction tiles [61], each U-Net was applied using 5% tile overlap with four different orientations of the dataset (original, X-flipped, Y-flipped, Z-flipped). The results were then returned to their original orientations before averaging to create an ensemble result. Image classification by a human image analyst in the same figure was conducted on two consecutive days and required approximately 10 minutes to complete.

3D Distance Probes

To simulate the effect of result variation on membrane contact site measurements, all static distance probes were assigned in accordance with standard criteria for designation of a contact site (< 35 nm separation; [10]). Initial coordinates for the 3D distance probes were chosen by running a stochastic simulation of diffusion limited aggregation in Mathematica (version 9). These coordinates were annotated in a blank TIFF stack with identical dimensions to the test dataset, by converting the voxel located at each indicated coordinate from black (value of 0) to white (value of 255). Probe coordinates further than 35nm from (or inside) a membrane were removed from the collection by binarizing the membranes shown in Fig. 1C, dilating them by 35 nm, then subtracting the un-dilated membranes from the result; any 3D probes outside the boundaries of this result were deleted. This process was repeated until 30 unique probe locations were identified (Fig.2H).

Analyses of mitochondrial membranes (Fig.3G-M) required an additional set of criteria for the distance probes, since the original 3D probes were generated without relevance to any specific organelle (Fig.1H) and most were located further than 35 nm from a mitochondrion. 3D probe coordinates for mitochondrial membranes were therefore generated by using the same 35nm membrane boundary mask described for general membranes above in combination with an additional mask representing the mitochondria, as follows; All mitochondrial class labels shown in Fig.3F were averaged, binarized, dilated by 35 nm, then applied in a boolean AND operation with the 35 nm membrane mask used to define the first probe positions (Fig.1H). The resulting mask was then used to assign new random points as described above, until thirty new random points were defined within 35nm of a membrane and within 35 nm of the mitochondrial class label

(Fig.4K). Value range for distance measurements (Fig.2I; Fig.3L) were calculated by subtracting the

minimum recorded distance from the maximum recorded distance for one probe. Relative

absolute deviations (aka. relative mean absolute difference) in measurements from each probe

(Fig.2J; Fig.3M) were calculated by dividing the mean absolute difference (difference between

each measurement and the mean value) of measurements by the arithmetic mean of all

measurements from that probe, then averaging the result.

Nuclear pore analyses

Nuclear pores were initially chosen by defining the approximate 3D coordinates of each pore in

the raw dataset. A cuboidal region surrounding each coordinate (300 nm on each axis) was then

extracted from the raw (Fig.S1F), VE (Fig.3N), AIE (Fig.3O), and AIVE (Fig.3P) image stacks. These

pores were then arranged into an identical orientation as follows; All four representations of data

(raw, VE, AIE, & AIVE) for each pore were averaged to create datasets for inter-pore alignment,

without biasing toward any particular method. Using the Fijiyama plugin (V4.0.11) for FIJI (v1.53),

Each averaged pore was manually aligned to the next pore in the sequence to generate a

preliminary alignment between all 18 pores, which was then refined via automatic block-matching,

also through Fijiyama [62]. Individual transform operations describing this alignment were then

composed into one rigid transformation per nuclear pore, allowing each pore to be aligned into

the same orientation across all four processing methods. Feret diameters of the pores were

measured binarizing the pore, then applying the built-in "analyse particles" function in FIJI (v1.53)

at the centrally aligned slice of the pore. All depictions of the individual nuclear pores are shown

post-alignment (Fig.3N-P; Fig.S1F), to account for the variable orientation of some pores in the

original dataset (Fig.S1E).

3D rendering and numerical analyses

All 3D surfaces were generated using the marching cubes algorithm with octree binning (256

points per leaf) in ParaView (v5.7). Distance measurements and mapping for visualization were

conducted in ParaView, by calculating the signed distance field for one polygonal mesh then

measuring the field value at each vertex of a second polygonal mesh. The signed distance fields

were calculated using a custom ParaView filter proposed in a blog post by Cory Quammen

(Technical leader in Scientific Computing Team; Kitware); the filter has been provided. All 3D

renderings were generated using the Cycles render engine in Blender (v2.93). All other

morphometric analyses were conducted using the MorpholibJ plugin for FIJI [63]. Difference in

voxel values (ΔV) were calculated using the built-in "image calculator" function in FIJI (v1.53).

Mathematica (version 9) was used for all other statistical analyses and charts, including the

Buffon's needle simulation in Fig.8.

Figure Legends:

Figure 1: Variable AI predictions predominantly occur at object boundaries during segmentation.

(A-E) 3D test data acquired via FIB-SEM imaging, showing an **(A)** overview and **(B)** inset 2D slice from dataset with **(C)** corresponding predictions generated by a random forest classifier trained to detect membranes, alongside **(D)** 2D binary segmentation and **(E)** 3D surfaces based on the membrane predictions. **(F)** 2D slices depicting raw membrane predictions for test dataset (from Fig.1B) with values (V) shown as percentage of their total dynamic range (0 to 1), made using six different machine learning algorithms (also provided as 3D renders in **Fig.S2A**); RF (Random Forest), J48, MLP (Multi-Layer Perceptron), DT (Decision Table), JRip, and PART (Projective Adaptive Resonance Theory). **(G)** Visual depiction of the absolute difference in voxel values (ΔV) between each prediction from F, showing the average value difference for each 2-way value comparison as a percentage of the total dynamic range (0 to 1).

Figure 2: AI-directed voxel extraction (AIVE) minimizes the influence of variable AI predictions.

(A-D) Flow chart depicting the technical basis of AI-directed voxel extraction (AIVE), in which **(A)** raw 3D data from the FIB-SEM is used to **(B)** train and apply a classifier model to generate probability maps depicting the location of cellular membranes (dynamic range, 0 to 1), while **(C)** intensity values from the original dataset are normalized via CLAHE (dynamic range, 0 to 255); **(D)** values from B are then multiplied by values from C (see also **Fig.S1A**), which isolates the voxel values belonging to a membrane to generate raw AIVE data (application of AIVE to other classes of signal provided in **Fig.S1B**). **(E)** 2D slices depicting raw AIVE data generated using each algorithm shown in Fig.1F (3D renders of the same data provided in **Fig.S2C**), and **(F)** averaged ΔV values for each 2-way comparison between one AIVE output and each of the other outputs (both shown as a percentage of 8-bit dynamic range; 0 to 255). **(G)** Quantification of the average ΔV per slice for each 2-way comparison made between the model predictions (Fig.1F-G) and between AIVE outputs (E-F) generated using those same model predictions (quantitation and 3D renders of blurred AI predictions shown in **Fig.S1C** & **Fig.S2B** respectively). **(H)** 3D rendering of membrane surfaces from the test dataset (from Fig.1E) viewed from the front and above, indicating the

positions of the static 3D probes used to benchmark membrane distance measurements

(membranes nearest to each 3D probe are indicated by red shading; see Video 1 for rotation

animation depicting probe positions). **(I)** Quantification of the distance between each 3D probe

(shown in H) and the nearest membranes detected by the six models shown in Fig.1F (raw

distance measurements provided in **Fig.S2A**), after segmentation of the raw predictions, blurred AI

predictions, or AIVE outputs, showing **(I)** the range (min-max) for values reported by each probe

between models, and **(J)** the relative absolute deviation in measurements from each probe. Plot

markers in I-J indicate value for individual 3D probes. Interquartile range indicated by box, median

indicated by horizontal lines, minimum and maximum indicated by whiskers in G, I & J.

Figure 3: Object classification does not alter the precision or accuracy of

segmentation by AIVE.

(A-E) Flow chart depicting object classification during AIVE, in which **(A)** organelle labels and **(B)**

membrane predictions are multiplied to generate **(C)** a contextually masked membrane prediction,

(D) which is then combined with the normalized EM data to generate **(E)** AIVE data belonging to a

specific organelle. **(F)** 2D examples of mitochondrial classification labels provided by one human

on different days, or two 3D U-Nets with different initialization seeds. **(G-J)** 3D rendered results

classified using the labels shown in Fig.3F, after classification via **(G)** conventional instance

segmentation, or by using the binary labels **(H)** for selective voxel extraction (VE), **(I)** AI extraction

(AIE), or **(J)** AIVE. **(K-M)** Quantitative analysis of **(K)** the distance between mitochondrial

membranes and each mitochondrial 3D probe (probe locations shown in **Fig.S1D**; also see Video 2

for rotation animation depicting probe positions), **(L)** the range of values recorded by each 3D

probe, and **(M)** the relative absolute deviation in 3D probe measurements, for all datasets shown

in Fig.3H-J. **(N-P)** 2D maximum intensity projections of 18 individual nuclear pores, which were

spatially aligned prior processing and segmentation via **(N)** VE, **(O)** AIE, or **(P)** AIVE (Unprocessed

pores shown and pore locations shown in **Fig.S1E & F** respectively). **(Q)** Feret's diameter

measurements of the central transport channel from each nuclear pore after segmentation via VE,

AI or AIVE. (black arrowhead indicates diameter reported by Schuller *et al.* [25]; hollow arrowhead

range indicates diameters reported by Zimmerli *et al.* [26] indicated by). **(R-T)** 2D projections and

3D renderings showing the median averaged data generated via by **(Q)** VE, **(R)** AI, or **(S)** AIVE

(purple), with spatial alignment to the nuclear pore complex structure (orange) published in 2021

by Schuller *et al.* [25]. Plot markers in **K-L** indicate values for individual 3D probes. Plot markers in **T** indicate measurement for individual nuclear pores. Interquartile range indicated by box, median indicated by horizontal lines, minimum and maximum indicated by whiskers. Scalebars; **N-S**, 100nm.

Figure 4: The compositional analysis of membranes in bulk cellular volumes via AIVE.

(A-D) 3D renderings of the three FIB-SEM datasets processed via AIVE for quantitative analysis, showing **(A)** the test data used in Fig.1-3 alongside overviews of **(B)** the dataset it originated from (Set #1; leading lines indicate location), and **(C-D)** two additional datasets processed via AIVE (Set #2 & Set #3 respectively). All three overviews are shown at the same scale (See Video 3 for animated cutaways of the data). **(F-H)** Compositional analysis of membrane volumes belonging to non-nuclear cytoplasmic organelles in **(F)** Set #1, **(G)** Set #2, and **(G)** Set #3. See Video 3 for animated cutaways of overview data. Scalebars; A, bounding box is 2400nm wide; B-D, markers on scale grid for overviews are 1000nm apart.

Figure 5: The morphometric analysis of individual mitochondria via AIVE.

(A-G) Mitochondria of interest that were detected in the AIVE-processed datasets shown in Fig.4, showing the membranes of each mitochondrion (Red) alongside its matrix (Blue), and the binary skeleton calculated for its bulk morphology (Black; longest-optimal path endings shown in yellow); these include **(A)** the smallest mitochondrion, **(B)** the largest mitochondrion, **(C)** a nanotunnelling mitochondrion, **(D)** the most elongated mitochondrion, **(E)** the mitochondrion with the longest continuous tubule, **(F)** the most spherical mitochondrion, and **(G)** the mitochondrion closest to the average values detected by all morphological metrics. For animations of these data, see Video 4 (for **A & B**), Video 5 (for **C, F & G**), and Video 7 (for **D & E**). **(H-L)** Quantitative analysis of key morphological metrics all mitochondria in each set; **(H)** box whisker chart for the total volume of each mitochondrion, **(I)** a scatter chart of showing the relationship between the volume membrane (x-axis) and volume of matrix (y-axis) in each mitochondrion, with **(J-L)** box whisker charts showing additional metrics for **(J)** mitochondrial sphericity, **(K)** mitochondrial length, and **(L)** mitochondrial elongation. See Video 4 Plot markers indicate value for individual mitochondria.

Hollow plot markers indicate mitochondria that were intersected by the dataset boundary; these

mitochondria were excluded from analyses sensitive to incomplete mitochondria (Fig.5 A, E, D, G,

J-L). Interquartile range for mitochondrial population indicated by box, median indicated by

horizontal lines, minimum and maximum indicated by whiskers. Scalebars; A-G, markers on scale

grid are 1000nm apart.

Figure 6: AIVE for the analysis of the mitochondrial interactome.

(A-F) Quantification of the closest approach distance separating each individual mitochondrion

from **(A)** the rest of the mitochondrial network, **(B)** the ER, **(C)** early endosomes, **(D)** late

endosomes, **(E)** lipid droplets, or **(F)** cytosolic vesicles. **(G)** Average percentage of all mitochondria

detected within 35nm of indicated organelle per set. **(H-J)** 3D rendering of distance-mapped

mitochondrial surfaces from a sub-region in Set #1. The mitochondrial surfaces are shaded with a

colorimetric scale indicating surface distances from **(H)** other mitochondria, **(I)** the ER, or **(J)**

cytosolic vesicles. **(K-M)** Surface area histograms for the mitochondria located within 100nm of

indicated organelle, depicting the average percentage of a mitochondrial surface located within

indicated distance of **(K)** other mitochondria, **(L)** the ER, or **(M)** cytosolic vesicles (Arrowheads in K

& L indicate the anomalous peak, explained in **Fig.S4A-D**). Plot markers in A-F indicate value for

individual mitochondria. Interquartile range indicated by box, median indicated by horizontal lines,

minimum and maximum indicated by whiskers. Data in G are mean ± S.D calculated for the

averages in each dataset. Histogram data in K-L are mean surface area percentages binned into

5nm distance brackets relative to the target organelle; mean values for each set are also shown.

Figure 7: Mitochondrial intrusion sites facilitate membrane contact with the ER

and other mitochondria.

(A-C) Examples of ER-to-Mitochondrion intrusion sites detected in **(A)** Set #1, **(B)** Set #2, and **(C)**

Set #3, displayed as partial (50nm) 2D averaged projections of raw EM and AIVE data (left panels),

rendered in 3D (centre panels), and with a colorimetric distance-map relative to the ER (right

panels). For animations of these data, see Video 8. Additional examples are provided in Fig.S5A-C.

(D-E) Examples of Mitochondrion-to-Mitochondrion intrusion sites detected in **(D)** Set #1 and **(E)**

Set #2, displayed as partial (50nm) 2D averaged projections of raw EM and AIVE data (left panels),

rendered in 3D (centre panels), and with a colorimetric distance-map relative to other

mitochondria (right panels). For animations of these data, see Video 9. Additional examples are

provided in Fig.S5D-E. **(F)** 3D rendered examples of empty intrusion sites detected in each dataset.

Scalebars, 200nm; bounding dimensions of each 3D dataset, 1.5 x 1.5 x 1.5 μm.

Figure 8: Volume EM is essential for the analysis of mitochondrial intrusion sites

(A) 3D rendered examples of mitochondria containing intrusion sites, with their interior cavities

colourized by the category of intrusion. **(B)** Quantitative analysis comparing the membrane

volume mitochondria in the presence or absence of mitochondrial intrusions. **(C)** Bar chart

showing percentage of the mitochondrial population by their number of intrusions, with plot

markers indicating expected values calculated by a Poisson distribution of same average. **(D-F)**

Quantitative analyses of each intrusion cavity by category, showing their **(D)** total interior volume,

(E) radius of the largest inscribed sphere they could enclose, and **(F)** the length of their major

ellipsoidal axis. **(G)** Statistical modelling of Buffon's needle problem, showing the probability of

detecting an intact mitochondrial intrusion (solid line) or only half an intrusion (dashed line) when

an intrusion is present, as a function of the Z-depth ("d") in a hypothetical FIB-SEM dataset or TEM

section. Plot markers in B indicates value for individual mitochondria, in D-F indicates value for

individual intrusion cavities. Interquartile range for each chart indicated by box, median indicated

by horizontal lines, minimum and maximum indicated by whiskers. Scalebars; markers on scale

grid are 1000nm apart.

Figure S1: Supplementary AIVE descriptions and demonstrations.

(A) Signal matrix depicting how the input values for a voxel contribute to the output value during

AIVE (red line indicates approximate location of threshold used for 3D rendering and

segmentation). **(B)** 3D examples of other various categories of electron signal processed via AIVE,

as detected in the test dataset shown in Fig.1. **(C)** Quantification of the average ΔV per slice for

each 2-way comparison made between the model predictions from **Fig.1F** after the application of

a 10nm Gaussian blur (3D renders of corresponding data also provided in **Fig.S2B**). **(D)** 3D

renderings indicating the position of each static 3D probe used for mitochondrial distance

measurements in Fig.3, viewed from the front and above (membranes nearest to each 3D probe are indicated by red shading). See Video 2 for rotation animation showing probe positions. **(E)** Maximum value projection of AIVE processed data for the nucleus used in Fig.3, depicting the position of each nuclear pore used for analysis in Fig.3N-T. **(F)** Averaged value projections of the unprocessed EM data after spatial alignment for analyses shown in **Fig.3N-T**. Interquartile range for chart indicated by box, median indicated by horizontal lines, minimum and maximum indicated by whiskers. Scalebars; C, 1000nm; D, 100nm.

Figure S2: Data at different stages of AIVE processing.

(A-C) 3D depictions of 2D data shown in figures 1 & 2, showing the 3D surfaces generated using **(A)** raw predictions, **(B)** blurred predictions, or **(C)** AIVE with one of six different trained models (RF, J48, MLP, DT, JRip, and PART). **(D)** Raw distance measurements between the 3D probes (shown in Fig.2H) and membrane surfaces under every condition shown in **A-C**. **(E-G)** raw AIVE data from mitochondrion #33 (set 1) displayed as a sum projection from **(E)** the front and **(F)** above (Also see Video 5), with **(G)** a orthoslice montage of EM and AIVE data through the mitochondrial nanotunnel (position shown by box in F). Scalebars; 500nm.

Figure S3: AIVE of mouse skeletal muscle.

(A-C) AIVE reconstruction of mouse skeletal muscle tissue revealing mitochondrial nanotunnel connections. **(A)** 3D rendering depicting myofibrils (red), mitochondria (purple) and sarcoplasmic reticulum (gold) within a volume measuring 8 x 8 x 6.2 μm FIB-SEM volume, with **(B)** a rendered cutaway displaying the interconnectivity of mitochondria within the tissue (mitochondrial nanotunnels indicated by black arrow-heads). **(C)** Representative example of one nanotunnel linking the mitochondrial network around a myofibril. See Video 6 for animated depiction of data in A-C. Scalebar distances as indicated in figure.

Figure S4: Simulated intrusive contacts.

(A-B) 3D renderings of simulated contact between two spheres (128 voxel diameter) with colorimetric distance scale indicating the distance of Sphere #1 from Sphere #2, **(A)** when both

structures are simple spheres, or **(B)** when Sphere #2 intrudes into Sphere #1. **(C-D)** Surface area histograms depicting the surface area percentage of Sphere #1 that is located within the indicated distances of Sphere #2, **(C)** when both structures are simple spheres, or **(D)** when Sphere #2 intrudes into Sphere #1 (Arrowhead in D indicates the anomalous peak discovered in Fig.6). **(E)** Illustrative description of "Buffon's Needle Problem" with corresponding equations, which describe the probability of a randomly distributed needle of known length ("L") overlapping a line, if those lines are regularly spaced by a known distance ("d"). Histogram data in C-D are surface area percentages binned into 5 voxel distance brackets relative to Sphere #2. Scalebars in A-B, 10 voxels.

Figure S5: Additional examples of intrusive contact sites.

(A-C) Additional examples of ER-to-Mitochondrion intrusion sites, displayed as partial (50nm) 2D averaged projections of raw EM and AIVE data (left panels), rendered in 3D (centre panels), and with a colorimetric distance-map relative to the ER (right panels). **(D-E)** Additional examples of Mitochondrion-to-Mitochondrion intrusion sites, displayed as partial (50nm) 2D averaged projections of raw EM and AIVE data (left panels), rendered in 3D (centre panels), and with a colorimetric distance-map relative to other mitochondria (right panels). Scalebars, 200nm; bounding dimensions of each 3D dataset, 1.5 x 1.5 x 1.5 µm.

Video Legends:

Video 1: General membrane distance probes

Animated rotation of the test volume, depicting the 3D membrane probe positions (from Fig.2H) used for membrane distance measurements in Fig.2-3. General membranes are shown in diffuse navy blue, while membranes nearest to a 3D probe are indicated by red shading.

Video 2: Mitochondrial membrane distance probes

Animated rotation of the test volume, depicting the 3D mitochondrial membrane probe positions (from Fig.S1D) used for mitochondrial membrane distance measurements in Fig.4. Mitochondrial membranes are shown in diffuse navy blue, while membranes nearest to a 3D probe are indicated by red shading.

Video 3: Overview cutaways

Animated cutaways of the data presented in Fig.4B-D. The video begins with all organelles shown, before cutting away the endoplasmic reticulum (ER) and nuclear envelope (Nu), followed by the Golgi and vesicles (VE), and finally, the early endosomes (EE), late endosomes (LE), and lipid droplets (LD). The video concludes by revealing these organelles in reverse order. Render key is provided within the animation.

Video 4: Smallest and largest mitochondria

Animation depicting the structure of the smallest (Fig.5A) and largest (Fig.5B) volumed mitochondria from figure 5. Render key is provided within the animation.

Video 5: Spherical, nanotunneling, and unremarkable mitochondria

Animation depicting the structure of the most spherical mitochondrion (Fig.5F), a nanotunneling mitochondrion (Fig.5C), and the "least remarkable" mitochondrion (Fig.5G) from figure 5. Render key is provided within the animation.

Video 6: AIVE of mouse skeletal muscle

Animation of rendered data shown in Fig.S5, displaying an AIVE reconstruction of FIB-SEM data from mouse skeletal muscle tissue, measuring 8 x 8 x 6.2 µm. As the video begins, myofibrils (red) and sarcoplasmic reticulum (gold) are cut away to reveal the mitochondrial network (purple). Black arrowheads then appear to highlight the mitochondrial nanotunnels, as the camera pans

961 around the dataset. The video concludes by focussing on a specific mitochondrial nanotunnel

962 which is tightly wrapped around a myofibril, thus connecting mitochondria on opposing sides of

963 the myofibril.

Video 7: Most elongated and longest mitochondria

965 Animation depicting the structure of the most elongated (Fig.5D) and longest (Fig.5E)

966 mitochondria from figure 5. Render key is provided within the animation.

Video 8: Intrusive contact by ER

968 Animation of the ER intrusions (Fig.7A-C) depicted in figure 7. The video begins with the full ROI

969 centred around the intrusion sites, with colourations indicated by render key. As the intruding ER

970 fades away, the mitochondrial surface is transitioned to the distance-map (0 to 100nm) colour

971 scale shown in figure 7, before rotation of the camera vertically over the datasets. The video

972 concludes by returning to the initial shading.

Video 9: Intrusive contact by mitochondria

974 Animation of the mitochondrial intrusions (Fig.7D-E) depicted in figure 7. The video begins with

975 the full ROI centred around the intrusion sites, with colourations indicated by render key. As the

976 intruding mitochondrion fades away, the mitochondrial surfaces are transitioned to the distance-

977 map (0 to 100nm) colour scale shown in figure 7, before rotation of the camera vertically over the

978 datasets. The video concludes by returning to the initial shading.

References

1. Voeltz, G.K., et al., *Making the connection: How membrane contact sites have changed our view of organelle biology.* Cell, 2024. **187**(2): p. 257-270.

2. Scorrano, L., et al., *Coming together to define membrane contact sites.* Nature Communications, 2019. **10**(1).

3. Flis, V.V. and G. Daum, *Lipid Transport between the Endoplasmic Reticulum and Mitochondria.* Cold Spring Harbor Perspectives in Biology, 2013. **5**(6): p. a013235-a013235.

4. Schrader, M., et al., *The different facets of organelle interplay—an overview of organelle interactions.* Frontiers in Cell and Developmental Biology, 2015. **3**.

5. Shai, N., M. Schuldiner, and E. Zalckvar, *No peroxisome is an island - Peroxisome contact sites.* Biochim Biophys Acta, 2016. **1863**(5): p. 1061-9.

6. Coleman, R.A., *It takes a village: channeling fatty acid metabolism and triacylglycerol formation via protein interactomes.* J Lipid Res, 2019. **60**(3): p. 490-497.

7. Wang, Y., et al., *Nanoparticle-Driven Controllable Mitochondrial Regulation through Lysosome-Mitochondria Interactome.* ACS Nano, 2022. **16**(8): p. 12553-12568.

8. Venditti, R., M.C. Masone, and M.A. De Matteis, *ER-Golgi membrane contact sites.* Biochem Soc Trans, 2020. **48**(1): p. 187-197.

9. Valm, A.M., et al., *Applying systems-level spectral imaging and analysis to reveal the organelle interactome.* Nature, 2017. **546**(7656): p. 162-167.

10. Jing, J., et al., *A molecular toolbox for interrogation of membrane contact sites.* The Journal of Physiology, 2019. **598**(9): p. 1725-1739.

11. Huang, X., et al., *Current and Emerging Approaches for Studying Inter-Organelle Membrane Contact Sites.* Frontiers in Cell and Developmental Biology, 2020. **8**.

12. Peng, L., et al., *The Advances and Challenges of Deep Learning Application in Biological Big Data Processing.* Current Bioinformatics, 2018. **13**(4): p. 352-359.

13. Xu, C.S., et al., *Transforming FIB-SEM Systems for Large-Volume and*, in *Volume Microscopy*. 2020. p. 221-243.

14. Gallusser, B., et al., *Deep neural network automated segmentation of cellular structures in volume electron microscopy.* Journal of Cell Biology, 2023. **222**(2).

15. Liu, L., et al., *DeepContact: High-throughput quantification of membrane contact sites based on electron microscopy imaging.* Journal of Cell Biology, 2022. **221**(9).

16. Nguyen, T.N., et al., *ATG4 family proteins drive phagophore growth independently of the LC3/GABARAP lipidation system.* Mol Cell, 2021. **81**(9): p. 2013-2030 e9.

17. Richard, M.D. and R.P. Lippmann, *Neural Network Classifiers Estimate Bayesian a posteriori Probabilities.* Neural Computation, 1991. **3**(4): p. 461-483.

18. Rumelhart, D.E., G.E. Hinton, and R.J. Williams, *Learning representations by back-propagating errors.* Nature, 1986. **323**(6088): p. 533-536.

19. Wang, Z., E. Wang, and Y. Zhu, *Image segmentation evaluation: a survey of methods.* Artificial Intelligence Review, 2020. **53**(8): p. 5637-5674.

20. Lee, R.G., et al., *Quantitative subcellular reconstruction reveals a lipid mediated inter-organelle biogenesis network.* Nat Cell Biol, 2024. **26**(1): p. 57-71.

21. Gibson, S.F.F., *Constrained elastic surface nets: Generating smooth surfaces from binary segmented data*, in *Medical Image Computing and Computer-Assisted Intervention — MICCAI'98*. 1998. p. 888-898.

22. Lorensen, W.E. and H.E. Cline, *Marching cubes: A high resolution 3D surface construction algorithm.* ACM SIGGRAPH Computer Graphics, 1987. **21**(4): p. 163-169.

23. Kim, S., et al., *A novel interpolation scheme for dual marching cubes on octree volume fraction data.* Computers & Graphics, 2017. **66**: p. 169-178.

24. Thevathasan, J.V., et al., *Nuclear pores as versatile reference standards for quantitative superresolution microscopy.* Nature Methods, 2019. **16**(10): p. 1045-1053.

25. Schuller, A.P., et al., *The cellular environment shapes the nuclear pore complex architecture.* Nature,
2021. **598**(7882): p. 667-671.
26. Zimmerli, C.E., et al., *Nuclear pores dilate and constrict in cellulo.* Science, 2021. **374**(6573).
27. Drozdz, M.M., et al., *Formation of a nucleoplasmic reticulum requires de novo assembly of nascent*
phospholipids and shows preferential incorporation of nascent lamins. Scientific Reports, 2017. **7**(1).
28. Vincent, A.E., et al., *Mitochondrial Nanotunnels.* Trends Cell Biol, 2017. **27**(11): p. 787-799.
29. Friedman, J.R., et al., *ER sliding dynamics and ER–mitochondrial contacts occur on acetylated*
microtubules. Journal of Cell Biology, 2010. **190**(3): p. 363-375.
30. Robertson, J.M. and A.F. Siegel, *Designing Buffon's Needle for a Given Crossing Distribution.* The
American Mathematical Monthly, 2018. **93**(2): p. 116-119.
31. Lebovitz, S., N. Levina, and H. Lifshitz-Assa, *Is AI Ground Truth Really True? The Dangers of Training*
and Evaluating AI Tools Based on Experts' Know-What. MIS Quarterly, 2021. **45**(3): p. 1501-1526.
32. Woodhouse, I.H., *On 'ground' truth and why we should abandon the term.* Journal of Applied
Remote Sensing, 2021. **15**(04).
33. Zick, M., R. Rabl, and A.S. Reichert, *Cristae formation-linking ultrastructure and function of*
mitochondria. Biochim Biophys Acta, 2009. **1793**(1): p. 5-19.
34. Linda, S., et al., *Mitochondrial Diseases and Myopathies: A Series of Muscle Biopsy Specimens with*
Ultrastructural Changes in the Mitochondria. Ultrastructural Pathology, 2009. **16**(3): p. 263-275.
35. Vincent, A.E., et al., *The Spectrum of Mitochondrial Ultrastructural Defects in Mitochondrial*
Myopathy. Scientific Reports, 2016. **6**(1).
36. Jenkins, B.C., et al., *Mitochondria in disease: changes in shapes and dynamics.* Trends Biochem Sci,
2024. **49**(4): p. 346-360.
37. Siegmund, S.E., et al., *Three-Dimensional Analysis of Mitochondrial Crista Ultrastructure in a Patient*
with Leigh Syndrome by In Situ Cryoelectron Tomography. iScience, 2018. **6**: p. 83-91.
38. Ryu, K.W., et al., *Cellular ATP demand creates metabolically distinct subpopulations of mitochondria.*
Nature, 2024.
39. Huang, X., et al., *Kissing and nanotunneling mediate intermitochondrial communication in the heart.*
Proc Natl Acad Sci U S A, 2013. **110**(8): p. 2846-51.
40. Lavorato, M., et al., *Increased mitochondrial nanotunneling activity, induced by calcium imbalance,*
affects intermitochondrial matrix exchanges. Proc Natl Acad Sci U S A, 2017. **114**(5): p. E849-E858.
41. Miyazono, Y., et al., *Uncoupled mitochondria quickly shorten along their long axis to form indented*
spheroids, instead of rings, in a fission-independent manner. Scientific Reports, 2018. **8**(1).
42. Giacomello, M., et al., *The cell biology of mitochondrial membrane dynamics.* Nature Reviews
Molecular Cell Biology, 2020. **21**(4): p. 204-224.
43. Lackner, L.L., *The Expanding and Unexpected Functions of Mitochondria Contact Sites.* Trends in Cell
Biology, 2019. **29**(7): p. 580-590.
44. Hertz, H., *Ueber die Berührung fester elastischer Körper.* Journal für die reine und angewandte
Mathematik, 1882. **1882**(92): p. 156-171.
45. Scorrano, L., et al., *Coming together to define membrane contact sites.* Nat Commun, 2019. **10**(1): p.
1287.
46. Yeo, H.K., et al., *Phospholipid transfer function of PTPIP51 at mitochondria-associated ER*
membranes. EMBO Rep, 2021. **22**(6): p. e51323.
47. Freyre, C.A.C., et al., *MIGA2 Links Mitochondria, the ER, and Lipid Droplets and Promotes De Novo*
Lipogenesis in Adipocytes. Mol Cell, 2019. **76**(5): p. 811-825 e14.
48. Kim, H., et al., *Structural basis for mitoguardin-2 mediated lipid transport at ER-mitochondrial*
membrane contact sites. Nat Commun, 2022. **13**(1): p. 3702.
49. Naon, D., et al., *Splice variants of mitofusin 2 shape the endoplasmic reticulum and tether it to*
mitochondria. Science, 2023. **380**(6651): p. eadh9351.
50. De Vos, K.J., et al., *VAPB interacts with the mitochondrial protein PTPIP51 to regulate calcium*
homeostasis. Hum Mol Genet, 2012. **21**(6): p. 1299-311.
51. Stoica, R., et al., *ER-mitochondria associations are regulated by the VAPB-PTPIP51 interaction and*
are disrupted by ALS/FTD-associated TDP-43. Nat Commun, 2014. **5**: p. 3996.

52.   Csordas, G., et al., *Imaging interorganelle contacts and local calcium dynamics at the ER-*
mitochondrial interface. Mol Cell, 2010. **39**(1): p. 121-32.
53.   Seligman, A.M., H.L. Wasserkrug, and J.S. Hanker, *A new staining method (OTO) for enhancing*
contrast of lipid--containing membranes and droplets in osmium tetroxide--fixed tissue with
osmiophilic thiocarbohydrazide(TCH). J Cell Biol, 1966. **30**(2): p. 424-32.
54.   Silva, M.T., F.C. Guerra, and M.M. Magalhaes, *The fixative action of uranyl acetate in electron*
microscopy. Experientia, 1968. **24**(10): p. 1074.
55.   Walton, J., *Lead aspartate, an en bloc contrast stain particularly useful for ultrastructural*
enzymology. J Histochem Cytochem, 1979. **27**(10): p. 1337-42.
56.   Arganda-Carreras, I., et al., *Consistent and Elastic Registration of Histological Sections Using Vector-*
Spline Regularization, in *Computer Vision Approaches to Medical Image Analysis*. 2006. p. 85-95.
57.   Rueden, C.T., et al., *ImageJ2: ImageJ for the next generation of scientific image data.* BMC
Bioinformatics, 2017. **18**(1).
58.   Frank, E., et al., *Data mining in bioinformatics using Weka.* Bioinformatics, 2004. **20**(15): p. 2479-
2481.
59.   Arganda-Carreras, I., et al., *Trainable Weka Segmentation: a machine learning tool for microscopy*
pixel classification. Bioinformatics, 2017. **33**(15): p. 2424-2426.
60.   Schindelin, J., et al., *Fiji: an open-source platform for biological-image analysis.* Nat Methods, 2012.
9(7): p. 676-82.
61.   Huang, B., et al., *Deep Convolutional Segmentation of Remote Sensing Imagery: A Simple and*
Efficient Alternative to Stitching Output Labels, in *IGARSS 2018 - 2018 IEEE International Geoscience*
and Remote Sensing Symposium. 2018. p. 6899-6902.
62.   Fernandez, R. and C. Moisy, *Fijiyama: a registration tool for 3D multimodal time-lapse imaging.*
Bioinformatics, 2021. **37**(10): p. 1482-1484.
63.   Legland, D., I. Arganda-Carreras, and P. Andrey, *MorphoLibJ: integrated library and plugins for*
mathematical morphology with ImageJ. Bioinformatics, 2016. **32**(22): p. 3532-3534.

1111

Figure 1

A — *Backscattered Electron (BSE) Image*

Value (V) = Detected signal (8-bit)

0   64   128   192   256

B — *BSE Image (inset)*

C — *Membrane predictions*

V = Probability (P: 0-100%)

0%   25%   50%   75%   100%

D — 2D membrane segmentation

E — 3D membrane isosurfaces

Machine learning algorithm

| RF | J48 | MLP | DT | JRip | PART |

F — *Raw predicted values (V)*

V:
100%
75%
50%
25%
0%

G — *Value difference (ΔV)*

ΔV:
50%
25%
0%

Figure 2

A Raw EM voxel values — V_{EM} — 255 / 128 / 0

B AI predictions + Gaussian — P_{memb} — 100% / 50% / 0%

C Normalized voxels + Median — $N(V_{EM})$ — 255 / 128 / 0

D *AI-Directed Voxel Extraction (AIVE)* — V_{AIVE} — 255 / 128 / 0

$$V_{AIVE} = N(V_{EM}) \times P_{memb}$$

Machine learning algorithm

RF   *J48*   *MLP*   *DT*   *JRip*   *PART*

AIVE outputs

E Raw values (V) — V: 100% / 75% / 50% / 25% / 0% — [0,255]

F Value difference (ΔV) — ΔV: 50% / 25% / 0%

G Mean value difference (ΔV) per slice (%dynamic range)

Raw AI predictions   AI-directed Voxel Extraction (AIVE)

H *3D distance probes (front view)*   *3D distance probes (top view)*

I Value range measured by probe (nm)

J Relative absolute deviation of probe measurements (%probe mean)

Figure 3

A Mitochondrial Classification

Mask ☐ 1 ■ 0

B Raw AI predictions

100% 50% 0%

Multiply

C AI predictions + Gaussian

100% 50% 0%

(Mitochondrial membranes only)

D Normalized voxels

Multiply

E AIVE of mitochondrion

255 128 0

F Class labels by Human (2D) / Class labels by U-Net (2D)

Human (Day 1) | Human (Day 2) | U-Net 1 (1337) | U-Net 2 (1338)

G Class labels by Human (3D) / Class labels by U-Net (3D)

Human (Day 1) | Human (Day 2) | U-Net 1 (1337) | U-Net 2 (1338)

Raw Labels

H VE

I AIE

J AIVE

K Probe-to-membrane separation distances (nm)

VE: Human 1, Human 2, U-Net 1, U-Net 2
AIE: Human 1, Human 2, U-Net 1, U-Net 2
AIVE: Human 1, Human 2, U-Net 1, U-Net 2

L Value range measured by probe (nm)

VE, AIE, AIVE

M Relative absolute deviation of probe measurements (%probe mean)

VE, AIE, AIVE

Individual nuclear pores

Projections (2D)

N VE

O AIE

P AIVE

Q Feret's diameter of central transport channel within pore (nm)

VE, AI, AIVE

Median averaged nuclear pores

Projections (2D) | *Front view (3D)* | *Cutaway (3D)*

R VE

S AIE

T AIVE

90°

Figure 4

A

Colour Legend

- ⬤ Endoplasmic Reticulum (ER)
- ⬤ Mitochondria (mt)
- ⬤ Early Endosomes (EE)
- ⬤ Late endosomes & Lysosomes (LE)
- ⬤ Golgi apparatus
- ⬤ Nuclear envelope
- ⬤ Lipid droplets (LD)
- ⬤ Vesicles (VE)

B *Dataset #1*

C *Dataset #2*

D *Dataset #3*

1000 nm

1000 nm

E *Dataset #1*
(Non-nuclear organellar membranes)

Mt (32.35%)
Golgi (5.78%)
EE (2.41%)
LE (3.8%)
VE (4.34%)
LD (1.69%)
ER (49.62%)

F *Dataset #2*
(Non-nuclear organellar membranes)

Mt (34.14%)
EE (0.91%)
LE (2.04%)
VE (5.18%)
LD (0.82%)
ER (56.92%)

G *Dataset #3*
(Non-nuclear organellar membranes)

Mt (27.33%)
Golgi (5.00%)
EE (1.61%)
LE (2.42%)
VE (4.62%)
LD (1.17%)
ER (57.84%)

Figure 5

Mitochondria of interest

A *Smallest volume*
Set3 Mito#14
(V = 0.035 μm³)

B *Largest volume*
Set2 Mito#12
(V = 3.468 μm³)

C *Nanotunnelling mito*
Set1 Mito#33

D *Most elongated*
Set1 Mito#57
(Aspect Ratio = 11.89)

E *Longest optimal path*
Set1 Mito#12
(11.229 μm)

F *Most spherical*
Set3 Mito#11
(Sphericity = 0.860)

G *Most average*
Set1 Mito#74
(all metrics)

Render Key
Detected membranes
Imputed mitochondrial matrix
Binary skeleton
Longest path endpoints

1000 nm
1000 nm
1000 nm

H Total volume per mitochondrion (μm³)
set1 set2 set3

I $V_{matrix} = 0.8164 \times V_{memb} - 0.0054$
Volume of matrix (μm³)
Volume of membrane (μm³)
Set1 (R = 0.9947)
Set2 (R = 0.9980)
Set3 (R = 0.9903)

J Mitochondrial sphericity
set1 set2 set3

K Longest optimal path through mitochondria (μm)
set1 set2 set3

L Elliptoid aspect ratio (major/minor axis)
set1 set2 set3

Figure 6

Figure 7

Intrusive contact by ER

2D EM & AIVE **3D AIVE** **3D distance maps**

Cutaway | Full ROI | Original orientation | Alternative orientation

A Set #1 - Mito 44   −50 nm   0 nm   +50 nm

B Set #2 - Mito 16

C Set #3 - Mito 9

Intrusive contact by mitochondrion

2D EM & AIVE **3D AIVE** **3D distance maps**

Cutaway | Full ROI | Original orientation | Alternative orientation

D Set #1 - Mito 89   −50 nm   0 nm   +50 nm

E Set #2 - Mito 59

F Empty intrusions

Set #1 - Mito 12 | Set #2 - Mito 23 | Set #3 - Mito 7

Render Key

- Mitochondria
- ER
- Endolysosomes
- Vesicles
- Nucleus
- Intruding organelle
- Distance map scale (nm)

0   20   40   60   80   100

Figure 8

A

Set #1 Mito 11 · *Set #1 Mito 12* · *Set #1 Mito 19* · *Set #1 Mito 20* · *Set #1 Mito 32* · *Set #1 Mito 33* · *Set #1 Mito 36* · *Set #1 Mito 44* · *Set #1 Mito 54* · *Set #1 Mito 58* · *Set #1 Mito 59* · *Set #1 Mito 61* · *Set #1 Mito 71* · *Set #1 Mito 85* · *Set #1 Mito 87* · *Set #1 Mito 89* · *Set #2 Mito 12* · *Set #2 Mito 16* · *Set #2 Mito 21* · *Set #2 Mito 23* · *Set #2 Mito 24* · *Set #2 Mito 35* · *Set #2 Mito 37* · *Set #2 Mito 41* · *Set #2 Mito 42* · *Set #2 Mito 56* · *Set #2 Mito 57* · *Set #3 Mito 07* · *Set #3 Mito 09* · *Set #3 Mito 12* · *Set #3 Mito 13* · *Set #3 Mito 16* · *Set #3 Mito 18* · *Set #3 Mito 24* · *Set #3 Mito 27* · *Set #3 Mito 31*

Render Key
- Mitochondrial Membranes
- Category 1 (Empty)
- Category 2 (Mito)
- Category 3 (ER)

1000 nm / 1000 nm

B mitochondrial membrane volume (μm³) — Intrusions: − +

C Mitochondrial population by number (*n*) of intrusions (% all mitochondria) — *n* = 0 1 2 3 4 5

D Interior volume of intrusion cavity (μm³) — Empty (1), Mito (2), ER (3)

E Radius of largest inscribed sphere (nm) — Empty (1), Mito (2), ER (3)

F Major ellipsoidal axis length of cavity (nm) — Empty (1), Mito (2), ER (3)

G Probability (P) of an intrusion being fully intact in a detectable orientation — Z-depth (d) of dataset (nm) — $L = d$ — *Half intrusion* (dashed), *Whole intrusion* (solid)