

CoralLite: μ CT Reconstruction of Coral Colonies from Individual Corallites

Jess Jones¹, Leonardo Bertini^{1,2,3}, Kenneth Johnson³,
Erica Hendy¹, and Tilo Burghardt¹

¹ University of Bristol {jess.jones,e.hendy}@bristol.ac.uk;
tilo@cs.bris.ac.uk

² University of Liverpool leonardo.bertini@liverpool.ac.uk

³ Natural History Museum k.johnson@nhm.ac.uk

Abstract. The life history of an individual coral is archived within the accreting skeleton of the colony. While reef-forming coral colonies (e.g. massive *Porites* sp.) may live for hundreds of years and deposit calcareous structures many metres in height and width, their living tissue is a thin outer surface layer comprised of asexually-dividing polyps that only survive a few years. To understand the rate and timing of polyp division and the consequences for colony skeletal growth, scientists need to track the skeletal corallite deposited around each polyp. Here we propose CoralLite, an annotated μ CT scan dataset of entire calcareous skeletons and an associated, first corallite deep learning reconstruction baseline. CoralLite combines fully quantified volumetric segmentations with cross-slice linking for visualisations of 3D models for each corallite up to colony scale. For segmentation, we propose and evaluate in detail a hybrid V-Trans-UNet architecture applicable to segmenting tiled μ CT virtual slabs of *Porites* sp. colonies. The model is pre-trained on weakly annotated data and topology-aware fine-tuned using fully annotated slice sections with 8k+ manual corallite region annotations. On unseen slices of the same colony, the resulting model reaches 0.94 topological accuracy at mean Dice scores of 0.77 on the same colony and projection axis, and 0.63 mean Dice scores on a different, biologically unrelated specimen. Whilst our experiments are limited in scale and context, our results show for the first time that visual machine learning can effectively support full 3D individual corallite modelling from μ CT scans of coral skeletons alone. For reproducibility and as a baseline for future research we publish our full dataset of 697 μ CT slices, 37 partial or full slice annotations, and all network weights and source code with this paper.

Keywords: corallite reconstruction · μ CT · volumetric segmentation · topology-aware loss · animal biometrics · imageomics · computer vision

1 Introduction

Scientific Motivation. *Scleractinia*, or stony corals, can form colonies by asexual division of individual polyps. Each polyp deposits a surrounding protective tube of aragonite ($CaCO_3$), known as a corallite, which forms a porous wall between neighbouring polyps and collectively creates the lasting skeletal structure of the colony [7, 9]. For example, an individual massive *Porites* sp. colony can grow large enough to form a micro-atoll, or alternatively a dome-shaped




Fig. 1. Overview of CoralLite. The proposed pipeline comprises three main stages: (1) *Novel Dataset*: Sequential 2D cross-section slices from 3D μ CT scans of a *Porites* sp. coral colony are annotated and processed into 5x224x224 tiled snippets. (2) *Volumetric Segmentation*: A 3D-context-aware Trans-UNet pipeline with a hybrid ResNet50-ViT backbone (details in Fig. 4) generates corallite segmentations using a topological loss for accurate boundary prediction. (3) *Corallite Reconstruction*: Segmented tiles are stacked into a volume to generate a 3D corallite-level colony model.

skeletal structure that can be over 5 metres in diameter, and encompass over 500 years of growth [11, 14]. Fossils of *Scleractinia* date back to the middle Triassic period and have therefore been a feature of ocean ecosystems for more than 230 million years. Since the Miocene, the genus *Porites* has been functionally important in shallow, tropical and sub-tropical waters as one of the key reef-framework builders of warm-water coral reefs [2]. Massive *Porites* sp. are also the most commonly studied coral for palaeoclimate reconstructions and growth rate variability associated with environmental change [17]. In μ CT scans, the structure of these colonies appears as large stacks of cross-sectional slices containing many small interlocked corallites as shown in Figure 1. Manual tracing of the internal space within individual corallites through these volumes is possible and practised today for fine-grained skeletal analysis, however this approach is time-consuming and laborious [15]. *Porites* sp. serve as an excellent case study for analysis automation with direct applicability to vital coral reef research.

Research Aims and Contributions. This paper considers whether such identity-preserving reconstructions of individual corallites in 3D from μ CT data alone can be automated using deep learning for animal biometrics [13]. The approach is outlined in Figure 1. In contrast to most existing coral reconstruction research, we attempt fully automated segmentation of individual corallites of

a scanned colony and quantify baselines of this task here for the first time. Our aim is not to present a complete biological model of coral growth at this stage since expert annotations are limited and expensive to produce at scale. Instead, our aim is to assess whether a reasonably sized, annotated dataset and volumetric learning together are sufficient to support an automated colony reconstruction at corallite level which is useful to practitioners today and sets baselines for the community to build on this work. Recent progress in medical and volumetric image segmentation in particular presented a useful starting point for our work [4, 10]. Procedurally, we trained 3D-aware segmentation models to identify corallite regions encoded in binary segmentation maps, and then apply cross-slice linking procedures supported by topological learning constraints to assemble predicted regions into per-corallite, identity-preserving 3D models. In line with UNet architecture tests for corallites in [5], we build on hybrid Video-TransUNet architectures and research [3, 22] that combines local texture modelling with broader contextual reasoning. Our ResNet-ViT baselines follow this paradigm and operate on stacks of adjacent coral μ CT slices such that cross-slice aggregation can exploit continuity in 3D data. Our contributions are:

- **Annotated Dataset:** 697 μ CT scan slices plus a set of 8,412 manual corallite ground truth segmentations in full slices of massive *Porites* complemented further by point-annotation corallite priors across 33 of the slices;
- **Topology-guided 3D Reconstruction Pipeline:** a volumetric segmentation pipeline based on a V-TransUNet model with cross-slice spatial feature context learning, together with a topological loss for discouraging merged regions with full release of key code and network weights;
- **Baseline Results and Evaluation:** baselines with a first task evaluation across same-colony, cross-axis, and cross-specimen test settings; and
- **Visualisation Bridge:** a simple cross-slice tracing and geometry generation procedure for reconstructing individual corallites as 3D objects suitable for manual inspection by marine scientists using standard graphics packages;

2 Related Work

3D Coral Reconstruction in Marine Biology. Recent research has shifted toward volumetric 3D imaging techniques, such as computed axial tomography (CT) and micro-CT (μ CT), to precisely measure coral extension rates and corallite demography [21, 15]. These high-resolution reconstructions allow for a more accurate, non-destructive assessment of individual polyp longevity and skeletal growth trajectories compared to traditional 2D analysis [1, 7]. By resolving the 3D spatial organisation of coral modules, these methods hold rich, previously inaccessible information about colonial developmental dynamics [15].

Corallite Reconstruction Automation. Automated computer-vision pipelines for individual corallite reconstruction from whole-colony μ CT remain scarce: most coral μ CT work still emphasises manual processing, rather than learned, 3D corallite models. A notable recent exception [5] applies deep semantic segmentation (U-Net family models, including Attention U-Net) to μ CT of scleractinians such as *Montastrea cavernosa* and *Porites astreoides*, separating skeleton from pores to support quantitative 3D analyses without reconstructing corallites.

Cross-Domain Architectural Consideration. The segmentation of corallites in μ CT data shares several characteristics with medical image segmentation: structures are small relative to the image, neighbouring slices are strongly correlated, and dense annotation is difficult to obtain [12, 20]. For such settings, CNNs remain a common baseline, while hybrid CNN–Transformer models have shown that local feature extraction and longer-range context can be combined effectively [3]. Transformers are attractive in this setting because they can model broader spatial relationships than standard convolution alone, but they are also data-hungry and often benefit from pre-training [8]. Hybrid architectures help bypass the heavy data requirements of standard Transformers by using a convolutional backbone to handle initial feature extraction. We therefore adopt a ResNet–ViT backbone design for the UNet layout similar to Trans-UNet [3].

Temporal Contexts and Topological Guidance. A second relevant direction is the use of temporal or sequential context modules. Zeng et al. [22] introduce a temporal context module (TCM) for blending information across neighbouring video frames in X-ray segmentation. Although our input is not video, adjacent μ CT slices also contain short-range continuity. We use this idea to encode relationships across slices in a volumetric window. Finally, our downstream task places additional demands on the segmentation output. Standard region-overlap metrics such as Dice are not always sufficient when neighbouring instances must remain separated for corallite reconstruction. This motivates an auxiliary loss designed to penalise merged or over-expanded predicted regions. Similar topologically-motivated constraints have recently proven effective in medical image segmentation; for instance, the architecture in [18] utilises a binary topological interaction loss to rapidly enforce exclusion and containment relationships among intricate anatomical structures. [19] proposed cIDice, a topology-preserving loss function for tubular structures that is calculated on the intersection of segmentation masks and their morphological skeletons.

3 Dataset and Setup

Acquisition and Dataset Details. We use high-resolution μ CT data from three *Porites* specimens from the Natural History Museum (London, UK) and Naturalis Biodiversity Center (Leiden, NL) collections. Specimens were scanned at the Natural History Museum using a Nikon XT H 225 ST system. Reconstructions were exported as 16-bit volumes and further processed into virtual slabs using Avizo 2021.2. The data comprise both weakly annotated (corallite volume centre points) and fully annotated slices (full manual binary segmentations of corallite volumes). With oversight from coral experts, six annotators produced weak annotations for the Naturalis 6785 specimen. In addition, we provide four fully annotated slices comprising 8,412 unique corallite interior segments across two *Porites* specimens and two projection settings. Inter-annotator agreement was not assessed within the current scope and remains an objective for future work. The full dataset with all ground truth annotations will be made publicly available with this paper. Figure 2 visualises 5 of the 697 2D slices of full-colony specimen next to associated corallite segmentation results. Datasets contained in this study are listed in Tab. 1 and detailed in following four slice sets:




Fig. 2. Dataset Sample Slices next to Segmentation Results. Shown are 5 of the 697 full-colony slices from the dataset next to segmentation outputs using inference performed on the Naturalis 6781 colony. This specific coral colony was originated from a single coral polyp forming a single corallite. The coral grew over a five year period into a lumpy hemispheric colony of thousands of corallites. Depicted are samples from both the Growth projection axis, and the Ortho projection axis. These examples illustrate the task’s complexity and provide qualitative result visualisations across an entire colony volume and different μ CT projection axes.

- **Naturalis 6785:** 60 non-contiguous slices from a growth-axis projection, 33 of which are partially annotated. These data are used for pre-training.
- **Naturalis 6781 Growth:** 248 slices from the growth-axis projection of a different colony (registration number ZMA.Coel.6781). Slice #1279 is used for fine-tuning and slice #2499 for evaluation.
- **Naturalis 6781 Ortho:** 388 slices from an orthogonal projection of the same colony. Slice #1003 is used only for evaluation.
- **Astraeiformis:** single fully annotated evaluation slice of different *Porites* specimen from the Natural History Museum collection (cat.#1883.11.8.9).

Segmentation Task and Evaluation Settings. The segmentation target is defined straight forwardly as binary corallite volume detection: each pixel is classified as corallite or background. Ground truth dataset slices are represented as 8-bit binary images with 255 and 0 pixel values, respectively. Evaluation is designed to characterise 3 generalisation levels for learned segmentation models:

- **(Task A) same colony, same axis:** Naturalis 6781 Growth #2499;
- **(Task B) same colony, different axis:** Naturalis 6781 Ortho #1003;
- **(Task C) entirely different specimen:** *Astraeiformis*.

Table 1. Summary of Datasets. Our dataset contribution contains four different slice sets with varying levels of annotation. Partial annotation stipulates corallite centres, whilst full annotation provides exact masks. We use selections from these to perform pre-training, fine-tuning, and different generalisation performance evaluations.

Specimen	Axis	Slices	Annotation	Slice Usage
Naturalis 6785	Growth	60	33 partial	Pre-train
Naturalis 6781	Growth	248	Full: #1279, #2499	Fine-tune or Test
Naturalis 6781	Ortho	388	Full: #1003	Test
Astraeiformis	–	1	Full slice	Test

4 Methods and Experiments

4.1 Volumetric Input Generation

Sliding-Window Tile Extraction. Each μ CT slice has high spatial resolution, while individual corallites occupy only a small number of pixels. To avoid down-sampling and fit network tensor dimensions, we process each slice using tiled crops of size 224×224 (see Figure 3). Tiles are generated by a sliding window with step size k . During pre-training on Naturalis 6785 we use $k = 224$, while during fine-tuning on Naturalis 6781 Growth #1279 we use $k = 50$ for overlap.

Cross-Slice Volumetric Stacking. To exploit continuity across adjacent slices, each training sample is constructed as a snippet of neighbouring tiles: $\mathbf{x} \in \mathcal{R}^{D \times H \times W}$, where $H = W = 224$, D is odd, and the annotation corresponds to the centre slice. We adopt a $D = 5$ regime following the methodology of [22] to capture local volumetric context. This depth provides a balance between capturing the longitudinal trajectory of corallites and minimising the memory footprint during training. Figure 3 illustrates the tiling process and resulting volumetric inputs.




Fig. 3. Volumetric Input Tensor Generation. Tiled volumetric stacking provides local context from adjacent slices for segmentation. Left to right: Illustrative example of raw image tiling with window 224×224 and step size 224; An example composite tile with annotation overlay in red, from the Naturalis 6785 dataset; An example snippet of dimensions $5 \times 224 \times 224$. Note that slice n is situated in the centre of the volume.




Fig. 4. Full Architecture Diagram. Our pipeline uses a Video-TransUNet [22] design but applied here to *spatial* volume analysis. It utilises CNN feature encoding with short-range volumetric learning and complex spatial context modelling via a ViT before DCNN decoding: **(a)**: $5 \times 224 \times 224$ input tensors are provided to **(b)**: a pre-trained ResNet-50 CNN encoder with 3x skip connections. **(c)**: Feature encodings are blended across the volumetric dimension before being passed to **(d)**: a pre-trained 12 layer ViT with Multi-head Attention and **(f)**: MLP transformer block components. **(e)**: Finally, through cascaded DCNN up-sampling incorporating skip information a final segmentation output of size 224×224 is produced in **(g)**. (adapted Video Trans-UNet of [22])

4.2 Segmentation Pipeline

Hybrid Trans-UNet Architecture. We adapt the Video-TransUNet design of [22] to snippets formed as tiled μ CT slabs. We symmetrise and reinterpret Video-TransUNet’s temporal dimension as spatial depth. Our model then combines a ResNet-50 encoder, a temporal context module (TCM) [22], a 12-layer ViT, and a cascaded decoder with skip connections as shown in Figure 4.

Stage-by-stage Feature Processing. Given an input stack $\mathbf{x} \in \mathcal{R}^{D \times H \times W}$, the ResNet encoder extracts per-slice features. The attention-based context module then considers features across the depth dimension, allowing neighbouring slices to contribute local volumetric context. The resulting representation is passed to the ViT encoder, which models longer-range spatial dependencies before decoding using a deconvolutional DCNN to yield a segmentation mask output $f(\mathbf{x}) \in \mathcal{R}^{H \times W}$ for the processed tile, i.e. the centre slice of $\mathbf{x}$.

4.3 Topology-aware Loss Regime

Motivation for Connectivity Penalty. For downstream single corallite reconstruction, it is important that neighbouring corallites remain separated in the predicted masks. Standard overlap losses such as binary cross-entropy and Dice do not directly penalise some important failure modes, such as thinly merged neighbouring corallite volumes or over-expanded false positive regions.




Fig. 5. Topological Error Maps. Two example depictions of connectivity-driven error maps used for training loss calculation highlighting incorrectly connected or generated corallite regions. The topological error map $\mathcal{E}$ is presented on the left of each image pair, while composite ground-truth (blue) and prediction (red) maps are given on the right of each pair. The grayscale error map demonstrates how the topological loss can specifically target both false positives (isolated red pixel groups) and merged regions (bridged by pixels shown in red), encouraging structural separation specifically, which is vital for accurate downstream 3D tracing of individual corallites.

Component Penalty Formulation. Using basic connected-component analysis on binary annotation and prediction maps via `skimage.measure`, we obtain the set of predicted regions P and labelled regions L . For each predicted connected component $p \in P$ with area A_p , we identify its nearest labelled region $l \in L$ with area A_l and define the component penalty:

$$p_\epsilon = \min \left(\left| 1 - e^{\frac{A_p - A_l}{A_p}} \right|, 1 \right). \quad (1)$$

Error Map and Topological Loss Stipulation. Each pixel belonging to a region p is assigned the corresponding penalty if it does not belong to the matched ground-truth region l . This yields an error map $\mathcal{E}(n)$ that highlights for each pixel n its contribution towards incorrect corallite connections, false positive corallites and over-expanded regions (see Figure 5). This map is used to calculate a topological loss by averaging evidence over all N pixels of a tile during training as:

$$\mathcal{L}_{\text{Topo}} = \frac{1}{2N} \sum_{n=1}^N \sigma(\mathcal{E}(n)) - 0.5. \quad (2)$$

4.4 Definition of Learning Objectives

Pre-Training. Since weakly annotated pre-training data (see Tab. 1) lacks sufficient fine-grained information for connectivity analysis we utilise standard Binary Cross-Entropy [6] and Dice [16] losses without topological awareness as bulk pre-training objective:

$$\mathcal{L} = \frac{1}{2} (\mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{Dice}}). \quad (3)$$

Training Loss Mixture. During main fine-tuned training of the network—where manually annotated pixel-specific ground truth is available—we add the topological loss with $T = 0.1$ after a warm-up period of 50 epochs:

$$\mathcal{L} = \frac{1}{2} (\mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{Dice}}) + T\mathcal{L}_{\text{Topo}}. \quad (4)$$




Fig. 6. Topological Contribution Tuning, Training and Validation Plots. Performance development during main fine-tuning of our model. *(left)* A comparison of different topology coefficients T confirms stable performance around 0.1-1.0. *(right)* With peak validation accuracy of 0.96 and saturated loss reduction at 320 epochs at a stable plateau, the depicted model (FT VLrg 0.1xT + GN in Tab. 2) has learned to segment unseen validation corallite slice data without significant over-fitting.

4.5 Network Training

Two-Stage Optimisation with Augmentation. Learning is happening in two stages. First, the segmentation model is pre-trained on the weakly annotated Naturalis 6785 data where ResNet-50 + ViT are initialised with standard ImageNet weights. Secondly, the pipeline is fine-tuned in a main training run on a fully annotated slice from the target colony, Naturalis 6781 Growth #1279. All experiments are conducted with Adam, learning rate 5×10^{-4} , and batch size 5. In both stages, 10% of the available training data are held out for validation. The remaining training tiles are augmented by random rotation, horizontal and vertical flips, and, in later experiments, additive Gaussian noise. We use a snippet depth of 5, fix context module parameters after pre-training, and set the topological loss coefficient to $T = 0.1$. Figure 6 depicts the full training curves.

5 Segmentation Results and Evaluation

5.1 Summary of Performance Baselines

Segmentation Metrics for Quantification. Segmentation performance is evaluated using the mean Dice Similarity Coefficient (mDSC) [16] as standard primary segmentation metric alongside a task-specific Topology Score (TS) defined as $1 - \mathcal{L}_{\text{Topo}}$ for approximating reproduction success of connected structures essential for downstream corallite tracing.

Segmentation Performance. Table 2 summarises the main quantitative corallite segmentation results across all settings. Fine-tuning substantially improves all over the pre-trained baseline across all metrics. Performance differences based on hyperparameter choices are fully documented, but are relatively small, showing the stability and suitability of the adapted V-Trans-UNet for the task at hand. The overall best-performing configuration in our experiments combines dense tiling ($k = 50$) during fine-tuning, topology-aware loss at ($T = 0.1$), and additional augmentation with Gaussian noise. On the target tasks, this model reaches mDSC scores of 0.77 on **(Task A)** Naturalis 6781, 0.71 on **(Task B)** Naturalis 6781, and still cross-species performance 0.63 on **(Task C)** *Astraeiformis*.

Table 2. Corallite Segmentation Performance Baselines. Full testing performance baselines for all models across the three generalisation tasks (see Sec. 3) using **mDSC**: Mean Dice similarity, and **Topo**: Topological accuracy: $1 - \mathcal{L}_{Topo}$. The full model with maximum tile overlap, topological loss utilisation and Gaussian Noise augmentation generalises best and offers best, or joint best, performance across all testing domains. Best or shared best performance with the full model is highlighted in bold.

Model	Parameters	Task A (same axis) Naturalis 6781 Growth #2499		Task B (same species) Naturalis 6781 Ortho #1003		Task C (diff. species) <i>Astraeiformis</i>	
		mDSC	Topo	mDSC	Topo	mDSC	Topo
Pre-training only	n/a	0.008	0.66	0.01	0.75	0.01	0.68
FT (fine-tuned)	$k = 224$	0.75	0.93	0.67	0.92	0.61	0.88
	$k = 112$	0.76	0.94	0.69	0.93	0.57	0.80
	$k = 50$	0.76	0.94	0.71	0.93	0.56	0.79
FT + $\mathcal{L}_{Topo}$ no tile overlap ($k = 224$)	$T = 0.01$	0.76	0.93	0.65	0.92	0.59	0.85
	$T = 0.1$	0.76	0.93	0.64	0.92	0.59	0.85
	$T = 1$	0.76	0.93	0.62	0.90	0.60	0.85
	$T = 10$	0.76	0.93	0.65	0.92	0.59	0.85
	$T = 100$	0.77	0.93	0.62	0.90	0.60	0.85
FT + $\mathcal{L}_{Topo}$ max tile overlap ($k = 50$)	$T = 0.01$	0.77	0.94	0.70	0.93	0.51	0.80
	$T = 0.1$	0.77	0.94	0.70	0.93	0.52	0.80
	$T = 1$	0.77	0.94	0.70	0.93	0.51	0.79
	$T = 10$	0.76	0.94	0.69	0.93	0.48	0.80
	$T = 100$	0.76	0.94	0.69	0.93	0.53	0.81
FT + $\mathcal{L}_{Topo}$ max tile overlap ($k = 50$) plus noise augmentation	$T = 0.01$	0.76	0.94	0.70	0.92	0.61	0.87
	$T = 0.1$	0.77	0.94	0.71	0.93	0.63	0.87
	$T = 1$	0.77	0.94	0.70	0.93	0.60	0.86
	$T = 10$	0.77	0.94	0.70	0.93	0.63	0.89

5.2 Detailed Segmentation Evaluation

Fine-Tuning is Critical. Without fine-tuning, the pre-trained model performs poorly on all evaluation slices. Expectedly, only detailed training annotations that reflect the exact corallite shape are sufficient to learn how to segment them given the visual complexity of coral colony scans. However, fine-tuning on one fully annotated slice is sufficient to recover strong performance on unseen, same colony data, and transfers reasonably to a different projection axis and specimen.

Qualitative Effect of Topology-aware Loss & Augmentation. Adding the topological loss does not produce large changes in mDSC, but improves the separation and shape of neighbouring predicted regions. Whilst the footprint in the overall Topo quantification shown in Tab. 2 is minimal due to the averaging nature of the overall loss shown, the practical effect is visible in qualitative comparisons (e.g. see Figure 7), where connected false positives between adjacent corallites are reduced. For instance, when testing for Task B at $k = 50$, utilising $\mathcal{L}_{Topo}$ reduces the absolute corallite count error by 2% (measured with ($T = 0.1$)). The effect is most consistent for densely tiled fine-tuning settings at values of $T \leq 1$). Performance on the *Astraeiformis* slice testing cross-species performance is lower than on the Naturalis 6781 slices throughout. Adding Gaussian noise augmentation alone improves *Astraeiformis* mDSC by up to 8% (compare final two row blocks of Tab. 2), while preserving performance on same-colony test slices. Qualitative inspection suggests that false positives in regions with scanner noise or echos are particularly improved.




Fig. 7. Connectivity Impact of Topological Loss. Four composite example pairs of predictions (red) and ground-truth data (blue) superimposed on coral tiles, that is (left) without and (right) with topological loss utilised ($k = 50, T = 1$). Note the improvement in structural corallite separation, particularly in areas highlighted by the red bounding box. **Bottom right pair:** Note that scanning noise and xray echoes, here on a detail of the *Astraeiformis* slice, lead to major out-of-domain (OOD) misinterpretation and false positives independent of topological loss considerations. Detection of OOD content is one key direction of future work in this domain and wider AI research.

5.3 3D Corallite Reconstruction

Basic Cross-Slice Tracing. In order to turn 2D corallite segmentation into 3D corallite tube reconstructions suitable for visualisation of entire colonies and for manual inspection, for each connected region R_i in slice S_n we search for a matching counterpart region R_j in the previous slice S_{n-1} . A match is accepted based on centroid distance and intersection-over-union (IoU) satisfying $d(R_i, R_j) < \gamma$ and $\text{IoU}(R_i, R_j) > \beta$. Thresholds can be altered for task-specific visualisations – visuals in this paper are generated with $\gamma = 0.3$ and $\beta = 0.3$. If a match is found, R_i is assigned to the same corallite object as R_j ; otherwise, a new object trace is started.

3D Colony Visualisation based on Corallite Objects. Each traced region is represented by its centroid, major and minor axis lengths, and in-plane orientation (see Figure 8, left). These parameters define an ellipse in 3D space at a depth corresponding to slice index. Ellipses belonging to the same track are then connected along the slice axis to form corallite object models, which can then be visually inspected as individual corallites (see Figure 8, middle) or in large groups up to the context of a full colony volume (see Figure 8, right). The resulting reconstructions capture corallite tube structures, alignment and growth directions within the colony and many other features. Whilst such visualisations at colony scale have never been attempted before, we emphasise that this is for inspection only. A full scientific 3D quantification against manually traced ground truths remains a subject for future work, as discussed in Section 6.




Fig. 8. Corallite Reconstruction and 3D Visualisation. (*left*) Per-slice corallite representation based on centroid, major and minor axis lengths, and in-plane orientation shown as vector pair. (*middle*) Corallite #87 as viewed along the X-axis where centroid tracing has translated stacks of 2D segments into a coherent 3D structure for visualisation that captures the natural geometry, curvature, split and termination as well as scale of the corallite. (*right*) Model of reconstructed colony viewed along the Z-axis via corallites #0 - #20,000. Note how this first individual corallite visualisation at colony scale highlights key coral features such as directional alignment of local corallite groups, dense packing of tubes, and the overall outward growth of the colony.

6 Brief Discussion

First Proof-of-Concept Pipeline. The results suggest that volumetric segmentation via deep learning architectures can indeed provide a workable basis for automatic, individual-level corallite reconstruction from μ CT data alone. In particular, pre-training on weak annotations and fine-tuning on just one fully annotated slice from the target colony were sufficient to obtain useful segmentation quality on unseen slices from the same colony. The resulting masks were adequate for generating interpretable 3D corallite object models at colony scale performed and visualised here for the first time.

Assumptions, Constraints and Future Work. Since our fully annotated dataset of manually labelled 8,412 corallite regions is relatively limited in size, reported results only provide performance baselines, but imply by no means performance caps on the tasks investigated. Furthermore, the full 3D reconstruction is not scientifically quantified here and is generated for first interpretable visualisations only. Further, our cross-specimen results suggest that out-of-domain application remains a practical challenge. Thus, this study provides first datasets and baselines to enable and inspire research on full scientific coral colony reconstructions – it does not yet claim to provide an imageomics application that can reliably estimate corallite demography indicators such as budding rate variability and polyp lifespan as a function of colony age. This paper lays down the first foundational work for this via public datasets, annotations, and comparison baselines for the community. That is, further research is required to completely automatically produce fully accurate and scientifically quantified 3D corallite-level reconstruction of coral colonies from CT scans alone.




Fig.9. Further Segmentation Results. Examples from Naturalis 6781 colony; *(left)* Annotation example, slice 1279 (Growth). *(right)* Inference outputs, clockwise from top-right: slice 2043 (Ortho), 2718 (Ortho), 1379 (Growth), and 0119 (Growth).

7 Conclusion

We presented CoralLite, a new μ CT coral dataset, with individual corallite annotations and a V-Trans-UNet inspired deep learning pipeline for corallite segmentation and reconstruction, together with full segmentation baselines. We produced a first approximate colony-scale coral visualisation made from generated individual corallite objects. The underlying segmentation model was fully evaluated on manually annotated test data. Our hybrid V-Trans-UNet inspired architecture combined short-range cross-slice feature blending with a topological loss aimed at discouraging merged neighbouring regions. In our experiments, fine-tuning on one fully annotated slice from a target colony was shown sufficient to produce useful segmentation quality on unseen data. The resulting masks supported an automatic 3D corallite reconstruction process at colony scale. The presented system establishes a proof-of-concept that fully automated reconstructions of corallites are feasible at colony level. It provides useful visualisations of colonies and marks a first step towards imageomic automation for understanding the rate and timing of polyp division and the consequences for colony skeletal growth. We publish the full CoralLite dataset of 697 μ CT slices, 37 partial or full slice annotations with 8,412 manual corallite region labellings, and all network weights and pipeline source code with this paper. We hope that this can enable and inspire the computer vision for conservation community to engage in furthering research in this area in order to understand corals better and find new ways to safeguard the remaining reefs of the planet.

Acknowledgements

J.J. is supported by FARSCOPE EPSRC CDT. J.J. and T.B. would like to thank and acknowledge the support of the University of Bristol HPC facilities. E.H. and L.B. gratefully acknowledge the dedicated labelling of corallites on scan slices completed by an enthusiastic team of University of Bristol Environmental Geoscience undergraduates; Arisha Fancy, Poppy Hammond, Sharifah Nor Qistina Syed Izuan, Louisa Lou, Emma Mody, and Khalil Imran Abdul Rahman, who were funded through the NERC GW4 DTP Undergraduate Research Experience Placement scheme and the University of Bristol School of Earth Sciences Summer Internship programme. Support for L.B. and this study is also acknowledged through 4D-REEF, a Marie Skłodowska-Curie Innovative Training Network funded by European Union Horizon 2020 research and innovation programme under grant agreement number 813360. We are also grateful for support from SYNTHESIS+ Transnational Access project NL-TAF-TA3-009, the Natural History Museum Science Investment Fund and the University of Bristol Covid Recovery Fund.

References

1. Benedetti, M.C., Bramanti, L., Priori, C., Erra, F., Iannelli, M., Bulleri, F., Santangelo, G.: Polyp longevity in a precious gorgonian coral: hints toward a demographic approach to polyp dynamics. *Coral Reefs* **39**(4), 1125–1136 (2020)
2. Cabioch, G., Davies, P., Done, T., Gischler, E., Macintyre, I., Wood, R., Woodroffe, C.: *Encyclopedia of modern coral reefs: structure, form and process*. Springer Science & Business Media (2010)
3. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306* (2021)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 424–432. Springer (2016)
5. Coronel-Zegarra, A., Knaub, J.L., Merk, V., Pandya, A.: Leveraging deep learning semantic segmentation for imaging coral skeletons. *Journal of Structural Biology* p. 108313 (2026)
6. Cox, D.R.: The regression analysis of binary sequences. *Journal of the Royal Statistical Society. Series B (Methodological)* **20**(2), 215–242 (1958)
7. Darke, W., Barnes, D.: Growth trajectories of corallites and ages of polyps in massive colonies of reef-building corals of the genus *porites*. *Marine Biology* **117**(2), 321–326 (1993)
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
9. Drake, J.L., Varsano, N., Mass, T.: Genetic basis of stony coral biomineralization: History, trends and future prospects. *Journal of structural biology* **213**(4), 107782 (2021)

10. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 574–584 (2022)
11. Hendy, E., Gagan, M., Lough, J.: Chronological control of coral records using luminescent lines and evidence for non-stationary enso teleconnections in northeast australia. *The Holocene* **13**(2), 187–199 (2003)
12. Jonsson, T.: Micro-ct and deep learning: Modern techniques and applications in insect morphology and neuroscience. *Frontiers in Insect Science* **3**, 3 (2023)
13. Kühn, H.S., Burghardt, T.: Animal biometrics: quantifying and detecting phenotypic appearance. *Trends in Ecology & Evolution* **28**(7), 432–441 (2013)
14. Lough, J., Barnes, D.: Several centuries of variation in skeletal extension, density and calcification in massive porites colonies from the great barrier reef: a proxy for seawater temperature and a background of variability against which to identify unnatural change. *Journal of Experimental Marine Biology and Ecology* **211**(1), 29–67 (1997)
15. Medellín-Maldonado, F., López-Pérez, A., Ruiz-Huerta, L., Carricart-Ganivet, J.P.: Understanding corallite demography to comprehend potential bias in sclerochronology: Analysis of coral modular growth by micro-computed tomography. *Limnology and Oceanography* (2022)
16. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV). pp. 565–571 (2016). <https://doi.org/10.1109/3DV.2016.79>
17. Sadler, J., Webb, G.E., Leonard, N.D., Nothdurft, L.D., Clark, T.R.: Reef core insights into mid-holocene water temperatures of the southern great barrier reef. *Paleoceanography* **31**(10), 1395–1408 (2016). <https://doi.org/https://doi.org/10.1002/2016PA002943>, <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016PA002943>
18. Shi, P., Guo, X., Yang, Y., Ye, C., Ma, T.: Nextou: Efficient topology-aware u-net for medical image segmentation. arXiv preprint arXiv:2305.15911 (2023)
19. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P., Bauer, U., Menze, B.H.: cldice-a novel topology-preserving loss function for tubular structure segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16560–16569 (2021)
20. Würfl, T., Ghesu, F.C., Christlein, V., Maier, A.: Deep learning computed tomography. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part III 19. pp. 432–440. Springer (2016)
21. Yudelman, E.A., Slowey, N.C.: Coral extension rate analysis using computed axial tomography. *Coral Reefs* **41**(4), 973–985 (2022)
22. Zeng, C., Yang, X., Mirmehdi, M., Gambaruto, A.M., Burghardt, T.: Video-transunet: temporally blended vision transformer for ct vfss instance segmentation. In: Fifteenth International Conference on Machine Vision (ICMV 2022). vol. 12701, pp. 98–105. SPIE (2023)