



Segmentation of Individual Foraminifera in X-ray Microtomography Volumes

Alexandros Sopasakis · Najmeh Abiri · Constance Choquel · Jonas Donnenfield · Behnaz Pirzamanbin · Helena L. Filipsson

Lund University · Malmö University · Université d'Angers / Nantes Université / Le Mans Université / CNRS · University of Washington


CVNH Workshop · ECCV 2026 · Malmö

FUNDING & INFRASTRUCTURE

Crafoord · Carl Tryggers · ASF · Lund University MicroMorph · InfraVis · NAISS / VR 2022-06725 · VR-TVAR 107231 · eSSENCE 156938 · FORMAS 153657

1 Why separation matters

Synchrotron micro-CT reveals the internal 3D structure of foraminiferal tests. To save beamtime, several specimens are scanned together and often touch.




Touching specimens become one connected foreground object

The practical problem

- A merged pair gives the wrong specimen count.
- Morphometrics require one isolated test at a time.

Morphology after correct isolation



Goal: recover each specimen as a distinct 3D instance while preserving its chambers, walls and pores.

How performance is measured

$$\text{Dice} = 2 |P \cap G| / (|P| + |G|)$$

Foreground overlap: are the correct voxels selected?

$$\text{F1} = 2PR / (P + R)$$

Instance detection: are specimens counted separately?

A prediction is matched to an expert instance when $\text{IoU} \geq 0.5$. Dice can remain high even when two touching specimens are merged; instance F1 exposes that topological error.

2 Training-free 3D watershed

The proposed pipeline operates directly on the reconstructed volume.


- Otsu thresholding**
separates bright calcite from background
- EDT-seeded watershed**
grows candidate specimens from interior peaks
- Saddle-height merge**
rejoins false fragments of one shell
- Solidity-based split**
separates non-convex touching pairs
- Graph-cut refinement**
places shared boundaries using image evidence

Solidity requires a 3D convex hull

$$S(R) = V(R) / V(\text{CH}_3D(R))$$

Compute the 3D convex hull $\text{CH}_3D(R)$ of the candidate region R , then divide the occupied voxel volume by the hull volume.

For every proposed child R_k , compute a new 3D convex hull and a new solidity.



$$\Delta S = \min_k S(R_k) - S(R_{\text{parent}})$$

Accept the coarsest split only if $\Delta S > \delta_s = 0.05$

The least-solid child must improve, so every resulting specimen is more plausible than the fused parent.

Detailed held-out results

Metric		3D watershed
Precision	0.830	0.956
Recall	0.912	0.956
Matched-instance IoU	0.939	0.948
Detected / missed / spurious	83 / 8 / 17	87 / 4 / 4
Count bias	+0.30	0.00

3 Same footprint, better topology

Fine-tuned SAM and 3D watershed recover almost identical foreground masks.

Foreground Dice

$$0.981 = 0.981$$

The difference appears when identities and counts matter.

Instance detection

Instance F1

$$0.869 \rightarrow 0.956$$

SAM 3D watershed

Counting accuracy

Slices with exact count

$$67\% \rightarrow 93\%$$

SAM 3D watershed

Held-out evaluation

30 slices · 91 annotated instances


Volume 005_AN1950_A never used for training, validation or parameter selection

Error reduction

Merge / split errors: 4 / 6 $\rightarrow$ 2 / 2

Count MAE: 0.50 $\rightarrow$ 0.07 specimens per slice



Observed solidity gain



Example: $\min(0.87, 0.86) - 0.66 = 0.20 > 0.05$, so the fused pair is split. Forced splits of genuine single tests give $\Delta S < 0$.

4 Where slice-wise prediction breaks

The key challenge is assigning touching shells to the correct instances.



Take-home message

For compact, high-contrast specimens and scarce annotations, explicit 3D geometry preserved instance topology better than slice-wise learning.

Scope and interpretation

Best suited to compact, approximately convex, high-contrast tests.

The comparison uses one deliberately held-out volume and slice-level annotations.

Ambiguous 3D contacts are flagged for expert review, but no result was manually corrected.