


Segmentation of Individual Foraminifera in X-ray Microtomography Volumes

Alexandros Sopasakis¹ · Najmeh Abiri² · Constance Choquel^{3/4} · Jonas Donnenfield^{3/5} · Behnaz Pirzamanbin⁶ · Helena L. Filipsson³

¹ Lund University, Mathematics ² Malmö University, Computer Science & Media Tech. ³ Lund University, Earth & Environmental Sciences

⁴ Université d'Angers, LPG UMR 6112 ⁵ University of Washington, Oceanography ⁶ Lund University, Statistics



Five contacting tests, automatically separated

MOTIVATION

Morphometrics needs one specimen at a time

1

Foraminifera are marine protists whose calcite tests archive past ocean conditions — key proxies for palaeoclimate reconstruction.




Elphidium clavatum: surface rendering, pore-bearing ultimate chamber, pores as cylindrical canals, isolated pore network — each requires the specimen to be isolated first.

MOTIVATION

Morphometrics needs one specimen at a time

- 1 Foraminifera are marine protists whose calcite tests archive past ocean conditions — key proxies for palaeoclimate reconstruction.
- 2 Every descriptor of interest — test volume, wall thickness, chamber arrangement, pore networks — assumes a correctly isolated individual.




Elphidium clavatum: surface rendering, pore-bearing ultimate chamber, pores as cylindrical canals, isolated pore network — each requires the specimen to be isolated first.

MOTIVATION

Morphometrics needs one specimen at a time

- 1 Foraminifera are marine protists whose calcite tests archive past ocean conditions — key proxies for palaeoclimate reconstruction.
- 2 Every descriptor of interest — test volume, wall thickness, chamber arrangement, pore networks — assumes a correctly isolated individual.
- 3 SR- μ CT at the ANATOMIX beamline (SOLEIL) images them non-destructively at $0.325\text{ }\mu\text{m}$ isotropic voxels.



Elphidium clavatum: surface rendering, pore-bearing ultimate chamber, pores as cylindrical canals, isolated pore network — each requires the specimen to be isolated first.

THE PROBLEM

Beamtime is scarce, so the shells touch


To use beamtime efficiently, specimens are mounted and scanned together — so tests end up in mutual contact.

THE PROBLEM

Beamtime is scarce, so the shells touch

To use beamtime efficiently, specimens are mounted and scanned together — so tests end up in mutual contact.

Thresholding plus connected components then fuses them into a single bright object and cannot recover the count.



Maximum-intensity projection: one bright, mutually contacting mass


THE PROBLEM

Beamtime is scarce, so the shells touch

To use beamtime efficiently, specimens are mounted and scanned together — so tests end up in mutual contact.

Thresholding plus connected components then fuses them into a single bright object and cannot recover the count.

A merge or a split does not just cost pixels: it propagates into specimen counts and into every downstream measurement.



Maximum-intensity projection: one bright, mutually contacting mass

THE PROBLEM

Beamtime is scarce, so the shells touch

To use beamtime efficiently, specimens are mounted and scanned together — so tests end up in mutual contact.

Thresholding plus connected components then fuses them into a single bright object and cannot recover the count.

A merge or a split does not just cost pixels: it propagates into specimen counts and into every downstream measurement.

Held-out volume: **4 specimens, 3 in mutual contact**

30 evaluation slices · 91 annotated instances



Maximum-intensity projection: one bright, mutually contacting mass

Three routes to instance separation

1

2

3

All three are compared as complete end-to-end pipelines, against the same expert annotations from one deliberately held-out volume.

Three routes to instance separation

1

Cellpose

fine-tuned

Generalist flow-field model, fine-tuned on the same 100 expert-annotated slices. At this annotation budget it produced no stable instance masks.

→ qualitative baseline only

2

3

All three are compared as complete end-to-end pipelines, against the same expert annotations from one deliberately held-out volume.

Three routes to instance separation

1

Cellpose

fine-tuned

Generalist flow-field model, fine-tuned on the same 100 expert-annotated slices. At this annotation budget it produced no stable instance masks.

→ qualitative baseline only

2

Fine-tuned SAM

decoder only

Image and prompt encoders frozen; only the mask decoder is trained, for prompt-free inference. Predictions are made slice by slice, then linked into 3D instances across a ± 20 -slice window.

→ accurate footprints, fragile topology

3

All three are compared as complete end-to-end pipelines, against the same expert annotations from one deliberately held-out volume.

Three routes to instance separation

1

Cellpose

fine-tuned

Generalist flow-field model, fine-tuned on the same 100 expert-annotated slices. At this annotation budget it produced no stable instance masks.

→ qualitative baseline only

2

Fine-tuned SAM

decoder only

Image and prompt encoders frozen; only the mask decoder is trained, for prompt-free inference. Predictions are made slice by slice, then linked into 3D instances across a ± 20 -slice window.

→ accurate footprints, fragile topology

3

3D watershed

training-free

Otsu foreground → EDT-seeded 3D watershed → saddle-height merge → solidity-based split → graph-cut boundary refinement. No annotations, no learned parameters, no GPU.

→ our proposal

All three are compared as complete end-to-end pipelines, against the same expert annotations from one deliberately held-out volume.

KEY IDEA

The geometric prior that does the work

One test is compact and nearly convex; two in contact form a dumbbell.


So split a watershed region only when the minimum solidity of the split pieces exceeds that of the parent by $\delta s = 0.05$.

KEY IDEA

The geometric prior that does the work

One test is compact and nearly convex; two in contact form a dumbbell.

So split a watershed region only when the minimum solidity of the split pieces exceeds that of the parent by $\delta s = 0.05$.




The fused pair is detected as re-entrant and cut, recovering the correct count of five.

KEY IDEA


The geometric prior that does the work

One test is compact and nearly convex; two in contact form a dumbbell.

So split a watershed region only when the minimum solidity of the split pieces exceeds that of the parent by $\delta s = 0.05$.



The fused pair is detected as re-entrant and cut, recovering the correct count of five.



All 13 candidate labels in one volume.

Fused pair $\Delta = +0.20$ · every genuine single test $\Delta < -0.10$ · a ~ 0.3 margin with no per-sample tuning

Same footprint accuracy, better topology

Instance F_1 (IoU ≥ 0.5)

0.869 → 0.956

fine-tuned SAM → 3D watershed

Slices with exact count

67% → 93%

fine-tuned SAM → 3D watershed

Same footprint accuracy, better topology

Instance F_1 ($\text{IoU} \geq 0.5$)

0.869 → 0.956

fine-tuned SAM → 3D watershed

Slices with exact count

67% → 93%

fine-tuned SAM → 3D watershed

Metric	Fine-tuned SAM	3D watershed
Dice (footprint)	0.981	0.981
IoU, semantic	0.963	0.963
IoU, matched instances	0.939	0.948
Precision / Recall	0.83 / 0.91	0.96 / 0.96
Merges / splits	4 / 6	2 / 2
Count MAE (per slice)	0.50	0.07


n = 30 slices, 91 annotated instances.

Cellpose excluded — no usable instance accuracy at this annotation budget.

QUALITATIVE

Where slice-wise prediction breaks

red = prediction **cyan** = expert annotation




Fine-tuned SAM — merge

The two upper tests are combined into a single predicted instance: the mask runs straight across the shared boundary.

Where slice-wise prediction breaks

red = prediction cyan = expert annotation




Fine-tuned SAM — merge

The two upper tests are combined into a single predicted instance: the mask runs straight across the shared boundary.

QUALITATIVE


Where slice-wise prediction breaks

red = prediction cyan = expert annotation



Fine-tuned SAM — merge

The two upper tests are combined into a single predicted instance: the mask runs straight across the shared boundary.



3D watershed — correct

All four specimens are recovered, including the contacting pair, with the cut following the true chamber walls.

Preserve topology, not just overlap

- In high-contrast SR- μ CT, foreground recovery is effectively solved (Dice ≈ 0.98 for both methods). The operative challenge is assigning contacting structures to the right specimen.
- When the morphology is compact and annotations are scarce, an explicit 3D geometric prior outperforms slice-wise learned segmentation — and needs no labels and no GPU.
- Scope: one held-out volume, one taxon, slice-level metrics; constants fixed on separate development data. Next: volumetric object-level validation, further taxa, and hybrid geometry + learned boundary refinement.