Centre for Computational Personalised
Medicine International
Research Foundation
www.sano.science

MultiGeoMed: geometryczne uczenie głębokie w obrazowaniu medycznym, od reprezentacji kształtu po naprawę topologii naczyń

Gniewosz Drwięga, PhD
Sano science – Centre for Computational Personalised
Medicine International Research Foundation, Kraków
Postdoctoral Researcher at
https://mwodzinski.edu.pl/multigeomed/team/

Our team

MultiGeoMed

sano



Marek Wodziński, PhD
Principal Investigator &
Deep Learning Expert



Wojciech Szymański
PhD Candidate &
Deep Learning Expert



Gniewosz Drwięga, PhD
Postdoctoral Researcher &
Deep Learning Expert

https://mwodzinski.edu.pl/multigeomed/

NATIONAL SCIENCE CENTRE
POLAND

MultiGeoMed: geometry beyond the voxel grid

MULTIGEOMED · THE PROJECT

- Medical-imaging AI mostly works on **voxel grids**, the format the scanner produces.
- Halve the voxel size and you get **8× more voxels**: memory forces cropping and downsampling.
- **Surface meshes, point clouds and implicit functions** describe anatomy differently, yet most studies use only one of them.
- MultiGeoMed asks: **which representation helps which task**, and **how to combine them**?

Geometric deep learning: neural networks that learn directly from meshes, graphs, point clouds or continuous functions, not only from voxel grids.

MultiGeoMed

Multimodal Geometric Deep Learning in Medical Image Analysis

Funding: NCN SONATA 20

Grant: No. 2024/55/D/ST6/02081

Duration: Oct 2025 – Sep 2028

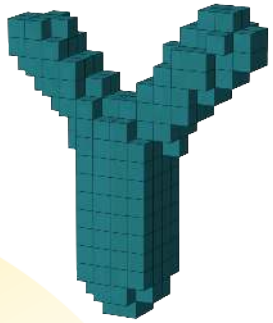
Host · PI: Sano, Kraków · Marek Wodziński

Target tasks: shape analysis & completion, image registration, classification

NATIONAL SCIENCE CENTRE POLAND

One vessel, four representations

MULTIGEOMED · SHAPE REPRESENTATIONS



Voxel grid

+ native scanner format; CNN-friendly (nnU-Net)

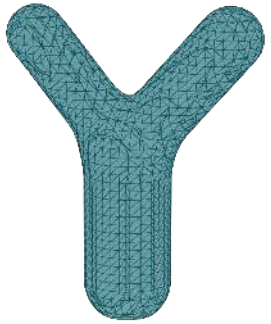
– memory grows as n^3; no explicit surface or topology



Point cloud

+ light and flexible; copes with incomplete data

– no connectivity, no ordering



Surface mesh

+ explicit surface and connectivity; topology control

– hard to change topology; usually template-based



Implicit function

+ continuous: query f(x) at any point and resolution

– usually fitted per shape; slower to query

No single winner → MultiGeoMed studies how to combine them

Three experiments today

MULTIGEOMED · ROADMAP OF THIS TALK

1 ONGOING

Shape reconstruction from incomplete data

Wojciech Szymański · PhD project



Representation: Voxels + point clouds
Task: Shape completion
Why: Completing damaged or missing anatomy (e.g. skull defects)

2 ONGOING

INR for whole-slide histology registration

Gniewosz Drwięga · started Oct 2026



Representation: Implicit neural (INR)
Task: Image registration
Why: Combining H&E and immunostains from serial sections

3 FINISHED

Mind the Gap: mesh-guided repair of broken vessels

Drwięga, Szymański, Wodziński



Representation: Voxel mask + mesh
Task: Shape analysis → topology repair
Why: Centrelines, branch labels and flow simulation need connected vessels

Shared tooling: our in-house multigeomed PyTorch library loads, converts and saves voxel grids, point clouds and meshes.

EXPERIMENT 1 OF 3

Shape reconstruction from incomplete data

Wojciech Szymański · PhD project (AGH University of Krakow & Sano)

ONGOING

Multimodal Geometric Deep Learning

Research focus

Developing geometric deep learning methods to reconstruct medical structures from incomplete data.

Studying how representation choice affects reconstruction and generalisation.

Multimodal architecture

Representation-specific encoders learn complementary geometric features.

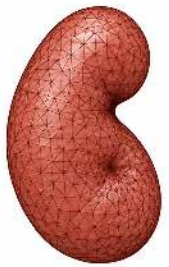
Feature fusion combines these features for anatomical reconstruction.



Voxel grids



Point clouds

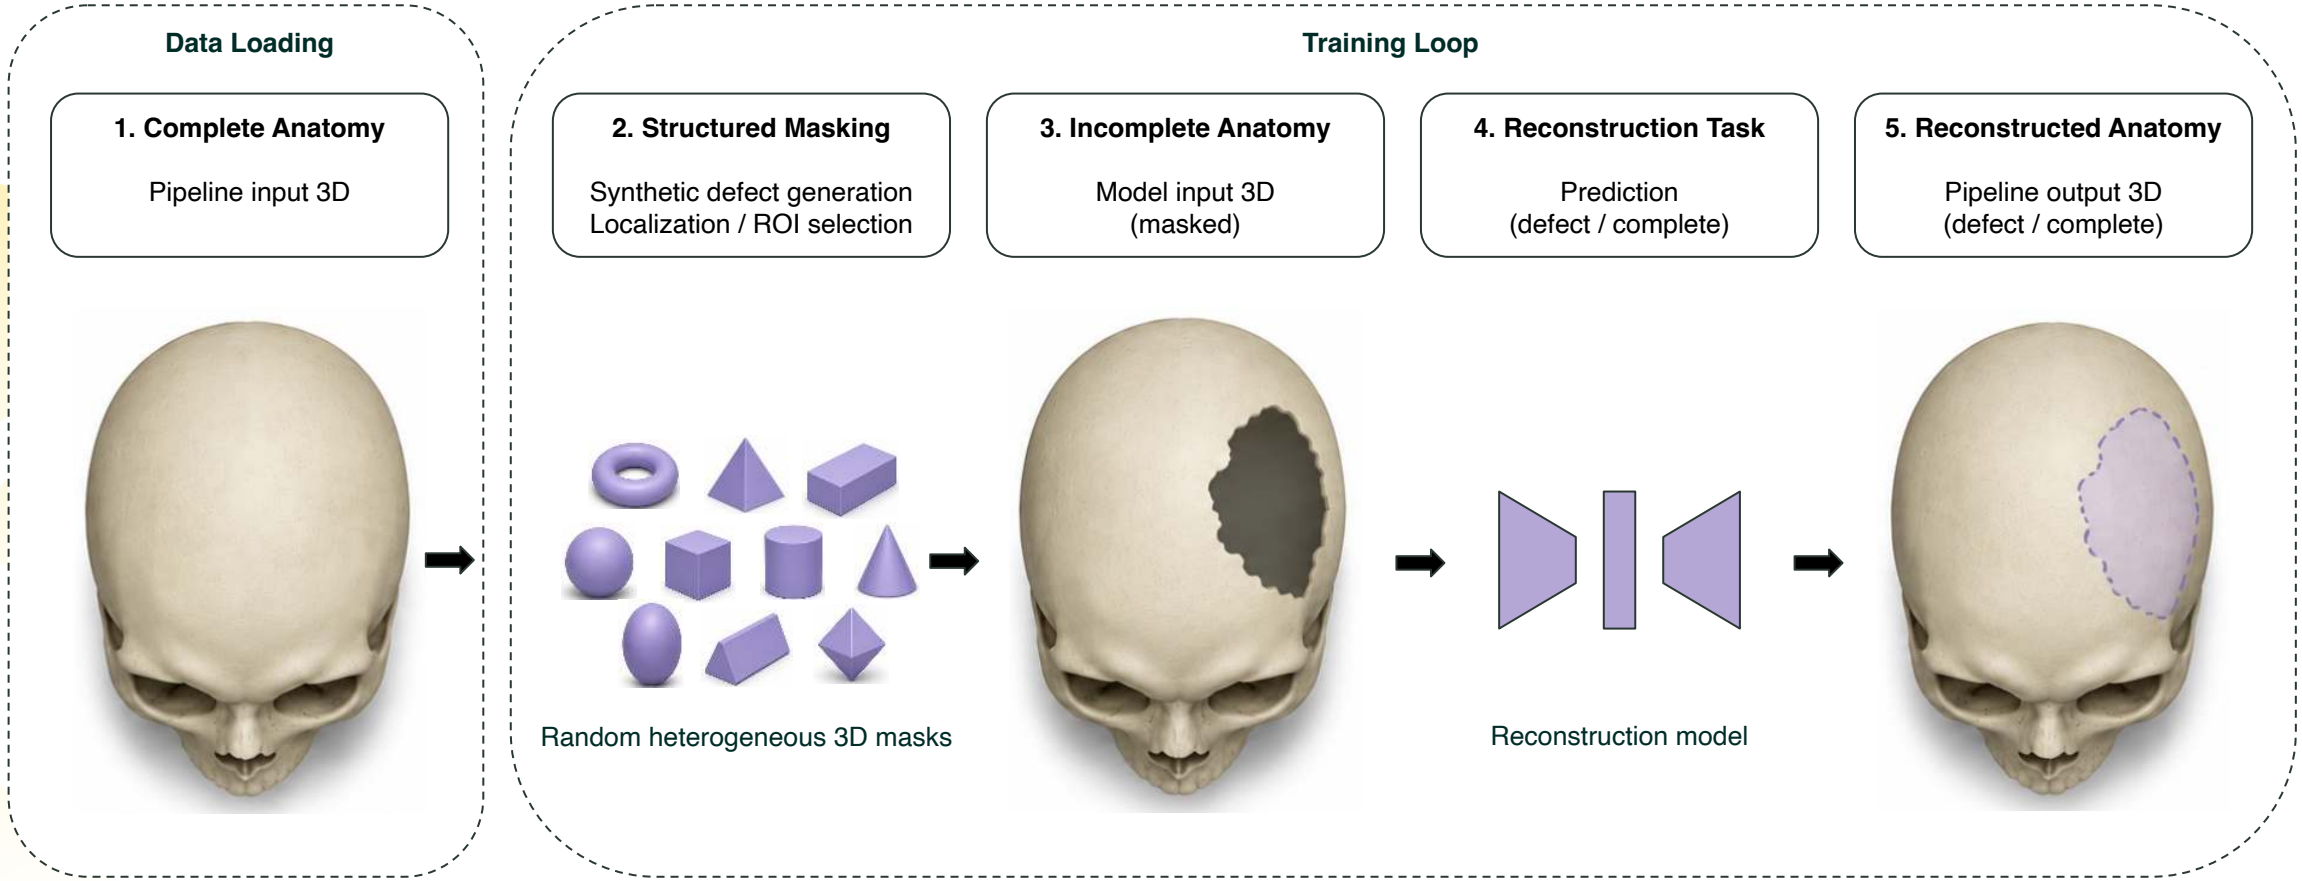


Surface meshes



Implicit functions

Robustness with Self-Supervised Learning



Medical Shape Reconstruction

Reconstructing complete anatomy or the missing fragment from incomplete geometry

Voxel grids

nnU-Net

Volumetric reconstruction with a custom trainer

Synthetic structural defects generated on-the-fly during training

Point clouds

Direct PTv3	Single-pass coordinate prediction
Diffusion	Iterative denoising in point space
Flow matching	Generation via a learned velocity field
Flat transformers	Other general-purpose transformer backbones
Foundation models	Pretrained completion models evaluated without fine-tuning

Evaluation separates surface fidelity from completion of the missing region

EXPERIMENT 2 OF 3

Implicit neural representations for histology registration

Gniewosz Drwięga · started October 2026 · plan and first baseline

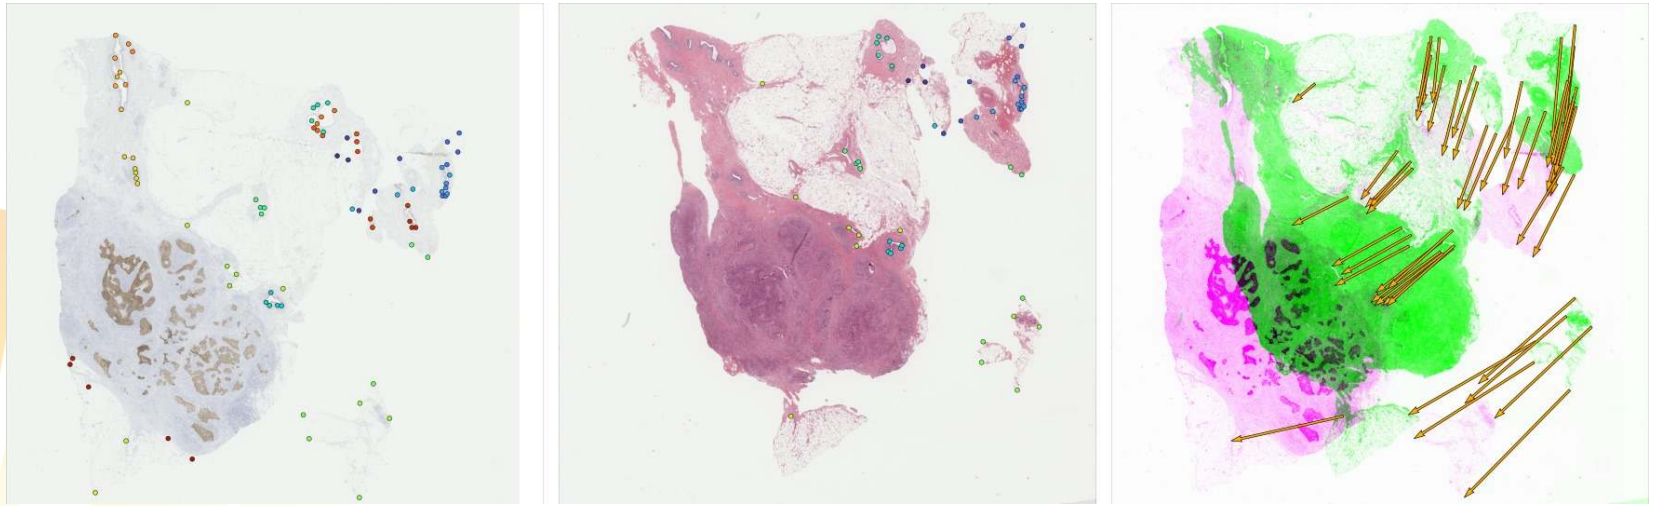
Same question, now for a deformation: a vector per pixel, or one continuous function?

ONGOING

Serial sections do not line up: registration on ANHIR

EXPERIMENT 2 · INR HISTOLOGY REGISTRATION · ONGOING

HER2 (source) | H&E (target) | Overlay, unregistered



ANHIR breast_2. Magenta = HER2, green = H&E, dark = both. Arrows: 50 landmark offsets, median rTRE 0.117.

230 training pairs, landmarks public on both slides

8 · 18 tissue types · stains, results reported per tissue

rTRE landmark error ÷ image diagonal (lower is better)

- Each section is cut, mounted and stained on its own: rotated, shifted, stretched, sometimes torn.
- Registration maps one slide onto the other, so the same structures overlap (e.g. routine H&E with an immunostain such as HER2).

Baseline: DeeperHistReg, running in our pipeline

EXPERIMENT 2 · INR HISTOLOGY REGISTRATION · FIRST RESULT (PRELIMINARY)

Preprocess
Grayscale, contrast boost, resize to 4,096 px

→

Initial alignment
Matched keypoints → rotation + scale + shift: 4 numbers

→

Non-rigid
A shift vector for every pixel (~44 M numbers)

INR replaces this

→

Output
One displacement field warps images and landmarks

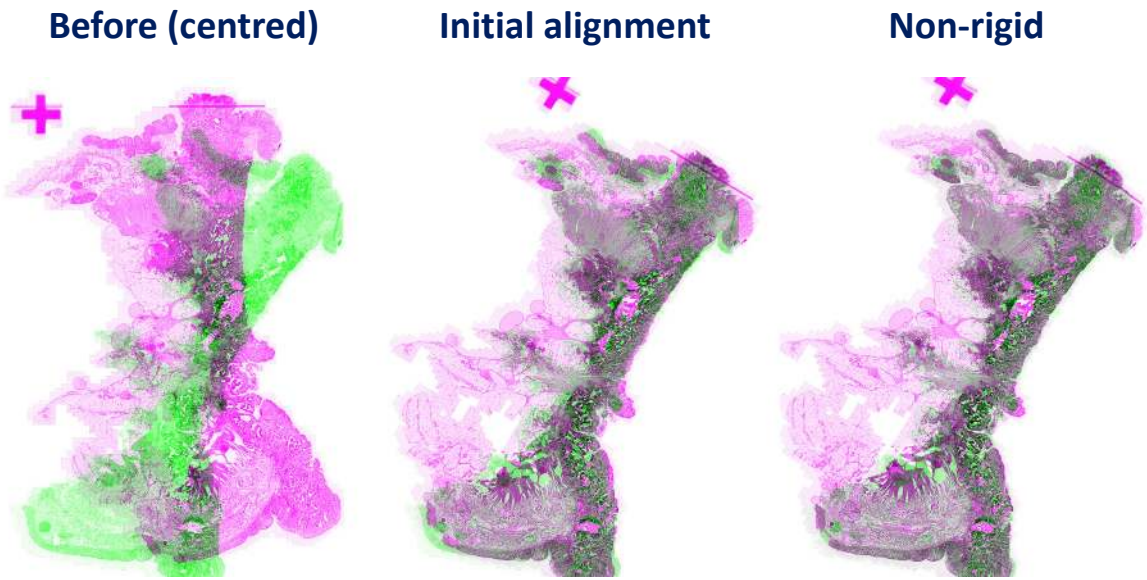
Before registration
0.082

After initial alignment
0.010

After non-rigid
0.0049

- 99 % of landmarks end up closer
- warp folds tissue onto itself: 0.2-0.8 % of pixels
- ~2 min per pair

Median rTRE, 5 colon-cancer (COAD_01) training pairs: 8.2 % → 0.5 % of the image diagonal. Preliminary: one tissue set.

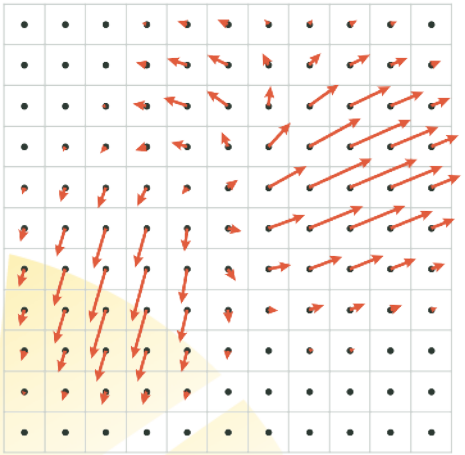


Pair 4: magenta = source, green = target, dark = overlap; magenta left after registration = tissue absent from the target.

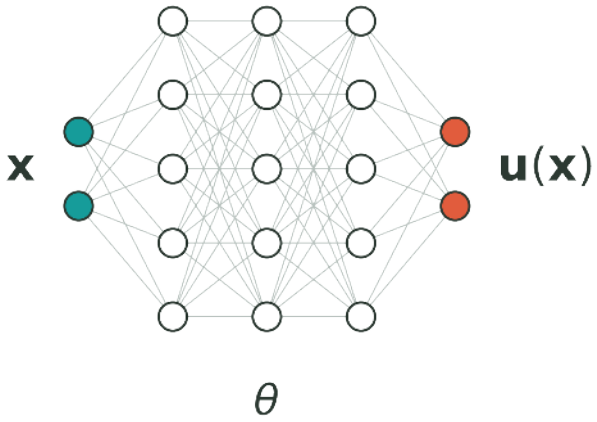
Wodziński et al., DeeperHistReg (2024): open framework whose methods (RegWSI) won ACROBAT 2023; default configuration.

Idea and plan: replace the dense field with an INR

EXPERIMENT 2 · INR HISTOLOGY REGISTRATION · THE PLAN



Today: a vector stored at every pixel



INR: u(x) = MLP(x; θ), a small network

- **Continuous:** evaluate u at any point (e.g. landmarks)
- **Compact:** ~10^5–10^6 weights for a plain MLP, independent of slide size
- **Scales by sampling:** fit on random coordinate batches
- **Smoother?** may fold tissue less: a hypothesis to test

Gap: for whole slides, the only INR method we found is NEMESIS (ACROBAT 2022: 5th of 8, fitted on ~10× downsampled images; Weitz et al., MedIA 2024). We found no INR results on ANHIR or HyReCo.

NOW — Baseline on all 230 ANHIR pairs → NEXT — Per-pair INR: accuracy, folding, resolution, per tissue → LATER — One network across many pairs (optional)

EXPERIMENT 3 OF 3

Mind the Gap: mesh-guided repair of broken vessels

G. Drwięga, W. Szymański, M. Wodziński · ShapeMI 2026 @ MICCAI, Strasbourg

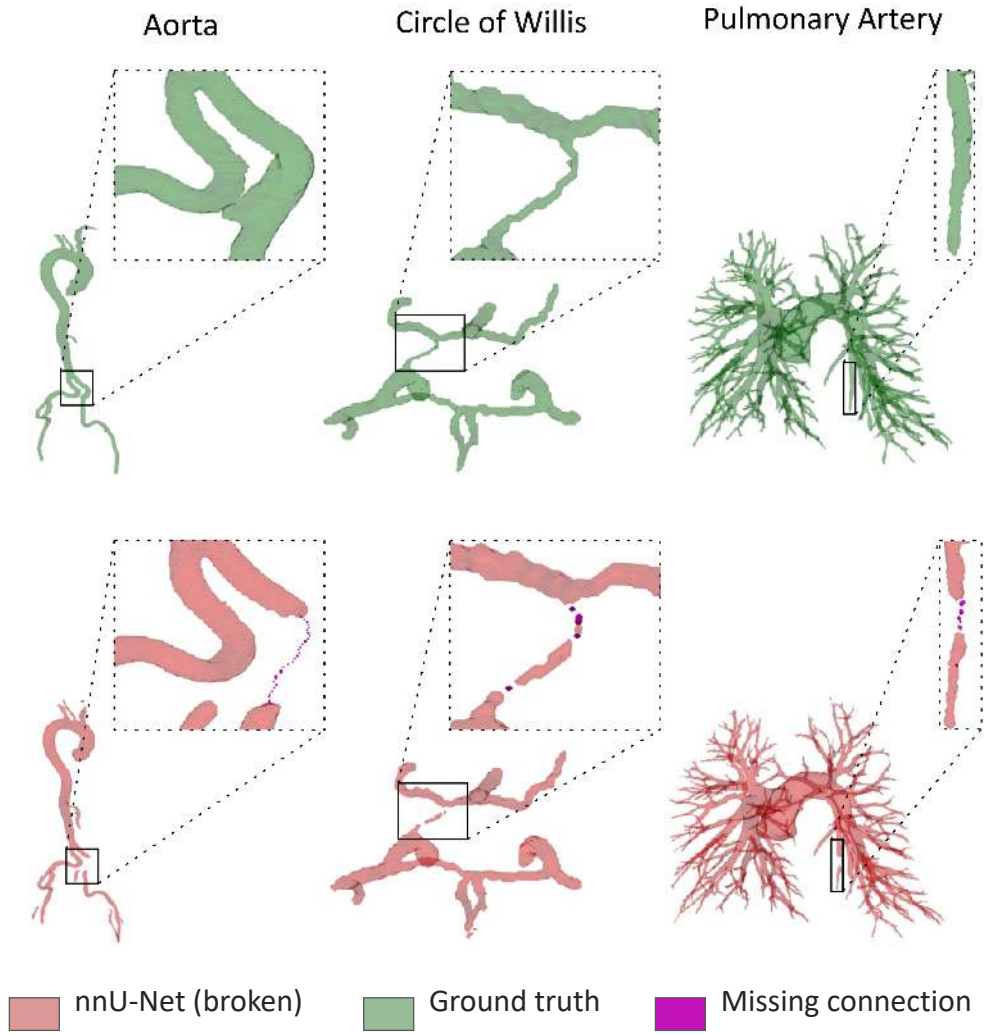
FINISHED · BEST POSTER AWARD

High Dice, broken vessels

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

- Vessel segmentation is trained and scored voxel by voxel.
- One missing bridge of a few voxels splits a vessel in two, and Dice barely moves.
- Everything downstream needs a connected vessel: centrelines, branch labels, diameters, flow simulation.
- Goal: repair the connectivity of an existing mask and keep the rest of the prediction.

Dice stays high, and the vessel is still broken.



Why Dice misses it: measuring connectivity

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

Reference: one vessel



Prediction: one tiny gap



	Reference	Prediction
Dice	1.00	≈ 0.99
β0 (pieces)	1	**2**
ccDice	1.00	**0.67**

Illustrative example.

Dice
voxel overlap: 2|P∩G| / (|P| + |G|)

Betti-0 (β0)
number of connected pieces; ideal = the reference count (usually 1)

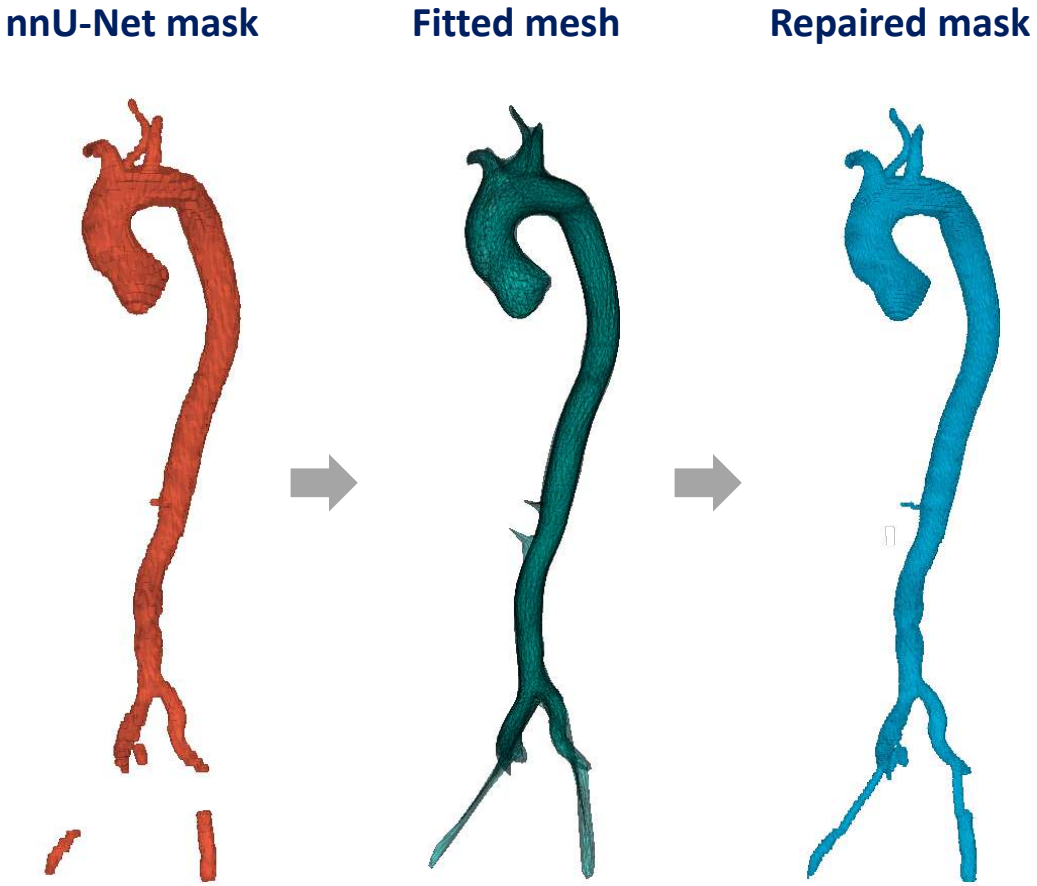
ccDice
Dice computed over connected components instead of voxels

False-branch fraction
share of the predicted centreline outside the reference; guards against spurious bridges

Idea: use a mesh as a geometric scaffold

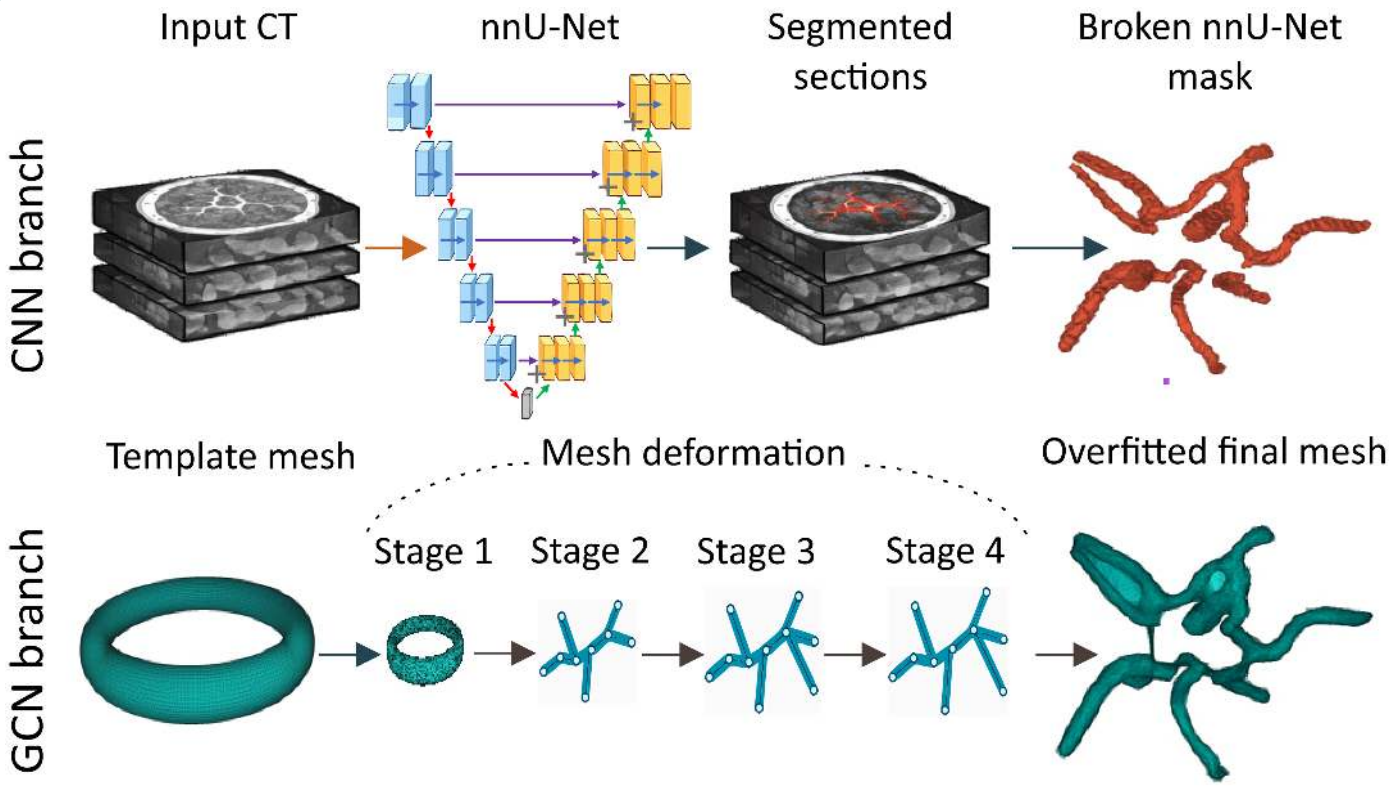
EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

- Post-processing for nnU-Net vessel masks: no retraining, no image intensities.
- Fit a deformable template mesh to each predicted mask.
- The mesh is never the output: it only proposes or validates thin bridges.
- The result stays a voxel mask: thin bridges added, spurious islands removed, everything else kept.



Method: fit a deformable template to every mask

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)



Paper Fig. 2 (cropped): the broken nnU-Net mask (top) is the target; a template mesh is deformed onto its surface in four stages (bottom).

- Template per anatomy: tube (aorta), ring (Circle of Willis), sphere (pulmonary arteries).
- A graph-convolutional network (GCN) deforms it in 4 coarse-to-fine stages.
- Objective: Chamfer distance (how far mesh and mask surface are apart) + smoothness terms.
- Fitted separately for every case (6000 steps); the surface keeps the scan's physical geometry and is normalised for fitting.

Mesh-guided repair: conservative bridges

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

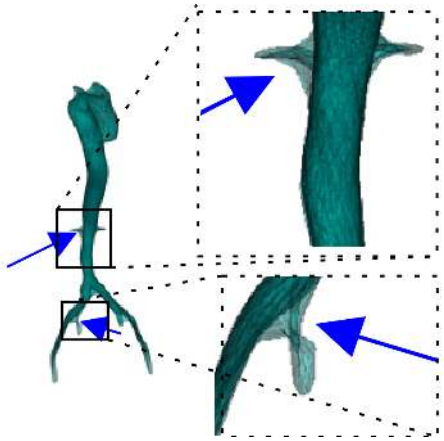
Mesh-graph path · aorta, Circle of Willis

Anchor each component to nearby mesh vertices (≤ 3 mm); bridge along the shortest path on the mesh graph.

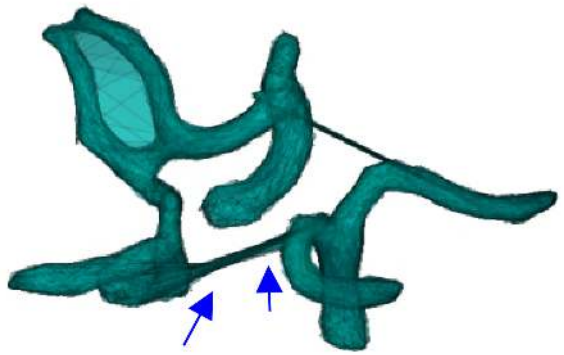
Endpoint + mesh validation · pulmonary arteries

Short bridges (≤ 12 mm) between component boundaries, kept only if ≥ 50 % supported by the mesh (≤ 3 mm).

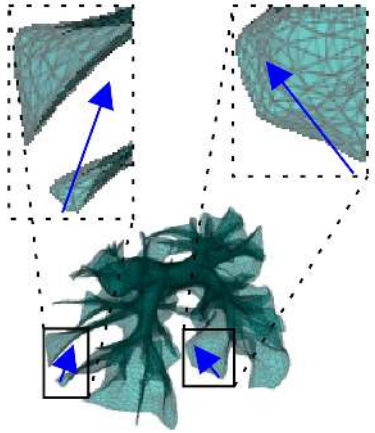
Acceptance rule: every bridge is a thin tube, kept only if it merges components and adds ≤ 3 % / 15 % / 6 % of the mask (aorta / CoW / pulmonary).



Aorta



Circle of Willis



Pulmonary arteries

Why not voxelise the mesh?

Fitted meshes have flaps, filled loops and webbing. The mesh guides; the mask decides.

Blue arrows: the repaired breaks (paper Fig. 3).

Experimental setup

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

Anatomy	Dataset	Cases
Aorta	SEGA + AortaSeg24	145
Circle of Willis	TopCoW	125
Pulmonary arteries	PARSE	100

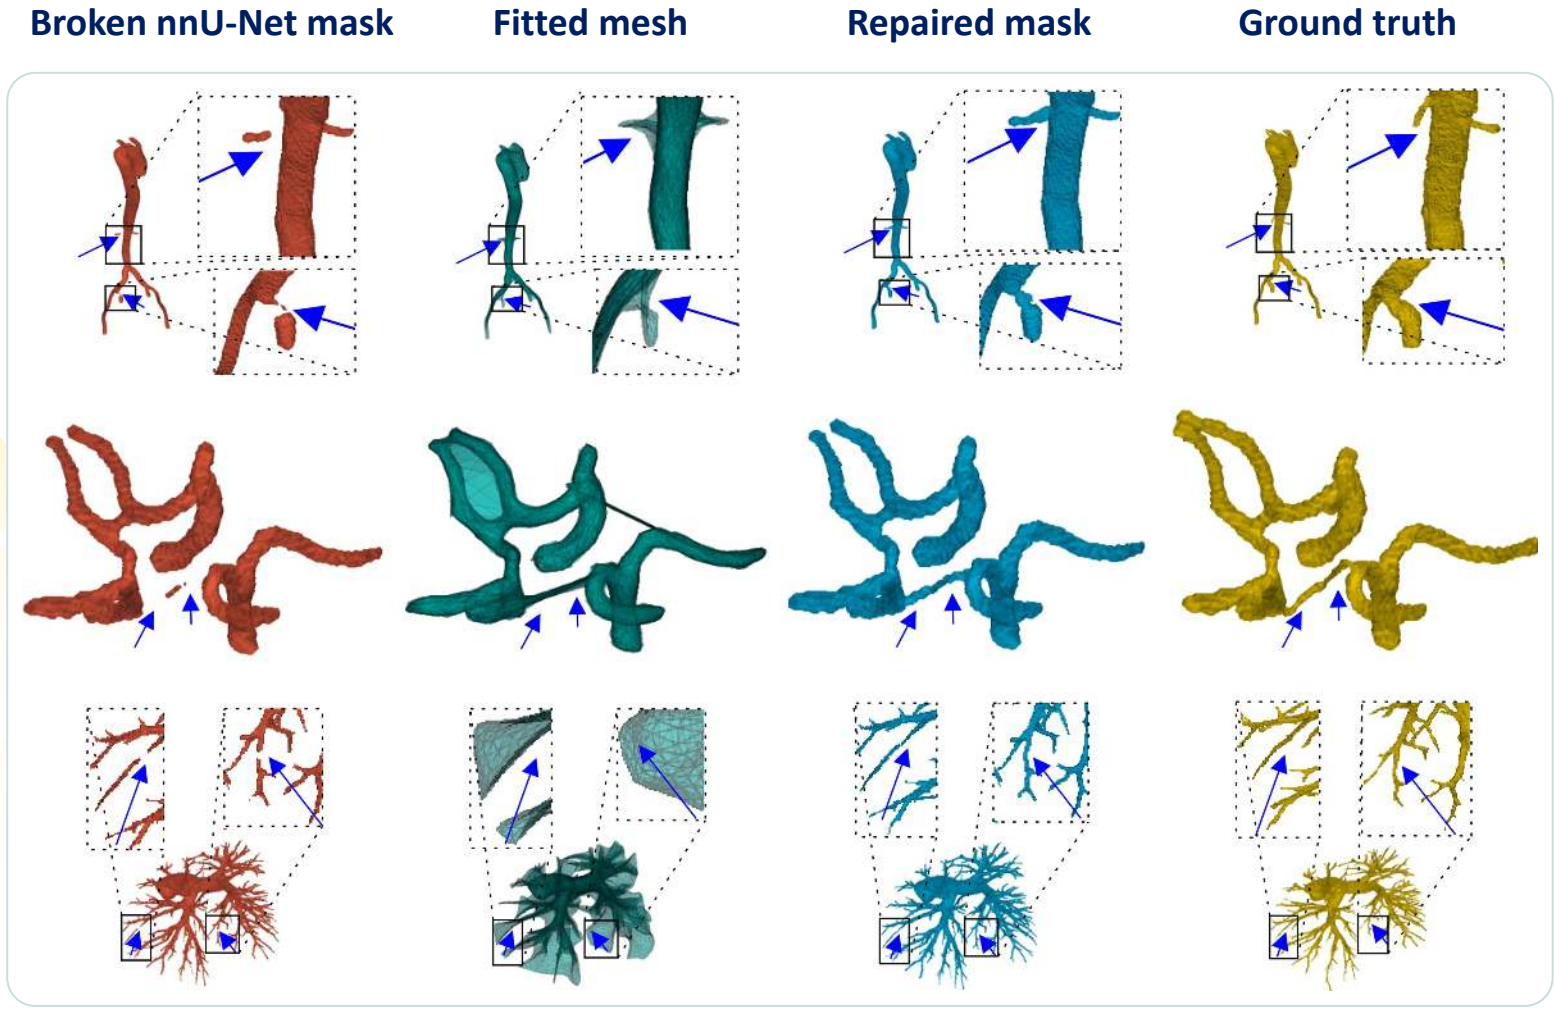
Repair sees only the binary mask

no CT intensities, no probability maps

- Baseline: nnU-Net v2 (a self-configuring U-Net), 3D full resolution; all vessel labels merged into one class.
- Five-fold cross-validation: every repaired case was segmented by a model that never saw it in training.
- Metrics: Dice, ccDice, β0 and false-branch fraction.
- Fitting and repair on one RTX 5090 workstation; nnU-Net trained on PLGrid (Helios).

Qualitative results: local bridges, minimal other changes

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)



- Blue arrows: breaks in the nnU-Net mask.
- Repair adds thin local connectors; elsewhere only spurious islands are removed.
- Here the mesh artefacts (filled loop, extra strand) do not reach the output.

Connectivity restored, Dice unchanged

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

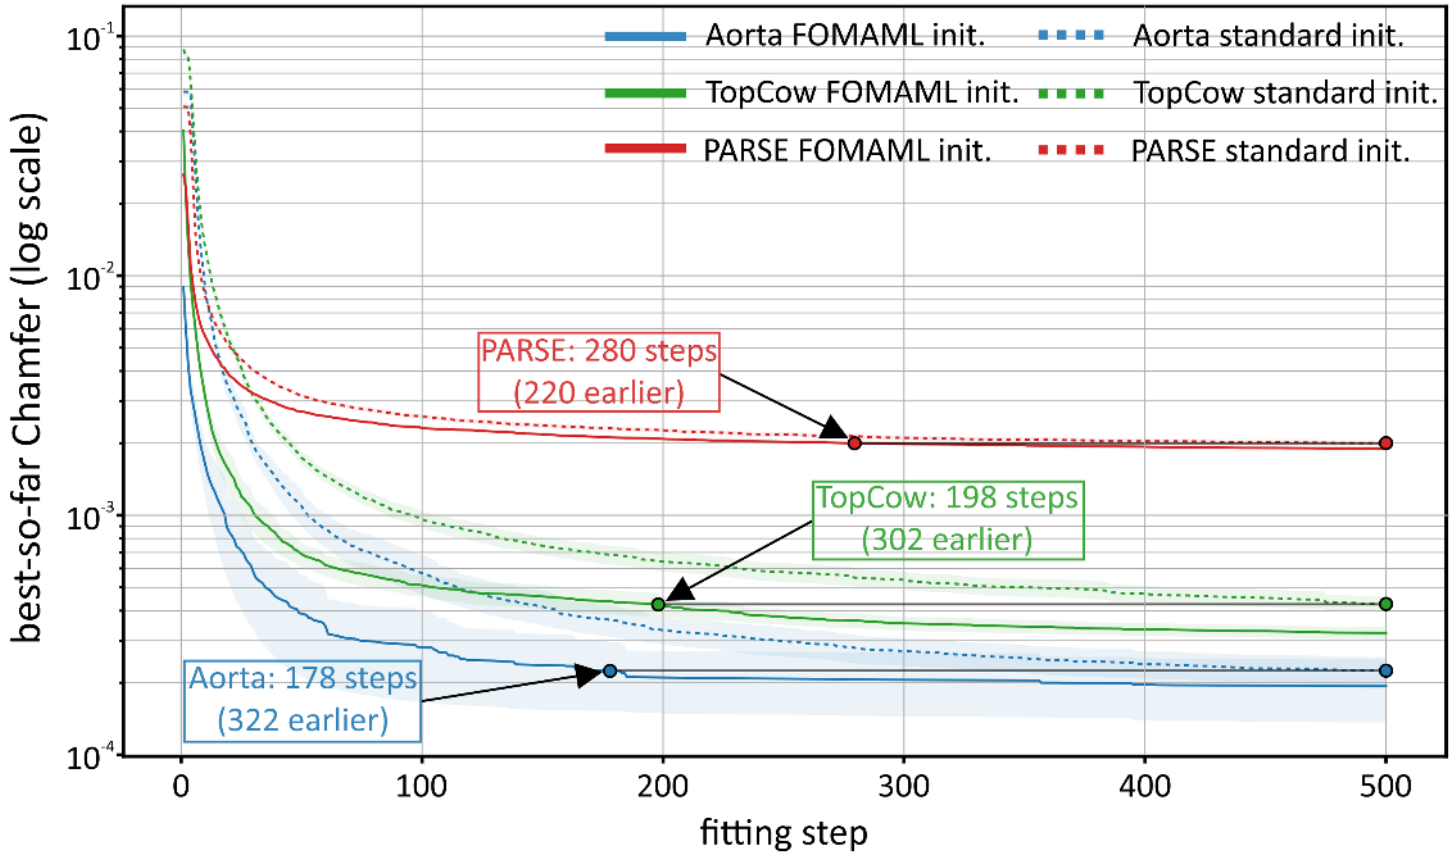
raw → repaired	Dice ↑	ccDice ↑	β0 (pieces) ↓	False-branch fraction ↓
Aorta (n = 145)	0.934 → 0.934	0.596 → **0.992**	3.35 → **1.01**	0.056 → 0.055
Circle of Willis (n = 125)	0.870 → 0.867	0.722 → **0.835**	2.58 → **1.02**	0.040 → **0.063**
Pulmonary arteries (n = 100)	0.877 → 0.876	0.028 → **0.862**	87.46 → **1.56**	0.130 → 0.132

Means over all cases (95 % confidence intervals in the paper); out-of-fold nnU-Net predictions.

- Dice changes by at most 0.003 on any dataset.
- Pulmonary arteries: from ~87 pieces per case to ~1.6.
- Trade-off on the Circle of Willis: more false branches (0.040 → 0.063), and 1.02 pieces vs 1.45 in the reference: some anatomically incomplete circles get bridged.

Faster fitting with a meta-learned start (FOMAML)

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)



Best-so-far Chamfer distance (log scale); solid = FOMAML, dotted = standard init.

- Meta-learning (first-order MAML) finds a shared starting point from which each new case fits quickly.
- Reaches the standard 500-step quality after 178 (aorta), 198 (CoW) and 280 (pulmonary) steps.
- The previous slide used full 6000-step fits (~20 min per aortic case). With FOMAML, ~300 steps + repair take ≈ 30 s (nnU-Net inference: 44 s).

Limitations and next steps

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)

Limitations

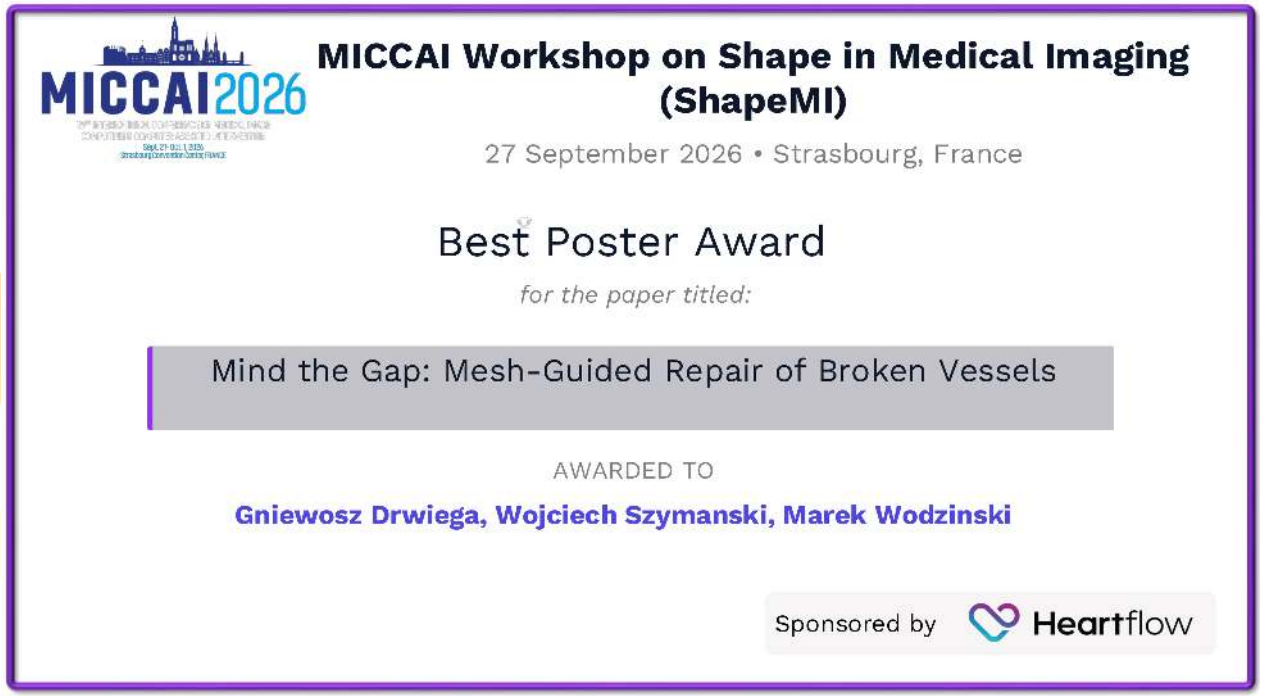
- False positives can mislead the mesh into a wrong bridge, or survive the filtering.
- Templates and thresholds are dataset-specific.
- Binary masks only: no image intensities or probabilities.

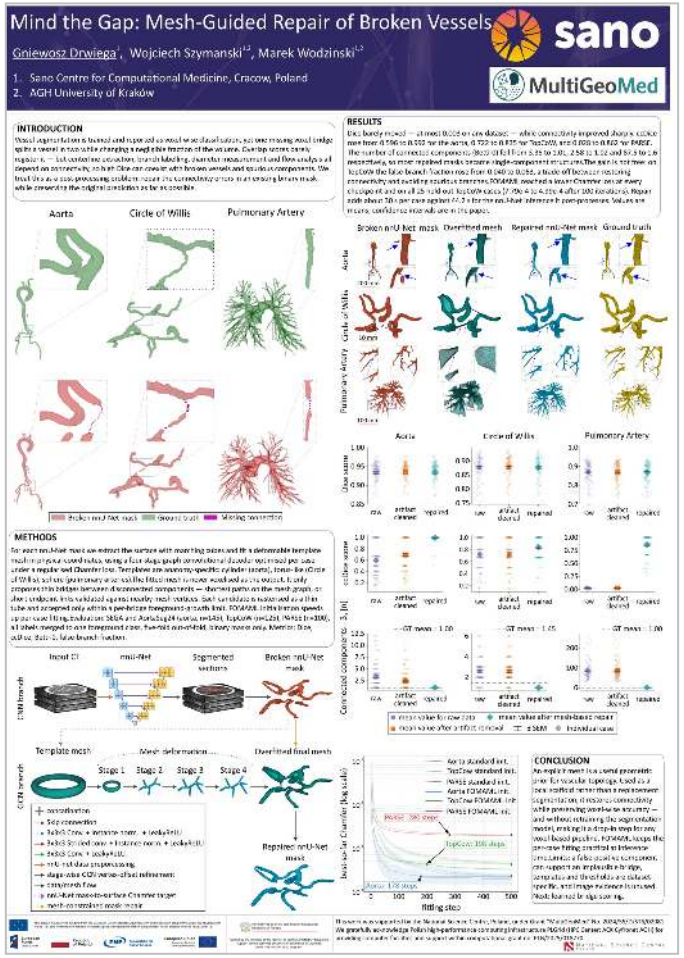
Next steps

- Learned bridge scoring using image evidence.
- Anatomical constraints against shortcuts.
- Hemodynamic simulations: do small topology fixes change the simulated flow?

Best Poster Award, ShapeMI 2026

EXPERIMENT 3 · MIND THE GAP · FINISHED (SHAPEMI 2026)





Project page

code + paper

Code: github.com/SanoScience/mind-the-gap

Project page: sanoscience.github.io/mind-the-gap-page

Paper: openreview.net/forum?id=BVB2SExfid

From shape representation to vessel topology repair

SUMMARY



Shape reconstruction

ONGOING

Voxel grids and point clouds: which representation reconstructs missing anatomy best, and how to fuse them.



INR histology registration

ONGOING

Baseline running in our pipeline (5 pairs, median rTRE 0.082 → 0.0049, preliminary); next, a continuous and compact INR deformation field.



Mind the Gap

FINISHED

A fitted mesh as a scaffold restores vessel connectivity (pieces per case → 1.0–1.6) while Dice is essentially unchanged (Δ ≤ 0.003).

Next in MultiGeoMed: fusing representations in one model · foundation models for 3D shape reconstruction · INRs in histopathology.

Thank you for your attention

sano
Centre for Computational Medicine

We use computing to find new ways to keep people healthy

www.sano.science

Acknowledgements

NATIONAL SCIENCE CENTRE POLAND

This work was supported by the National Science Centre, Poland, under Grant "MultiGeoMed" No. 2024/55/D/ST6/02081.

We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018770 and grant no. PLG/2026/020104.

MultiGeoMed mwodzinski.edu.pl/multigeomed

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857533 and from the International Research Agendas Programme of the Foundation for Polish Science No MAB PLUS/2019/13.

European Funds Smart Growth
Republic of Poland
FNP Foundation for Polish Science
European Union European Regional Development Fund

Ministry of Science and Higher Education Republic of Poland

Funded by the initiative of the Minister of Science and Higher Education "Support for the activities of Centers of Excellence established under the Horizon 2020 program"