Local convolutional segmentation and geometric error — review sheet

Use this after attempting the worked examples and practice. Explain each claim in words before calculating.

Key ideas

- A shared 3×3 layer learns ten parameters including bias.
- Loss gradients can be checked numerically before trusting a training run.
- The split unit should match the intended unit of generalization.
- Overlap, boundary distance, area and connectivity describe different failures.
- Synthetic improvement is evidence about the declared generator, not clinical validation.

Common errors and repairs

Call the small scaffold U-Net. It has one trained convolution and sigmoid; the U-Net reading is an architecture comparison.

Tune a threshold after inspecting test masks. Choose with training/validation information, freeze it, then evaluate the test protocol.

Treat a missing predicted object as zero boundary error. Mark the nonempty/empty distance unavailable and report the failure.

Accept high Dice as proof that topology is preserved. Check the feature relevant to the downstream question, such as connected components.

Retrieve without your notes

- Why can area error be zero with an incorrect boundary?
- What does a central finite-difference check compare?
- How does this assignment fit the course policy?

Readiness check

- Distinguish a fixed intensity baseline from a trained shared-convolution segmenter.
- Trace a ten-parameter model from local image patches through a loss and gradient update.
- Evaluate complete held-out images using overlap, physical boundary distances and connectivity.
- Combine geometry provenance and a sensitivity counterexample in a short Assignment 3 report.