

Synthetic gastric histopathology for anomaly-detection training

Executive summary | SGH Synthetic Histopathological Image Generation Challenge | 11 September 2026

We developed a generation and evaluation pipeline for gastric H&E tissue, producing 4096 x 2048 pixel fields at a configured 0.25 micrometres per pixel. It combines explicit tissue construction, pathology diffusion, targeted organism rendering and quantitative selection. The submission contains ten selected synthetic images, a seventy-image synthetic set containing those ten, and seventy public real reference fields.

Method

A procedural three-dimensional mucosal model controls glands, pits, lumina, epithelium, stroma and section angle. Virtual sections produce tissue maps. PixCell-1024 supplies cellular appearance using UNI2-h features from public gastric reference fields. We compared direct structure adapters, single-field conditioning and spatially matched appearance features from multiple reference fields. Overlapping windows share a latent canvas during generation. A final procedural stage places *H. pylori*-like curved rods in selected mucus and luminal locations; coordinates and counts are recorded.

The experiments separate tissue appearance from structural control. In a paired forty-layout comparison, spatial reference matching reduced median morphometric distance from 6.30 to 3.86 and increased outputs inside the reference band from 14/40 to 35/40. Adapter comparisons showed stronger retention of the intended gland arrangement. The final selection combines these complementary generation routes.

Selected images

The ten comprise three HP-with-gastritis targets, two intestinal-metaplasia targets, two co-occurrence targets and three normal controls. We ranked the combined eighty-image pool with common tissue measurements, then used gland topology and field-level variety to choose the showcase. The normal selection includes a mucosal surface and gland profiles; all three satisfy the original normal ring threshold. The wider synthetic set contains 21/14/14/21 images in the same category order.

The final ten have mean morphometric distance 3.189, minimum recorded embedding-copy margin 0.0532, and maximum recorded pixel NCC 0.8219. Within the seventy-image set, the maximum same-category SSIM across 602 pairs is 0.7053 at 1024-pixel width, below the 0.75 screen threshold. The five HP-target showcase images contain 390 procedural organisms in total.

Public reference data

The real-image package contains seventy fields from nineteen HiESD whole slides, with no more than four fields per slide. Each was extracted at 4166 x 2083 native pixels (0.2458 micrometres per pixel) and downsampled to the required raster. The public region annotations identify chronic gastritis, intestinal metaplasia and normal glands. HP organism status is unrecorded in HiESD; real-image metadata therefore states source annotations explicitly. Coordinates, overlap information, attribution and per-file checksums accompany the images.

Evaluation and development

Our evaluation records nuclear and tissue statistics, gland-lumen topology, distance from category-specific reference distributions, layout retention, embedding similarity and targeted pixel correlation. Selection also includes full-field views, native-pixel crops and inspection of recorded organism locations. These measurements make changes between generation methods and candidate sets traceable.

The intended endpoint is useful synthetic augmentation for SGH's anomaly detector. The evaluation sequence is morphology and label agreement, source-similarity analysis, then a real-data comparison of detector training with and without synthetic augmentation. The latter comparison measures the intended downstream benefit; it is separate from generation-quality scores.

Licensing and a dedicated public-data model

The current prototype uses PixCell-1024 and UNI2-h, whose non-commercial research terms create a licensing dependency for commercial development. Our proposed replacement is full generator-weight domain adaptation of the Apache-2.0 FLUX.2 Klein Base 4B model, using public, compatible histopathology data and its native text and image conditioning, with no PixCell or UNI2-h weights or embeddings in the replacement pipeline.

We will use calibrated HiESD fields for gastric architecture and metaplasia, expand with suitable public gastric whole slides, and use the 245,196 small GasHisSDB patches as an auxiliary texture stream at their native information scale. Source descriptions and annotations will be converted into structured captions by automated image review, preserving measured pixel spacing and source labels. Organism supervision will come from explicitly annotated public H&E examples, with unknown findings kept separate from negative labels.

Training begins by checking image reconstruction through the autoencoder. We then update all generator weights with flow-matching on captioned gastric fields, extend the same model to reference-plus-caption conditioning, and train spatial control from source-aligned tissue maps. The curriculum progresses from local fields to overlapping windows in a shared 4096 x 2048 latent canvas. Generation uses a category, physical sampling, tissue layout, optional reference examples and fresh noise, followed by the same morphology, diversity and source-similarity selection. Finalist compute will be benchmarked and allocated across adaptation, evaluation and generation; larger allocations extend the corpus and field-scale training.

References

HiESD: Ge et al., Scientific Data (2025), doi:10.1038/s41597-025-05679-1. Public dataset doi:10.6084/m9.figshare.28919840, CC BY 4.0.

GasHisSDB: Hu et al., doi:10.6084/m9.figshare.15066147; arXiv:2106.02473.

PixCell: arXiv:2506.05127; huggingface.co/StonyBrook-CVLab/PixCell-1024. UNI2-h: huggingface.co/MahmoodLab/UNI2-h.

Replacement model: huggingface.co/black-forest-labs/FLUX.2-klein-base-4B.