Generative AI for Designing Topology-Optimized Metamaterial Unit Cells

•••

Summer 2026

Dataset

- Generated synthetic binary unit cells
  - Six families (holes, lattice, cross, ring, bars, re-entrant) – spanned useful range of volume fraction and stiffness
  - ~50% of the cells are symmetric.
  - Cells with degenerate volume fraction (near-empty and/or near-solid) are regenerated.
  - Script wasn't able to generate cells with negative poisson's ratio, therefore trained generative AI model cannot target auxetic behavior
- Scored each unit cell on mechanical behavior
  - Stiffness is relative to the solid material (solid E = 1)
  - Homogenization process:
    - For each macroscopic unit strain (xx, yy, shear) it solves the periodic equilibrium problem on a Q4 pixel mesh and integrates the energy to get the effective stiffness tensor, from which E, ν, K, and G follow.
    - Validate by comparing against known values

Metrics collected

- volume_fraction — fraction of solid pixels.
- symmetry_class — square / biaxial / uniaxial / diagonal / none.
- symmetry_score — 0–1, best mirror match (1.0 = perfectly symmetric).
- mirror_x, mirror_y, diagonal — individual symmetry booleans.
- n_solid_components — number of disconnected solid pieces (8-connectivity; 1 ideal).
- largest_component_fraction — share of solid in the largest piece.
- percolates_x, percolates_y — solid spans edge to edge in x / y.
- periodic_connected_x, periodic_connected_y — solid joins its tiled copies (the correct connectivity test for a repeating unit cell).
- min_feature_px — the minimum printable width used for the checks (default 3).
- thin_solid_fraction — fraction of solid in members thinner than that.
- thin_void_fraction — fraction of void in gaps narrower than that.
- floating_solid_fraction — solid not attached to the main structure.
- trapped_void_fraction — enclosed voids (bad for powder/casting processes).
- est_min_solid_thickness_px — estimated thinnest member, from a distance transform.
- manufacturable — true when thin / floating / trapped fractions are all small.
- E_eff_x_rel, E_eff_y_rel — effective Young's modulus in x and y.
- poisson_eff — effective Poisson's ratio.
- G_eff_rel — effective shear modulus.
- K_eff_rel — effective 2D (area) bulk modulus.
- auxetic — true when Poisson's ratio is negative (below −0.01).
- specific_stiffness_rel — stiffness-to-weight = mean effective modulus ÷ volume fraction.

VAE

- Trained on 480 cell dataset
- Convolutional encoder/decoder, latent dimension 32, ~1.07M parameters, 64×64.
- Encoder: four stride-2 conv layers (32→64→128→128), then linear heads for the latent mean and log-variance. Decoder mirrors it with transposed convolutions.
- Loss: binary cross-entropy reconstruction + KL divergence, with a KL warmup over the first third of training so the model learns to reconstruct before the latent is regularized.
- Reconstruction IoU 0.984
- Volume-fraction conditioning: requests 0.3 / 0.5 / 0.6 produce ~0.22 / 0.41 / 0.61.
- Stiffness conditioning lands the requested tercile in roughly 70–100% of samples depending on the target.
- Symmetry conditioning through the label alone failed (0 of 16), so symmetry is enforced at sampling time by mirror-averaging the decoder's probability field before thresholding (the --symmetrize flag). With it, symmetry is exact.
- Nearest-training-neighbour IoU averages ~0.68 (max ~0.90), so the model interpolates

Diffusion model

- Upgrade from VAE
- Trained on 8040 cell dataset
- Small U-Net noise predictor (~2.0M params), 32×32, sinusoidal time embedding.
- DDPM with 400 noise steps. Sampled with DDIM (40–50 steps) for speed.
- The condition is dropped 12% of the time in training; at sampling, guidance w (default 3) pushes toward the target.
- EMA of the weights. Samples come from the EMA copy.
- Trained up to epoch 29, but epoch 25 had better controllability
- Requests 0.2/0.4/0.6 give achieved ~0.37/0.44/0.51. Reliable target range ~0.37–0.57
- Low vs high separates by about 2.4× in effective modulus at a fixed volume fraction
- 22–24 of 24 samples periodically connected.
- Rejection yield: ~40% at vf 0.40 and ~72% at vf 0.50 (up from 8% and 25% for the original 480-cell model)