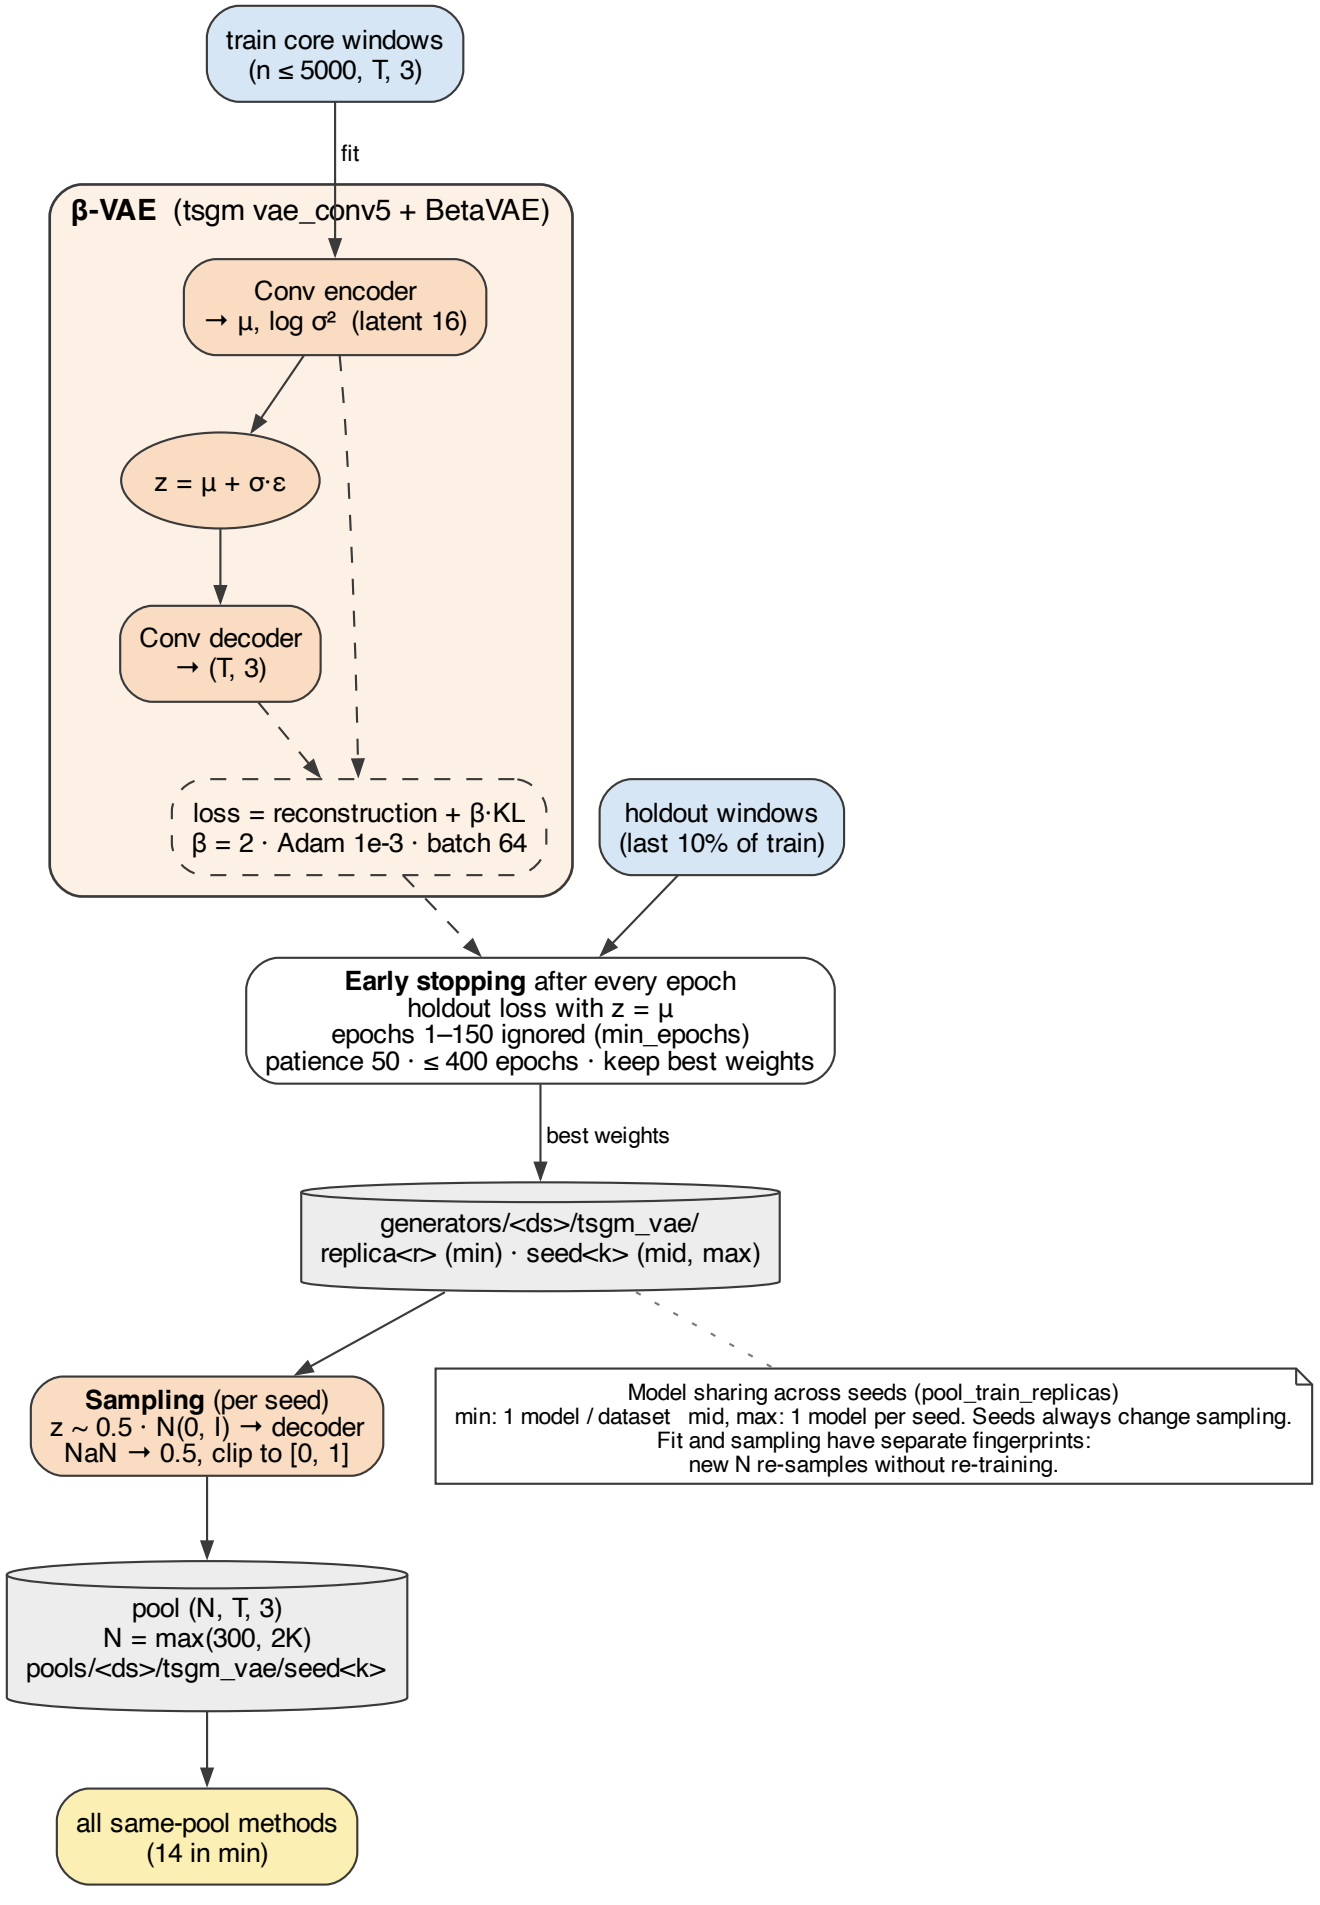

train core windows
(n ≤ 5000, T, 3)
fit
β-VAE (tsgm vae_conv5 + BetaVAE)
Conv encoder
→ μ, log σ² (latent 16)
z = μ + σ·ε
Conv decoder
→ (T, 3)
loss = reconstruction + β·KL
β = 2 · Adam 1e-3 · batch 64
holdout windows
(last 10% of train)
Early stopping after every epoch
holdout loss with z = μ
epochs 1–150 ignored (min_epochs)
patience 50 · ≤ 400 epochs · keep best weights
best weights
generators/<ds>/tsgm_vae/
replica<r> (min) · seed<k> (mid, max)
Sampling (per seed)
z ~ 0.5 · N(0, I) → decoder
NaN → 0.5, clip to [0, 1]
Model sharing across seeds (pool_train_replicas)
min: 1 model / dataset mid, max: 1 model per seed. Seeds always change sampling.
Fit and sampling have separate fingerprints:
new N re-samples without re-training.
pool (N, T, 3)
N = max(300, 2K)
pools/<ds>/tsgm_vae/seed<k>
all same-pool methods
(14 in min)