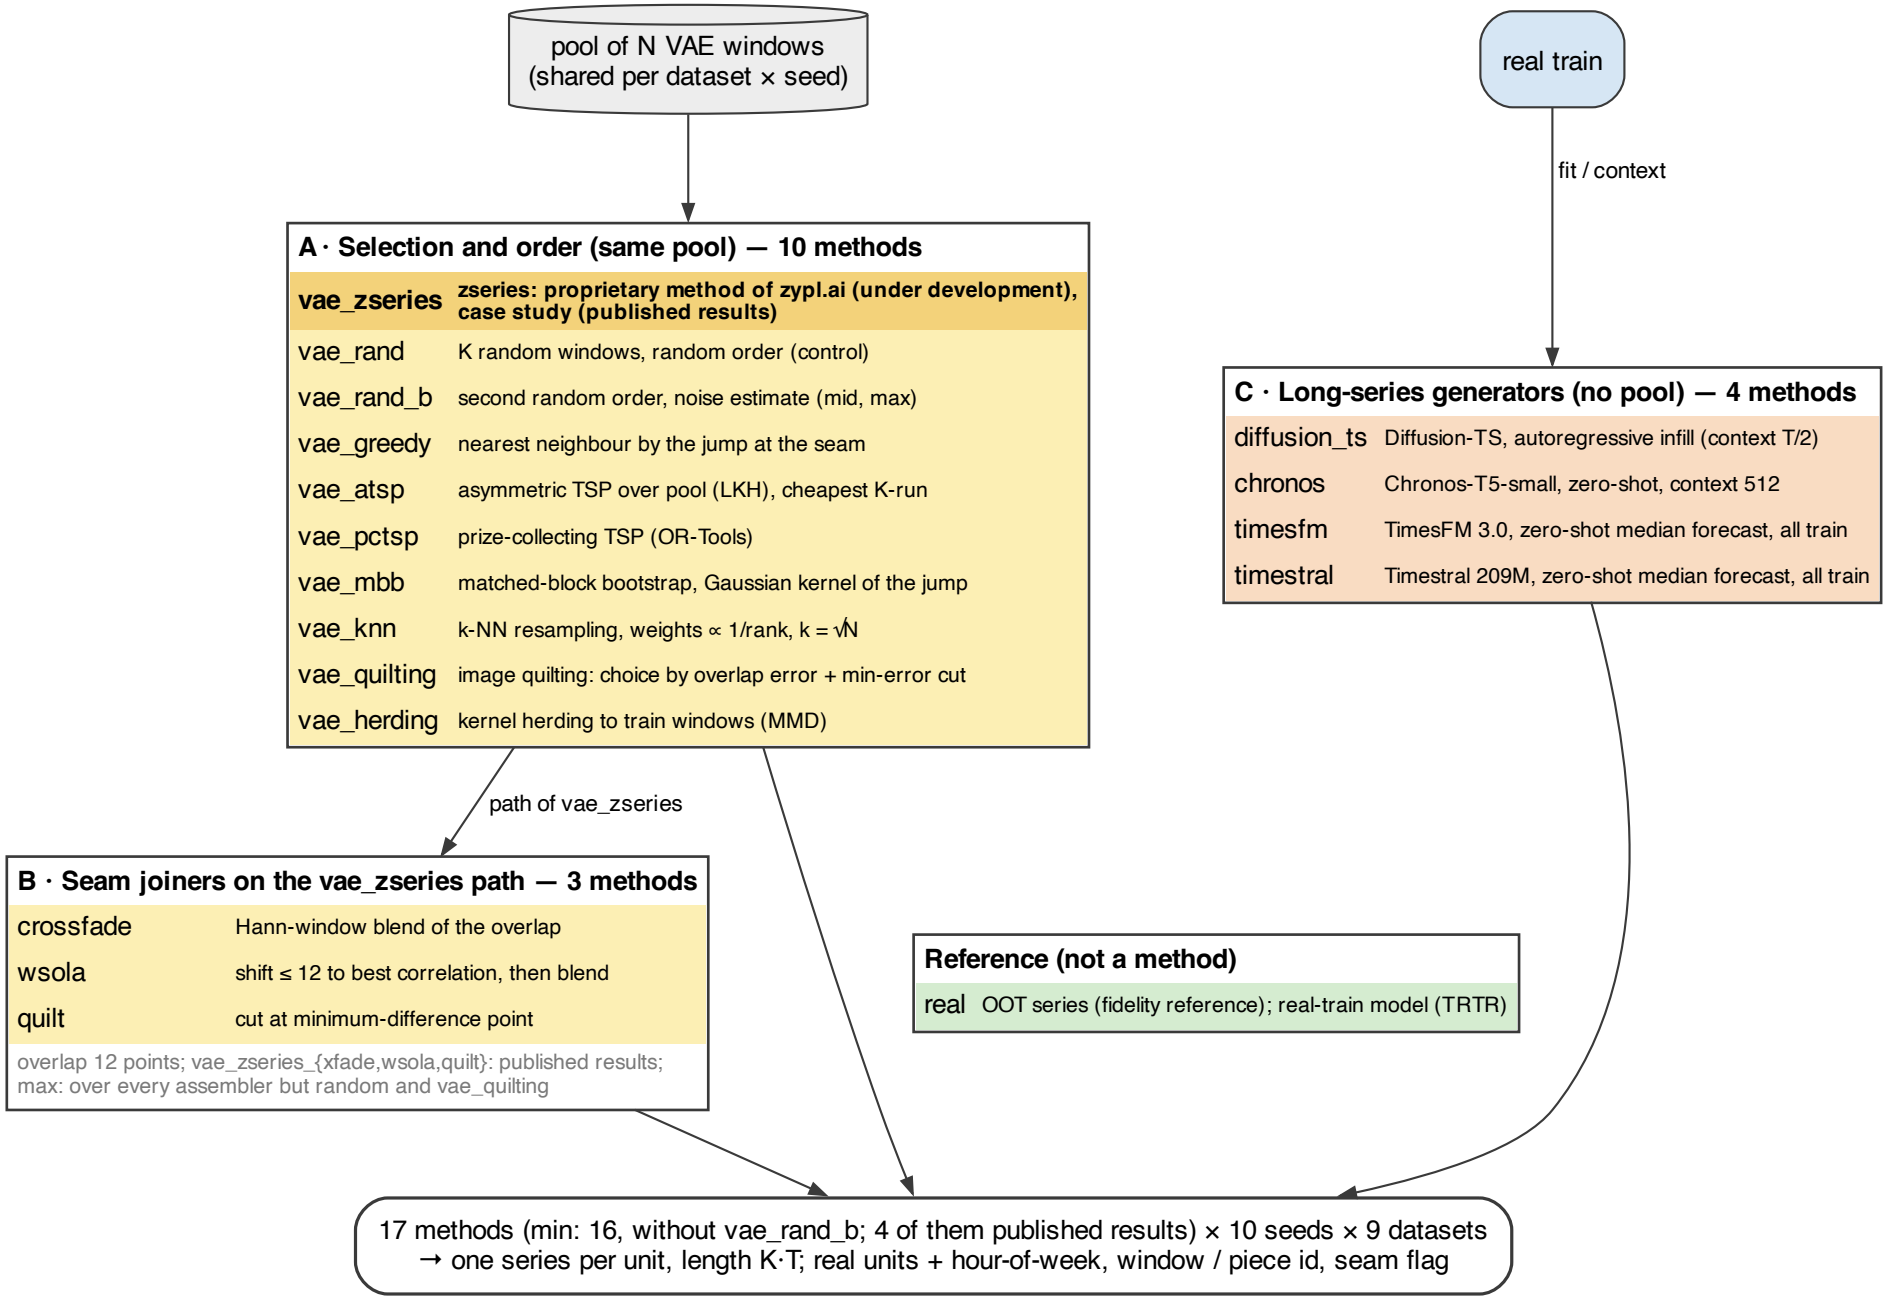

pool of N VAE windows
(shared per dataset × seed)
real train
fit / context
A · Selection and order (same pool) — 10 methods
vae_zseries
zseries: proprietary method of zypl.ai (under development), case study (published results)
vae_rand
K random windows, random order (control)
vae_rand_b
second random order, noise estimate (mid, max)
vae_greedy
nearest neighbour by the jump at the seam
vae_atsp
asymmetric TSP over pool (LKH), cheapest K-run
vae_pctsp
prize-collecting TSP (OR-Tools)
vae_mbb
matched-block bootstrap, Gaussian kernel of the jump
vae_knn
k-NN resampling, weights ∝ 1/rank, k = √N
vae_quilting
image quilting: choice by overlap error + min-error cut
vae_herding
kernel herding to train windows (MMD)
C · Long-series generators (no pool) — 4 methods
diffusion_ts
Diffusion-TS, autoregressive infill (context T/2)
chronos
Chronos-T5-small, zero-shot, context 512
timesfm
TimesFM 3.0, zero-shot median forecast, all train
timestral
Timestral 209M, zero-shot median forecast, all train
path of vae_zseries
B · Seam joiners on the vae_zseries path — 3 methods
crossfade
Hann-window blend of the overlap
wsola
shift ≤ 12 to best correlation, then blend
quilt
cut at minimum-difference point
overlap 12 points; vae_zseries_{xfade,wsola,quilt}: published results;
max: over every assembler but random and vae_quilting
Reference (not a method)
real
OOT series (fidelity reference); real-train model (TRTR)
17 methods (min: 16, without vae_rand_b; 4 of them published results) × 10 seeds × 9 datasets
→ one series per unit, length K·T; real units + hour-of-week, window / piece id, seam flag