OFFICIAL NEURAL MOCO RANKINGS ARE EXTRACTOR-CONDITIONED: A FINITE-K PROTOCOL AUDIT

Anonymous authors
Paper under double-blind review

ABSTRACT

Neural multi-objective combinatorial optimization is used for multi-attribute route planning and resource allocation. Published evaluations report a finite approximate Pareto set extracted by preference-based scalarization: one survivor per preference vector. That official archive is a legitimate score of the advertised preference-to-solution map. The same tables, however, report set-level hypervolume, so a printed ranking mixes per-preference quality with set attainment. We re-score frozen official decoded pools of five open-source constructors (WE-CA, PMOCO, CNH, GIMF-P, and NHDE-P) with a zero-search set-level extractor, Preference-Agnostic Selection (PAS): classical greedy hypervolume subset selection over the entire pool, with no extra neural pass and no local search. PAS strictly exceeds official scalarized extraction on 53 of 56 finite-K settings. Mean-rank reversals occur on close pairs with official margins of 4–5×10^{-4}; published 0.001–0.003 gaps at larger n need not reverse, because the PAS increment is often similar across a pair. We recommend reporting official and PAS hypervolume in parallel as two legitimate targets. Unmoved-then-Align (U2N) is a post-hoc local refinement on top of PAS. Official finite-K rankings are therefore extractor-conditioned.

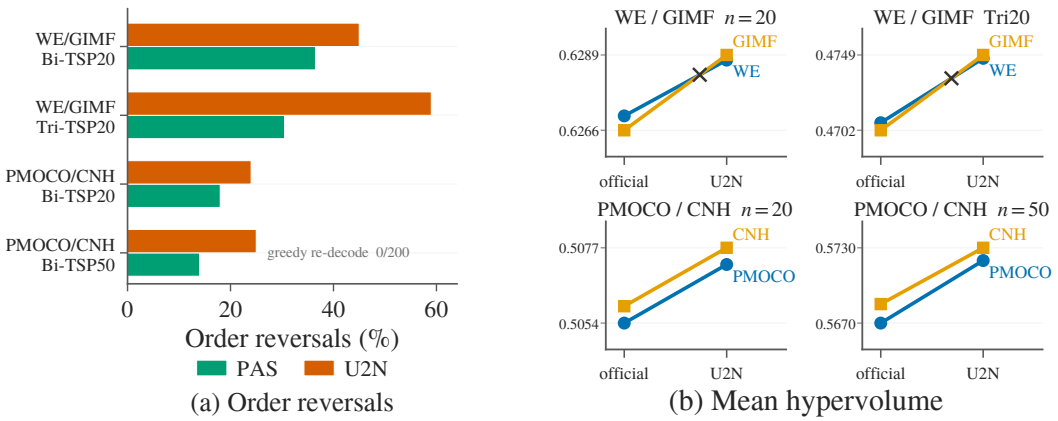
1 INTRODUCTION

Neural multi-objective combinatorial optimization is used for multi-attribute route planning and resource allocation. The returned object is a finite set of trade-offs, and official neural MOCO evaluations score that set by hypervolume—the volume covered toward a reference (Zitzler et al., 2003; Fleischer, 2003). A network cannot emit a continuum of solutions in one pass, so constructors inherit the MOEA/D template (Zhang & Li, 2007). They query one preference vector, decode a path, move the vector, repeat (multi-start decoders such as POMO (Kwon et al., 2020; Lin et al., 2022; Chen et al., 2023; Fan et al., 2025a; Chen et al., 2025a;b), using n starts on an n-city or n-item instance). Using a one-dimensional projection to *search* is reasonable; using it as the only scored K-archive confounds a $\lambda \to x$ query with set-level hypervolume. Those evaluations still report set-level hypervolume, yet each paper's extractor keeps one weighted-sum (WS), Tchebycheff (TCH), or PBI survivor per preference vector and discards mutually non-dominated solutions that sit between those discrete vectors. The mismatch is the same as scoring an archery target by coverage while keeping only arrows that landed on the cardinal axes.

A constructor is the generative policy that emits a decoded pool P; an extractor returns the archive A that is actually scored. The audit scores extractors on frozen official pools, without changing parameters or decoding budgets. Those are two legitimate targets on the same decode, per-preference map fidelity scored by A_{off} (Lin et al., 2022) and set-level attainment scored by PAS. Mean-rank reversals under PAS sit on official gaps of 4–5×10^{-4}; published 0.001–0.003 gaps at larger n need not reverse, because the PAS increment is often shared across a pair (Table 1; Appendix Table A.38). Zero-search PAS flips $73/200$ Bi-TSP20 with a mean-rank change (Figure 1(a)). The $118/200$ WE-CA versus GIMF-P Tri-TSP20 count is a refinement-track flip under U2N. Finite-K leaderboards are therefore extractor-conditioned.

The decoded pool is both a search trace along preference vectors and a set of trade-offs scored by hypervolume. Using the training projection as the extraction map induces two extraction gaps and

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.



(a) Order reversals (b) Mean hypervolume

Figure 1: (a) Instance-order reversals under PAS versus U2N on identical frozen pools (extractor-induced). A second greedy A_{off} pass leaves that pairing unchanged $(0/200)$. (b) Mean hypervolume from the official extractor to U2N; y-ticks mark each panel's HV range; a crossing marks a mean-rank reversal under U2N.

one training-side observation (Figure 2).

- M1, pool truncation. Preference-wise argmin keeps one survivor per preference vector, so a globally non-dominated solution that is not that vector's best scalarized cost is discarded from the archive.
- M2, local inefficiency. A retained solution need not be locally efficient under a feasible combinatorial move $\mathcal{N}(x)$.
- M3, exploratory training-side observation. Policy gradients minimize $\mathcal{L}_{\text{WS}}(\theta) = \mathbb{E}[\lambda^\top f(x)]$ rather than $J_K(P)$. A continued-$\mathcal{L}_{\text{WS}}$ recipe on one official WE-CA basin can collapse onto extreme preference vectors (Appendix F.3), reflecting training-objective divergence on that constructor basin.

Extractor hierarchy. The audit places the official extractor and zero-search PAS on each frozen pool; U2N is M2 refinement on top of PAS. PAS is the greedy hypervolume subset of that pool, so it cannot exceed the best in-pool K-set; it isolates pool truncation (M1) with no extra decode and no local search.

Contributions.

1. **Official finite-K rankings are extractor-conditioned.** Official finite-K evaluations score constructors jointly with a preference-based scalarization extractor, confounding a $\lambda \to x$ map with set-level hypervolume. Re-extracting identical pools can reverse close pairs ($g \sim 4$–5×10^{-4}); larger published gaps need not (Table 1). The recommended primary track is A_{off} with PAS.
2. **Zero-search PAS is the set-level extractor.** Preference-Agnostic Selection returns a classical greedy hypervolume subset of the frozen decoded pool, with no extra neural pass and no local search. It keeps between-vector non-dominated solutions that scalarized extraction discards, and strictly exceeds the official archive on 53 of 56 finite-K settings. Those 53 gains isolate pool truncation (M1) as expected subset-of-superset geometry.
3. **U2N is a feasible refinement on top of PAS.** Unmoved-then-Align applies a dominance-monotone neighborhood move, then residual alignment only on unmoved points, then the same PAS selector. On unconstrained Bi-TSP the increment over PAS grows with n (PAS leftover share $92.9\% \to 19.3\%$ at n=20 $\to$ 200 on official WE-CA).

A K=2 count makes the drop explicit (Figure 2). Two preference vectors keep the extremes $x_1 = (1, 10)$ and $x_3 = (10, 1)$ because the compromise $x_2 = (3, 3)$ wins neither one-dimensional contest (HV 19 at $r = (11, 11)$). Zero-search PAS keeps $\{x_1, x_2\}$ (HV 66); x_3 is unselected at K=2. A dominance-monotone move then advances x_2 to $(2.5, 2.5)$ (HV 73.75) without lengthening either tour.

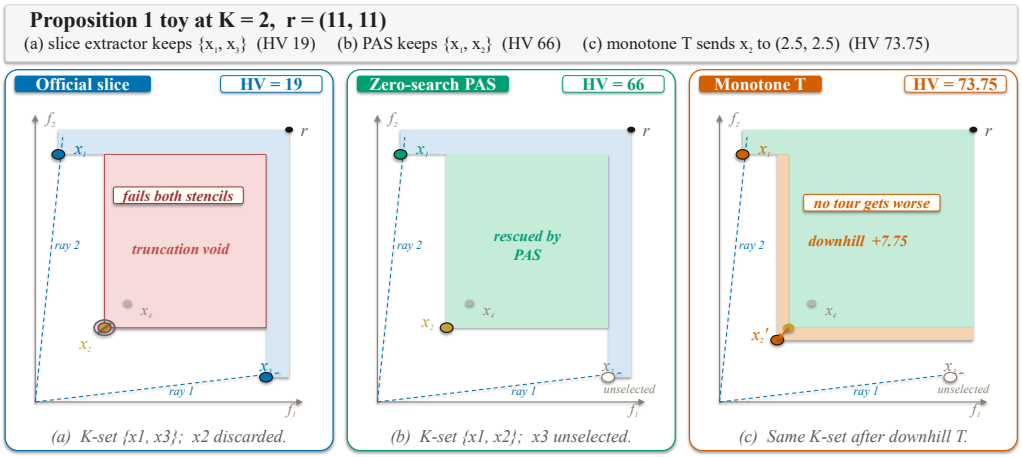


Figure 2: Proposition 1 toy at K=2, $r = (11, 11)$. (a) Slice extractor keeps $\{x_1, x_3\}$ (HV 19); the non-dominated compromise $x_2 = (3, 3)$ fails both one-dimensional stencils (M1). (b) Zero-search PAS keeps $\{x_1, x_2\}$ (HV 66). (c) A dominance-monotone move sends x_2 to $(2.5, 2.5)$ (HV 73.75) without lengthening either tour (M2). Continued $\mathcal{L}_{\text{WS}}$ (M3) is Appendix F.3.

2 Related Work

Neural MOCO and multi-objective decomposition. Learning-based MOCO instantiates the MOEA/D template (Zhang & Li, 2007) with preference-conditioned models for multi-attribute route planning and resource allocation (Li et al., 2021; Kool et al., 2019; Lin et al., 2022; Chen et al., 2023; Fan et al., 2025a; Chen et al., 2025a;b; Fan et al., 2025b; Wu et al., 2026). On TSP, CVRP, and knapsack, finite-K evaluations of WE-CA, PMOCO, CNH, and GIMF-P keep one scalarization survivor per weight and then score that K-archive by hypervolume. Those evaluations therefore measure the constructor together with a scalarized extractor, not generative capacity alone; A_{off} is also the per-preference score those models advertise (Lin et al., 2022).

Scalarized projections, HSSP, and the EMO precedent. Weighted sums and Tchebycheff (Zhang & Li, 2007) score a one-dimensional projection; Pareto ranking (Deb et al., 2002) and hypervolume (Zitzler et al., 2003; Fleischer, 2003) score a set. EMO already uses an unbounded external archive and two-scenario subset selection (Tanabe et al., 2017; Ishibuchi et al., 2018; Bringmann et al., 2014; Kuhn et al., 2016; Guerreiro et al., 2016); the contribution here is to quantify that extractor gap on official neural tables, not to introduce subset selection. Official evaluations still report set-level hypervolume while extracting a finite argmin $A_{\text{off}}(P)$ (Definition 2), so mutually non-dominated points that fail a discretized supporting hyperplane are discarded (Figure 2). Linear weighted sums cannot support non-convex segments (Das & Dennis, 1997; Miettinen, 1999).

In-constructor selection versus official extraction. Two-phase Pareto local search (Lust & Teghem, 2010) updates an archive during search; CDE (Lu et al., 2024) rewires that constructor's dual-inference stack; NHDE MPO (Chen et al., 2023) stores in-constructor memories. Those changes sit inside the constructor—how the policy searches $\mathcal{X}$. The audit replaces only the official extraction map A_{off} on the already-decoded $K \times n$ set; the network and decode protocol remain those of the constructor under evaluation.

Two extraction pipelines. Section 4 defines zero-search PAS (in-pool greedy hypervolume; Guerreiro et al. (2016)) and U2N (M2 refinement). The recommended reporting track reports A_{off} with PAS for per-preference fidelity versus set-level attainment. Official finite-K rankings remain extractor-conditioned.

3 Preliminaries

Sets, attainment, and operating budgets. Hypervolume is the volume of objective space a set covers toward an admissible reference r. Let $\mathcal{X}$ be the discrete feasible set and $f : \mathcal{X} \to \mathbb{R}^m_{\geq 0}$ the

m-objective additive cost. Write $x \preceq y$ when $f(x) \le f(y)$ componentwise, and $x \prec y$ when $x \preceq y$ and $f(x) \neq f(y)$. For finite $S \subset \mathcal{X}$, drop duplicate costs and set $\mathrm{ND}(S) = \{x \in S : \text{no } y \in S \text{ has } y \prec x\}$. A reference r is admissible if $f(x) \le r$ on S (official WE-CA points). The attainment region and hypervolume

$$\mathcal{B}_r(S) := \bigcup_{x \in S} \{q \in \mathbb{R}^m : f(x) \le q \le r\}, \qquad \mathrm{HV}_r(S) := \mathrm{vol}_m\big(\mathcal{B}_r(S)\big) \tag{1}$$

coincide on S and $\mathrm{ND}(S)$ (Zitzler et al., 2003).

Candidate pool and official archive. The pool P is the decoded cloud across preference vectors and starts. A_{off} is the official stencil, one one-dimensional winner per preference vector. Write the preference simplex $\Sigma^{m-1} = \{\lambda \in \mathbb{R}^m_{\ge 0} : \|\lambda\|_1 = 1\}$. A constructor π_θ is queried on $\Lambda = \{\lambda_1, \ldots, \lambda_K\} \subset \Sigma^{m-1}$ with $K = 101$ (105 when $m = 3$) and n POMO starts on an instance of size n (Kwon et al., 2020). *Constructor* is that generative policy; *extractor* is A_{off} returning a K-archive from P. The audit replaces only the extractor. Weights, grids, and references stay fixed. WE-CA / GIMF-P / NHDE-P use $(20, 20)/(35, 35)/(65, 65)$ at $n \le 100$ and that release's r at $n > 100$ and KroAB; PMOCO / CNH use that paper's r^* (Section 5). Reported hypervolume follows each source paper's official normalization.

Definition 1 (Candidate set)**.** The candidate set is the multiset $P = \{x(\lambda_k, s)\}$ with $|P| = Kn$, start index s, and slices $P_k = \{x(\lambda_k, s)\}_{s=1}^n$.

Definition 2 (Finite WS-argmin)**.** Official evaluation keeps one minimizer per slice. Under official WS extraction,

$$A_{\mathrm{WS}}(P) = \big\{\arg\min_{x \in P_k} \lambda_k^\top f(x)\big\}_{k=1}^K, \tag{2}$$

ties are broken by the official implementation's first index and $A_{\mathrm{off}} = A_{\mathrm{WS}}$. PMOCO / CNH evaluations report a TCH or PBI K-archive $A_{\mathrm{off}}(P)$ under that paper's r^*. The same decoded pool supports the cardinality-K hypervolume subset

$$J_K(P) = \max\{\mathrm{HV}_r(A) : A \subseteq \mathrm{ND}(P),\ |A| \le K\}. \tag{3}$$

Official evaluations report $\mathrm{HV}_r(A_{\mathrm{off}}(P))$, not $J_K(P)$.

Zero-search PAS is a greedy in-pool selector, so it cannot exceed the best K-subset of $\mathrm{ND}(P)$:

$$\mathrm{HV}_r(A_{\mathrm{PAS}}(P)) \le J_K(P) \le \mathrm{HV}_r(\mathrm{ND}(P)). \tag{4}$$

On Table A.2 Panel A, $\mathrm{HV}_r(A_{\mathrm{off}}) \le \mathrm{HV}_r(A_{\mathrm{PAS}})$ holds empirically (three ties). Γ is scored beside (4) as remaining M2, not as a bound on $J_K(P)$.

Three mismatches. M1 is the official per-weight stencil (2); M2 is local inefficiency under a feasible move; M3 is that training targets $\lambda^\top f$ rather than $J_K(P)$ (Appendix F.3).

Definition 3 (Recovered hypervolume)**.** Write $A_{\mathrm{off}}(P)$ for that paper's official archive (finite WS / TCH / PBI, or unbounded EP on NHDE-P). The recovered hypervolume of a named K-set A is $\Delta(A, A_{\mathrm{off}}) = \mathrm{HV}_r(A) - \mathrm{HV}_r(A_{\mathrm{off}})$; we write leftover for the per-instance mean of Δ. Under official WS extraction, $A_{\mathrm{off}} = A_{\mathrm{WS}}$.

Proposition 1 (Finite WS-argmin discards non-dominated candidates)**.** *There exist candidate sets P with $|A_{\mathrm{WS}}(P) \cap \mathrm{ND}(P)| < |\mathrm{ND}(P)|$, so A_{WS} need not attain $J_K(P)$ (unsupported-point geometry of Das & Dennis (1997)).*

Proof. Take $K{=}2$, $\lambda_1 = (0.9, 0.1)^\top$, $\lambda_2 = (0.1, 0.9)^\top$, $f(x_1) = (1, 10)$, $f(x_2) = (3, 3)$, $f(x_3) = (10, 1)$, and $r = (11, 11)^\top$ (the $K{=}2$ count in the introduction; Figure 2(a) is the same stencil). Both slice inner products equal $1.9 < 3.0$, so $A_{\mathrm{WS}} = \{x_1, x_3\}$ and $\mathrm{HV}_r(A_{\mathrm{WS}}) = 19 < 66 \le J_K(P)$.

4 AUDIT PROTOCOL

The primary contribution is empirical and methodological. We score evaluation-time extractors on already-decoded official pools. The statements below are elementary geometric facts used to justify the two maps. A differentiable surrogate of $J_K(P)$ at training time is left open. Two maps act on the same frozen $K \times n$ pool. Zero-search PAS ($\mathrm{GreedyHV}_{K,r} \circ \mathrm{Floor}_K \circ \mathrm{ND}$) is the set-level extractor and the recommended reporting companion to A_{off}. It is the classical greedy hypervolume K-subset

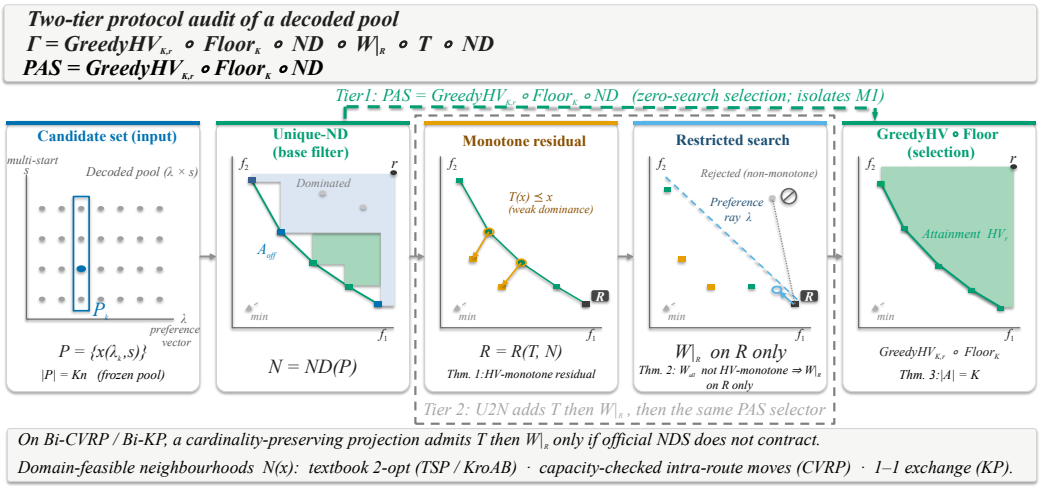


Figure 3: **Two-tier protocol audit of a decoded pool.** Tier 1 (PAS) drops dominated points, restores the official count, then greedy hypervolume (zero-search; isolates M1). Tier 2 (dashed U2N) applies dominance-monotone T, then residual alignment only on unmoved points, then the same PAS selector. On Bi-CVRP / Bi-KP, a cardinality-preserving projection blocks domination shadowing. T is hypervolume-monotone (Lemma 2).

of P. U2N is a post-hoc local refinement we propose on top of PAS, obtained by inserting monotone T then residual $W|_R$ before that same selector (M2; Lemma 2 for T). Both maps reuse the official decode; only A_{off} is replaced. Read the map left to right as four jobs. ND drops clearly dominated candidates. T never worsens any objective, so the attainment of $T(S)$ cannot shrink (Lemma 2). $W|_R$ then pushes only points that T left unmoved, because unrestricted alignment need not preserve hypervolume (Proposition 2). Floor_K restores the official count with copies that occupy no extra volume; greedy hypervolume is the same PAS selector.

$$\Gamma = \text{GreedyHV}_{K,r} \circ \text{Floor}_K \circ \text{ND} \circ W\big|_R \circ T \circ \text{ND}, \tag{5}$$

applied to P. Those four jobs are the maps in (5). T-moved points stay fixed under residual alignment; Floor_K restores $|A| = K$ (Theorem 1). Γ scores remaining M2 on the same frozen pool (Figure 3; Appendix Algorithm A.1); Table A.2 reports the PAS extractor and the proposed refinement Γ versus that paper's extractor, with Γ_{hold} on Bi-CVRP / Bi-KP. $\mathcal{N}(x)$ is the feasible move that makes T well-defined (Section 4.1).

4.1 Dominance-monotone local refinement

Write $\mathcal{N}(x)$ for a feasible neighbourhood of x and $\mathcal{N}_{\preceq}(x) = \{x' \in \mathcal{N}(x) : x' \preceq x\}$ for its dominance-monotone restriction. A first-improvement oracle on $\mathcal{N}_{\preceq}$ returns $T(x) \in \mathcal{N}_{\preceq}(x) \cup \{x\}$, so $T(x) \preceq x$ for every x.

Family-feasible neighbourhood. T must stay in $\mathcal{X}$, using textbook 2-opt on Bi-TSP / KroAB; capacity-feasible intra-route 2-opt or or-opt on Bi-CVRP (inter-route IMVND only inside Γ_{hold} at n=20; Appendix Table A.51); 1–1 item-exchange on Bi-KP. On Bi-TSP / KroAB the operator is P2O (Pareto 2-opt).

Definition 4 (Unmodified residual)**.** A map $T : \mathcal{X} \to \mathcal{X}$ is dominance-monotone if $T(x) \preceq x$ for every $x \in \mathcal{X}$. For finite $N \subset \mathcal{X}$ the unmodified residual is $R(T, N) = \{x \in N : \|f(T(x)) - f(x)\|_\infty \le \varepsilon\}$ with $\varepsilon = 10^{-12}$ (objective-space round-off). A point of R is locally f-efficient under $\mathcal{N}_{\preceq}$, so first-improvement found no accepted dominating neighbour. Appendix Algorithm A.1 stores that residual as an index set.

Lemma 1 (Attainment inclusion)**.** *If T is dominance-monotone and S is finite and admissible for r, then $\mathcal{B}_r(S) \subseteq \mathcal{B}_r(T(S))$. Moreover $\mathcal{B}_r(S) = \mathcal{B}_r(\text{ND}(S))$, so $\text{HV}_r(S) = \text{HV}_r(\text{ND}(S))$ (Equation (1)). The argument is Appendix C.*

Lemma 2 (Pointwise weak dominance is hypervolume-monotone)**.** *Let T be dominance-monotone and $S \subset \mathcal{X}$ finite and admissible for r. Then $\text{HV}_r(\text{ND}(T(S))) \ge \text{HV}_r(\text{ND}(S))$.*

A pointwise weak-dominance move expands the attainment region, so hypervolume cannot fall (Appendix C). Lemma 2 admits T inside Γ; recovered hypervolume of the later K-archive is experimental (Table A.2).

4.2 Residual preference alignment

T only accepts weak dominators and therefore cannot advance a trade-off along an inclined preference ray. Residual preference alignment is the source-λ operator W_λ. It admits any $x' \in \mathcal{N}(x)$ with $\lambda^\top f(x') < \lambda^\top f(x)$ and returns $W_\lambda(x)$. Write W_{all} for applying W_λ component-wise to every candidate in its input set.

Lemma 3 (Source-λ search is not dominance-monotone)**.** *There exist a candidate x, a weight $\lambda \in \Sigma^{m-1}$, and a neighbour $x' \in \mathcal{N}(x)$ such that $\lambda^\top f(x') < \lambda^\top f(x)$ and $x' \not\preceq x$. The counter-example is Appendix C.*

Proposition 2 (Unconstrained source-λ search need not be hypervolume-monotone)**.** *There exist a finite set $S \subset \mathcal{X}$, a weight $\lambda \in \Sigma^{m-1}$, and an admissible reference r such that* $\mathrm{HV}_r(W_{\text{all}}(S)) < \mathrm{HV}_r(S)$.

Restriction to the unmodified residual. The argument is Appendix C. Proposition 2 shows that greedy scalarized descent along an inclined ray can erode an orthogonal attainment region and lower global hypervolume. To keep every T-moved point invariant, U2N applies W_λ only on the unmodified residual $R(T, \mathrm{ND}(P))$. $W|_R(x) = W_\lambda(x)$ if $x \in R(T, \mathrm{ND}(P))$ and $W|_R(x) = x$ otherwise. By construction, $W|_R$ leaves every T-moved point unchanged. T moves only inside the weak-dominance cone $\mathcal{N}_\preceq$, so it cannot expand orthogonal trade-offs; unconstrained W_{all} has directional reach but is not hypervolume-monotone (Proposition 2; Figure 4(b)). The residual $R(T, \mathrm{ND}(P))$ is the set of T-unimproved points—the locus where a source-λ step can add a lateral trade-off without overwriting a dominance-monotone update.

4.3 Cardinality preservation and hypervolume selection

Official protocols require $|A| = K$ ($K = 101$ when $m = 2$; $K = 105$ when $m = 3$), so they compare equal-cardinality archives. The unique non-dominated set need not reach K (official WE-CA Bi-TSP20 unique-ND 24 at K=101; Table A.2).

Definition 5 (Fixed-cardinality K-archive)**.** An archive A is a multiset of size K. Official evaluations compare $|A| = K$, while $|\mathrm{ND}(A)|$ may be smaller. Duplicate pads add no Lebesgue measure, $\mathrm{HV}_r(A) = \mathrm{HV}_r(\mathrm{ND}(A))$.

Theorem 1 (Padding restores $|A| = K$, not unique non-dominated cardinality)**.** *Let $N_0 = \mathrm{ND}(P)$ and let Q be the multiset after T and $W|_R$. Write $N_1 = \mathrm{ND}(Q)$ and $\mathcal{R}_{\text{cand}} = \{y \in N_0 \cup Q : \nexists z \in N_1,\ z \prec y\}$. Write $\kappa = \min(K, |N_0|)$. If $|N_1| < \kappa \le |\mathcal{R}_{\text{cand}}|$, restoring $\kappa - |N_1|$ points of $\mathcal{R}_{\text{cand}} \setminus N_1$ and padding yields F with $|F| = K$. The unique non-dominated set need not reach K. Any restored candidate that was updated by T is required to remain weakly dominating its pre-move cost, so Lemma 2 still applies to those restorations.*

Padding matches the official $|A| = K$ rule; duplicate pads add no Lebesgue measure. The last stage of Γ is the same greedy hypervolume selection as PAS (Guerreiro et al., 2016), now after T and $W|_R$. On official n=100, exact two-objective hypervolume DP exceeds the greedy selector by 6.6×10^{-6}; residual maps after that stage recover 0.00166, about $250\times$ the greedy–exact discrepancy, so the audit is not an artefact of greedy HSSP approximation. Bi-CVRP / Bi-KP additionally apply (6).

4.4 Family instantiations and the cardinality-preserving projection

The guarantees of U2N do not name a particular $\mathcal{N}(x)$; each family supplies one so T stays in $\mathcal{X}$ (Section 4.1). Write A_{PAS} for the PAS archive of P, Γ_T for the T-only truncation of (5), and $n_{\text{off}}(A)$ for the official unique non-dominated count of A (`Pareto_sols` on Bi-CVRP; else that of the reported K-set). On Bi-CVRP and Bi-KP the canonical evaluation mapping is the instance-wise dual-criterion cascade

$$\Gamma_{\text{hold}}(P) = \begin{cases} \Gamma(P), & n_{\text{off}}(\Gamma(P)) \ge n_{\text{off}}(A_{\text{PAS}}), \\ \Gamma_T(P), & n_{\text{off}}(\Gamma_T(P)) \ge n_{\text{off}}(A_{\text{PAS}}), \\ A_{\text{PAS}}, & \text{otherwise.} \end{cases} \tag{6}$$

Constrained families expose a Pareto-breadth versus hypervolume trade-off. A deeper feasible move can dominate two marginal trade-offs (domination shadowing), raising HV while the official unique-ND count falls. The cascade uses unique-ND counts, not leftover Δ, and is identity on unconstrained permutations (Appendix G).

Drop-in PAS extractor. Algorithm 1 is the zero-search PAS extractor (Guerreiro et al., 2016); Γ is Appendix Algorithm A.1.

Algorithm 1 Zero-search PAS extractor

Input: Decoded costs of P, cardinality K, reference r
Output: Archive A_{PAS} with $|A| = K$
$N \leftarrow \mathrm{UniqueND}(P)$
if $|N| \leq K$ **then**
  return $\mathrm{Floor}_K(N)$ ▷ pad duplicates; no HV change
end if
$A \leftarrow \emptyset$
for $k = 1$ **to** K **do**
  $x^\star \leftarrow \arg\max_{x \in N \setminus A} \left(\mathrm{HV}_r(A \cup \{x\}) - \mathrm{HV}_r(A)\right)$
  $A \leftarrow A \cup \{x^\star\}$
end for
return A

5 Experiments

Table 1 reports official versus PAS gaps on the same-r^* constructor pairs that decide published order. The full per-setting hypervolume archive, including the proposed refinement A_{U2N}, is Appendix Table A.2. Re-extracting those frozen pools can reverse close constructor pairs (Section 5.3; Figure 1). Finite-K claims cite the 56 official finite-K settings (Appendix Table A.3).

5.1 Problems, testdata, and protocols

Evaluation uses the official test suites of Chen et al. (2025a). Bi-TSP/Tri-TSP minimize two/three Euclidean tour lengths on n cities; Bi-CVRP trades length against makespan under capacity; Bi-KP is a two-value 0–1 knapsack; KroAB is TSPLIB KroA/KroB (n=1). Sizes are Bi-TSP $n \leq 200$, Bi-CVRP / Tri-TSP $n \leq 100$, Bi-KP $n \in \{50, 100, 200\}$. Checkpoints are the frozen official files of WE-CA, PMOCO, NHDE-P, CNH, and GIMF-P. K=101 (105 if m=3); instance counts are 200 except official PMOCO/CNH seed1234 sets (100). PAS reuses the already-decoded $K \times n$ pool; monotone T then residual $W|_R$ re-extract that same pool (no second neural pass). Reported recovered hypervolume uses that paper's extractor, r^*, grid, and testdata; paired Wilcoxon signed-rank is one-sided greater (α=0.01) on A_{U2N} versus A_{off}, with paired effect size $d_z = \bar{\Delta}/s_\Delta$. The identification is the frozen pool. Parameters, grids, r^*, and the paid $K \times n$ decode stay those of the constructor under evaluation. Official WS/TCH/PBI extractors remain those of the paper under evaluation; the audit replaces only the scored K-set. Ranking uses the sign of the same-r^* instance-mean gap (Table 1). Instance-bootstrap 95% CIs of Δ_{PAS} use 10^4 resamples (seed 2026).

5.2 Does a finite official extractor attain the pool?

Zero-search PAS strictly exceeds A_{off} on $53/56$ and ties on three PMOCO Bi-KP settings; A_{U2N} is strictly above A_{off} on all 56 finite-K settings (Table A.2, Panel A). Those three ties occur where the unique non-dominated set is already sparse, so Tchebycheff fills every selectable slot and greedy subset selection cannot add volume. The 53 strict gains isolate pool truncation (M1) on a frozen decode: a hypervolume subset of a superset is expected to rise. Ranking evidence is Δ_{PAS} relative to g_{off} (Section 5.3; Appendix Table A.38). Panel A is the finite-K census on each frozen $K \times n$ decode. Panel B (NHDE-P EP) records $\mathrm{HV}_r(A_{\mathrm{PAS}}) \leq \mathrm{HV}_r(\mathrm{EP})$ empirically: a cardinality-K selector cannot match an archive that is allowed to grow. On official WE-CA Bi-TSP, leftover after PAS is 0.00156 at n=20 and 0.00126 at n=100 (Table 3). On capacitated routing, knapsack, and three-objective TSP, official finite-K extraction remains lossy after neighbourhood and cardinality constraints (Eq. (6); leftover $+3.4$–3.6×10^{-4} on WE-CA Bi-CVRP; $|\mathrm{ND}|$ 796 $\rightarrow$ 3241 on Tri-TSP, Appendix Table A.47). The hold cascade on those suites compares unique-ND size, not leftover Δ, and returns A_{PAS} if unique-ND would contract.

5.3 DO OFFICIAL RANKINGS REVERSE ON IDENTICAL POOLS?

Table 1: Constructor gaps on identical frozen pools (five-decimal means from Appendix Table A.2). Signed $g_{\text{off}} = \text{HV}(A_{\text{off}}) - \text{HV}(B_{\text{off}})$ and $g_{\text{PAS}} = \text{HV}(A_{\text{PAS}}) - \text{HV}(B_{\text{PAS}})$ use the pair order. Δ_{PAS} is the differential extractor increment. Instance-bootstrap 95% CIs (10^4, seed 2026) of Δ_{PAS} are Bi-TSP20 $[-0.00062, -0.00039]$ and Bi-TSP100 $[-0.00004, +0.00004]$. Full pair CIs are Appendix Table A.39. Mean-rank flip when the PAS winner differs from the official winner. The $118/200$ Tri-TSP20 count is a U2N refinement-track flip (official/PAS winner WE-CA). KroAB is n=1 (no CI).

Pair	Problem	g_{off}	g_{PAS}	Δ_{PAS}	Inst. flip	Mean
WE-CA–GIMF	Bi-TSP20	$+0.00044$	-0.00008	-0.00052	$73/200$	Flip
	Bi-TSP50	$+0.00161$	$+0.00150$	-0.00011	$13/200$	Pres.
	Bi-TSP100	$+0.00326$	$+0.00325$	-0.00001	$0/200$	Pres.
	Bi-TSP150	$+0.00330$	$+0.00323$	-0.00007	$0/200$	Pres.
	Bi-TSP200	-0.00029	-0.00029	0	$16/200$	Pres.
	Tri-TSP20	$+0.00047$	$+0.00027$	-0.00020	$61/200$	Pres.
	KroAB200	$+0.00009$	-0.00004	-0.00013	$1/1$	Flip
PMOCO–CNH	Bi-TSP20	-0.00052	-0.00058	-0.00006	$18/100$	Pres.
	Bi-TSP50	-0.00150	-0.00125	$+0.00025$	$14/100$	Pres.
	Bi-TSP100	-0.00169	-0.00149	$+0.00020$	$4/100$	Pres.

Table 1 asks whether set-level extraction changes published superiority on identical decoded pools. The ranking object is the sign of the same-r^* instance-mean gap. That sign moves when $|\Delta_{\text{PAS}}|$ exceeds $|g_{\text{off}}|$.

Close official margins. On WE-CA versus GIMF-P Bi-TSP20 the published gap is $+0.00044$. Zero-search PAS reverses the mean winner ($73/200$ instances change pairing). The differential extractor increment $\Delta_{\text{PAS}} = -0.00052$ has a 95% instance-bootstrap CI $[-0.00062, -0.00039]$ (10^4, seed 2026). A second greedy A_{off} pass on the same decode reproduces the locked lists ($0/200$), so the mean-rank reversal is extractor-induced rather than a second-sample artefact. Tri-TSP20 keeps the official mean winner under PAS ($61/200$); the $118/200$ count is a U2N refinement-track flip.

Wider official gaps. The same pair at n=$100/150$ has $g_{\text{off}} \approx 0.003$ and $0/200$ flips, because the PAS gains are essentially shared ($\Delta_{\text{PAS}} \approx 0$; Bi-TSP100 CI $[-0.00004, +0.00004]$ includes 0). PMOCO versus CNH never changes the mean winner. Those rankings survive because the truncation penalty is nearly symmetric across the pair.

What the numbers establish. Close published margins of order 4–5×10^{-4} can be smaller than $|\Delta_{\text{PAS}}|$. On WE-CA versus GIMF-P Bi-TSP20, $g_{\text{off}} = +0.00044$ and $\Delta_{\text{PAS}} = -0.00052$, so the mean winner reverses on the same decode. Remaining 0.001–0.003 gaps after PAS, as at n=$100/150$, keep the official winner because the increment is shared. Constructor order is the sign of the same-r^* instance-mean gap. All 46 leftover Wilcoxon tests of A_{U2N} versus A_{off} reject, including pairs whose official winner is preserved under PAS.

5.4 POOL TRUNCATION VERSUS LOCAL INEFFICIENCY

Leftover under PAS and U2N could be pool-wide subset recovery (M1) or unused local search (M2). Two controls separate those sources.

Local search on the reported archive. If leftover were mainly local roughness, 2-opt on the official K-set would recover it without inspecting discarded points. Table 2 shows that T applied only to A_{WS} recovers 2.5%, 9.0%, and 22.3% of leftover at n=$20/50/100$. At those sizes the bottleneck is mechanical truncation between rays (M1), not roughness of the reported tours (M2). At n=$150/200$ the shares rise to $40.9\%/69.4\%$ and $T(A_{\text{WS}})$ exceeds PAS, so local inefficiency of the reported set then dominates (Table 3).

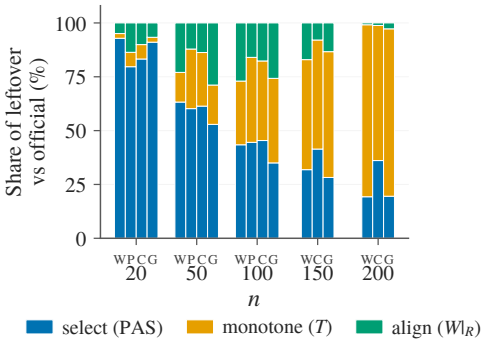
Zero-search PAS versus extra compute. PAS makes no local move and no extra neural forward pass; it is an unparameterized K-subset selector on the already-decoded pool. Logged runtime of 10–11 min/100 is an unoptimized general-dimension loop; two-objective greedy HSSP is $O(|N| \log |N| + K|N|)$ (Guerreiro et al., 2016). Doubling the preference budget ($101 \rightarrow 201$) still leaves PAS 1.39×10^{-3} below U2N at 101 (Figure A.5). Extra neural samples enlarge P and do not close leftover at the official K. Continued $\mathcal{L}_{\text{WS}}$ is empirical (M3; Appendix F.3).

Table 2: T on A_{WS} only (official WE-CA Bi-TSP). Share of leftover vs. U2N; WTL 0–0–200 at n=50/100.

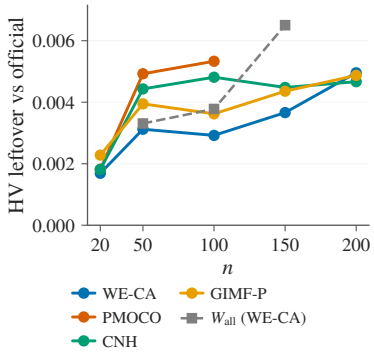
n	WS	$T(A_{\mathrm{WS}})$	PAS	U2N	share
20	0.62704	0.62708	0.62860	**0.62873**	2.5%
50	0.63905	0.63933	0.64103	**0.64217**	9.0%
100	0.70386	0.70451	0.70512	**0.70678**	22.3%
150	0.69912	0.70062	0.70028	**0.70278**	40.9%
200	0.72638	0.72982	0.72733	**0.73134**	69.4%

Table 3: Official WE-CA Bi-TSP staircase and leftover share PAS / T / $W|_R$; W_{all} is unconstrained source-λ search.

	official WE-CA HV					share			
n	WS	PAS	T	U2N	W_{all}	WE-CA	PMOCO	CNH	GIMF-P
20	0.62704	0.62860	0.62864	0.62873	—	92.9/2.2/4.9	79.7/6.7/13.6	83.3/6.7/10.1	91.1/2.3/6.7
50	0.63905	0.64103	0.64146	0.64217	0.64236	63.3/13.8/23.0	60.2/27.7/12.1	61.4/24.9/13.7	52.9/18.3/28.8
100	0.70386	0.70512	0.70599	0.70678	0.70764	43.4/29.6/27.0	44.6/39.5/16.0	45.4/37.0/17.7	35.0/39.3/25.7
150	0.69912	0.70028	0.70216	0.70278	0.70562	31.9/51.1/17.0	—	41.4/50.6/7.9	28.3/58.4/13.3
200	0.72638	0.72733	0.73129	0.73134	—	19.3/79.9/0.8	—	36.1/62.7/1.2	19.5/77.7/2.7



(a) Leftover share



(b) Leftover vs official extractor

Figure 4: (a) Leftover share of PAS / T / $W|_R$. Bars denoted W, P, C, and G represent WE-CA, PMOCO, CNH, and GIMF-P, respectively, and sum to 100% of that constructor's U2N leftover. (b) Leftover versus official extractor on Bi-TSP; dashed W_{all} on WE-CA (Proposition 2).

6 Conclusion

Official finite-K neural MOCO tables rank constructors jointly with a lossy slice-wise extractor, mixing directional ray alignment with set-level attainment. Close published margins (g∼4–5 $\times 10^{-4}$) can reverse under set-level extraction on identical decodes; wider rankings survive when the truncation penalty is shared (Figure 1, Table 1). The 53/56 PAS gains isolate pool truncation. Future finite-K tables should follow three reporting rules.

Dual-track reporting (Table 4). Primary tables report A_{off} and A_{PAS} side by side. A claim of set-level attainment must raise the PAS score; an A_{off}-only gain is per-preference ray alignment.

Local-inefficiency diagnostic. Ablations report A_{U2N} as leftover M2, separating generative quality from neighbourhood roughness.

Protocol and cardinality parity. Keep finite-K archives separate from unbounded external archives (EP), and report unique non-dominated size with hypervolume on constrained suites.

A differentiable set-level surrogate that would close M3, and scaling zero-search selection to $m \geq 4$, remain open.

Table 4: Illustrative four-way reporting template (rounded Table A.2 Bi-TSP100 magnitudes, not a new experiment). Primary track is A_{off} and A_{PAS}. Official-only gain is per-preference quality; PAS-only gain is hidden pool capacity. Table A.2 also reports A_{U2N}.

Constructor	$\mathrm{HV}_r(A_{\mathrm{off}})$	$\mathrm{HV}_r(A_{\mathrm{PAS}})$	Verdict
Previous model	0.7006	0.7019	baseline
Official-only gain	0.7039 ↑	0.7019 =	per-preference only
Masked capacity	0.7006 =	0.7039 ↑	hidden pool
Genuine progress	0.7039 ↑	0.7051 ↑	Pareto expansion

AI USE STATEMENT

Generative AI tools were used only to assist with writing and editing the manuscript. The authors reviewed the text and take responsibility for the final content.

ETHICS STATEMENT

All authors adhere to the ICLR Code of Ethics (`https://iclr.cc/public/CodeOfEthics`). The study uses public synthetic Euclidean routing archives and TSPLIB Kro files from official WE-CA testdata (Section 5); it involves no human subjects or personal data. Checkpoints are re-scored under each paper's extractor and r^*. LLM assistance is disclosed in the AI use statement. Sponsorship, if any, appears after review.

REPRODUCIBILITY STATEMENT

This statement points to the artifacts that make the claims checkable. The fixed-cardinality archive map, dominance-monotone T, residual alignment $W|_R$, and padding to $|A| = K$ are Appendix Algorithm A.1 and Section 4. The unmodified residual R is defined in Section 4.1. Recovered hypervolume is Definition 3. The unconstrained-W_{all} ablation is Appendix Figure A.2; the continued-training diagnostic is Appendix F.3 and Appendix Figure A.5.

The evaluation protocol (each paper's official extractor and r^*; published testdata as in Section 5.1; size-matched versus zero-shot checkpoints; paired Wilcoxon signed-rank of A_{U2N} versus that official extractor, one-sided greater, α=0.01, with $d_z = \bar{\Delta}/s_\Delta$) is Section 5. The primary ranking columns are Table 1; gap bootstrap CIs are Appendix Table A.39; the full HV archive is Appendix Table A.2; A_{U2N} win–tie–loss is Table A.4. The finite-K census is Appendix Table A.3. PMOCO/NHDE-P Bi-KP evaluation stays in Appendix Table A.49; CNH officially releases MOTSP/MOCVRP only (no Bi-KP/Tri-TSP checkpoints; Appendix H). The official mix constructor is Appendix D. Leftover shares are Figure 4; the reported staircase and the unconstrained-W_{all} ablation versus restricted $W|_R$ are Table 3 and Figure 4; cross-constructor leftover shares are Appendix Tables A.33–A.41; mix-constructor PAS order reversals are Table A.6; official constructor order reversals are Figure 1 and Table A.7; T-stage order reversals are Table A.8; extractor-to-gap ratios are Tables A.38–A.40; hold choices are Table A.52; $T(A_{\text{WS}})$ is Table 2, with $W_{\text{all}}(A_{\text{WS}})$ in Appendix Table A.20; $A_{\text{off}}(T(P))$ is Appendix Table A.21; $T(A_{\text{off}})$ versus $A_{\text{off}}(T(P))$ is Appendix Table A.22. Computational resource and training-budget comparisons versus paper† WS-LKH and paper† POCCO-W are Appendix E.

Official WE-CA, PMOCO, NHDE-P, CNH, and GIMF-P checkpoints and the WE-CA testdata, including official `kro*.pt` transforms, are the public releases cited in Section 5 and Appendix H. Evaluation scripts, canonical JSON tables, and figure code are bundled with the OpenReview supplement for anonymous review (Appendix H); a camera-ready repository follows acceptance. KroAB official rows are n=1 and have no Wilcoxon.

REFERENCES

Karl Bringmann, Tobias Friedrich, and Patrick Klitzke. Generic postprocessing via subset selection for hypervolume and epsilon-indicator. In *Parallel Problem Solving from Nature – PPSN XIII*, volume 8672 of *Lecture Notes in Computer Science*, pp. 518–527. Springer, 2014. doi: 10.1007/978-3-319-10762-2_51.

Jinbiao Chen, Zizhen Zhang, Zhiguang Cao, Yaoxin Wu, Yining Ma, Te Ye, and Jiahai Wang. Neural multi-objective combinatorial optimization with diversity enhancement. In *Advances in Neural Information Processing Systems*, volume 36, pp. 39176–39188, 2023. doi: 10.52202/075280-1702. URL `https://proceedings.neurips.cc/paper_files/paper/2023/hash/7b5ae891000049b91b3b62de596b1560-Abstract-Conference.html`.

Jinbiao Chen, Zhiguang Cao, Jiahai Wang, Yaoxin Wu, Hanzhang Qin, Zizhen Zhang, and Yue-Jiao Gong. Rethinking neural multi-objective combinatorial optimization via neat weight embedding. In *International Conference on Learning Representations*, 2025a. URL `https://openreview.net/forum?id=GM7cmQfk2F`.

Jinbiao Chen, Jiahai Wang, Zhiguang Cao, and Yaoxin Wu. Neural multi-objective combinatorial optimization via graph-image multimodal fusion. In *International Conference on Learning Representations*, 2025b. URL https://openreview.net/forum?id=4sJ2FYE65U.

Indraneel Das and J. E. Dennis. A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Structural Optimization*, 14(1):63–69, 1997. doi: 10.1007/BF01197559.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi: 10.1109/4235.996017.

Mingfeng Fan, Yaoxin Wu, Zhiguang Cao, Wen Song, Guillaume Sartoretti, Huan Liu, and Guohua Wu. Conditional neural heuristic for multiobjective vehicle routing problems. *IEEE Transactions on Neural Networks and Learning Systems*, 36(3):4677–4689, 2025a. doi: 10.1109/TNNLS.2024.3371706.

Mingfeng Fan, Jianan Zhou, Yifeng Zhang, Yaoxin Wu, Jinbiao Chen, and Guillaume Sartoretti. Preference-driven multi-objective combinatorial optimization with conditional computation. In *Advances in Neural Information Processing Systems*, volume 38, pp. 12931–12964, 2025b. doi: 10.52202/085713-0389. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/10fc83943b4540a9524af6fc67a23fef-Abstract-Conference.html.

Mark Fleischer. The measure of Pareto optima: Applications to multi-objective metaheuristics. In *International Conference on Evolutionary Multi-Criterion Optimization*, volume 2632 of *Lecture Notes in Computer Science*, pp. 519–533. Springer, 2003. doi: 10.1007/3-540-36970-8_37.

Andreia P. Guerreiro, Carlos M. Fonseca, and Luís Paquete. Greedy hypervolume subset selection in low dimensions. *Evolutionary Computation*, 24(3):521–544, 2016. doi: 10.1162/EVCO_a_00188.

Hisao Ishibuchi, Ryo Imada, Yu Setoguchi, and Yusuke Nojima. How to specify a reference point in hypervolume calculation for fair performance comparison. *Evolutionary Computation*, 26(3):411–440, 2018. doi: 10.1162/evco_a_00226.

Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ByxBFsRqYm.

Tobias Kuhn, Carlos M. Fonseca, Luís Paquete, Stefan Ruzika, Miguel M. Duarte, and José Rui Figueira. Hypervolume subset selection in two dimensions: Formulations and algorithms. *Evolutionary Computation*, 24(3):411–425, 2016. doi: 10.1162/EVCO_a_00157.

Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. POMO: Policy optimization with multiple optima for reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/f231f2107df69eab0a3862d50018a9b2-Abstract.html.

Kaiwen Li, Tao Zhang, and Rui Wang. Deep reinforcement learning for multiobjective optimization. *IEEE Transactions on Cybernetics*, 51(6):3103–3114, 2021. doi: 10.1109/TCYB.2020.2977661.

Xi Lin, Zhiyuan Yang, and Qingfu Zhang. Pareto set learning for neural multi-objective combinatorial optimization. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=QuObT9BTWo.

Yongfan Lu, Zixiang Di, Bingdong Li, Shengcai Liu, Hong Qian, Peng Yang, Ke Tang, and Aimin Zhou. Context-aware diversity enhancement for neural multi-objective combinatorial optimization. *arXiv preprint arXiv:2405.08604*, 2024. URL https://arxiv.org/abs/2405.08604.

Thibaut Lust and Jacques Teghem. Two-phase Pareto local search for the biobjective traveling salesman problem. *Journal of Heuristics*, 16(3):475–510, 2010. doi: 10.1007/s10732-009-9103-9.

Kaisa Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, 1999. doi: 10.1007/978-1-4615-5563-6.

Ryoji Tanabe, Hisao Ishibuchi, and Akira Oyama. Benchmarking multi- and many-objective evolutionary algorithms under two optimization scenarios. *IEEE Access*, 5:12097–12126, 2017. doi: 10.1109/ACCESS.2017.2717482.

Xuan Wu, Jinbiao Chen, Yang Li, Lijie Wen, Chunguo Wu, Yuanshu Li, Yubin Xiao, Chunyan Miao, You Zhou, and Di Wang. WeCon: An efficient weight-conditioned neural solver for multi-objective combinatorial optimization problems. *arXiv preprint arXiv:2605.22876*, 2026. URL `https://arxiv.org/abs/2605.22876`.

Qingfu Zhang and Hui Li. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007. doi: 10.1109/TEVC.2007.892759.

Eckart Zitzler, Lothar Thiele, Marco Laumanns, Carlos M. Fonseca, and Viviane Grunert da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003. doi: 10.1109/TEVC.2003.810758.

A NOTATION

Table A.1: Main symbols.

Symbol	Meaning
P, P_k, s	decoded $K \times n$ pool, its λ_k slice, start index
n	instance size (cities or items) and matching POMO starts
f, $\mathcal{X}$, m	additive cost; discrete feasible set; number of objectives
$\mathcal{N}(x)$	family-feasible neighbourhood (distinct from $N = \mathrm{ND}(P)$)
A_{off}, A_{WS}, A_{PAS}	official archive; weighted-sum archive; PAS archive
r, r^*	admissible hypervolume reference; that paper's official r
Σ^{m-1}	preference simplex $\{\lambda \geq 0 : \|\lambda\|_1 = 1\}$
$\mathcal{L}_{\mathrm{WS}}(\theta)$, $J_K(P)$	scalarized training loss $\mathbb{E}[\lambda^\top f]$; max HV_r of a K-subset of $\mathrm{ND}(P)$
PAS	recommended set-level extractor: $\mathrm{GreedyHV}_{K,r} \circ \mathrm{Floor}_K \circ \mathrm{ND}$
U2N, Γ	post-hoc local refinement (Eq. (5))
T, P2O	dominance-monotone refinement; TSP instance is Pareto 2-opt
W_λ, $W\|_R$, W_{all}	source-λ search; residual restriction; unconstrained ablation
$R(T, N)$	unmodified residual of T on N (distinct from $\mathcal{R}_{\mathrm{cand}}$)
Floor_K	cardinality restoration before greedy HV (Theorem 1)
$\mathcal{R}_{\mathrm{cand}}$	Floor restoration candidates (distinct from R)
Γ_{hold}	cascade over Γ that preserves unique-ND count relative to A_{PAS} (Eq. (6))
Δ, leftover	recovered $\Delta(A, A_{\mathrm{off}})$; Trunc. $= \mathrm{HV}_r(A_{\mathrm{PAS}}) - \mathrm{HV}_r(A_{\mathrm{off}})$
NDS	instance-mean unique-ND of scored A_{off}/A_{U2N}, not $\|\mathrm{ND}(P)\|$
WTL, d_z	win–tie–loss of A_{U2N} vs. A_{off}; paired effect size
Proto	official extractor of that row: WS / TCH / PBI / EP

Algorithm A.1 The audit map U2N Γ (Figure 3; Eq. (5))

Input: Candidate set $P = \{(x(\lambda_k, s), \lambda_k)\}$, reference r, K (101 if m=2; 105 if m=3), $\varepsilon = 10^{-12}$
Output: Archive A with $|A| = K$ (padding does not change hypervolume)

$N \leftarrow \mathrm{UniqueND}(P)$; $N_0 \leftarrow N$; $R \leftarrow \emptyset$ ▷ R: index set of locally unmodified points; each x keeps $\lambda(x)$
for $i = 1, \ldots, |N|$ **do**
  $x' \leftarrow \textsc{FirstImprovement}(N_i)$ ▷ accept iff $x' \preceq N_i$; else keep N_i
  if $\|f(x') - f(N_i)\|_\infty \leq \varepsilon$ **then**
    $R \leftarrow R \cup \{i\}$ ▷ locally f-efficient under $\mathcal{N}_\preceq$
  end if
  $N_i \leftarrow x'$ ▷ inherit decode-source $\lambda(N_i)$
end for
for $i \in R$ **do**
  $N_i \leftarrow \textsc{SourceLambdaSearch}(N_i, \lambda(N_i))$ ▷ $W_{\lambda(N_i)}$ on R only
end for
$Q \leftarrow N$ ▷ multiset after T and $W|_R$
$N \leftarrow \mathrm{UniqueND}(Q)$
if $|N| < \min(K, |N_0|)$ **then**
  $N \leftarrow \textsc{FloorRestore}(N, N_0, Q)$ ▷ Theorem 1: exclusive-HV restore from $\mathcal{R}_{\mathrm{cand}}$
end if
return $\textsc{GreedyHV}(N, K, r)$

Constrained Bi-CVRP / Bi-KP rows apply Γ_{hold} (Eq. (6)) to that returned archive.

B FINITE-K EVALUATION CENSUS

Table A.2: Full protocol-audit HV archive (same official protocol; complement of Table 1). Panel A (56 finite-K). Trunc. $= \Delta(A_{\mathrm{PAS}}, A_{\mathrm{off}})$; other HV absolute; A_{U2N} is the proposed refinement. Panel B (unbounded EP) is not a finite-K ranking claim. † $\mathrm{HV}_r(A_{\mathrm{PAS}}) - \mathrm{HV}_r(\mathrm{EP})$. NDS is the instance-mean unique-ND of the scored archives (A_{off}/A_{U2N}), not $|\mathrm{ND}(P)|$; KroAB (n=1) excluded from 46 Wilcoxon tests (Table A.4).

Ctor	Problem	Proto	A_{off}	A_{PAS}	Trunc.	A_{U2N}	NDS
Panel A: finite-K extractors.							
WE-CA	Bi-TSP20	WS	0.62704	0.62860	+0.00156	0.62873	24/47
	Bi-TSP50	WS	0.63905	0.64103	+0.00198	0.64217	55/101
	Bi-TSP100	WS	0.70386	0.70512	+0.00126	0.70678	77/101
	Bi-TSP150	WS	0.69912	0.70028	+0.00116	0.70278	85/101
	Bi-TSP200	WS	0.72638	0.72733	+0.00095	0.73134	86/101
	KroAB100	WS	0.69585	0.69718	+0.00133	0.69915	81/101
	KroAB150	WS	0.68982	0.69135	+0.00153	0.69387	87/101
	KroAB200	WS	0.72324	0.72418	+0.00094	0.72789	94/101
	Bi-CVRP20	WS	0.42941	0.42959	+0.00018	0.42975	9/14
	Bi-CVRP50	WS	0.40904	0.40922	+0.00018	0.40940	11/20
	Bi-CVRP100	WS	0.40723	0.40737	+0.00014	0.40758	12/23
	Bi-KP50	WS	0.35608	0.35667	+0.00059	0.35669	19.9/47.5
	Bi-KP100	WS	0.45333	0.45445	+0.00112	0.45446	32.1/93.9
	Bi-KP200	WS	0.36013	0.36097	+0.00084	0.36097	47.5/100.9
	Tri-TSP20	WS	0.47067	0.47446	+0.00379	0.47467	75.3/105
	Tri-TSP50	WS	0.43893	0.44212	+0.00319	0.44537	104.1/105
	Tri-TSP100	WS	0.49759	0.49952	+0.00193	0.50477	104.9/105
PMOCO	Bi-TSP20	TCH	0.50536	0.50681	+0.00145	0.50718	32/48
	Bi-TSP50	TCH	0.56704	0.57001	+0.00297	0.57197	59/101
	Bi-TSP100	TCH	0.66640	0.66878	+0.00238	0.67173	75/101
	KroAB100	TCH	0.65850	0.66114	+0.00264	0.66471	75/101
	Bi-CVRP20	TCH	0.21447	0.21561	+0.00114	0.21721	10/11
	Bi-CVRP50	TCH	0.25082	0.25202	+0.00120	0.25285	10/14
	Bi-CVRP100	TCH	0.23815	0.23935	+0.00120	0.24020	10/14
	Bi-KP50	TCH	0.75382	0.75382	0	0.75840	12/12
	Bi-KP100	TCH	0.95033	0.95033	0	0.95197	19/19
	Bi-KP200	TCH	0.68080	0.68080	0	0.68153	21/21
	Tri-TSP20	PBI	0.32044	0.32610	+0.00566	0.32788	71/105
	Tri-TSP50	PBI	0.33649	0.34145	+0.00496	0.34966	103/105
	Tri-TSP100	PBI	0.43965	0.44246	+0.00281	0.45275	105/105
CNH	Bi-TSP20	TCH	0.50588	0.50739	+0.00151	0.50770	33/50
	Bi-TSP50	TCH	0.56854	0.57126	+0.00272	0.57297	59/101
	Bi-TSP100	TCH	0.66809	0.67027	+0.00218	0.67290	73/101
	Bi-TSP150	TCH	0.71597	0.71783	+0.00186	0.72045	80/101
	Bi-TSP200	TCH	0.74376	0.74544	+0.00168	0.74842	83/101
	KroAB100	TCH	0.66093	0.66290	+0.00197	0.66563	73/101
	KroAB150	TCH	0.70850	0.71076	+0.00226	0.71359	84/101
	KroAB200	TCH	0.74025	0.74199	+0.00174	0.74525	91/101
	Bi-CVRP20	TCH	0.22574	0.22635	+0.00061	0.22689	10/13
	Bi-CVRP50	TCH	0.26444	0.26498	+0.00054	0.26519	9/18
	Bi-CVRP100	TCH	0.25095	0.25163	+0.00068	0.25198	8/18
GIMF-P	Bi-TSP20	WS	0.62660	0.62868	+0.00208	0.62888	21/50
	Bi-TSP50	WS	0.63744	0.63953	+0.00209	0.64139	49/101
	Bi-TSP100	WS	0.70060	0.70187	+0.00127	0.70422	73/101
	Bi-TSP150	WS	0.69582	0.69705	+0.00123	0.70018	81/101
	Bi-TSP200	WS	0.72667	0.72762	+0.00095	0.73154	84/101
	KroAB100	WS	0.69360	0.69473	+0.00113	0.69719	70/101
	KroAB150	WS	0.68777	0.68902	+0.00125	0.69226	77/101
	KroAB200	WS	0.72315	0.72422	+0.00107	0.72827	89/101
	Bi-CVRP20	WS	0.42877	0.42905	+0.00028	0.42928	7/12
	Bi-CVRP50	WS	0.40768	0.40793	+0.00025	0.40826	9/16
	Bi-CVRP100	WS	0.40515	0.40531	+0.00016	0.40571	10/17
	Bi-KP50	WS	0.35604	0.35664	+0.00060	0.35665	18.9/46.0
	Bi-KP100	WS	0.45333	0.45446	+0.00113	0.45447	32.8/94.1
	Tri-TSP20	WS	0.47020	0.47419	+0.00399	0.47488	73.8/105
	Tri-TSP50	WS	0.43536	0.43872	+0.00336	0.44412	103.0/105
Panel B: unbounded EP; Trunc. $^{\dagger} = \mathrm{HV}_r(A_{\mathrm{PAS}}) - \mathrm{HV}_r(\mathrm{EP})$.							
NHDE-P	Bi-TSP20	EP	0.62863	0.62863	0	0.62879	55/101
	Bi-TSP50	EP	0.63886	0.63885	−0.00001	0.64046	128/101
	Bi-TSP100	EP	0.70059	0.70045	−0.00014	0.70296	191/101
	Bi-TSP150	EP	0.69547	0.69522	−0.00025	0.69867	230/101
	Bi-TSP200	EP	0.71782	0.71764	−0.00018	0.72431	216/101
	KroAB100	EP	0.69327	0.69316	−0.00011	0.69574	177/101
	KroAB150	EP	0.68569	0.68546	−0.00023	0.68913	211/101

Table A.2: Full protocol-audit HV archive (continued).

Ctor	Problem	Proto	A_{off}	A_{PAS}	Trunc.	A_{U2N}	NDS
	KroAB200	EP	0.71463	0.71443	-0.00020	0.72143	226/101
	Bi-CVRP20	EP	0.42851	0.42851	0	0.42886	12/101
	Bi-CVRP50	EP	0.40504	0.40504	0	0.40562	14/101
	Bi-CVRP100	EP	0.39340	0.39340	0	0.39442	10/101
	Bi-KP50	EP	0.35535	0.35535	0	0.35559	12/101
	Bi-KP100	EP	0.45204	0.45204	0	0.45235	22/101
	Tri-TSP20	EP	0.47579	0.47442	$-0.00137^{\dagger}$	0.47507	889/105
	Tri-TSP50	EP	0.45055	0.44095	$-0.00960^{\dagger}$	0.44530	2549/105
	Tri-TSP100	EP	0.51105	0.49629	$-0.01476^{\dagger}$	0.50379	4978/105

Table A.3 enumerates the 56 official finite-K settings of Table A.2. PAS versus the official archive is already strictly positive on 53 settings, and a tie on the three sparse non-dominated PMOCO Bi-KP settings. U2N leftover versus A_{off} is strictly positive on all 56 settings; residual is the majority of leftover on $32/56$. At five-decimal Table A.2 precision, A_{U2N} equals PAS on WE-CA Bi-KP200. On the 46 multi-instance non-Kro settings, one-sided Wilcoxon U2N>PAS has $p<0.01$ on all 46 (WE-CA Bi-KP200 is a five-decimal mean tie with $p=6.5\times10^{-7}$). PAS>Off. has $p<0.01$ on every nonzero-PAS setting ($43/46$); the three PMOCO Bi-KP settings are all ties (p undefined). U2N>Off. has $p<0.01$ on all 46. U2N leftover is strictly positive on all 56.

Table A.3: Census of the 56 official finite-K settings. PAS>Off. on $53/56$; U2N>Off. on all 56; residual majority on $32/56$.

Ctor	Proto	Cells	#	PAS>Off.	U2N>Off.
WE-CA	WS	Bi-TSP 20–200, KroAB 100–200, Bi-CVRP 20–100, Bi-KP 50–200, Tri-TSP 20–100	17	17	17
GIMF-P	WS	Bi-TSP 20–200, KroAB 100–200, Bi-CVRP 20–100, Bi-KP 50/100, Tri-TSP 20/50	15	15	15
PMOCO	TCH	Bi-TSP 20–100, KroAB100, Bi-CVRP 20–100, Bi-KP 50–200	10	7	10
PMOCO	PBI	Tri-TSP 20–100	3	3	3
CNH	TCH	Bi-TSP 20–200, KroAB 100–200, Bi-CVRP 20–100	11	11	11
Total			56	53	56

Table A.4: A_{U2N} vs. A_{off}: paired d_z and win–tie–loss moved off Table A.2. WE-CA and GIMF-P multi-instance settings are 200–0–0 (d_z as in the previous main table). KroAB (n=1) has no test. Only non-unanimous settings are listed.

Ctor	Cell	d_z	WTL
PMOCO	Bi-TSP20	1.66	97–0–3
	Bi-CVRP20	0.88	99–1–0
	Bi-KP50/100/200	0.84/0.98/1.27	139–61–0 / 160–40–0 / 188–12–0
CNH	Bi-TSP20/50	2.09/2.91	99–0–1
NHDE-P	Bi-TSP20	0.47	148–0–52
	Bi-CVRP20/50/100	0.62/0.51/0.62	185–14–1 / 191–9–0 / 192–8–0
	Bi-KP50/100	0.80/0.92	154–46–0 / 185–15–0
	Tri-TSP20/50/100†	$-1.35/-4.57/-7.99$	18–0–182 / 0–0–200 / 0–0–200

C Proofs of theoretical guarantees

The arguments below record the attainment-region consequences of dominance-monotone refinement and of unrestricted source-weight search. They locate the finite WS-argmin interface of neural MOCO relative to classical hypervolume geometry (Zitzler et al., 2003; Bringmann et al., 2014; Das & Dennis, 1997). Throughout, $\mathcal{X}$ is finite and discrete, $f:\mathcal{X}\to\mathbb{R}^m$ is the vector cost, and r is an admissible reference with $f(x)\le r$ entrywise on every evaluated set. For finite $S\subset\mathcal{X}$, $\mathcal{B}_r(S)=\bigcup_{x\in S}\{q: f(x)\le q\le r\}$ and $\mathrm{HV}_r(S)=\mathrm{vol}_m(\mathcal{B}_r(S))$ (Equation (1)).

C.1 Attainment-region inclusion

Proof of Lemma 1. Let $q\in\mathcal{B}_r(S)$. By Equation (1) there exists $x\in S$ with $f(x)\le q\le r$ entrywise. Dominance-monotonicity gives $T(x)\preceq x$, hence $f_i(T(x))\le f_i(x)\le q_i\le r_i$ for every coordinate $i\in\{1,\dots,m\}$. Thus $q\in\{q': f(T(x))\le q'\le r\}\subseteq\mathcal{B}_r(T(S))$, which is $\mathcal{B}_r(S)\subseteq\mathcal{B}_r(T(S))$.

Now let $y\in S$ be dominated: some $x^*\in S$ satisfies $x^*\prec y$, so $f(x^*)\le f(y)$ with at least one strict coordinate. The closed box $\{q: f(y)\le q\le r\}$ is then contained in $\{q: f(x^*)\le q\le r\}$. Removing dominated points therefore leaves both the union of boxes and its Lebesgue measure unchanged,

$$\mathcal{B}_r(S)=\bigcup_{x\in S}\{q: f(x)\le q\le r\}=\bigcup_{x\in\mathrm{ND}(S)}\{q: f(x)\le q\le r\}=\mathcal{B}_r(\mathrm{ND}(S)),$$

so $\mathrm{HV}_r(S) = \mathrm{HV}_r(\mathrm{ND}(S))$.

C.2 Hypervolume monotonicity of pointwise dominance

Proof of Lemma 2. Lemma 1 yields $\mathcal{B}_r(S) \subseteq \mathcal{B}_r(T(S))$. Monotonicity of Lebesgue measure ($A \subseteq B$ implies $\mathrm{vol}_m(A) \leq \mathrm{vol}_m(B)$) then gives

$$\mathrm{HV}_r(S) = \mathrm{vol}_m(\mathcal{B}_r(S)) \leq \mathrm{vol}_m(\mathcal{B}_r(T(S))) = \mathrm{HV}_r(T(S)).$$

The second claim of Lemma 1 replaces each side by its unique non-dominated set without changing the measure:

$$\mathrm{HV}_r(\mathrm{ND}(S)) = \mathrm{HV}_r(S) \leq \mathrm{HV}_r(T(S)) = \mathrm{HV}_r(\mathrm{ND}(T(S))).$$

Hence pointwise weak dominance is hypervolume-monotone. The inequality is non-decreasing, not strictly increasing: T may return x unchanged.

C.3 Source-weight search is not dominance-monotone

Proof of Lemma 3. Take $m = 2$, $f(x) = (1, 1)^\top$, $\lambda = (0.9, 0.1)^\top$, and a neighbour $x' \in \mathcal{N}(x)$ with $f(x') = (0.8, 2.0)^\top$. Then $\lambda^\top f(x) = 1$ and $\lambda^\top f(x') = 0.9 \cdot 0.8 + 0.1 \cdot 2.0 = 0.92 < 1$, so W_λ accepts x'. The second coordinate rises ($2.0 > 1$), hence $x' \not\preceq x$ and $\mathcal{N}_\preceq$ rejects x'. Source-weight search is therefore not dominance-monotone.

C.4 Unconstrained search need not be hypervolume-monotone

Proof of Proposition 2. Reuse the costs of Lemma 3: $f(x) = (1, 1)^\top$, $\lambda = (0.9, 0.1)^\top$, and a neighbour x' with $f(x') = (0.8, 2.0)^\top$, so W_λ accepts x'. Take $S = \{x\}$ and $r = (3, 3)^\top$. Then $\mathrm{HV}_r(S) = (3 - 1)(3 - 1) = 4$ and $\mathrm{HV}_r(\{x'\}) = (3 - 0.8)(3 - 2.0) = 2.2 < 4$. Hence W_{all} need not be hypervolume-monotone. The same operator can raise volume on other instances (Table 3); the claim is existence of a drop, not a universal decrease.

C.5 Cardinality restoration by invariant padding

Proof of Theorem 1. Write $N_0 = \mathrm{ND}(P)$ and let Q be a finite multiset ($Q = N_0$ on PAS; after T and $W|_R$ on U2N). Set $N_1 = \mathrm{ND}(Q)$ and $\mathcal{R}_{\mathrm{cand}} = \{y \in N_0 \cup Q : \nexists z \in N_1,\ z \prec y\}$. Write $\kappa = \min(K, |N_0|)$. Then $N_1 \subseteq \mathcal{R}_{\mathrm{cand}}$, so the restoration pool is $\mathcal{R}_{\mathrm{cand}} \setminus N_1$. By construction, every restored point of $\mathcal{R}_{\mathrm{cand}} \setminus N_1$ that was moved by T remains weakly dominating its pre-move cost, so Lemma 1 keeps $\mathcal{B}_r(\mathcal{R}_{\mathrm{cand}}) \supseteq \mathcal{B}_r(N_0)$ when the whole of $\mathcal{R}_{\mathrm{cand}}$ is restored. If $|N_1| < \kappa \leq |\mathcal{R}_{\mathrm{cand}}|$, restoring $\kappa - |N_1|$ points of $\mathcal{R}_{\mathrm{cand}} \setminus N_1$ yields a base set F_{base} with $|F_{\mathrm{base}}| = \kappa$. When more candidates are available, the implementation restores in exclusive-hypervolume order at r (official first-index ties). That step restores cardinality toward $\min(K, |N_0|)$; it does not restore the full attainment of $\mathcal{R}_{\mathrm{cand}}$. If $\kappa < K$, pad F_{base} by $K - \kappa$ duplicate objective vectors. A duplicate x_{dup} has the same closed box as its source, and $A \cup A = A$, so

$$\mathcal{B}_r(F) = \mathcal{B}_r(F_{\mathrm{base}}).$$

Duplicate padding therefore contributes no Lebesgue measure and enforces $|F| = K$. The unique non-dominated set need not reach K: official WE-CA Bi-TSP20 retains unique-ND size 24 with $|A| = K$ (Table A.2).

D Extended benchmark evaluations

Diagnostic tables in this appendix are empirical stress-tests of already-decoded candidate pools under the official evaluations.

Table A.5 evaluates archive recovery on the official WE-CA-U mix constructor on the official WE-CA testdata. Reported leftover follows the official WE-CA WS map: the audit map Γ on Bi-TSP / Tri-TSP20/50, the cardinality-preserving projection Γ_{hold} on Bi-CVRP, and Γ on Tri-TSP100. That map raises same-pool hypervolume versus mix WS under the unified 101λ+WS benchmark protocol. The size-specific constructor remains the primary column: mix Bi-TSP100 U2N 0.70641 sits below the official WE-CA setting of Table A.2 0.70678. Table A.5 reports leftover of the official mix constructor under the official WE-CA WS protocol. On mix Bi-CVRP the reported method is Γ_{hold}: leftover $+0.00049/+0.00037/+0.00042$ (U2N / T / PAS $181/4/15$, $160/0/40$, $183/0/17$; non-dominated size held or enlarged; win–tie–loss 200–0–0). Table A.6 records the same-testdata ranking of the size-specific constructor against the mix constructor. Official WS ranks stay almost identical (Spearman ≈ 0.998 at n=20). Selecting an archive via U2N reverses instance-wise mix-versus size-specific order in $78/200$ cases at n=20 (39%) under that benchmark protocol. Re-extracting the

same pools from WS to U2N is therefore a same-pool re-ranking protocol: per-weight WS conceals structural variation inside the candidate pools even when mean ranks stay correlated.

Table A.5: Official WE-CA-U mix. Method = Γ on Bi-TSP / Tri-TSP20/50, Γ_{hold} on Bi-CVRP, and Γ on Tri-TSP100.

Problem	Mix WS	Method	Δ	C / occ
Bi-TSP20	0.62703	**0.62886**	+0.00183	0.392 / 88.2%
Bi-TSP50	0.63921	**0.64227**	+0.00306	0.483 / 19.4%
Bi-TSP100	0.70344	**0.70641**	+0.00297	0.706 / 0.7%
Bi-TSP150	0.70083	**0.70417**	+0.00335	0.780 / 0
Bi-TSP200	0.73462	**0.73769**	+0.00306	0.805 / 0
Tri-TSP20	0.47069	**0.47488**	+0.00419	—
Tri-TSP50	0.43888	**0.44536**	+0.00648	—
Tri-TSP100	0.49748	**0.50472**	+0.00725	—
Bi-CVRP20	0.42903	**0.42952**	+0.00049	0.500 / 62.0%
Bi-CVRP50	0.40898	**0.40935**	+0.00037	0.626 / 39.9%
Bi-CVRP100	0.40691	**0.40732**	+0.00042	0.718 / 29.1%

Table A.6: Official WE-CA testdata, official WS extractor: size-specific constructor versus mix on identical already-decoded pools. PAS / T / U2N order reversal = instances whose mix-versus-size-specific order reverses from WS to that stage. T uses seed-matched mix evaluation runs that record intermediate T hypervolume. At n=20, PAS already reverses $67/200$ and the aggregate winner; T is $69/200$; U2N is $78/200$.

n	WS	PAS	T	U2N	PAS flip	T flip	U2N flip	Spearman
20	official	mix	mix	mix	$67/200$	$69/200$	$78/200$	0.998
50	mix	mix	mix	mix	$52/200$	$54/200$	$54/200$	0.996
100	official	official	official	official	$21/200$	$31/200$	$31/200$	0.994
150	mix	mix	mix	mix	$1/200$	$1/200$	$2/200$	0.993
200	mix	mix	mix	mix	$0/200$	$0/200$	$0/200$	0.978

Table A.7: Official-protocol constructor ranking on shared testdata and r^*. WE-CA–GIMF-P uses official WS and the WE-CA testdata (Bi-TSP/CVRP/KP and leftover Tri n=20/50/100; Tri100 leftover+T is Appendix Table A.36, not a census row). PMOCO–CNH uses official TCH, seed1234, and that paper's official r^*; there is no official PMOCO n>100 pair. KroAB is n=1 (Spearman undefined). Order reversal = instances whose pairwise order reverses from the official extractor to PAS / U2N.

Pair	Cell	Off. win.	PAS win.	U2N win.	PAS flip	U2N flip	Spearman
WE-CA–GIMF	Bi-TSP20	WE-CA	GIMF	GIMF	$73/200$	$90/200$	0.998
	Bi-TSP50	WE-CA	WE-CA	WE-CA	$13/200$	$29/200$	0.995
	Bi-TSP100	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	0.993
	Bi-TSP150	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	0.991
	Bi-TSP200	GIMF	GIMF	GIMF	$16/200$	$36/200$	0.981
	Bi-CVRP20	WE-CA	WE-CA	WE-CA	$25/200$	$45/200$	0.999
	Bi-CVRP50	WE-CA	WE-CA	WE-CA	$13/200$	$20/200$	0.999
	Bi-CVRP100	WE-CA	WE-CA	WE-CA	$4/200$	$27/200$	0.999
	Bi-KP50	WE-CA	WE-CA	WE-CA	$80/200$	$77/200$	1.000
	Bi-KP100	GIMF	GIMF	GIMF	$76/200$	$78/200$	1.000
	Tri-TSP20	WE-CA	WE-CA	GIMF	$61/200$	$118/200$	0.999
	Tri-TSP50	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	0.998
	Tri-TSP100	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	0.997
	KroAB100	WE-CA	WE-CA	WE-CA	$0/1$	$0/1$	—
	KroAB150	WE-CA	WE-CA	WE-CA	$0/1$	$0/1$	—
	KroAB200	WE-CA	GIMF	GIMF	$1/1$	$1/1$	—
PMOCO–CNH	Bi-TSP20	CNH	CNH	CNH	$18/100$	$24/100$	0.999
	Bi-TSP50	CNH	CNH	CNH	$14/100$	$25/100$	0.993
	Bi-TSP100	CNH	CNH	CNH	$4/100$	$6/100$	0.992
	Bi-CVRP20	CNH	CNH	CNH	$1/100$	$3/100$	0.990
	Bi-CVRP50	CNH	CNH	CNH	$1/100$	$1/100$	0.995
	Bi-CVRP100	CNH	CNH	CNH	$0/100$	$0/100$	0.993
	KroAB100	CNH	CNH	CNH	$0/1$	$0/1$	—

Table A.9 repeats the mix re-extraction on Bi-CVRP, Tri-TSP, and KroAB. Table A.10 splits those reversals into PAS reversals that persist through T/U2N versus reversals that appear only later. At mix Bi-TSP20, all 67 PAS reversals persist to T and 55 persist to U2N; U2N adds 23 new. Mix Tri U2N

Table A.8: Official-protocol T-stage ranking on evaluation runs that record intermediate T statistics. Order reversal = official extractor → T / T →U2N. WE-CA Bi-KP T is the seed-matched recompile of the cardinality-preserving projection archive; U2N remains that hold archive. Bi-TSP$100/150$ have leftover but no close-pair order reversal at any stage. Tri100 leftover+T is not a 56-setting census row.

Pair	Cell	Off. win.	T win.	U2N win.	T flip	T →U2N	U2N flip
WE-CA–GIMF	Bi-TSP20	WE-CA	GIMF	GIMF	$76/200$	$40/200$	$90/200$
	Bi-TSP50	WE-CA	WE-CA	WE-CA	$15/200$	$26/200$	$29/200$
	Bi-TSP100	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	$0/200$
	Bi-TSP150	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	$0/200$
	Bi-TSP200	GIMF	GIMF	GIMF	$37/200$	$13/200$	$36/200$
	Bi-CVRP20	WE-CA	WE-CA	WE-CA	$32/200$	$35/200$	$45/200$
	Bi-CVRP50	WE-CA	WE-CA	WE-CA	$19/200$	$7/200$	$20/200$
	Bi-CVRP100	WE-CA	WE-CA	WE-CA	$22/200$	$9/200$	$27/200$
	Bi-KP50	WE-CA	WE-CA	WE-CA	$82/200$	$9/200$	$77/200$
	Bi-KP100	GIMF	GIMF	GIMF	$77/200$	$7/200$	$78/200$
	Tri-TSP20	WE-CA	WE-CA	GIMF	$64/200$	$104/200$	$118/200$
	Tri-TSP50	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	$0/200$
	Tri-TSP100	WE-CA	WE-CA	WE-CA	$0/200$	$0/200$	$0/200$
	KroAB100	WE-CA	WE-CA	WE-CA	$0/1$	$0/1$	$0/1$
	KroAB150	WE-CA	WE-CA	WE-CA	$0/1$	$0/1$	$0/1$
	KroAB200	WE-CA	GIMF	GIMF	$1/1$	$0/1$	$1/1$
PMOCO–CNH	Bi-TSP20	CNH	CNH	CNH	$20/100$	$16/100$	$24/100$
	Bi-TSP50	CNH	CNH	CNH	$13/100$	$22/100$	$25/100$
	Bi-TSP100	CNH	CNH	CNH	$5/100$	$5/100$	$6/100$
	Bi-CVRP20	CNH	CNH	CNH	$1/100$	$2/100$	$3/100$
	Bi-CVRP50	CNH	CNH	CNH	$1/100$	$0/100$	$1/100$
	Bi-CVRP100	CNH	CNH	CNH	$0/100$	$0/100$	$0/100$
	KroAB100	CNH	CNH	CNH	$0/1$	$0/1$	$0/1$

Table A.9: Mix re-extraction on the remaining official-WE-CA families in Table A.5 (same testdata, official WS). Order reversal = size-specific versus mix order from WS to PAS / T / U2N. T uses seed-matched mix evaluation runs. KroAB is n=1.

Cell	Off.	PAS	T	U2N	PAS flip	T flip	U2N flip
Bi-CVRP20	official	official	official	official	$28/200$	$32/200$	$47/200$
Bi-CVRP50	official	official	official	official	$13/200$	$17/200$	$21/200$
Bi-CVRP100	official	official	official	official	$9/200$	$17/200$	$22/200$
Tri-TSP20	mix	mix	mix	mix	$78/200$	$80/200$	$97/200$
Tri-TSP50	official	official	official	official	$67/200$	$62/200$	$90/200$
Tri-TSP100	official	official	official	official	$46/200$	$61/200$	$80/200$
KroAB100	official	official	official	official	$0/1$	$0/1$	$0/1$
KroAB150	mix	mix	mix	mix	$0/1$	$0/1$	$0/1$
KroAB200	mix	mix	mix	mix	$0/1$	$0/1$	$0/1$

adds 48–49 new reversals at every size. Official same-r^* persistence is Table A.11: WE-CA–GIMF-P KP is already decided at PAS (72–73 of 76–80 persist); Tri20 U2N adds 84 new. KroAB n=1 is omitted (no multi-instance persist). On the same decoded pool, the official, PAS, and U2N columns distinguish three signatures (Table A.12): a joint rise of A_{off} and A_{PAS} is generative capacity; a rise only on A_{off} is alignment to the discrete preference vectors Λ; a rise only on A_{PAS} or $\Gamma(P)$ is hidden pool capacity that the scalarized extractor truncated.

D.1 TSPLIB KRO PAIRS

Table A.13 repeats official KroAB150/200 next to the ten-pair family. Table A.14 details pairwise evaluations across all $\binom{5}{2} = 10$ TSPLIB Kro-100 bi-objective pairs under the same unit-max-axis reconstruction as official `kro100.pt` (Chen et al., 2025a), including the official KroAB100 row of Table A.2. Across the ten urban topologies, U2N records a consistent hypervolume improvement over official WS (mean $\Delta = +0.00312$; paired Wilcoxon $10/10$, $p = 9.8 \times 10^{-4}$).

Table A.10: Mix re-extraction persistence on seed-matched evaluation runs. Persist = a PAS order reversal that remains at T / U2N. New = a reversal that is absent at PAS and present at that stage. n_{test}=200. KroAB n=1 omitted.

Cell	PAS flip	persist T	persist U2N	new T	new U2N
Bi-TSP20	67/200	67/200	55/200	2/200	23/200
Bi-TSP50	52/200	42/200	33/200	12/200	21/200
Bi-TSP100	21/200	16/200	12/200	15/200	19/200
Bi-TSP150	1/200	1/200	1/200	0/200	1/200
Bi-TSP200	0/200	0/200	0/200	0/200	0/200
Bi-CVRP20	28/200	26/200	23/200	6/200	24/200
Bi-CVRP50	13/200	9/200	10/200	8/200	11/200
Bi-CVRP100	9/200	6/200	6/200	11/200	16/200
Tri-TSP20	78/200	69/200	49/200	11/200	48/200
Tri-TSP50	67/200	47/200	42/200	15/200	48/200
Tri-TSP100	46/200	34/200	31/200	27/200	49/200

Table A.11: Official same-r^* persistence. Pairs and evaluation runs match Tables A.7–A.8. WE-CA KP T is the seed-matched recompile. Tri100 leftover+T is not a 56-setting census row.

Pair	Cell	PAS	persist T	persist U2N	new T	new U2N
WE-CA–GIMF	Bi-TSP20	73/200	69/200	60/200	7/200	30/200
	Bi-TSP50	13/200	12/200	8/200	3/200	21/200
	Bi-TSP200	16/200	10/200	9/200	27/200	27/200
	Bi-CVRP20	25/200	25/200	18/200	7/200	27/200
	Bi-CVRP100	4/200	4/200	2/200	18/200	25/200
	Bi-KP50	80/200	78/200	72/200	4/200	5/200
	Bi-KP100	76/200	73/200	73/200	4/200	4/200
	Tri-TSP20	61/200	56/200	34/200	8/200	84/200
PMOCO–CNH	Bi-TSP20	18/100	17/100	14/100	3/100	10/100
	Bi-TSP50	14/100	10/100	7/100	3/100	18/100

Table A.12: Same-pool reading of official, PAS, and U2N columns.

Signature	Diagnosis	Reading
Off. and PAS/U2N rise	generative gain	capacity enlarged
Off. rises, PAS/U2N flat	alignment tuning	ray fit, not J_K
Off. flat, PAS/U2N rise	hidden pool capacity	pool truncation

Table A.13: Unified-WS Δ (diagnostic) under a common evaluation protocol on zero-shot KroAB150/200 from the official n=100 file (same settings as Table A.2). Δ is recovered hypervolume Ours−WS; NDS is non-dominated size $|\mathrm{ND}(\cdot)|$ (WS / Ours). KroAB is n=1.

Ctor	Problem	WS	PAS	Ours	Δ / NDS
WE-CA	KroAB150	0.68982	0.69135	**0.69387**	+0.00405 / 87/101
	KroAB200	0.72324	0.72418	**0.72789**	+0.00465 / 94/101
PMOCO	KroAB150	0.68561	0.68731	**0.69089**	+0.00527 / 77/101
	KroAB200	0.72417	0.72562	**0.72948**	+0.00531 / 79/101
NHDE-P	KroAB150	0.67234	0.67337	**0.68005**	+0.00771 / 73/101
	KroAB200	0.70794	0.70889	**0.71665**	+0.00871 / 81/101
CNH	KroAB150	0.68713	0.68897	**0.69211**	+0.00498 / 82/101
	KroAB200	0.72596	0.72749	**0.73094**	+0.00498 / 80/101
GIMF-P	KroAB150	0.68777	0.68902	**0.69226**	+0.00449 / 77/101
	KroAB200	0.72315	0.72422	**0.72827**	+0.00513 / 89/101

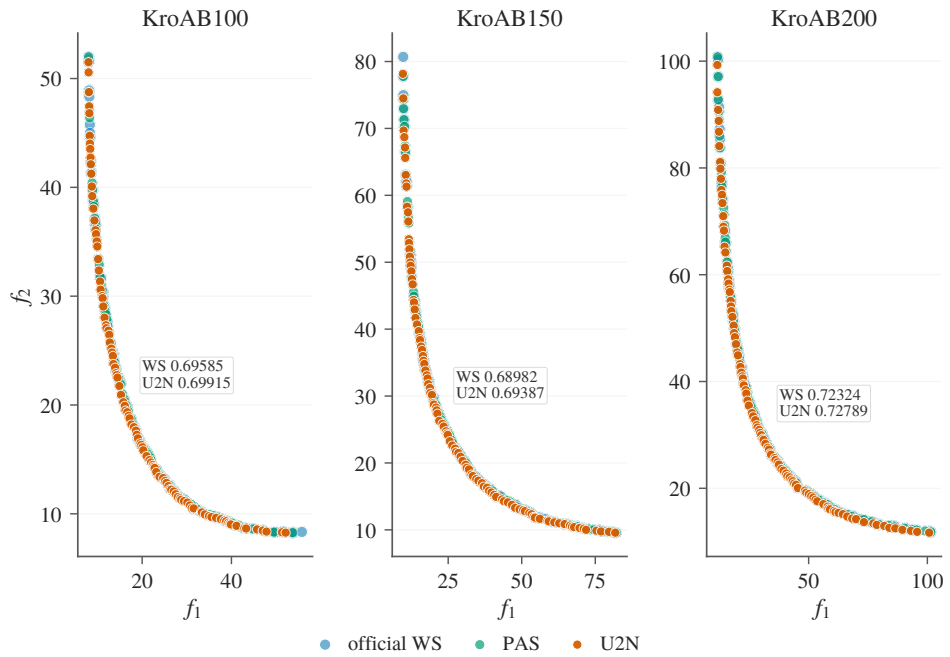


Figure A.1: Official KroAB fronts (n=1) on `kro*.pt` (Table A.2 has no Wilcoxon). Ten-pair Wilcoxon 10/10 is Table A.14.

Table A.14: Ten TSPLIB Kro-100 pairs under the official unit-max-axis reconstruction, official WE-CA+$U2N$ under a common evaluation protocol. KroAB100 is the official evaluation setting.

Pair	WS	U2N	Δ
KroAB100	0.69585	**0.69915**	+0.00329
KroAC100	0.69416	0.69724	+0.00308
KroAD100	0.70326	0.70624	+0.00298
KroAE100	0.70285	0.70571	+0.00286
KroBC100	0.69554	0.69872	+0.00318
KroBD100	0.69362	0.69713	+0.00351
KroBE100	0.68610	0.68912	+0.00302
KroCD100	0.71081	0.71383	+0.00303
KroCE100	0.70136	0.70431	+0.00295
KroDE100	0.69666	0.69998	+0.00332

E RESOURCE BUDGETS AND EVALUATION-PROTOCOL WALLS

U2N is a CPU post-process of an already-decoded pool: it adds no second neural pass and no extra training epoch. The official extractor is the already-incurred GPU decode ($9\,\mathrm{s}/100$); it is not a computational peer of subset selection. PAS-only is zero-search (10–$11\,\mathrm{min}/100$) and is the adoption default; it already exceeds that extract on $53/56$ settings. Full U2N (19–$23\,\mathrm{min}/100$) is the recovery of remaining M2 by residual $T{+}W|_R$ on tri-objective manifolds and sparse KP, where selection gain can be zero. Those minutes are evaluation-protocol post-processing on a frozen suite, not online-inference parity with the official extractor. A second greedy A_{off} pass on size-matched WE-CA and GIMF-P Bi-TSP20 reproduces the locked lists ($0/200$). An off-protocol softmax two-seed diagnostic flips $49/200$ official pairings; that is not the Figure 1(a) null. Spending U2N's extra 8–$12\,\mathrm{min}/100$ on more samples is a GPU decode, not CPU 2-opt: PAS $101 \to 201$ is $4\,\mathrm{min}/100$ and still $+1.39 \times 10^{-3}$ below U2N at 101 (Figure A.5). Official 64-fold extra starts $n{\to}4n$ on that pair take $13\,\mathrm{s}$ (WE-CA) and $22\,\mathrm{s}$ (GIMF-P) per 200 instances and lock at n ($73/200$); WS@$4n$ recovers 11% of leftover, and PAS@$4n$ can exceed locked U2N because it is extra decode. The $4n$ PAS constructor gap is 3×10^{-6} and is not a mean-rank claim. Table A.15 reports that same-pool split per 100 instances on the RTX 4090 node of Table A.2.

Table A.15: Per-100-instance wall-clock on the Table A.2 node (RTX 4090). Decode† is the official GPU pass, already required by the protocol. PAS-only and $\Delta(T{+}W|_R)$ are the same-pool stage-wall split (rounded; exact $11.47{+}11.76{=}23.22$, $10.45{+}8.49{=}18.94$, $10.57{+}9.17{=}19.74$, $10.86{+}9.47{=}20.33$).

Pipeline	Decode†	PAS-only	$\Delta(T{+}W\|_R)$	U2N total
WE-CA Bi-TSP100	9 s	11 min	+12 min	23 min
PMOCO Bi-TSP100	9 s	10 min	+9 min	19 min
CNH Bi-TSP100	9 s	11 min	+9 min	20 min
GIMF-P Bi-TSP100	9 s	11 min	+9 min	20 min

Residual $T{+}W|_R$ adds 8–$12\,\mathrm{min}/100$ and is 78.5% of official PMOCO Tri-TSP100 leftover. Those minutes are not a direct, head-to-head runtime benchmark against paper† WS-LKH. Published WS-LKH† is 6 h (Fan et al., 2025b) on a different node. Tables A.16–A.17 and A.18 are the 200-instance stage totals on the same node. The zero-search PAS stage is a greedy hypervolume subset on the already-decoded $K \times n$ pool and adds no second neural pass. A separate PAS-only evaluation pipeline (decode + unique non-dominated set + greedy HV, no 2-opt) is $21/46\,\mathrm{min}$ at $n{=}100/150$ on this node; PAS-only checkpoints at $n{=}50/200$ were not logged. The GPU decode ($5/17/49/81\,\mathrm{s}$ per 200 instances at $n{=}50/100/150/200$) is the official decode budget and is already inside both the PAS-only and U2N walls. The U2N walls $19/46/66/89\,\mathrm{min}$ on Bi-TSP50/100/150/200 (KroAB100: $19\,\mathrm{s}$) include decode and compile (the $n{=}50$ wall is the logged evaluation checkpoint; the size-matched checkpoint used for Table A.2 HV is $42\,\mathrm{min}$); comparing those checkpoints, U2N $-$ PAS-only is about $26/20\,\mathrm{min}$ at $n{=}100/150$. The 200-instance WE-CA totals remain Table A.17; Table A.15 uses the same-pool stage-wall split on official Bi-TSP100 pools. Per 100 instances the WE-CA wall is $23\,\mathrm{min}$; PMOCO/CNH/GIMF-P leftover-compile walls on this node (19–$20\,\mathrm{min}/100$) are the same compiler on those constructors, not the official extractor. Official Bi-CVRP Γ_{hold} on the same machine class is $20/10/20\,\mathrm{min}$ at $n{=}20/50/100$ ($n{=}20$ includes IMVND). Published WS-LKH times on the same-size Bi-TSP suites are 6–22 hours (Fan et al., 2025b); those hours are a different solver class, not a peer of the official extractor, and we do not use them as a runtime claim for PAS or U2N. Published POCCO-W† means at $n \in \{150, 200\}$ use 200 extra training epochs on 10^5 new instances (Fan et al., 2025b). The official extractor sits below those means.

Table A.16: Hypervolume and wall-clock versus paper† WS-LKH (Fan et al., 2025b). Decode† is a same-λ GPU pass already required by the official protocol and is included in both later columns. PAS-full is decode+greedy HV with no 2-opt, from a separate evaluation pipeline ($n{=}100/150$ only). U2N wall‡ is decode+PAS+$T{+}W|_R$ on one RTX 4090 node, 200 instances. LKH† hours are cited from that paper (A100+EPYC) and use a different node.

Problem	U2N HV	LKH† HV	Decode†	PAS-full	U2N wall‡	LKH†	Extra train
Bi-TSP100	0.7068	0.7090	17 s	21 min	46 min	6 h	0
Bi-TSP150	0.7028	0.7149	49 s	46 min	66 min	13 h	0
Bi-TSP200	0.7313	0.7490	81 s	—	89 min	22 h	0
KroAB100	0.6991	0.7022	—	—	19 s	2.3 min	0

Table A.17: Same-node leftover compile versus the official extractor pass. Rows are decode+compile on the RTX 4090 node that produced Table A.2. PMOCO/CNH use the 100-instance seed1234 evaluation; NHDE-P and WE-CA use 200 instances. Per-100 renormalizes the 200-instance walls. Tri-TSP $n \geq 50$ merge-only clocks are omitted.

Ctor	Problem	n_{test}	Wall	Per 100 inst.	Protocol
PMOCO	Bi-TSP100	100	19 min	19 min	official TCH
CNH	Bi-TSP100	100	20 min	20 min	official TCH
NHDE-P	Bi-TSP100	200	23 min	11 min	official EP
WE-CA	Bi-TSP100	200	46 min	23 min	official WS
PMOCO	Bi-TSP100	200	38 min	19 min	unified WS

Table A.18: Training budget of the official extractor. Extra training is 0 epochs. Published POCCO-W† uses 200 epochs $\times 10^5$ instances (Fan et al., 2025b); those paper means are not reported here and are not settings from the same evaluation protocol. Three official-init WE-CA Bi-TSP100 seeds drop the freeze from 0.70386 to mean 0.70315.

Problem	Frozen official U2N	Extra train
Bi-TSP150	0.7028	0 ep.
Bi-TSP200	0.7313	0 ep.
KroAB150	0.6939	0 ep.
KroAB200	0.7279	0 ep.

F Diagnostics on unconstrained search and training divergence

Table 3 and Figure A.2 record source-weight search without the R-restriction (unconstrained W_{all}). The stage staircase and leftover shares are Figure 4.

On official WE-CA Bi-TSP100, unconstrained W_{all} raises hypervolume from 0.70678 to 0.70764 while remaining at NDS 101. It lacks a monotone guarantee (Proposition 2) and can contract Pareto breadth. On n=50, unique non-dominated size under W_{all} is 99.43/82.52/76.84 on official WE-CA/PMOCO/NHDE-P (Table A.46): the search contracts Pareto breadth on the latter two constructors in exchange for a single projection. Restricting W_λ to the unmodified residual R therefore freezes candidates already advanced by T, so W does not apply scalarized moves to those points. Local refinement of A_{WS} and continued training under the same weighted-sum loss are recorded below.

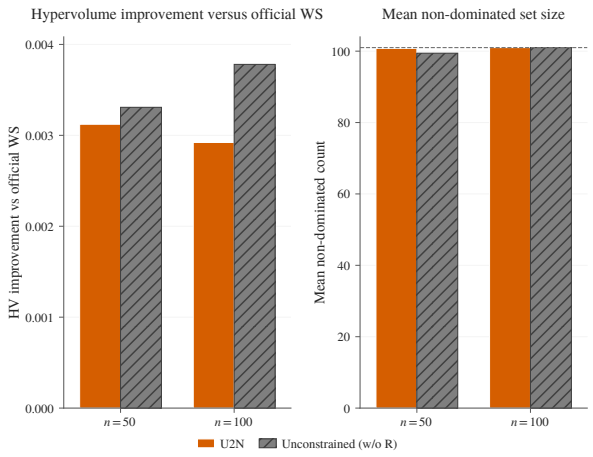


Figure A.2: Unconstrained W_{all} versus the cardinality-K archive on official WE-CA Bi-TSP (Proposition 2). Numbers are the W_{all} column of Table 3.

F.1 Reported-set residual

Table A.20 applies the same first-improvement operator that instantiates T, but only to the official WS survivors. The hypothesis is that leftover of P can be closed without reading the discarded unique-ND candidates. $T(A_{\text{WS}})$ recovers 2.5%, 9.0%, 22.3%, 40.9%, and 69.4% of the U2N leftover

at n=20/50/100/150/200 (paired Wilcoxon 0–0–200 against U2N at n=50/100). $W_{\text{all}}(A_{\text{WS}})$ is source-λ search on the same 101 reported solutions: at n=100 it exceeds PAS yet remains below U2N (mean 0.70619 versus 0.70678; WTL 12–0–188); at n=150/200 the means are 0.70395 and 0.73324 against U2N 0.70278 and 0.73134. Official WS already kept only 77 of 276 unique-ND candidates on Bi-TSP100, so the same 2-opt on A_{WS} cannot replace subset selection of P. Writing $T(\text{ND}(P))$ back into the decoded pool and re-running the official extractor, $A_{\text{off}}(T(P))$, recovers 2.1% of leftover on official WE-CA Bi-TSP20 and 82.2% at n=200; GIMF-P is 3.0%/80.1%. The four-constructor n=100 range is 31.7–43.0% (Table A.21). Table A.22 is the same leftover ruler for reported-set $T(A_{\text{off}})$ versus write-T-then-re-extract $A_{\text{off}}(T(P))$. At n=100, $A_{\text{off}}(T)$ is 34.2% on WE-CA and 43.0% on GIMF-P against $T(A_{\text{off}})$ 22.3% and 33.2%; on official TCH, $T(A_{\text{off}})$ is 39.3% against 33.8% (PMOCO) and 34.4% against 31.7% (CNH). At n=200, $T(A_{\text{off}})$ is 69.4% on WE-CA and 70.3% on GIMF-P against $A_{\text{off}}(T)$ 82.2% and 80.1%; CNH is 59.2% against 61.3%. KroAB n=1 repeats the ruler (Table A.23): official WS write-back is larger (31.0–79.8% versus 24.8–66.1%); official TCH is tied or reported-set-led at $n\leq$150, then write-back at CNH200 (60.8% versus 55.3%). Official Tri leftover at $n\geq$50 is not a reported-set 2-opt gap (Table A.24): $T(A_{\text{off}})$ is 2.9–10.2% and $A_{\text{off}}(T)$ is 8.2–19.3%. On Tri-20 the same duals are 0.4–2.3% and 2.1–5.5%, so monotone T on the official K-set cannot be the official WE-CA–GIMF-P aggregate order reversal. Constrained leftover is not one dual (Tables A.25–A.26). On official PMOCO Bi-KP (official HV), PAS ties Off.; $T(A_{\text{off}})$ is 70.3/111.8/108.7% at n=50/100/200 (item-exchange on the reported K-set; $T(A_{\text{off}})$>U2N at $n\geq$100), while $A_{\text{off}}(T)$ is 0. Hold WE-CA/GIMF-P KP leftover is PAS-dominated; both duals are 0.4–7.9%. On Bi-CVRP the two duals are close at n=20; at $n\geq$50 official TCH $T(A_{\text{off}})$ can exceed U2N (PMOCO 122.5/160.0%), matching the already reported T>U2N note on official PMOCO CVRP50. A share above 100% is not a partition of the U2N increment: it is the ratio of two archives against the same Off. baseline (Table A.27). Official TCH on PMOCO Bi-KP places a unique-ND cost on only 39.2/39.0/35.2% of the K=101 slots (|ND| 12.5/19.0/20.7). U2N extracts that ND set; $T(A_{\text{off}})$ polishes the official K-archive, including dominated slice winners, so the increment can exceed leftover. $A_{\text{off}}(T)$=0 because writing $T(\text{ND})$ back does not move those winners. Official TCH CVRP is the same mismatch as n grows (PMOCO in-ND 50.6→25.5%; T share 53.1→160.0%). Tri is the opposite: |ND| is hundreds to thousands and official costs sit on the front (60–99%), so both duals stay a minority of leftover. Official Bi-TSP is the unconstrained 2-objective middle (Table A.28): WS keeps 70.7–80.8% of the reported K on unique-ND (|ND| 58–323); official TCH is 46.1–69.4% on a similarly large front. KroAB n=1 repeats that split (Table A.29): official WS 73.3–90.1%, official TCH 42.6–71.3%. Neither is the sparse-ND TCH mismatch of PMOCO Bi-KP, and both duals stay below leftover (Table A.22). Hold official WS Bi-CVRP occupancy is Table A.30: in-ND falls 53.8→28.9% (WE-CA) and 62.2→33.0% (GIMF-P) while |ND| stays 12.8–22.6. Hold WS Bi-KP is Table A.31: |ND| 46.0–208.9, in-ND 38.1–49.1%, and both duals stay 0.4–7.9% (no GIMF-P KP200). Neither is the sparse-ND official TCH mismatch of Table A.27. Mix WE-CA-U official K on the same unique-ND test is Table A.19: Bi-TSP occupancy 70.5–84.3% tracks size-specific WS 70.7–80.8%; hold Bi-CVRP 55.3→28.3% tracks 53.8→28.9%; Tri 64.4→99.5% tracks 59.9→99.3%. The mix re-extraction is therefore not an occupancy collapse of official K.

Table A.19: Mix WE-CA-U extracted K versus unique-ND(P), aligned to leftover mix evaluation runs (10^{-6}). in-ND is unique-ND occupancy of mix official K. Bi-TSP/KroAB/CVRP K=101; Tri K=105. KroAB is n=1. Not a census row and not a new method.

Problem	n	|ND|	in-ND
Bi-TSP	20	59.5	70.5%
	50	189.5	75.0%
	100	298.8	79.7%
	150	350.5	82.5%
	200	375.0	84.3%
Bi-CVRP	20	14.0	55.3%
	50	20.2	38.4%
	100	23.1	28.3%
Tri-TSP	20	505.8	64.4%
	50	1853.3	95.4%
	100	3260.1	99.5%
KroAB	100	268	77.2%
	150	357	89.1%
	200	357	82.2%

Table A.20: Reported-set residual on official WE-CA, including $W_{\mathrm{all}}(A_{\mathrm{WS}})$. The main-text subset is Table 2. T is the dominance-monotone operator (2-opt on Bi-TSP); W_{all} is source-λ on the same reported solutions. Share is leftover of $T(A_{\mathrm{WS}})$ over leftover of U2N. T versus U2N is WTL 0–0–200 at n=50/100. W_{all} versus U2N is 0–0–200 at n=50 and 12–0–188 at n=100.

n	WS	$T(A_{\mathrm{WS}})$	$W_{\mathrm{all}}(A_{\mathrm{WS}})$	PAS	U2N	share
20	0.62704	0.62708	0.62717	0.62860	0.62873	2.5%
50	0.63905	0.63933	0.64047	0.64103	0.64217	9.0%
100	0.70386	0.70451	0.70619	0.70512	0.70678	22.3%
150	0.69912	0.70062	0.70395	0.70028	0.70278	40.9%
200	0.72638	0.72982	0.73324	0.72733	0.73134	69.4%

Table A.21: Official extractor after writing $T(\mathrm{ND}(P))$ back into the decoded Bi-TSP pool. $A_{\mathrm{off}}(T(P))$ is not $T(A_{\mathrm{off}})$. Share is leftover of $A_{\mathrm{off}}(T(P))$ over leftover of U2N. n=20/150/200 are evaluated on freshly decoded official solutions; n=50 WE-CA/GIMF-P uses the officially released candidate pools; n=100 is the four-constructor checkpoint. WE-CA $T(A_{\mathrm{WS}})$ share 22.3% at n=100 is Table A.20; side-by-side shares are Table A.22.

Ctor	n	Off.	$A_{\mathrm{off}}(T)$	U2N	share
WE-CA	20	0.62704	0.62707	0.62873	2.1%
	50	0.63905	0.63962	0.64217	18.2%
	100	0.70386	0.70485	0.70678	34.2%
	150	0.69912	0.70112	0.70278	54.8%
	200	0.72638	0.73045	0.73134	82.2%
GIMF-P	20	0.62660	0.62667	0.62888	3.0%
	50	0.63744	0.63833	0.64139	22.4%
	100	0.70060	0.70216	0.70422	43.0%
	150	0.69582	0.69854	0.70018	62.3%
	200	0.72667	0.73057	0.73154	80.1%
PMOCO	20	0.50536	0.50566	0.50718	16.4%
	50	0.56704	0.56825	0.57197	24.6%
	100	0.66640	0.66821	0.67173	33.8%
CNH	20	0.50588	0.50614	0.50770	14.4%
	50	0.56854	0.56942	0.57297	19.8%
	100	0.66809	0.66961	0.67290	31.7%
	150	0.71597	0.71806	0.72045	46.6%
	200	0.74376	0.74662	0.74842	61.3%

Table A.29: Official KroAB extracted K versus unique-ND(P), n=1. in-ND is the fraction of official K=101 costs on unique-ND. WE-CA/GIMF-P are official WS; PMOCO/CNH are official TCH. No official PMOCO KroAB150/200. Same direction as Table A.28: WS occupancy is high; TCH is lower.

Ctor	n	|ND|	in-ND
WE-CA	100	277	79.2%
	150	286	80.2%
	200	362	90.1%
GIMF-P	100	232	73.3%
	150	267	74.3%
	200	319	82.2%
PMOCO	100	239	42.6%
CNH	100	234	42.6%
	150	308	53.5%
	200	329	71.3%

Table A.30: Hold official WS Bi-CVRP extracted K versus unique-ND(P), n_{test}=200. in-ND is unique-ND occupancy of official K. T / $A(T)$ repeat Table A.25.

Ctor	n	|ND|	in-ND	T	$A(T)$
WE-CA	20	15.2	53.8%	31.1%	30.1%
	50	20.6	37.5%	110.7%	101.4%
	100	22.6	28.9%	111.4%	102.7%
GIMF-P	20	12.8	62.2%	30.5%	37.4%
	50	16.4	44.4%	102.9%	106.2%
	100	17.3	33.0%	108.6%	98.2%

Table A.22: Reported-set 2-opt $T(A_{\text{off}})$ versus official slice extract after writing $T(\text{ND}(P))$ back into the decoded Bi-TSP pool, $A_{\text{off}}(T(P))$. $T(A_{\text{off}})$ does not read discarded unique-ND. HV matches Table A.2. Share is leftover of that archive over leftover of U2N. WE-CA n=50/100 $T(A_{\text{off}})$ is Table A.20; GIMF-P n=50/100 use the unified official WS evaluation runs (official = WS).

Ctor	n	Off.	$T(A_{\text{off}})$	$A_{\text{off}}(T)$	U2N	T	$A(T)$
WE-CA	20	0.62704	0.62708	0.62707	0.62873	2.5%	2.1%
	50	0.63905	0.63933	0.63962	0.64217	9.0%	18.2%
	100	0.70386	0.70451	0.70485	0.70678	22.3%	34.2%
	150	0.69912	0.70062	0.70112	0.70278	40.9%	54.8%
	200	0.72638	0.72982	0.73045	0.73134	69.4%	82.2%
GIMF-P	20	0.62660	0.62665	0.62667	0.62888	2.1%	3.0%
	50	0.63744	0.63805	0.63833	0.64139	15.3%	22.4%
	100	0.70060	0.70180	0.70216	0.70422	33.2%	43.0%
	150	0.69582	0.69800	0.69854	0.70018	50.0%	62.3%
	200	0.72667	0.73010	0.73057	0.73154	70.3%	80.1%
PMOCO	20	0.50536	0.50588	0.50566	0.50718	28.5%	16.4%
	50	0.56704	0.56862	0.56825	0.57197	32.0%	24.6%
	100	0.66640	0.66850	0.66821	0.67173	39.3%	33.8%
CNH	20	0.50588	0.50632	0.50614	0.50770	24.1%	14.4%
	50	0.56854	0.56968	0.56942	0.57297	25.7%	19.8%
	100	0.66809	0.66974	0.66961	0.67290	34.4%	31.7%
	150	0.71597	0.71808	0.71806	0.72045	47.1%	46.6%
	200	0.74376	0.74652	0.74662	0.74842	59.2%	61.3%

Table A.23: KroAB n=1 reported-set 2-opt $T(A_{\text{off}})$ versus official extractor after writing $T(\text{ND}(P))$ back into the decoded pool. HV matches Table A.2. Share is leftover of that archive over leftover of U2N. No Wilcoxon (n=1). No official PMOCO TCH r^* for KroAB150/200.

Ctor	n	Off.	$T(A_{\text{off}})$	$A_{\text{off}}(T)$	U2N	T	$A(T)$
WE-CA	100	0.69585	0.69667	0.69687	0.69915	24.8%	31.0%
	150	0.68982	0.69124	0.69229	0.69387	35.0%	61.0%
	200	0.72324	0.72604	0.72695	0.72789	60.3%	79.8%
GIMF-P	100	0.69360	0.69482	0.69524	0.69719	34.0%	45.8%
	150	0.68777	0.68981	0.69028	0.69226	45.5%	55.9%
	200	0.72315	0.72654	0.72707	0.72827	66.1%	76.6%
PMOCO	100	0.65850	0.66146	0.66145	0.66471	47.6%	47.5%
CNH	100	0.66093	0.66288	0.66226	0.66563	41.5%	28.4%
	150	0.70850	0.71075	0.71059	0.71359	44.1%	41.0%
	200	0.74025	0.74302	0.74329	0.74525	55.3%	60.8%

Table A.31: Hold official WS Bi-KP extracted K versus unique-ND(P), n_{test}=200. in-ND is unique-ND occupancy of official K. T / $A(T)$ repeat Table A.26. No GIMF-P KP200.

Ctor	n	|ND|	in-ND	T	$A(T)$
WE-CA	50	47.5	38.1%	6.3%	7.9%
	100	110.7	45.7%	0.8%	2.4%
	200	208.9	49.1%	1.2%	3.4%
GIMF-P	50	46.0	42.8%	4.0%	5.0%
	100	114.7	44.7%	0.4%	1.8%

Table A.24: Official Tri-TSP reported-set dominance-monotone 2-opt $T(A_{\text{off}})$ versus official extract after writing $T(\text{ND}(P))$ back into the decoded $105\times n$ pool. HV matches Table A.2 (PMOCO official r^*). Share is leftover of that archive over leftover of U2N. GIMF-P Tri-100 is not a 56-setting census row.

Ctor	n	Off.	$T(A_{\text{off}})$	$A_{\text{off}}(T)$	U2N	T	$A(T)$
WE-CA	20	0.47067	0.47068	0.47075	0.47467	0.4%	2.1%
	50	0.43893	0.43911	0.43945	0.44537	2.9%	8.2%
	100	0.49759	0.49807	0.49862	0.50477	6.7%	14.4%
GIMF-P	20	0.47020	0.47026	0.47036	0.47488	1.4%	3.5%
	50	0.43536	0.43585	0.43635	0.44412	5.6%	11.4%
PMOCO	20	0.32044	0.32061	0.32085	0.32788	2.3%	5.5%
	50	0.33649	0.33730	0.33823	0.34966	6.2%	13.2%
	100	0.43965	0.44098	0.44218	0.45275	10.2%	19.3%

Table A.25: Official Bi-CVRP reported-set intra-route 2-opt $T(A_{\text{off}})$ versus official extract after writing $T(\text{ND}(P))$ back. HV Off. / U2N match Table A.2. PMOCO/CNH use official r^*, seed1234, n_{test}=100; WE-CA/GIMF-P use hold WS, n_{test}=200. T is capacity-feasible intra-route 2-opt, not IMVND. Share is leftover of that archive over leftover of U2N; it can exceed 100% when the archive exceeds U2N.

Ctor	n	Off.	$T(A_{\text{off}})$	$A_{\text{off}}(T)$	U2N	T	$A(T)$
WE-CA	20	0.42941	0.42951	0.42951	0.42975	31.1%	30.1%
	50	0.40904	0.40944	0.40940	0.40940	110.7%	101.4%
	100	0.40723	0.40762	0.40759	0.40758	111.4%	102.7%
GIMF-P	20	0.42877	0.42893	0.42896	0.42928	30.5%	37.4%
	50	0.40768	0.40828	0.40829	0.40826	102.9%	106.2%
	100	0.40515	0.40576	0.40570	0.40571	108.6%	98.2%
PMOCO	20	0.21447	0.21593	0.21591	0.21721	53.1%	52.6%
	50	0.25082	0.25331	0.25309	0.25285	122.5%	111.4%
	100	0.23815	0.24143	0.24123	0.24020	160.0%	150.1%
CNH	20	0.22574	0.22619	0.22607	0.22689	39.4%	28.4%
	50	0.26444	0.26501	0.26495	0.26519	76.5%	68.5%
	100	0.25095	0.25173	0.25171	0.25198	75.7%	74.3%

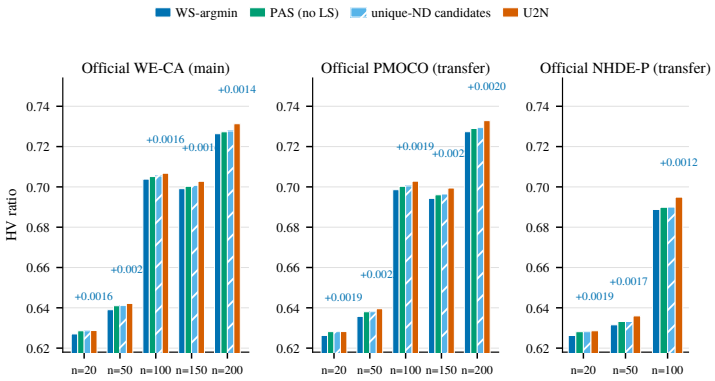


Figure A.3: Identical already-decoded pools before local search (Table A.2). Bars 1–3 are WS, PAS, and unique-ND. Bar 4 is leftover after U2N. WE-CA is the full suite; PMOCO / NHDE-P are Bi-TSP n=20/50/100 only.

Table A.26: Official Bi-KP reported-set 1–1 item-exchange $T(A_{\text{off}})$ versus official extract after writing $T(\text{ND}(P))$ back. HV Off. / U2N match Table A.2. PMOCO is official HV; WE-CA/GIMF-P are hold WS. T is first-improvement item-exchange, not TSP 2-opt. WE-CA KP U2N is evaluated on the frozen candidate pool. No official GIMF-P KP200. Share can exceed 100% when $T(A_{\text{off}})$>U2N. $A_{\text{off}}(T)$=0 on official PMOCO is exact.

Ctor	n	Off.	$T(A_{\text{off}})$	$A_{\text{off}}(T)$	U2N	T	$A(T)$
WE-CA	50	0.35608	0.35612	0.35613	0.35669	6.3%	7.9%
	100	0.45333	0.45334	0.45335	0.45446	0.8%	2.4%
	200	0.36013	0.36014	0.36015	0.36097	1.2%	3.4%
GIMF-P	50	0.35604	0.35606	0.35607	0.35665	4.0%	5.0%
	100	0.45333	0.45334	0.45335	0.45447	0.4%	1.8%
PMOCO	50	0.75382	0.75704	0.75382	0.75840	70.3%	0.0%
	100	0.95033	0.95216	0.95033	0.95197	111.8%	0.0%
	200	0.68080	0.68160	0.68080	0.68153	108.7%	0.0%

Table A.27: Official extracted K versus unique-ND(P) on the decoded pool (same decoded solutions as Tables A.25–A.24). $|\text{ND}|$ is the mean unique non-dominated size. in-ND is the fraction of official K slots whose cost lies on that set. T / $A(T)$ repeat leftover shares of $T(A_{\text{off}})$ and $A_{\text{off}}(T)$. A share >100% means that archive exceeds U2N. GIMF-P Tri-100 is not a census row.

Ctor	Prob	n	$\|\text{ND}\|$	in-ND	T	$A(T)$
PMOCO	KP	50	12.5	39.2%	70.3%	0.0%
	KP	100	19.0	39.0%	111.8%	0.0%
	KP	200	20.7	35.2%	108.7%	0.0%
	CVRP	20	11.9	50.6%	53.1%	52.6%
	CVRP	50	13.6	30.8%	122.5%	111.4%
	CVRP	100	14.5	25.5%	160.0%	150.1%
CNH	CVRP	20	13.8	47.7%	39.4%	28.4%
	CVRP	50	18.3	45.8%	76.5%	68.5%
	CVRP	100	17.9	34.4%	75.7%	74.3%
WE-CA	Tri	20	467.7	59.9%	0.4%	2.1%
	Tri	50	1875.6	95.3%	2.9%	8.2%
	Tri	100	3230.1	99.3%	6.7%	14.4%
GIMF-P	Tri	20	515.2	69.5%	1.4%	3.5%
	Tri	50	1644.6	95.7%	5.6%	11.4%
PMOCO	Tri	20	463.7	78.6%	2.3%	5.5%
	Tri	50	1498.9	96.6%	6.2%	13.2%
	Tri	100	2715.2	99.5%	10.2%	19.3%

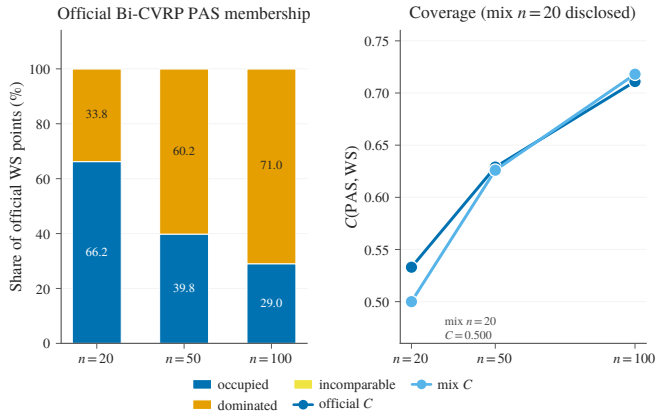


Figure A.4: Official Bi-CVRP under PAS only: occupancy falls and point-domination rises while incomparable stays 0; coverage C rises. Improvement is 1–2 $\times 10^{-4}$ covered volume, matching the official PAS settings.

The official WE-CA hypervolume staircase is Table 3.

Table A.28: Official Bi-TSP extracted K versus unique-ND(P). HV Off. matches Table A.2. T / $A(T)$ repeat Table A.22. WE-CA/GIMF-P are official WS; PMOCO/CNH are official TCH, seed1234, n_{test}=100. No official PMOCO TCH n>100.

Ctor	n	|ND|	in-ND	T	$A(T)$
WE-CA	20	58.0	70.7%	2.5%	2.1%
	50	186.2	78.2%	9.0%	18.2%
	100	277.7	79.2%	22.3%	34.2%
	150	321.5	80.8%	40.9%	54.8%
	200	312.8	79.7%	69.4%	82.2%
GIMF-P	20	60.4	75.2%	2.1%	3.0%
	50	149.0	78.2%	15.3%	22.4%
	100	238.8	76.4%	33.2%	43.0%
	150	291.6	78.4%	50.0%	62.3%
	200	315.1	79.4%	70.3%	80.1%
PMOCO	20	54.6	63.6%	28.5%	16.4%
	50	143.7	46.1%	32.0%	24.6%
	100	235.3	46.3%	39.3%	33.8%
CNH	20	57.5	69.4%	24.1%	14.4%
	50	155.8	54.9%	25.7%	19.8%
	100	240.1	53.1%	34.4%	31.7%
	150	290.5	53.7%	47.1%	46.6%
	200	326.5	55.3%	59.2%	61.3%

F.2 Cross-constructor factorial leftover

HV columns of Table 3 isolate PAS / T / $W|_R$ on official WE-CA Bi-TSP; share columns add official PMOCO/CNH/GIMF-P. The selection stage of U2N is PAS; T then $W|_R$ score remaining M2; constrained families apply Γ_{hold}. Leftover is versus the official extractor. Table A.33 reports which stage carries leftover under each official extractor. Shares on a common ruler (same r^*, or unified 101λ+WS) isolate extraction efficiency and leave Table A.2 ranks unchanged. Official checkpoints store Off/PAS/U2N and therefore the PAS share versus residual $T+W|_R$ except on official Bi-TSP PMOCO/CNH/GIMF-P, official PMOCO/WE-CA/GIMF-P Tri, and official TCH CVRP, where a logged T now splits that residual (Tables A.34–A.37). NHDE-P Tri-TSP leftover is negative and stays outside the 56-setting census.

On Bi-TSP, PAS share falls with n on every constructor that starts from a K-slice (WE-CA/PMOCO/CNH/GIMF-P). Official NHDE-P EP already occupies a large unique-ND set, so the PAS share is vacuous (0 or slightly negative) and the positive leftover is residual. WE-CA/GIMF-P Bi-KP leftover is almost all PAS (Γ_{hold} preserves NDS); PMOCO/NHDE-P Bi-KP leftover is all residual (PAS = Off; the increment is under Γ_{hold}). PMOCO Tri-TSP leftover at n=100 is 21.5% PAS, 18.9% T, and 59.6% $W|_R$ (+0.01028); n=20/50 are 76.0/5.0/18.9 and 37.7/12.2/50.1.

Table A.32 is the n=100 cross-family excerpt of Table 3; Table A.34 is the full official Bi-TSP T checkpoint. The T share grows with n on PMOCO/CNH/GIMF-P; the $W|_R$ share stays the smaller mid-scale term. Official CNH n=150/200 continues that shift (41.4/50.6/7.9; 36.1/62.7/1.2). Table A.38 scales the same increments by the official constructor gap on a shared r^*. On the shared WS ruler, the extractor-to-gap ratio $\rho_{\text{U2N}}(\text{WE}-\text{CA})$ first exceeds 1 at n=150 (1.11); at n=200 the official Off. values nearly coincide (g=0.00029), which amplifies ρ. Table A.41 repeats the split on the unified 101λ+WS ruler, where WE-CA/GIMF-P official = WS. The $W|_R$ share can be negative (PMOCO n=20 on that ruler; PMOCO CVRP50 and WE-CA KP200 in Tables A.36–A.37). That sign is a local extreme-ray trade-off before Γ_{hold}, so the reported map keeps W on R only. Unified CNH n=20/50/100 now log T (Table A.41). Hypervolumes on Table A.41 are not the official TCH/EP settings of Table A.2.

Table A.32: Module shares at n=100 (PAS / T / $W|_R$). HV matches Table A.2. A dash means no logged T; share is then PAS / residual.

Ctor	Problem	Off.	PAS	T	U2N	share
WE-CA	Bi-TSP	0.70386	0.70512	0.70599	0.70678	43.4/29.6/27.0
PMOCO	Bi-TSP	0.66640	0.66878	0.67088	0.67173	44.6/39.5/16.0
CNH	Bi-TSP	0.66809	0.67027	0.67205	0.67290	45.4/37.0/17.7
GIMF-P	Bi-TSP	0.70060	0.70187	0.70329	0.70422	35.0/39.3/25.7
WE-CA	Bi-CVRP	0.40723	0.40737	0.40753	0.40758	41.1/44.7/14.2
GIMF-P	Bi-CVRP	0.40515	0.40531	0.40564	0.40571	27.6/59.3/13.0
PMOCO	Bi-CVRP	0.23815	0.23935	0.24011	0.24020	58.8/37.0/4.2
CNH	Bi-CVRP	0.25095	0.25163	0.25190	0.25198	66.0/26.4/7.6
WE-CA	Bi-KP	0.45333	0.45445	0.45446	0.45446	99.1/0.9/0.0
PMOCO	Bi-KP	0.95033	0.95033	0.95189	0.95197	0.0/95.1/4.9
PMOCO	Tri-TSP	0.43965	0.44246	0.44494	0.45275	21.5/18.9/59.6
WE-CA	Tri-TSP	0.49759	0.49952	0.50053	0.50477	26.9/14.2/59.0

Table A.33: Official leftover shares PAS versus residual $T+W|_R$. HV columns match Table A.2. Share is leftover of PAS / leftover of residual over leftover of U2N, from the Table A.2 checkpoints. WE-CA Bi-TSP three-way shares remain Table 3. Constrained rows use Γ_{hold}. NHDE-P Tri-TSP is omitted (negative leftover; not in the 56-setting census).

Ctor	Problem	Off.	PAS	U2N	PAS / res.
WE-CA	Bi-TSP20	0.62704	0.62860	0.62873	92.9/7.1
	Bi-TSP50	0.63905	0.64103	0.64217	63.3/36.7
	Bi-TSP100	0.70386	0.70512	0.70678	43.4/56.6
	Bi-TSP150	0.69912	0.70028	0.70278	31.9/68.1
	Bi-TSP200	0.72638	0.72733	0.73134	19.3/80.7
	Bi-CVRP100	0.40723	0.40737	0.40758	40.0/60.0
	Bi-KP100	0.45333	0.45445	0.45446	99.1/0.9
	Tri-TSP100	0.49759	0.49952	0.50477	26.9/73.1
PMOCO	Bi-TSP20	0.50536	0.50681	0.50718	79.7/20.3
	Bi-TSP50	0.56704	0.57001	0.57197	60.2/39.8
	Bi-TSP100	0.66640	0.66878	0.67173	44.6/55.4
	Bi-CVRP100	0.23815	0.23935	0.24020	58.8/41.2
	Bi-KP100	0.95033	0.95033	0.95197	0.0/100.0
	Tri-TSP20	0.32044	0.32610	0.32788	76.0/24.0
	Tri-TSP50	0.33649	0.34145	0.34966	37.7/62.3
	Tri-TSP100	0.43965	0.44246	0.45275	21.5/78.5
CNH	Bi-TSP20	0.50588	0.50739	0.50770	83.3/16.7
	Bi-TSP50	0.56854	0.57126	0.57297	61.4/38.6
	Bi-TSP100	0.66809	0.67027	0.67290	45.4/54.6
	Bi-CVRP100	0.25095	0.25163	0.25198	66.0/34.0
NHDE-P	Bi-TSP20	0.62863	0.62863	0.62879	0.0/100.0
	Bi-TSP50	0.63886	0.63885	0.64046	−1.1/101.1
	Bi-TSP100	0.70059	0.70045	0.70296	−5.6/105.6
	Bi-CVRP100	0.39340	0.39340	0.39442	0.0/100.0
	Bi-KP100	0.45204	0.45204	0.45235	0.0/100.0
GIMF-P	Bi-TSP20	0.62660	0.62868	0.62888	91.1/8.9
	Bi-TSP50	0.63744	0.63953	0.64139	52.9/47.1
	Bi-TSP100	0.70060	0.70187	0.70422	35.0/65.0
	Bi-CVRP100	0.40515	0.40531	0.40571	28.6/71.4
	Bi-KP100	0.45333	0.45446	0.45447	99.1/0.9
	Tri-TSP50	0.43536	0.43872	0.44412	38.4/61.6

Table A.34: Official Bi-TSP T split (PAS / T / $W|_R$). HV columns match Table A.2. Share is leftover of PAS / T / $W|_R$ over leftover of U2N. WE-CA remains Table 3. T is the official Bi-TSP monotone residual. CNH n=200 is the frozen official checkpoint.

Ctor	n	Off.	PAS	T	U2N	share
PMOCO	20	0.50536	0.50681	0.50693	0.50718	79.7/6.7/13.6
	50	0.56704	0.57001	0.57137	0.57197	60.2/27.7/12.1
	100	0.66640	0.66878	0.67088	0.67173	44.6/39.5/16.0
CNH	20	0.50588	0.50739	0.50752	0.50770	83.3/6.7/10.1
	50	0.56854	0.57126	0.57236	0.57297	61.4/24.9/13.7
	100	0.66809	0.67027	0.67205	0.67290	45.4/37.0/17.7
	150	0.71597	0.71783	0.72010	0.72045	41.4/50.6/7.9
	200	0.74376	0.74544	0.74837	0.74842	36.1/62.7/1.2
GIMF-P	20	0.62660	0.62868	0.62873	0.62888	91.1/2.3/6.7
	50	0.63744	0.63953	0.64025	0.64139	52.9/18.3/28.8
	100	0.70060	0.70187	0.70329	0.70422	35.0/39.3/25.7
	150	0.69582	0.69705	0.69960	0.70018	28.3/58.4/13.3
	200	0.72667	0.72762	0.73141	0.73154	19.5/77.7/2.7

Table A.35: Official KroAB T split (PAS / T / $W|_R$), n=1. Share is leftover of PAS / T / $W|_R$ over leftover of U2N. WE-CA is evaluated under the frozen official KroAB checkpoint (T is the controlled hypervolume split). PMOCO/CNH/GIMF-P are evaluated under the controlled parameter configurations. No official PMOCO KroAB$150/200$ T split is logged. These HVs use each paper's official ruler, not the unified-WS settings of Table A.13.

Ctor	n	Off.	PAS	T	U2N	share
WE-CA	100	0.69585	0.69718	0.69814	0.69915	40.2/29.4/30.4
	150	0.68982	0.69135	0.69320	0.69387	37.8/45.6/16.6
	200	0.72324	0.72418	0.72788	0.72789	20.2/79.6/0.2
PMOCO	100	0.65850	0.66114	0.66399	0.66471	42.5/45.9/11.6
CNH	100	0.66093	0.66290	0.66491	0.66563	41.9/42.6/15.4
	150	0.70850	0.71076	0.71327	0.71359	44.5/49.1/6.4
	200	0.74025	0.74199	0.74515	0.74525	34.8/63.1/2.0
GIMF-P	100	0.69360	0.69473	0.69620	0.69719	31.3/41.1/27.6
	150	0.68777	0.68902	0.69157	0.69226	27.8/56.7/15.5
	200	0.72315	0.72422	0.72815	0.72827	21.0/76.5/2.4

Table A.36: Official leftover T split on Tri-TSP and official r^* CVRP (PAS / T / $W|_R$). HV matches the leftover-share archive of that paper. WE-CA Tri-20 is the frozen leftover-share evaluation run (already records T). PMOCO/CNH CVRP is official r^*, seed1234, not the hold $\sim$ 0.39 ruler. PMOCO CVRP50 has a negative $W|_R$ share (T>U2N): a source-λ step on the residual can trade orthogonal coverage before Γ_{hold}. WE-CA Tri-$20/50/100$ now records T. GIMF-P Tri leftover T is this evaluation run (n=100 is not a 56-setting census row). Hold Γ_{hold} CVRP/KP T is Table A.37.

Ctor	Setting	Off.	PAS	T	U2N	share
WE-CA	Tri20	0.47067	0.47446	0.47452	0.47467	94.7/1.7/3.6
	Tri50	0.43893	0.44212	0.44264	0.44537	49.5/8.0/42.4
	Tri100	0.49759	0.49952	0.50053	0.50477	26.9/14.2/59.0
PMOCO	Tri20	0.32044	0.32610	0.32647	0.32788	76.0/5.0/18.9
	Tri50	0.33649	0.34145	0.34305	0.34966	37.7/12.2/50.1
	Tri100	0.43965	0.44246	0.44494	0.45275	21.5/18.9/59.6
GIMF-P	Tri20	0.47020	0.47419	0.47432	0.47488	85.3/2.8/11.9
	Tri50	0.43536	0.43872	0.43969	0.44412	38.4/11.1/50.5
	Tri100	0.49268	0.49472	0.49641	0.50235	21.1/17.4/61.4
PMOCO	CVRP20	0.21447	0.21561	0.21616	0.21721	41.4/20.4/38.2
	CVRP50	0.25082	0.25202	0.25292	0.25285	$59.1/44.0/-3.1$
	CVRP100	0.23815	0.23935	0.24011	0.24020	58.8/37.0/4.2
CNH	CVRP20	0.22574	0.22635	0.22649	0.22689	52.7/12.3/35.0
	CVRP50	0.26444	0.26498	0.26516	0.26519	72.7/23.3/4.0
	CVRP100	0.25095	0.25163	0.25190	0.25198	66.0/26.4/7.6

Table A.37: Official leftover T split on hold Γ_{hold} CVRP/KP and official HV PMOCO KP (PAS / T / $W|_R$). HV matches Table A.2. n=100 WE-CA/GIMF-P CVRP and WE-CA/PMOCO KP remain Table A.32. For WE-CA Bi-KP, T is evaluated under the fixed-seed control configuration. PMOCO KP200 has $T\approx$U2N after five decimals. WE-CA KP200 has a -0.6% $W|_R$ share at five decimals: $T\approx$U2N after rounding, not a global alignment failure.

Ctor	Setting	Off.	PAS	T	U2N	share
WE-CA	CVRP20	0.42941	0.42959	0.42962	0.42975	52.5/10.8/36.6
	CVRP50	0.40904	0.40922	0.40937	0.40940	51.6/41.9/6.5
	KP50	0.35608	0.35667	0.35669	0.35669	97.7/2.2/0.2
	KP200	0.36013	0.36097	0.36097	0.36097	$99.9/0.7/-0.6$
GIMF-P	CVRP20	0.42877	0.42905	0.42913	0.42928	53.8/16.8/29.4
	CVRP50	0.40768	0.40793	0.40819	0.40826	44.1/44.2/11.7
	KP50	0.35604	0.35664	0.35665	0.35665	98.1/1.5/0.4
	KP100	0.45333	0.45446	0.45447	0.45447	99.1/0.6/0.3
PMOCO	KP50	0.75382	0.75382	0.75626	0.75840	0.0/53.1/46.9
	KP200	0.68080	0.68080	0.68153	0.68153	0.0/99.5/0.5

Table A.38: Same-pool extractor increment versus the official constructor gap g on a shared r^*. (A) and (B) are the first and second constructor of each pair (WE-CA–GIMF: A = WE-CA, B = GIMF-P; PMOCO–CNH: A = PMOCO, B = CNH). $\rho_S(A) = \Delta_S(A)/g$ for $S \in \{\mathrm{PAS}, T, \mathrm{U2N}\}$. WE-CA–GIMF-P share WS r^*; PMOCO–CNH share TCH r^* and seed1234. $\rho > 1$ means the extractor increment exceeds that pair's official Off. gap. At n=200, the marginal denominator (g=0.00029; WE-CA 0.72638 vs GIMF-P 0.72667) amplifies ρ. PMOCO–CNH has no official n>100 pair (no published PMOCO official r^*). A dash is a missing T checkpoint. $\rho_T(B)$ uses the same five-decimal protocol as the other ρ columns.

Pair	n	g	$\rho_{\mathrm{PAS}}(A)$	$\rho_T(A)$	$\rho_{\mathrm{U2N}}(A)$	$\rho_{\mathrm{PAS}}(B)$	$\rho_T(B)$	$\rho_{\mathrm{U2N}}(B)$
WE-CA–GIMF	20	0.00044	3.55	3.64	3.84	4.73	4.84	5.18
	50	0.00161	1.23	1.50	1.94	1.30	1.75	2.45
	100	0.00326	0.39	0.65	0.90	0.39	0.83	1.11
	150	0.00330	0.35	0.92	1.11	0.37	1.15	1.32
	200	0.00029	3.28	16.93	17.10	3.28	16.34	16.79
PMOCO–CNH	20	0.00052	2.79	3.02	3.50	2.90	3.15	3.50
	50	0.00150	1.98	2.89	3.29	1.81	2.55	2.95
	100	0.00169	1.41	2.65	3.15	1.29	2.34	2.85

Table A.39: Instance-bootstrap 95% CIs (10^4 resamples, seed 2026) for the Table 1 gaps, from the locked per-instance HV lists. Each replicate redraws the n_{test} instances with replacement. KroAB (n=1) omitted.

Pair	Cell	g_{off} CI	g_{PAS} CI	Δ_{PAS} CI
WE-CA–GIMF	Bi-TSP20	$[+0.00028, +0.00059]$	$[-0.00017, +0.00002]$	$[-0.00062, -0.00039]$
	Bi-TSP50	$[+0.00146, +0.00175]$	$[+0.00137, +0.00162]$	$[-0.00020, -0.00003]$
	Bi-TSP100	$[+0.00316, +0.00336]$	$[+0.00317, +0.00335]$	$[-0.00004, +0.00004]$
	Bi-TSP150	$[+0.00320, +0.00339]$	$[+0.00314, +0.00333]$	$[-0.00010, -0.00003]$
	Bi-TSP200	$[-0.00040, -0.00018]$	$[-0.00040, -0.00018]$	$[-0.00003, +0.00004]$
	Tri-TSP20	$[+0.00035, +0.00059]$	$[+0.00021, +0.00033]$	$[-0.00031, -0.00009]$
PMOCO–CNH	Bi-TSP20	$[-0.00077, -0.00028]$	$[-0.00078, -0.00040]$	$[-0.00019, +0.00007]$
	Bi-TSP50	$[-0.00175, -0.00125]$	$[-0.00146, -0.00105]$	$[+0.00012, +0.00037]$
	Bi-TSP100	$[-0.00188, -0.00149]$	$[-0.00168, -0.00131]$	$[+0.00011, +0.00028]$

Table A.40: Extractor-to-gap ratios on the remaining official same-r^* pairs (five-decimal g as in Table A.38). (A) and (B) follow the pair order as in Table A.38. WE-CA–GIMF-P KP100 is omitted (g=0 at five decimals). KP50 (g=0.00004) and KroAB200 (g=0.00009) are near-ties; large ρ is the thin denominator, not a method ranking. PMOCO–CNH CVRP g is $O(10^{-2})$, so $\rho \ll 1$: the extractor increment does not close that constructor gap. GIMF-P Tri100 is not a census row. No official PMOCO n>100.

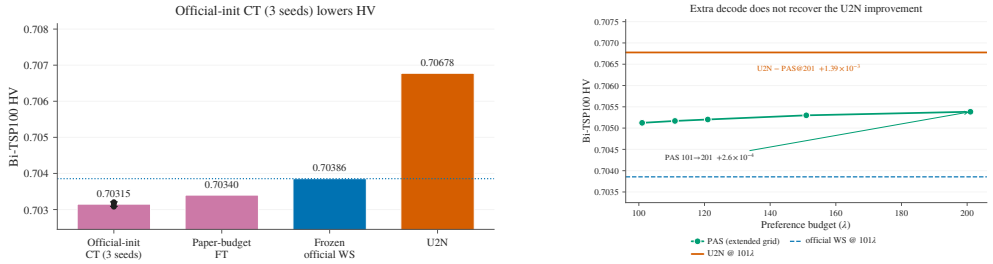
Pair	Cell	g	$\rho_{\mathrm{PAS}}(A)$	$\rho_T(A)$	$\rho_{\mathrm{U2N}}(A)$	$\rho_{\mathrm{PAS}}(B)$	$\rho_T(B)$	$\rho_{\mathrm{U2N}}(B)$
WE-CA–GIMF	CVRP20	0.00064	0.28	0.33	0.53	0.44	0.56	0.80
	CVRP50	0.00135	0.14	0.25	0.27	0.19	0.38	0.43
	CVRP100	0.00208	0.07	0.14	0.17	0.08	0.24	0.27
	KP50	0.00004	14.75	15.25	15.25	15.00	15.25	15.25
	Tri20	0.00047	8.06	8.19	8.51	8.49	8.77	9.96
	Tri50	0.00357	0.89	1.04	1.80	0.94	1.21	2.45
	Kro100	0.00225	0.59	1.02	1.47	0.50	1.16	1.60
	Kro150	0.00205	0.75	1.65	1.98	0.61	1.85	2.19
	Kro200	0.00009	10.44	51.56	51.67	11.89	55.56	56.89
PMOCO–CNH	CVRP20	0.01127	0.10	0.15	0.24	0.05	0.07	0.10
	CVRP50	0.01362	0.09	0.15	0.15	0.04	0.05	0.06
	CVRP100	0.01280	0.09	0.15	0.16	0.05	0.07	0.08
	Kro100	0.00243	1.09	2.26	2.56	0.81	1.64	1.93

Table A.41: Unified 101λ+WS T split (PAS / T / $W|_R$) on Bi-TSP. Share is leftover of PAS / T / $W|_R$ over leftover of U2N. WE-CA coincides with Table 3. A dash means that checkpoint does not record T; the two-number share is then PAS / residual. GIMF-P official $=$ WS, so its T matches Table A.34. These HVs are not official TCH leftover. A negative $W|_R$ share (PMOCO/CNH n=20) means T>U2N on that unified ruler: residual alignment can still trade orthogonal coverage, so the reported map keeps W on R only.

Ctor	n	WS	PAS	T	U2N	share
WE-CA	20	0.62704	0.62860	0.62864	0.62873	92.9/2.2/4.9
	50	0.63905	0.64103	0.64146	0.64217	63.3/13.8/23.0
	100	0.70386	0.70512	0.70599	0.70678	43.4/29.6/27.0
PMOCO	20	0.62630	0.62818	0.62827	0.62822	$97.9/4.5/-2.5$
	50	0.63575	0.63805	0.63913	0.63953	60.7/28.5/10.8
	100	0.69860	0.70028	0.70214	0.70283	39.6/43.9/16.5
CNH	20	0.62686	0.62855	0.62862	0.62854	$100.9/4.3/-5.2$
	50	0.63691	0.63915	0.64001	0.64032	65.6/25.3/9.2
	100	0.69992	0.70168	0.70319	0.70390	44.2/38.0/17.9
	150	0.69602	0.69781	0.70031	0.70069	38.3/53.7/8.0
NHDE-P	20	0.62627	0.62815	0.62827	0.62861	80.6/5.0/14.4
	50	0.63156	0.63320	0.63484	0.63598	37.2/37.0/25.8
	100	0.68879	0.68988	0.69423	0.69492	17.8/70.9/11.3
GIMF-P	20	0.62660	0.62868	0.62873	0.62888	91.1/2.3/6.7
	50	0.63744	0.63953	0.64025	0.64139	52.9/18.3/28.8
	100	0.70060	0.70187	0.70329	0.70422	35.0/39.3/25.7
	150	0.69582	0.69705	0.69960	0.70018	28.3/58.4/13.3

F.3 TRAINING-OBJECTIVE DIVERGENCE

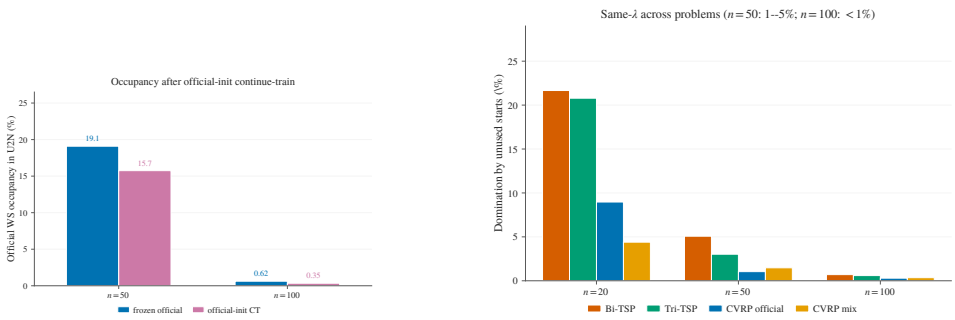
A linear scalarization $\min \lambda^\top f(x)$ is supported only at extreme points of a convex Pareto front (Das & Dennis, 1997; Miettinen, 1999). As POMO policy-gradient updates continue on a fixed preference grid, the rollout distribution concentrates on those supporting hyperplanes: entropy falls, and unsupported non-dominated coverage between adjacent rays is eroded. We interpret those drops as preference collapse on this constructor and recipe: continued scalarized updates prioritize supported extreme rays at the expense of intermediate coverage, so $\mathcal{L}_{\mathrm{WS}}$ and J_K pull in opposite geometric directions. This is a diagnostic on one constructor and recipe, not a claim that every training objective fails. Figure A.5 visualizes the $n{=}100$ bar. Three official-init WE-CA seeds (2024/2025/2026) continue the same $\mathcal{L}_{\mathrm{WS}}$ for 20 epochs and lower $\mathrm{HV}_r(A_{\mathrm{off}})$ from 0.70386 to mean 0.70315 ($0.70317/0.70320/0.70308$; s.d. $6{\times}10^{-5}$); all three sit below the freeze. The same recipe at $n{=}50$ (one seed) lowers 0.63905 to 0.63861; a paper-budget fine-tune at $n{=}100$ (0.70340) likewise sits below that freeze.



(a) (b)

Figure A.5: (a) M3 diagnostic, three official-init seeds: mean 0.70315 ($0.70317/0.70320/0.70308$), all below freeze 0.70386; paper-budget FT 0.70340 also below (same recipe $n{=}50$: $0.63905 \rightarrow 0.63861$). (b) Decode-budget ablation: PAS $101 \rightarrow 201$ is $+2.60 \times 10^{-4}$ and remains $+1.39 \times 10^{-3}$ below U2N at 101.

Leftover by official-extractor hypervolume tertile remains on the order of 10^{-3} on the high band except official NHDE-P Tri-TSP EP. On $30/33$ multi-instance official checkpoints the low/mid/high Off. HV tertile means are all positive; leftover-only-on-high is $0/33$. Every $n{=}200$ checkpoint keeps the same sign on the first 100 (PMOCO Tri-20 first 100 leftover $+0.00725$, distinct from an official online sample).



(a) (b)

Figure A.6: (a) Occupancy of official WS points after official-init continued training still falls with n; further λ-training leaves occupancy on the same declining trajectory. (b) Same-λ domination across Bi-TSP, Tri-TSP, and Bi-CVRP. At $n{=}50$ the rate is 1–5%; at $n{=}100$ it is below 1%: leftover is distinct from extra multi-start inference.

F.4 OBJECT IDENTITY AND CARDINALITY

Table A.42 lists evaluated objects on official WE-CA Bi-TSP100. U2N is the fixed-cardinality R-restricted map reported in Table A.2. Table A.43 separates $|A| = K$ from unique-ND: official WS Bi-TSP100 has $|A| = K$ and NDS 77; U2N Bi-TSP20 has $|A| = K$ with reported NDS 47. Table A.44 isolates the same M1 object on synthetic 2-D fronts ($K{=}11$, 40 reps): a finite WS-argmin is not a hypervolume selector.

Table A.42: Object identity on official WE-CA Bi-TSP100. NDS is $|\mathrm{ND}(\cdot)|$; $|A|$ is the reported cardinality. Extra λ is PAS $101 \to 201$. $T(A_{\mathrm{WS}})$ is reported-set residual. M3 HV is seed 2026; three-seed mean 0.70315.

Object	HV	NDS	$\lvert A\rvert$	$W\vert_R$	Role
official WS	0.70386	77	101	—	protocol
continued training	0.70308	—	101	—	M3
PAS 101	0.70512	101	101	—	stage one
PAS 201	0.70538	—	201	—	extra λ
$T(A_{\mathrm{WS}})$	0.70451	—	101	yes	reported-set
Unconstrained	0.70764	101	101	no	diagnostic
U2N	0.70678	101	101	yes	reported

Table A.43: Fixed cardinality is $|A| = K$ (padding allowed). Unique-ND need not reach K.

Object	$\lvert A\rvert$	NDS	$\lvert A\rvert{=}K$	reported object
official WS Bi-TSP100	101	77	yes	finite WS-argmin
U2N Bi-TSP20	101	47	yes	selected HV K-set
U2N Bi-TSP100	101	101	yes	selected HV K-set
Unconstrained Bi-TSP100	101	101	yes	unrestricted W_{all}
Unconstrained Bi-TSP50	101	99.43	yes	unrestricted W_{all}

Table A.44: Synthetic 2-D fronts: finite WS-argmin versus greedy K-subset and the full front. $K{=}11$, 40 reps, $r = (1.05, 1.05)$.

Front	WS vs pool	WS vs greedy	greedy $>$ WS	unique WS
linear	0.335	0.292	40/40	3.0
concave	0.207	0.181	40/40	2.0
convex	0.032	0.006	40/40	10.9
clustered	0.056	0.055	40/40	4.0

Table A.45 records that PAS is stage one: the exact–greedy gap is selector discrepancy; U2N − PAS is the post-selection residual.

Table A.45: PAS is stage one on official WE-CA Bi-TSP100. Shares are select / monotone / align of leftover.

Contrast	$n{=}100$	Interpretation
exact HSSP–PAS	$+6.6 \times 10^{-6}$	selector discrepancy
U2N–PAS	$+0.00166$	residual + align
leftover share	43.4/29.6/27.0	selection share
$n{=}200$ leftover share	19.3/79.9/0.8	residual dominated by T

Table A.46 reports unconstrained W_{all} versus each constructor's own U2N. The operator lacks a monotone guarantee (Proposition 2). Every row has $|A|{=}K$; NDS records $|\mathrm{ND}(\cdot)|$. Official $n{=}50$ NDS is $99.43/82.52/76.84$ on WE-CA/PMOCO/NHDE-P.

Table A.46: Unconstrained W_{all} versus each constructor's own U2N. Every row has $|A|{=}K$; NDS records $|\mathrm{ND}(\cdot)|$.

Constructor	n	U2N	Unconstrained	Unconstrained NDS	HV vs U2N
official WE-CA	50	0.64217	0.64236	99.43	higher
	100	0.70678	0.70764	101	higher
	150	0.70278	0.70562	101	higher
official PMOCO	50	0.63953	0.63887	82.52	lower
	100	0.70283	0.70374	101	higher
	150	0.69943	0.70171	101	higher
	200	0.73287	0.73379	101	higher
official NHDE-P	50	0.63598	0.63932	76.84	higher
	100	0.69492	0.70109	99.12	higher

G DIAGNOSTICS ON CONSTRAINED AND HIGH-DIMENSIONAL INSTANCES

The envelope of official WE-CA Tri-TSP is Table A.47: leftover versus own WS, unique non-dominated size, and retain/increase/dominate class of A_{WS}. Table A.48 records official Bi-KP

leftover under item-exchange T with the instance-wise NDS oracle (U2N is the reported leftover). Table A.49 is the constructor-matched counterpart.

Table A.47: Official WE-CA Tri-TSP (K=105) vs. own WS. class. ret./inc./dom. (%).

n	leftover	$\lvert\mathrm{ND}\rvert$	class.
20	+0.00400	796	77.5/20.8/1.7
50	+0.00645	1917	45.3/53.0/1.7
100	+0.00718	3241	46.4/53.3/0.3

G.1 Bi-KP item-exchange residual

Table A.48 repeats the leftover protocol on official WE-CA Bi-KP with T instantiated as first-improvement 1–1 item-exchange (capacity $12.5/25/25$ at n=50/100/200).

Table A.48: Official WE-CA Bi-KP leftover versus own WS. T is first-improvement 1–1 item-exchange. Ours = U2N under the instance-wise NDS oracle (U2N / T / PAS counts $70/130/0$, $70/130/0$, $199/1/0$). NDS is $\lvert\mathrm{ND}(\cdot)\rvert$ (WS / Ours) and is held. $\lvert R\rvert/\lvert N\rvert$ is the unmodified residual after T. These settings are the Table A.2 WE-CA Bi-KP rows.

n	WS	Ours	leftover	$\lvert R\rvert/\lvert N\rvert$	NDS	WTL
50	0.35608	0.35669	+0.00060	0.967	19.9/47.5	200–0–0
100	0.45333	0.45446	+0.00113	0.963	32.1/93.9	200–0–0
200	0.36013	0.36097	+0.00084	0.962	47.5/100.9	200–0–0

The unmodified residual after T is $96.7/96.3/96.2\%$. Unconstrained W_{all} shadows official NDS ($47.51 \to 41.31$, $93.76 \to 87.74$, $100.95 \to 100.10$). The tabulated setting applies (6) instance-wise: keep U2N if official NDS does not fall, else T, else A_{PAS}. No instance reverts to PAS ($70/130/0$, $70/130/0$, $199/1/0$). Leftover versus own WS is U2N $+0.00060/+0.00113/+0.00084$ (unique-ND $19.9 \to 47.5$, $32.1 \to 93.9$, $47.5 \to 100.9$; WTL 200–0–0). These settings are the Table A.2 WE-CA Bi-KP rows. A diagnostic 1–2 / 2–2 exchange on the same official pool leaves the post-PAS residual at $\sim 10^{-5}$, so reported leftover versus own WS is pool truncation of $\mathrm{ND}(P)$, not a residual local move; T remains the prescribed 1–1 neighbourhood.

Table A.49 repeats that protocol on the official PMOCO / NHDE-P / GIMF-P checkpoints, still versus each constructor's own WS and on the same WE-CA testdata (K=101). PMOCO and NHDE-P emit sparse unique-ND pools (12–26), so PAS HV equals own WS: the remaining unified-WS Δ is item-exchange polish of the selected archive, not latent-pool coverage of $\mathrm{ND}(P)$. PMOCO unified-WS Δ stays below the leftover residual of Table A.48 ($+0.00017/+0.00020/+0.00016$). NHDE-P unified-WS Δ (versus own WS, not versus EP) is $+0.00039/+0.00040$. GIMF-P leftover matches official WE-CA (leftover versus own WS $+0.00061/+0.00114$; versus PAS $\sim 10^{-5}$) and is the Table A.2 GIMF-P Bi-KP pair. PMOCO / NHDE-P stay in this appendix (unique-ND versus own WS contracts).

Table A.49: Constructor-matched Bi-KP Γ_{hold} versus own WS on official WE-CA testdata. T is first-improvement 1–1 item-exchange. unique-ND is WS / hold. hold counts are U2N / T / PAS. PMOCO leftover is residual leftover after item-exchange. PMOCO hold counts here are unified-WS; official HV/TCH counts are Table A.52. WTL is versus own WS on all 200 instances (ties are the PAS-hold cases). GIMF-P is also Table A.2. CNH is omitted: the official release is MOTSP/MOCVRP only (no Bi-KP checkpoint; Appendix H).

	n	WS	Ours	vs WS	vs PAS	unique-ND	hold	WTL
PMOCO	50	0.35467	0.35484	+0.00017	+0.00017	13.9/12.5	42/88/70	130–70–0
	100	0.45112	0.45132	+0.00020	+0.00020	20.2/19.0	7/155/38	162–38–0
	200	0.35820	0.35836	+0.00016	+0.00016	22.0/20.7	3/185/12	188–12–0
NHDE-P	50	0.35448	0.35487	+0.00039	+0.00039	15.8/13.8	114/46/40	160–40–0
	100	0.45131	0.45172	+0.00040	+0.00040	29.6/26.2	109/76/15	185–15–0
GIMF-P	50	0.35604	0.35665	+0.00061	+0.00001	18.9/46.0	25/41/134	200–0–0
	100	0.45333	0.45447	+0.00114	+0.00001	32.8/94.1	29/98/73	200–0–0

G.2 Bi-CVRP residual neighbourhoods and domination shadowing

Table A.50 evaluates three capacity-feasible combinatorial compilers on official Bi-CVRP20 across 200 instances, on the same decoded pool: intra-route 2-opt (T only), inter-route IMVND (Relocate, Swap, 2-opt*), and mixed-neighbourhood search. Leftover versus WS rises for every residual ($+0.00022/+0.00036/+0.00034$ versus PAS $+0.00018$). Inter-route IMVND records $+0.00036$ versus WS and $+0.00018$ versus PAS (WTL 178–1–21). The same operators trigger a *domination-shadowing* effect: locally advanced solutions push deeper into the objective interior and geometrically subsume adjacent marginal trade-offs. Official `Pareto_sols` NDS contracts from 14.01 to $13.78/13.52/13.49$, while unique-ND of the candidate pool stays at the PAS floor ($14.01/14.01/14.09/14.04$). Intra-route $W|_R$ is 0; IMVND-as-T instantiates a different neighbourhood from the Bi-TSP Γ; mixed is a third object. These unconstrained residual compilers are therefore a controlled diagnostic of why (6) is required: they raise hypervolume by sharpening the official archive rather than by enlarging unique-ND, and they contract the official non-dominated archive. The Table A.2 WE-CA Bi-CVRP setting is Γ_{hold} (NDS $14.07/20.32/22.52$), which applies residual search only when the official non-dominated archive does not contract (Appendix Table A.51: leftover $+0.00034/+0.00036/+0.00035$; WTL 200–0–0). Unconstrained compilers in Table A.50 isolate the contraction; Table A.51 reports the cardinality-preserving projection.

Table A.50: Controlled diagnostic ablation of residual combinatorial compilers on official Bi-CVRP20 (200 instances, identical decoded pool). Capacity-checked inter-route moves (IMVND) lift leftover versus PAS ($+0.00018$) while official archive NDS falls (14.01 to 13.52): a domination-shadowing contraction of `Pareto_sols`. Unique-ND of the pool remains 14.01. NDS is official `Pareto_sols`. WTL is versus PAS. Unconstrained rows stay a diagnostic of why (6) is required. The Table A.2 Bi-CVRP20 setting is Γ_{hold} in Table A.51.

Object	vs WS	vs PAS	NDS	$W\|_R$	WTL
PAS	$+0.00018$	—	14.01	—	—
intra (T only)	$+0.00022$	$+0.00004$	13.78	0	174–26–0
IMVND-as-T	$+0.00036$	$+0.00018$	13.52	28.4%	178–1–21
mixed	$+0.00034$	$+0.00016$	13.49	74.9%	191–2–7

Table A.51: Official WE-CA Bi-CVRP leftover versus own WS. Ours $= \Gamma_{\text{hold}}$ (U2N / T / PAS counts $168/2/30$, $162/0/38$, $188/0/12$). NDS is official `Pareto_sols` (WS / Ours) and is held or enlarged. Operators (2-opt, or-opt, IMVND) are neighbourhood tools. IMVND at n=20 only; this Γ_{hold} is the Table A.2 WE-CA Bi-CVRP setting. Wall-clock $20/10/20$ min at n=$20/50/100$.

n	WS	Ours	leftover	NDS	WTL
20	0.42941	0.42975	$+0.00034$	$8.9/14.1$	200–0–0
50	0.40904	0.40940	$+0.00036$	$11.2/20.3$	200–0–0
100	0.40723	0.40758	$+0.00035$	$11.8/22.5$	200–0–0

G.3 Tri-objective scaling

Table A.2 shows that U2N executes on all 200 instances at n=$20/50/100$ and fills official NDS 105. Post-PAS residual (U2N − PAS) is $+2.10 \times 10^{-4}$ at n=20 (WTL 159–0–41, $p = 7.8 \times 10^{-20}$), $+3.25 \times 10^{-3}$ at n=50 ($W|_R = 96.1\%$; WTL 200–0–0), and $+5.25 \times 10^{-3}$ at n=100 ($W|_R = 99.9\%$; WTL 200–0–0). Mix Tri-TSP20 lifts the post-PAS residual by $+4.09 \times 10^{-4}$ versus mix PAS at official NDS 105 (Table A.5 Method 0.47488). Mix Tri-TSP50 lifts the post-PAS residual by $+3.34 \times 10^{-3}$ versus mix PAS ($W|_R = 96.0\%$; WTL 200–0–0; Table A.5 Method 0.44536). Mix Tri-TSP100 lifts the post-PAS residual by $+5.37e-3$ versus mix PAS ($W|_R = 99.9\%$). At n=100 the candidate pool exhibits a heavy cardinality tail ($|\text{ND}| = 3241$ against K=105; Table A.47); the WE-CA Table A.2 setting is $+0.00718$ versus own WS. Main-table Tri leftover uses each row's official Proto (PMOCO PBI, NHDE-P EP); official NHDE-P Tri-TSP20/50/100 leftover is negative (Table A.2).

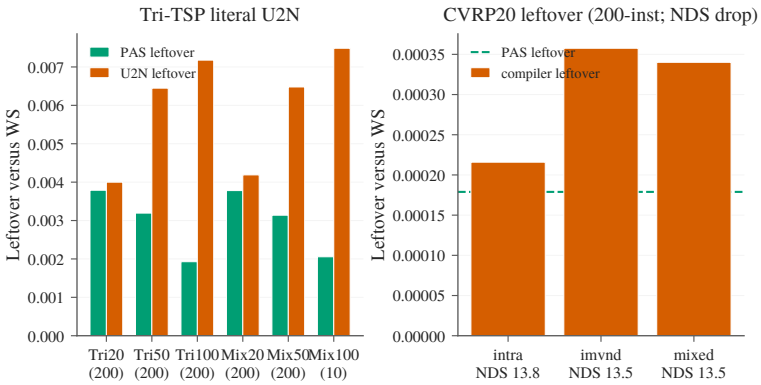


Figure A.7: Unified-WS Δ (diagnostic) versus own WS. Left: PAS versus U2N on official/mix Tri-TSP20/50/100. Right: Bi-CVRP20 leftover above the PAS line.

H Limitations, missing settings, and hold choices

Recovered hypervolume is the per-instance mean versus that paper's official extractor under the same official hypervolume, and is strictly positive on all 56 finite-K settings. The magnitude varies by domain—from large Tri-TSP gains to fine increments on constrained Bi-CVRP ($\Delta \approx 3$–6×10^{-4})—and Wilcoxon signed-rank tests of A_{U2N} versus A_{off} are significant on every multi-instance setting. Official PMOCO Bi-KP has high ties (PAS = Off.). Official NHDE-P Bi-TSP20 has $d_z = 0.47$ (WTL 148–0–52) and sits outside the census with the other 15 EP rows. On those 52 losses, PAS equals EP on every instance (WTL 0–52–0). T is hypervolume-monotone on its input set (Lemma 2); A_{U2N} can still fall below the unbounded EP because a later cardinality-K selector is compared with that archive (mean -1.56×10^{-4}). Official NHDE-P reports an unbounded EP: $\mathrm{PAS}(\mathrm{EP}) \subseteq \mathrm{EP}$, so a finite-$K$ selector cannot raise that reported archive. The negative Tri-TSP difference ($\Delta^{\dagger} < 0$) is that geometric boundary: a finite K-set versus an unbounded EP. Constrained official rows use Γ_{hold} (Table A.52). PMOCO/CNH Bi-CVRP official checkpoints apply the same projection and do not log per-instance arm counts. Evaluations use the officially released checkpoints and reference points r^* of each repository; unreleased configurations are omitted: PMOCO $n{>}100$ has no published official r^*; CNH releases MOTSP/MOCVRP only (no Bi-KP / Tri-TSP); GIMF-P has no official Bi-KP200 / Tri-TSP100 checkpoint. M3 is a diagnostic official-init continue-train on WE-CA Bi-TSP100 (three seeds, mean 0.70315) and Bi-TSP50 ($0.63905 \rightarrow 0.63861$), not a claim that every training objective fails. KroAB official rows are $n{=}1$ and have no Wilcoxon. On the twelve smallest-effect or high-tie checkpoints (WE-CA/GIMF-P Bi-CVRP, official PMOCO Bi-KP, official NHDE-P Bi-TSP20, official PMOCO/CNH Bi-TSP20), one-sided Wilcoxon signed-rank ($\alpha{=}0.01$, alternative greater) and paired instance-mean bootstrap 95% CIs (10^4, seed 2026) all sit strictly above 0. Official WE-CA Bi-TSP $n{=}20/50/100$ is $[0.00158, 0.00180]$, $[0.00303, 0.00321]$, $[0.00287, 0.00297]$; Bi-CVRP Γ_{hold} versus own WS is $[0.00029, 0.00039]$, $[0.00031, 0.00042]$, $[0.00031, 0.00039]$; none of those intervals contain 0. GIMF-P Bi-CVRP $n{=}20/50/100$ gives $[0.00044, 0.00058]$, $[0.00049, 0.00068]$, $[0.00046, 0.00069]$; official PMOCO Bi-KP $n{=}50/100/200$ gives $[0.00384, 0.00535]$, $[0.00142, 0.00188]$, $[0.00066, 0.00082]$; official NHDE-P Bi-TSP20 gives $[0.00011, 0.00020]$ (WTL 148–0–52); official PMOCO/CNH Bi-TSP20 give $[0.00160, 0.00203]$ and $[0.00165, 0.00199]$. Those intervals corroborate the Table A.2 signs.

Table A.52: Cardinality-preserving projection Γ_{hold} on constrained official settings that log an arm. Counts are U2N / T / PAS over n_{test}. PMOCO Bi-KP is the official HV checkpoint, not unified WS. NHDE-P Bi-CVRP is an official EP row (outside the 56-setting census).

Ctor	Problem	U2N / T / PAS	n_{test}
WE-CA	Bi-CVRP20	168/2/30	200
	Bi-CVRP50	162/0/38	200
	Bi-CVRP100	188/0/12	200
	Bi-KP50	70/130/0	200
	Bi-KP100	70/130/0	200
	Bi-KP200	199/1/0	200
GIMF-P	Bi-CVRP20	178/2/20	200
	Bi-CVRP50	179/0/21	200
	Bi-CVRP100	190/0/10	200
	Bi-KP50	25/41/134	200
	Bi-KP100	29/98/73	200
PMOCO	Bi-KP50	90/49/61	200
	Bi-KP100	11/149/40	200
	Bi-KP200	3/185/12	200
NHDE-P	Bi-CVRP20	193/1/6	200
	Bi-CVRP50	191/0/9	200
	Bi-CVRP100	198/0/2	200