NOT ALL PROGRESS TRANSFERS: ADAPTATION-STATE REUSE FOR GENERALIZABLE OBJECT REID

Anonymous authors
Paper under double-blind review

ABSTRACT

Generalizable object re-identification aims to learn identity-discriminative representations that remain effective across heterogeneous object categories and visual distributions. Learning must reconcile fine-grained specialization with unknown matching requirements, although seen-category supervision gives no direct guidance on which adaptation changes to retain for transfer. We identify Progressive Feature Repurposing (PFR): structure beyond an early discriminative span gains seen-category retrieval utility while its category-readout accessibility declines modestly even after refitting. Restoring earlier structure benefits some transfer settings but harms others, with heterogeneous query-level responses: later progress does not uniformly supersede historical utility. Motivated by PFR, we propose ReState (Reusing Adaptation States), treating adaptation history as a reusable representational resource rather than committing to one selected state. During adaptation, ReState retains historical states and composes their feature residuals relative to a pretrained anchor through input- and dimension-conditioned routing. A zero-adaptation option retains access to anchor content, while optimization- and supervision-level regulation shapes alternatives for reuse. Across ten heterogeneous retrieval tasks, ReState demonstrates strong performance against existing ReID methods, with particularly large gains on multi-category benchmarks. It also improves mAP over a strong baseline across all category-shifted tasks.

1 INTRODUCTION

Object re-identification (ReID) retrieves images of the same physical object across viewpoints, appearances, and sensing conditions. ReID has traditionally been developed through category-specific models (Ye et al., 2022; He et al., 2021). Extending this capability beyond predefined categories is important for reusable retrieval systems: encountering new object types and deployment conditions should not require a separately annotated and specialized ReID model. Recent work moves toward this goal through cross-category retrieval: VICP uses visual in-context examples (Huang & Liu, 2025), while Object-Generalized Re-Identification (OG-ReID) studies source-only transfer from seen to unseen object categories within a unified retrieval framework (Chen et al., 2026). Generalizable object ReID therefore asks one representation to preserve fine-grained identity discrimination across heterogeneous object types and visual distributions.

Pretrained visual encoders provide a broad starting point for such transfer (Oquab et al., 2023), but adaptation must learn fine-grained distinctions among seen identities without knowing which visual cues will distinguish identities in future categories. Prior work improves transfer by controlling adaptation through weight interpolation, selective updates, semantic anchoring, or prompt-based regularization (Wortsman et al., 2022b; Lee et al., 2023; Han et al., 2024; Wen et al., 2025), while generalizable ReID methods address source-specific variation in identity representations (Jin et al., 2020; Ni et al., 2023). These strategies provide effective mechanisms for robust adaptation. Yet observed-identity supervision gives no direct indication of which induced feature changes should be preserved, suppressed, or reused when the object category and visual distribution change. The central difficulty is reconciling fine-grained specialization with unknown future matching requirements, including how adaptation progress relates to those requirements.

To examine this gap, we trace the representation throughout a baseline ReID adaptation trajectory. Figure 1(a) first tests whether transfer utility evolves together with source discrimination: source

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.

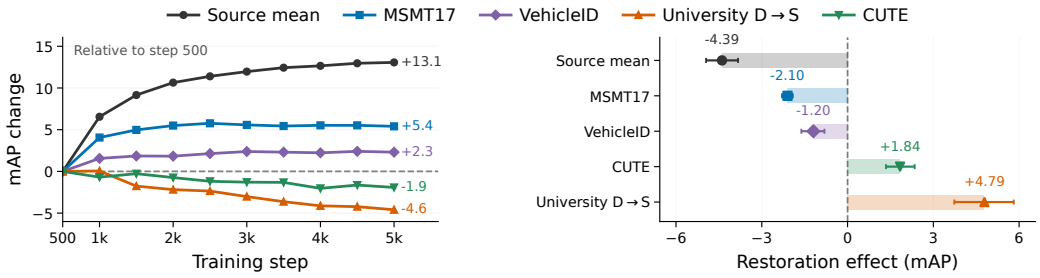

(a) Adaptation trajectories (b) Historical restoration effects

Figure 1: **Source progress does not uniformly supersede historical transfer utility.** (a) Source retrieval improves while representative transfer trajectories diverge. (b) Restoring earlier feature structure at the final checkpoint reduces source retrieval but benefits selected evaluation settings. Whiskers denote 95% confidence intervals.

retrieval continues to improve, while representative evaluation tasks follow different trajectories. Figure 1(b) then tests whether earlier feature structure remains useful after continued adaptation. Restoring that structure at the final checkpoint reduces source retrieval but improves some evaluation settings while degrading others. Together, these results show that later progress does not uniformly supersede historical transfer utility. We characterize the broader pattern as *Progressive Feature Repurposing* (PFR): continued identity learning develops source-discriminative structure while changing its measured functions, leaving historical structure whose utility depends on the retrieval setting. Section 3.1 develops this account through complementary discriminative development, category-readout change, and restoration responses across settings and queries. This motivates retaining adaptation history as reusable alternatives and learning how to combine them.

We propose **ReState** (Reusing Adaptation States), coupling state formation and reuse. To exploit historical utility without committing to one adaptation state, *Trajectory Residual Routing* (TRR) retains states and composes their feature residuals relative to a pretrained anchor per input and coordinate. Because these alternatives remain shaped by seen-identity learning, *Dual-Scale Trajectory Regularization* (DTR) moderates current updates and accumulated adaptation to temper specialization. *Dual-Space Relation Perturbation* (DSRP) instead targets rigid reinforcement of observed identity relations through selective perturbations to classification and metric supervision.

Our contributions can be summarized as follows:

- **We characterize PFR**, linking continued identity-discrimination development, changing category accessibility, and retained historical transfer utility across settings and queries.
- **We propose ReState**, which couples conditional reuse of trajectory states with optimization- and supervision-level regulation of their formation.
- **We improve a strong adaptation baseline** on all seven category-shifted tasks, alongside competitive complete-system results across 111 dataset-local categories.

2 Related Work

Generalizable object re-identification. Domain-generalizable ReID extends category-specific retrieval beyond training distributions (Cho et al., 2024) through meta-learning, normalization, expert aggregation, and local similarity (Song et al., 2019; Jin et al., 2020; Choi et al., 2021; Dai et al., 2021; Ni et al., 2023). These approaches primarily address category-consistent shifts. VICP extends retrieval across categories using visual support pairs (Huang & Liu, 2025), while OG-ReID formalizes source-only category generalization (Chen et al., 2026). We examine how representation utility evolves during adaptation in this setting.

Pretrained representation adaptation. Transferability depends on task relationships (Yosinski et al., 2014), and adaptation can distort features and impair out-of-distribution performance (Kumar et al., 2022). Preservation strategies regularize toward initialization (Li et al., 2018), select layers (Lee et al., 2023), interpolate weights (Wortsman et al., 2022b), or anchor pretrained semantics (Han et al., 2024). Foundation-model studies also identify adaptation bias under limited source diversity (Wen et al., 2025). Our analysis complements preservation methods by tracing changing utility throughout ReID adaptation.

Table 1: **Rapid early adaptation coexists with source-dependent pretrained accessibility.** mAP_0 is pretrained performance; G_{early} and G^{lin} denote early-adaptation and linear-readout gains; η_{src} compares source-specific readout gain with early gain; v_{early} and v_{late} are the corresponding adaptation rates.

	Pretrained / early		Linear accessibility			Adaptation rates	
Dataset	mAP_0	G_{early}	G^{lin}	$G^{\mathrm{lin}}_{\mathrm{src}}$	η_{src}	v_{early}	v_{late}
Market-1501	1.67	59.55	1.23	6.78	0.11	0.1191	0.0043
iPanda-50	18.51	10.03	-0.70	3.08	0.31	0.0201	0.0019
VeRi	4.95	47.90	4.78	11.86	0.25	0.0958	0.0040
ATRW	47.76	7.58	12.13	15.47	2.04	0.0152	0.0014
VesselReID	26.45	24.80	19.77	18.55	0.75	0.0496	0.0030

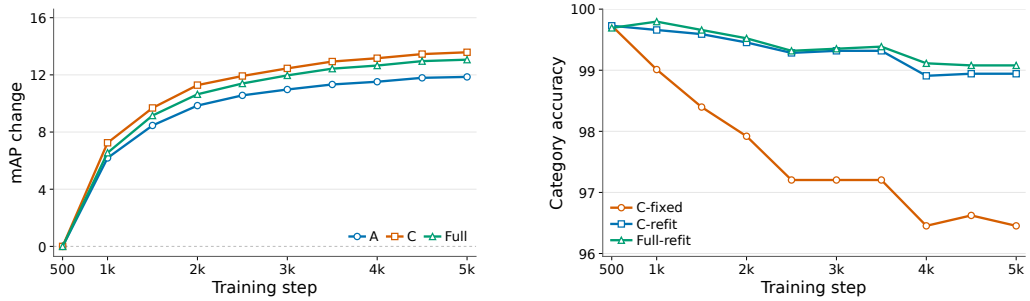

(a) Discriminative development along adaptation

(b) Category-readout accessibility along adaptation

Figure 2: **Continued adaptation develops complementary identity-discriminative utility while category-readout accessibility modestly declines.** (a) Five-source mean mAP changes in the early span A, complementary span C, and full representation, relative to step 500. (b) Category accuracy on the identity-disjoint CUTE-10 probe; C-fixed retains the early classifier, while C-refit and Full-refit refit each checkpoint.

Model combination. Snapshot Ensembles combine checkpoint predictions from one trajectory (Huang et al., 2017), Model Soups average independently fine-tuned weights (Wortsman et al., 2022a), and WiSE-FT interpolates pretrained/adapted weights (Wortsman et al., 2022b). Input-dependent approaches include VRF's sample-wise ensembling (Zhu et al., 2024), DaWin's uncertainty-based interpolation (Oh et al., 2025), and Twin-Merging's conditional knowledge integration (Lu et al., 2024). ReState builds on these ideas of historical reuse and conditional combination **within a single adaptation trajectory**, coupling the formation of adaptation states with their feature-level composition for generalizable object ReID.

3 From Progressive Feature Repurposing to ReState

3.1 Progressive Feature Repurposing

Progressive Feature Repurposing (PFR) couples continued identity-discrimination development with historical utility that later progress does not uniformly supersede. We trace a baseline ReID adaptation trajectory from a pretrained encoder, beginning with pretrained accessibility, examining how subsequent adaptation changes measured functions, and finally testing the retrieval value of earlier structure. All diagnostics use this trajectory without ReState (Appendix A).

Rapid early gains coexist with accessible identity structure. Table 1 shows that average mAP gain per step before the early reference checkpoint t_e, v_{early}, exceeds the later rate, v_{late}, on every source along this decreasing-learning-rate trajectory. Source-specific linear readouts also improve retrieval on all five sources with the pretrained encoder E_0 frozen, exposing identity distinctions without encoder updates. The ratio η_{src} compares these readout gains with early encoder gains and varies substantially: pretrained accessibility has no single relationship to early progress across sources (Appendix A.3). We use the early state as a reference for following subsequent development.

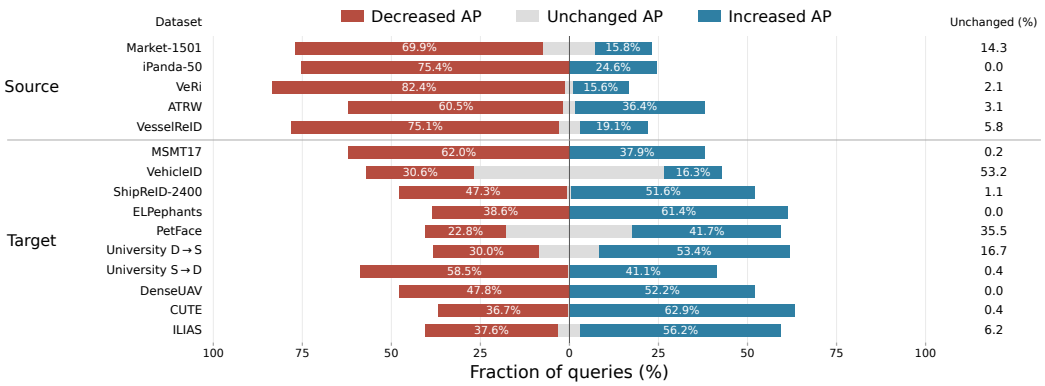

Figure 3: **Historical restoration produces heterogeneous query responses within and across retrieval settings.** Fractions of queries with decreased, unchanged, or increased AP after restoring early complementary structure in the final representation. Bars indicate response prevalence rather than effect magnitude.

3.1.1 Functional Change During Continued Adaptation

Later learning does more than refine early discriminative directions. Source identity scatter at t_e defines fixed projectors

$$P_A = Q_A Q_A^\top, \qquad P_C = I - P_A, \tag{1}$$

where Q_A is an orthonormal basis of the early discriminative span A, P_A projects onto A, P_C onto its complement C, and I is the feature-space identity matrix. Source-fitted orthogonal alignment places normalized features at adaptation step t, denoted u_t, in early coordinates.

Figure 2(a) shows sustained source-macro gains in A, C, and full features. Because the decomposition stays fixed, increasing utility in C reveals development outside the directions selected early, rather than a span refitted to follow each checkpoint. Later progress is therefore not confined to refining A: complementary structure becomes increasingly useful for retrieval, indicating that continued discrimination recruits structure beyond the early span. Each curve uses its own early baseline and separately normalized descriptors; gains are not additive (Appendices A.2 and A.4).

The same complement does not improve uniformly across functions. On identity-disjoint CUTE-10 splits, the early classifier retained on C (C-fixed) becomes less effective (Fig. 2(b)). Checkpoint-wise linear refitting (C-refit) recovers much of this decline, indicating reduced compatibility with the early decision rule. Yet C-refit and full-feature refitting (Full-refit) finish modestly below their early accuracies, with fluctuations near ceiling. Read together, the panels show growing complementary retrieval utility alongside reduced category accessibility even after refitting. This is measured functional change, not irreversible information loss (Appendix A.5). We next examine whether earlier structure remains useful for retrieval itself.

3.1.2 Historical Utility Across Settings and Queries

Earlier structure remains useful, but restoration is not uniformly beneficial. For an input image x, we retain current A and restore early C, using $\mathrm{normalize}(v) = v/\|v\|_2$ to obtain a unit descriptor:

$$f_t^{\mathrm{restore}\,C}(x) = \mathrm{normalize}\left(P_A u_t(x) + \kappa_C(t) P_C u_{t_e}(x)\right), \tag{2}$$

where $f_t^{\mathrm{restore}\,C}$ is the hybrid descriptor and $\kappa_C(t)$ matches restored component energy to its current counterpart using source calibration. Applied to both queries and galleries, final-checkpoint restoration has negative source point estimates but improves some targets and degrades others (Fig. 1(b)). Earlier structure thus retains transfer value after source discrimination improves.

Dataset averages conceal opposing query responses. Figure 3 shows improved and degraded queries coexisting, with predominantly degraded source queries and differing target profiles. Historical utility therefore varies within settings as well as across them, revealing competing preferences even within a single retrieval setting. These fractions measure response prevalence, not effect magnitude or optimal checkpoints for isolated queries. Reverse-A and matched-random controls show that historical benefits are not uniquely localized to the fitted complement (Appendix A.6).

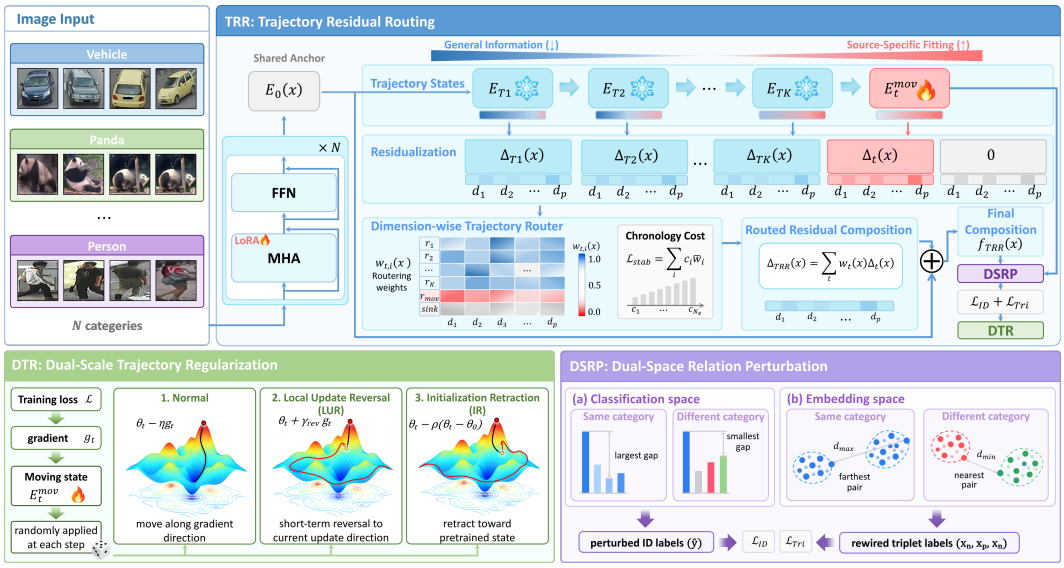

Figure 4: **Overview of ReState.** Trajectory Residual Routing (TRR) retains trajectory states, represents their changes relative to a pretrained anchor, and composes them through input- and dimension-wise routing. Dual-Scale Trajectory Regularization (DTR) and Dual-Space Relation Perturbation (DSRP) regulate state formation through optimization and supervision, respectively. Only TRR is used at inference.

Implication for adaptation-state reuse. Continued training develops seen-identity discrimination without making earlier structure uniformly obsolete for other matching requirements, yet restoring that structure helps some settings and harms others. These findings motivate retaining adaptation alternatives during continued learning rather than automatically replacing or indiscriminately restoring historical structure, and learning when different alternatives should contribute to retrieval.

3.2 ReState Overview

Historical utility alone does not determine how adaptation alternatives should be reused. Standard ReID supervision combines identity classification and metric learning, shaping adaptation around seen identities while providing no direct indication of the states' usefulness for future categories. ReState therefore addresses two coupled requirements: preserving access to different adaptation states and regulating the learning process that forms them, so that state reuse and state formation are learned together within one source-only learning process.

ReState jointly learns adaptation alternatives and their reuse. To avoid commitment to one state, *Trajectory Residual Routing* (TRR) retains history and conditionally composes its adaptation-induced changes. Because optimization shapes these alternatives, *Dual-Scale Trajectory Regularization* (DTR) intervenes in current updates and accumulated adaptation to moderate commitment to a source-driven path. *Dual-Space Relation Perturbation* (DSRP) varies classification and metric relations to temper repeated enforcement of observed identity assignments. Figure 4 summarizes these interventions in reuse, optimization, and supervision.

3.3 Trajectory Residual Routing

Retaining adaptation alternatives. Starting from the pretrained encoder E_0, a Moving State E_t^{mov} continues ReID adaptation while periodic LoRA snapshots form frozen Trajectory States sharing the pretrained backbone. Let $\mathcal{I}_t$ index the historical and moving states available at step t, with corresponding encoders E_s, $s \in \mathcal{I}_t$. Deploying one selected state would remove explicit access to these historical alternatives, whereas fixing an early state would forgo continued identity learning. TRR retains both opportunities by preserving the evolving trajectory for learned composition, rather than committing all inputs to one adaptation depth (Appendix B.1).

Residualized state representation. Retained states share a pretrained foundation but differ through adaptation. For state s, we expose its change from the Shared Anchor as $\Delta_s(x) = E_s(x) - E_0(x)$ and assign a coordinate-wise routing vector $w_s(x) \in \mathbb{R}^p$, where p is the descriptor dimension. Their composition is

$$f_{\mathrm{TRR}}(x) = E_0(x) + \sum_{s \in \mathcal{I}_t} w_s(x) \odot \Delta_s(x), \tag{3}$$

where $f_{\mathrm{TRR}}(x) \in \mathbb{R}^p$ is the composed feature and $\odot$ is elementwise multiplication. Residual conditioning makes adaptation changes explicit to the router around a common pretrained reference; the composition remains algebraically a coordinate-wise mixture of anchor and adapted features (Appendix B.1). Retrieval uses the normalized f_{TRR}.

Input- and dimension-wise routing. Heterogeneous restoration responses motivate input-dependent reuse, while subspace diagnostics motivate nonuniform composition within a descriptor. Coordinate-wise routing implements this flexibility in shared feature coordinates, without identifying them with the diagnostic subspaces. A shared MLP g encodes each residual, and a state-specific head h_s produces routing logits. With routing temperature $\tau_r > 0$,

$$\begin{aligned} r_s(x) &= h_s(g(\Delta_s(x))), \\ w_{s,j}(x) &= \frac{\exp(r_{s,j}(x)/\tau_r)}{1 + \sum_{v \in \mathcal{I}_t} \exp(r_{v,j}(x)/\tau_r)}, \qquad w_{0,j}(x) = 1 - \sum_{s \in \mathcal{I}_t} w_{s,j}(x), \end{aligned} \tag{4}$$

Here, j indexes feature coordinates, $r_{s,j}(x)$ is the corresponding state logit, and $w_{0,j}(x)$ represents the zero-adaptation weight. The denominator's constant 1 comes from a fixed zero-logit, zero-residual candidate, allowing retention of pretrained-anchor content.

Chronology-aware regularization. The source-trained router may favor states specialized to seen identities. A soft chronology prior counterbalances this pressure, charging more for later states while keeping them accessible, thereby discouraging premature collapse toward highly adapted states. Let c_i increase with state i's chronological position and let $\bar{w}_i$ be its routing mass averaged over images and coordinates. We use $\mathcal{L}_{\mathrm{chron}} = \sum_i c_i \bar{w}_i$, with zero cost for the zero-adaptation option. Chronology measures order, not parameter or feature displacement (Appendix B.1).

3.4 Dual-Scale Trajectory Regularization

The states available for reuse are shaped by source-driven optimization. Prior work shows that controlling adaptation relative to a pretrained solution can alter transfer behavior (Kumar et al., 2022; Wortsman et al., 2022b; Ilharco et al., 2023). This motivates regulating not only the endpoint, but also how updates accumulate along the trajectory producing reusable states.

Retaining history is useful only if it supplies alternatives worth reusing. DTR introduces controlled variation at two temporal scales: *Local Update Reversal* (LUR) occasionally opposes the current loss-gradient contribution, while *Initialization Retraction* (IR) contracts accumulated LoRA adaptation (Hu et al., 2022) toward the pretrained reference. Ordinary updates continue identity learning; together, these operations aim to reduce premature commitment to one source-driven path. Let g_t and $\tilde{g}_t$ denote the current loss gradient and optimizer input, respectively. For Moving State LoRA factor ℓ, $\theta_t^{(\ell)}$ is its value at step t, with $\theta_0^{(\ell)} = 0$ as the zero-functional-adaptation reference. LUR transforms the gradient input, whereas IR replaces the usual factor update with contraction:

$$\tilde{g}_t = \begin{cases} g_t, & \text{Normal}, \\ -\gamma_{\mathrm{amp}} g_t, & \text{LUR}, \end{cases} \qquad \theta_{t+1}^{(\ell)} = \theta_t^{(\ell)} - \rho_t^{(\ell)} (\theta_t^{(\ell)} - \theta_0^{(\ell)}) \quad \text{(IR)}, \tag{5}$$

where $\gamma_{\mathrm{amp}} \geq 1$ and $\rho_t^{(\ell)} \in [0, 1]$ control reversal amplification and retraction strength, respectively. LUR applies to all trainable gradients, including router and identity heads; IR contracts only Moving State factors. Normal steps predominate; probabilities, optimizer handling, and gradient-scale-matched retraction appear in Appendix B.2.

3.5 Dual-Space Relation Perturbation

Identity supervision fixes sample-to-class assignments and metric neighborhoods, but learning these source relations does not specify which distinctions will transfer across categories. Prior metric-learning studies show that sample relations contain structure beyond identity equivalence (Kim et al.,

2019; Wang et al., 2019). DSRP therefore complements DTR by selectively varying the supervision that shapes adaptation states. Label perturbation can regularize classification (Xie et al., 2016), whereas incorrect metric labels can disrupt many pairwise relations (Liu et al., 2021). DSRP uses category information and current model scores to guide quota-limited perturbations, retaining ordinary supervision elsewhere.

Classification-space perturbation. For sample i with identity y_i and detached class probability $p_{i,j}$, an alternative identity $j \neq y_i$ is scored by $g^{\text{cls}}_{i,j} = |p_{i,y_i} - p_{i,j}|$. DSRP prioritizes large within-category gaps and small cross-category gaps, changing at most N_q sample targets per operation while leaving the remaining labels unchanged.

Embedding-space perturbation. Classification relabeling does not directly alter the labels used for triplet mining, so DSRP also varies metric neighborhoods. Using Euclidean distance $d_{i,j} = \|u_i - u_j\|_2$ between normalized branch descriptors, it selects farthest same-category and nearest cross-category different-identity pairs. Swapping their triplet labels within quota N_q changes surrounding positive/negative relations while keeping each selected pair different-identity. The four operation quotas share N_q; identity protection and pair-selection constraints appear in Appendix B.3.

3.6 Training Objective and Inference

Each path uses the standard ReID objective $\mathcal{L}_{\text{ReID}} = \mathcal{L}_{\text{ID}} + \mathcal{L}_{\text{Tri}}$, combining identity classification and triplet metric learning with equal weights (Hermans et al., 2017; Luo et al., 2019). The composed path learns reuse while the Moving State produces future alternatives. Once history is available, both receive supervision:

$$\mathcal{L} = \mathcal{L}^{\text{TRR}}_{\text{ReID}} + \mathcal{L}^{\text{mov}}_{\text{ReID}} + \lambda_{\text{chron}}\mathcal{L}_{\text{chron}}, \tag{6}$$

The first two terms are the ReID losses of the composed and Moving State paths, respectively, while λ_{chron} weights chronology regularization. TRR forms the representation, DSRP constructs each path's targets, and DTR modifies subsequent updates.

At inference, the terminal state joins the frozen pool and the moving branch is removed. TRR produces normalized query/gallery descriptors. Shared backbone weights reduce storage but do not eliminate multi-state encoding cost (Appendix B.4).

4 Experiments

4.1 Datasets and Evaluation Metrics

Five training categories support ten evaluation tasks in three regimes: *category-aligned single-category*, *category-shifted single-category*, and *category-shifted multi-category*. We report mean Average Precision (mAP), Rank-1 (R1), and mean Inverse Negative Penalty (mINP). Dataset statistics, category groupings, and retrieval protocols appear in Appendix B.5.

4.2 Implementation Details

To isolate how identity supervision reshapes a generic visual representation, our experiments require a strong but minimally engineered reference whose pretraining prior introduces as little task-specific bias as possible. We therefore instantiate **Generic Visual Adaptation (GVA)** with DINOv2 ViT-B/14 (Oquab et al., 2023) and rank-128 Q/K/V LoRA (Hu et al., 2022) in all transformer blocks. DINOv2 provides a suitable visual foundation because its image-only self-supervised pretraining learns broadly transferable structure without explicit language alignment (Radford et al., 2021) or optimization toward a fixed supervised label taxonomy (Dosovitskiy et al., 2021), reducing additional semantic biases that could become entangled with the changes induced by ReID supervision. LoRA then enables network-wide identity adaptation while keeping the pretrained backbone fixed, preserving this visual reference throughout training. Together, these choices provide a strong, minimally engineered reference for isolating ReState's effects. GVA and ReState share the same backbone, adapters, image resolution, and 5,000-step training budget; ReState additionally caches states every 500 steps with $\gamma_{\text{amp}} = N_q = 4$. Further details are provided in Appendix B.4.

Table 2: Comparison with state-of-the-art methods across ten heterogeneous evaluation settings. Methods are grouped into category-specific, domain-generalizable, and object-general ReID.

Method	MSMT17			VehicleID			ShipReID-2400			ELPephants			PetFace		
	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
TransReID(He et al., 2021)	13.3	0.8	31.8	48.1	48.1	38.9	2.7	1.2	3.0	11.7	**2.9**	34.5	28.2	26.6	25.9
PASS-ReID(Zhu et al., 2022)	4.1	0.2	11.6	37.3	37.3	29.1	1.6	0.7	1.6	9.6	2.2	26.7	22.7	21.3	20.9
PHA(Zhang et al., 2023a)	8.4	0.5	21.5	44.5	44.5	35.3	3.1	1.4	3.1	10.8	2.4	32.0	25.6	24.1	23.6
CLIP-ReID(Li et al., 2023)	20.3	**1.6**	45.1	54.6	54.6	44.3	3.4	1.4	3.7	9.3	2.0	27.2	23.6	22.2	21.1
PAT(Ni et al., 2023)	15.5	0.9	37.5	49.4	49.4	40.6	4.0	1.6	5.4	12.1	2.7	33.8	36.8	35.1	33.7
BAU(Cho et al., 2024)	1.8	0.1	6.6	31.6	31.6	21.7	0.6	0.3	0.7	3.7	1.9	5.9	1.9	1.9	1.6
ReNorm(Nie et al., 2024)	7.0	0.2	24.1	53.7	53.7	45.4	1.7	0.5	2.1	6.8	2.1	16.2	18.2	17.3	16.8
DACS(Yang et al., 2024a)	5.0	0.3	14.5	37.0	37.0	29.9	2.6	1.2	2.7	10.7	2.3	30.8	28.4	26.9	25.9
PromptSG(Yang et al., 2024b)	17.7	1.1	40.4	56.4	56.4	47.0	4.0	1.6	5.7	9.9	2.8	25.4	21.1	19.7	18.9
Instruct-ReID(He et al., 2024)	14.0	0.5	37.6	48.0	48.0	38.4	0.8	0.3	1.1	7.4	2.0	20.0	19.3	18.2	17.8
CILP-FGDI(Zhao et al., 2025)	15.3	0.7	37.7	54.1	54.1	43.6	2.1	0.8	3.3	6.8	2.0	19.5	13.6	12.9	11.9
ADSR(Chen et al., 2025)	0.7	0.1	3.3	30.2	30.2	24.1	1.1	0.4	1.7	7.7	2.2	19.3	19.1	17.9	17.5
VICP(Huang & Liu, 2025)	15.6	0.7	40.2	**57.5**	**57.5**	**49.4**	4.2	1.6	6.2	12.0	2.6	33.5	30.6	29.1	27.8
OG-ReID(Chen et al., 2026)	16.7	1.1	38.7	51.3	51.3	42.9	4.5	**1.9**	5.7	**13.3**	2.8	**38.4**	39.8	38.1	36.6
Ours	**23.2**	1.5	**53.2**	55.6	55.6	47.7	**4.8**	**1.9**	**7.5**	13.2	2.4	36.4	**45.6**	**44.0**	**41.7**

Method	University D→S			University S→D			DenseUAV			CUTE			ILIAS		
	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
TransReID(He et al., 2021)	17.8	17.8	12.8	10.7	1.3	26.8	1.3	0.2	1.5	20.2	0.6	84.2	10.6	1.9	23.6
PASS-ReID(Zhu et al., 2022)	11.7	11.7	7.4	6.3	0.5	25.5	0.8	0.1	0.9	14.2	0.6	78.5	4.2	0.5	9.9
PHA(Zhang et al., 2023a)	15.7	15.7	11.1	8.5	1.0	22.5	0.8	0.1	0.9	20.1	0.6	84.8	9.3	1.5	22.5
CLIP-ReID(Li et al., 2023)	25.2	25.2	17.6	14.4	1.7	30.7	1.3	0.2	1.8	24.7	0.9	84.8	19.1	3.7	37.7
PAT(Ni et al., 2023)	22.0	22.0	15.7	12.8	1.5	**34.8**	1.5	0.2	1.6	25.0	0.7	85.2	17.2	3.1	37.3
BAU(Cho et al., 2024)	1.7	1.7	0.5	0.3	0.1	1.0	0.2	0.1	0.1	7.7	0.6	59.2	0.6	0.2	0.7
ReNorm(Nie et al., 2024)	6.7	6.7	4.0	2.9	0.1	20.8	0.4	0.1	0.4	16.4	0.6	79.7	3.1	0.4	7.5
DACS(Yang et al., 2024a)	19.8	19.8	13.8	12.1	1.2	31.1	1.5	0.2	2.1	18.4	0.6	80.4	11.5	2.4	24.7
PromptSG(Yang et al., 2024b)	27.0	27.0	19.0	16.4	2.9	**34.8**	1.2	0.2	1.9	23.9	0.7	85.7	18.4	4.3	36.3
Instruct-ReID(He et al., 2024)	3.5	3.5	1.8	1.3	0.3	12.3	0.6	0.1	1.0	10.4	0.6	69.6	2.0	0.4	4.5
CILP-FGDI(Zhao et al., 2025)	8.2	8.2	4.7	3.0	0.1	17.0	0.3	0.1	0.3	27.0	1.1	82.1	6.2	0.9	14.4
ADSR(Chen et al., 2025)	7.9	7.9	4.2	3.0	0.3	13.8	0.5	0.1	0.4	13.3	0.6	76.4	3.2	0.7	6.6
VICP(Huang & Liu, 2025)	17.6	17.6	11.2	8.9	1.0	22.1	3.0	0.6	3.5	29.1	1.3	89.6	26.5	6.3	52.2
OG-ReID(Chen et al., 2026)	23.5	23.5	17.1	15.6	3.0	34.4	2.8	0.5	3.0	31.0	2.1	87.9	29.4	8.2	54.3
Ours	**32.7**	**32.7**	**23.9**	**18.3**	**4.9**	32.0	**5.3**	**1.6**	**5.0**	**53.3**	**16.4**	**91.8**	**41.1**	**14.1**	**67.1**

Table 3: Component ablation of ReState across five selected evaluation settings.

Components			ShipReID-2400			University D→S			PetFace			CUTE			ILIAS		
TRR	DSRP	DTR	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
×	×	×	4.3	1.8	6.7	23.8	23.8	16.3	41.7	40.1	37.8	50.0	16.2	90.4	36.0	11.8	61.0
✓	×	×	4.5	1.8	6.8	27.8	27.8	19.5	43.3	41.7	39.4	51.2	16.1	90.9	38.2	12.7	63.7
×	✓	×	4.3	1.8	6.2	26.9	26.9	18.6	42.4	40.8	38.6	51.3	17.0	91.1	37.6	12.6	63.1
×	×	✓	4.6	**1.9**	**7.5**	28.9	28.9	20.8	44.2	42.6	40.3	53.1	**18.1**	91.6	39.9	13.4	65.8
✓	✓	✓	**4.8**	**1.9**	**7.5**	**32.7**	**32.7**	**23.9**	**45.6**	**44.0**	**41.7**	**53.3**	16.4	**91.8**	**41.1**	**14.1**	**67.1**

4.3 Comparison with SOTA ReID Methods

Table 2 compares complete ReID systems across ten heterogeneous evaluation settings. ReState leads mAP on eight tasks, with different advantages across the three transfer regimes.

Category-aligned single-category transfer. ReState retains an explicit pretrained reference during continued identity learning. Selective historical composition lets deployment use learned distinctions without inheriting one state's entire accumulated adaptation. Its improvement over CLIP-ReID (Li et al., 2023) on MSMT17 illustrates this formulation's value within familiar categories, although VICP remains stronger on VehicleID.

Category-shifted single-category transfer. ReState learns historical reuse from training identities and applies it directly to test images from unfamiliar categories and acquisition conditions. VICP (Huang & Liu, 2025) instead conditions retrieval on a small set of target-task visual examples. ReState's advantage on DenseUAV demonstrates useful cross-view transfer without requiring such task-specific visual support.

Category-shifted multi-category transfer. OG-ReID (Chen et al., 2026) learns unified cross-category representations; ReState additionally learns the formation and reuse of historical alternatives without assigning states to predefined target categories. Gains on CUTE and ILIAS and leading results on PetFace demonstrate practical value across heterogeneous identity cues.

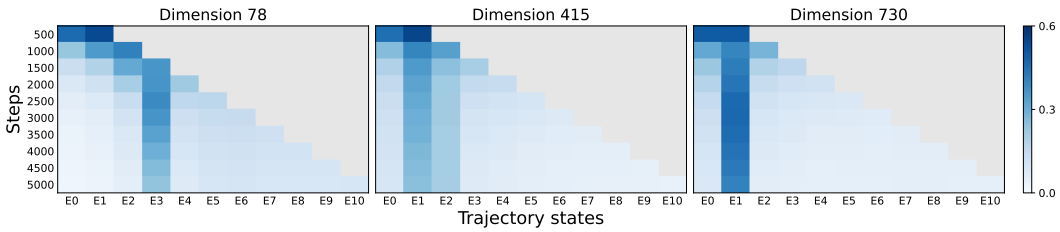

Figure 5: **Historical states remain active with dimension-specific temporal preferences.** Mean routing weights over ten evaluation settings for three randomly selected feature dimensions.

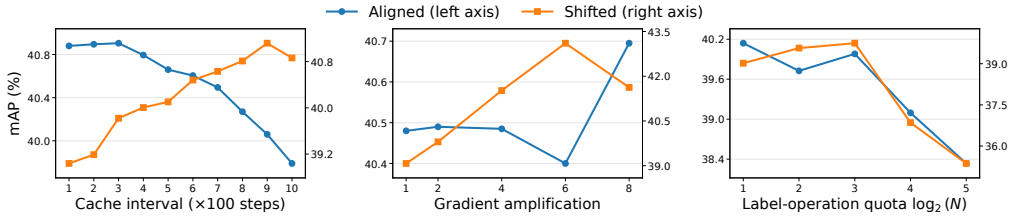

(a) TRR (b) DTR (c) DSRP

Figure 6: **Retention and amplification preferences differ across transfer groups, while large supervision-perturbation quotas degrade both.** Each panel adds only the indicated module to GVA, the component-free baseline. Aligned (left axis) is equal-task mean mAP over MSMT17 and VehicleID; Shifted (right axis) averages PetFace, University D→S, CUTE, and ILIAS.

4.4 Ablation Study

Component effectiveness. Table 3 shows complementary roles for reuse, optimization, and supervision. TRR improves mAP on all five tasks, demonstrating useful historical alternatives beyond a single adapted state. DTR gives the strongest standalone gains on the shifted tasks, highlighting the importance of shaping the trajectory that produces reusable states. DSRP also improves shifted retrieval, supporting variation in source supervision as another way to form transferable states. Full ReState achieves the highest mAP on all four shifted tasks shown, supporting joint regulation of state formation and reuse. Module-specific controls appear in Appendix B.7.

Trajectory-state routing. Figure 5 shows historical states remaining active after newer states enter the pool. The displayed coordinate averages favor different stages: some retain substantial early-state weight, while others emphasize intermediate states. Later checkpoints thus expand the available history without uniformly replacing earlier contributions. Appendix B.7 further compares state-wise and static coordinate-wise routing.

Parameter sensitivity. Figure 6 reveals group-dependent operating preferences. For TRR, shifted performance rises and then falls with longer cache intervals, while aligned performance eventually declines: useful historical coverage depends on the retention schedule. DTR favors moderate amplification for shifted transfer and stronger amplification for aligned retrieval, supporting controlled trajectory variation. Excessive DSRP quotas degrade both groups, supporting selective perturbation that retains ordinary identity supervision. These patterns favor balancing history retention, update variation, and supervision perturbation. Appendix B.6 details the sweeps.

5 Conclusion

PFR reveals that continued identity learning does not simply replace earlier representation structure with a uniformly better one. Instead, adaptation progressively reorganizes feature utility, allowing later states to gain source discrimination while historical structure can remain useful for different transfer settings and queries. This suggests that adaptation history itself can be a resource for generalizable retrieval rather than merely a sequence of intermediate checkpoints. ReState operationalizes this view by retaining trajectory alternatives while jointly regulating how they are formed and selectively composed for each input. Its consistent improvements over a strong baseline across all seven category-shifted tasks further show that generalization benefits from considering both the formation and reuse of adaptation states, rather than relying solely on the final adapted representation. In future, we will explore adaptive state retention beyond fixed intervals.

AI USE STATEMENT

We used generative AI tools for draft generation and writing polish, including clarifying descriptions of the methods and existing experimental results. AI was not used to implement the research methods, process datasets, or generate experimental data. The authors remain responsible for the accuracy of the manuscript and its claims.

ETHICS STATEMENT

This work studies object re-identification using existing research benchmarks and does not collect new human-subject data or release new identity data. ReID technology, particularly when applied to people, can raise privacy, surveillance, and bias concerns beyond controlled benchmark evaluation. Our experiments are intended for methodological research; real-world use should follow applicable requirements for data access, privacy, consent, fairness, and responsible deployment.

REPRODUCIBILITY STATEMENT

We provide the definitions, objectives, and training procedure needed to reproduce ReState in the main text and Appendix B. Appendices A and B further document the diagnostic protocol, uncertainty estimation, dataset and retrieval protocols, implementation details, parameter sweeps, and extended ablations. Experimental settings and matched baselines are specified alongside their corresponding analyses to facilitate reproducible comparison.

REFERENCES

Long Chen, De Cheng, Shizhou Zhang, Yinghui Xing, Di Xu, and Yanning Zhang. Amplitude-aware domain style replay for lifelong person re-identification. In *ACM MM*, pp. 2026–2035, 2025.

Shuoyi Chen, Yurui Wu, and Mang Ye. Object-generalized re-identification: A step towards universal instance perception. In *CVPR*, 2026.

Yoonki Cho, Jaeyoon Kim, Woo Jae Kim, Junsik Jung, and Sung eui Yoon. Generalizable person re-identification via balancing alignment and uniformity. In *NeurIPS*, 2024.

Seokeon Choi, Taekyung Kim, Minki Jeong, Hyoungseob Park, and Changick Kim. Meta batch-instance normalization for generalizable person re-identification. In *CVPR*, pp. 3425–3435, 2021.

Ming Dai, Enhui Zheng, Zhenhua Feng, Lei Qi, Jiedong Zhuang, and Wankou Yang. Vision-based uav self-positioning in low-altitude urban environments. *IEEE TIP*, 33:493–508, 2024. doi: 10.1109/TIP.2023.3346279.

Yongxing Dai, Xiaotong Li, Jun Liu, Zekun Tong, and Ling-Yu Duan. Generalizable person re-identification with relevance-aware mixture of experts. In *CVPR*, 2021.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

Jinwei Han, Zhiwen Lin, Zhongyisun Sun, Yingguo Gao, Ke Yan, Shouhong Ding, Yuan Gao, and Gui-Song Xia. Anchor-based robust finetuning of vision-language models. In *CVPR*, pp. 26919–26928, 2024.

Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *ICCV*, pp. 15013–15022, 2021.

Weizhen He, Yiheng Deng, Shixiang Tang, Qihao Chen, Qingsong Xie, Yizhou Wang, Lei Bai, Feng Zhu, Rui Zhao, Wanli Ouyang, Donglian Qi, and Yunfeng Yan. Instruct-reid: A multi-purpose person re-identification task with instructions. In *CVPR*, pp. 17521–17531, June 2024.

Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.

Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E. Hopcroft, and Kilian Q. Weinberger. Snapshot ensembles: Train 1, get m for free. In *ICLR*, 2017.

Zhizhong Huang and Xiaoming Liu. Generalizable object re-identification via visual in-context prompting. In *ICCV*, 2025.

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *ICLR*, 2023.

Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *CVPR*, pp. 3143–3152, 2020.

Sungyeon Kim, Minkyo Seo, Ivan Laptev, Minsu Cho, and Suha Kwak. Deep metric learning beyond binary supervision. In *CVPR*, pp. 2288–2297, 2019.

Giorgos Kordopatis-Zilos, Vladan Stojnić, Anna Manko, Pavel Šuma, Nikolaos-Antonios Ypsilantis, Nikos Efthymiadis, Zakaria Laskar, Jiří Matas, Ondřej Chum, and Giorgos Tolias. Ilias: Instance-level image retrieval at scale. In *CVPR*, 2025.

Matthias Körschens and Joachim Denzler. ELPephants: A fine-grained dataset for elephant re-identification. In *Int. Conf. Comput. Vis. Worksh.*, 2019.

Klemen Kotar, Stephen Tian, Hong-Xing Yu, Daniel L.K. Yamins, and Jiajun Wu. Are these the same apple? comparing images based on object intrinsics. In *NeurIPS*, 2023.

Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *ICLR*, 2022.

Yoonho Lee, Annie S. Chen, Fahim Tajwar, Ananya Kumar, Huaxiu Yao, Percy Liang, and Chelsea Finn. Surgical fine-tuning improves adaptation to distribution shifts. In *ICLR*, 2023.

Shuyuan Li, Jianguo Li, Hanlin Tang, Rui Qian, and Weiyao Lin. Atrw: a benchmark for amur tiger re-identification in the wild. *arXiv preprint arXiv:1906.05586*, 2019.

Siyuan Li, Li Sun, and Qingli Li. Clip-reid: exploiting vision-language model for image re-identification without concrete text labels. In *AAAI*, volume 37, pp. 1405–1413, 2023.

Xuhong Li, Yves Grandvalet, and Franck Davoine. Explicit inductive bias for transfer learning with convolutional networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2825–2834, 2018.

Baolong Liu, Roukai Huang, Xin Pan, Chuanhuang Li, Jie Sun, Jianfeng Dong, and Xun Wang. Advancing ship re-identification in the wild: The shipreid-2400 benchmark dataset and d2internet baseline method. In *SIGIR*, pp. 106–115, 2025.

Chang Liu, Han Yu, Boyang Li, Zhiqi Shen, Zhanning Gao, Peiran Ren, Xuansong Xie, Lizhen Cui, and Chunyan Miao. Noise-resistant deep metric learning with ranking-based instance selection. In *CVPR*, pp. 6811–6820, 2021.

Hongye Liu, Yonghong Tian, Yaowei Wang, Lu Pang, and Tiejun Huang. Deep relative distance learning: Tell the difference between similar vehicles. In *CVPR*, pp. 2167–2175, 2016.

Xinchen Liu, Wu Liu, Tao Mei, and Huadong Ma. Provid: Progressive and multimodal vehicle reidentification for large-scale urban surveillance. *IEEE TMM*, 20(3):645–658, 2018.

Zhenyi Lu, Chenghao Fan, Wei Wei, Xiaoye Qu, Dangyang Chen, and Yu Cheng. Twin-merging: Dynamic integration of modular expertise in model merging. In *NeurIPS*, volume 37, 2024.

Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *CVPRW*, 2019.

Hao Ni, Yuke Li, Lianli Gao, Heng Tao Shen, and Jingkuan Song. Part-aware transformer for generalizable person re-identification. In *ICCV*, pp. 11280–11289, 2023.

Ren Nie, Jin Ding, Xue Zhou, and Xi Li. Rethinking normalization layers for domain generalizable person re-identification. In *ECCV*, pp. 267–284. Springer, 2024.

Changdae Oh, Yixuan Li, Kyungwoo Song, Sangdoo Yun, and Dongyoon Han. Dawin: Training-free dynamic weight interpolation for robust adaptation. In *ICLR*, 2025.

Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pp. 8748–8763. PMLR, 2021.

Risa Shinoda and Kaede Shiohara. Petface: A large-scale dataset and benchmark for animal identification. In *ECCV*, pp. 19–36. Springer, 2024.

Jifei Song, Yongxin Yang, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. Generalizable person re-identification by domain-invariant mapping network. In *CVPR*, pp. 719–728, 2019.

Le Wang, Rizhi Ding, Yuanhao Zhai, Qilin Zhang, Wei Tang, Nanning Zheng, and Gang Hua. Giant panda identification. *IEEE TIP*, 30:2837–2849, 2021.

Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R. Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *CVPR*, pp. 5022–5030, 2019.

Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, pp. 79–88, 2018.

Changsong Wen, Zelin Peng, Yu Huang, Xiaokang Yang, and Wei Shen. Domain generalization in clip via learning with diverse text prompts. In *CVPR*, pp. 9559–9569, 2025.

Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, volume 162, pp. 23965–23998, 2022a.

Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo-Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. Robust fine-tuning of zero-shot models. In *CVPR*, pp. 7959–7971, 2022b.

Lingxi Xie, Jingdong Wang, Zhen Wei, Meng Wang, and Qi Tian. Disturblabel: Regularizing cnn on the loss layer. In *CVPR*, pp. 4753–4762, 2016.

Fengxiang Yang, Zhun Zhong, Zhiming Luo, Yifan He, Shaozi Li, and Nicu Sebe. Diversity-authenticity co-constrained stylization for federated domain generalization in person re-identification. In *AAAI*, volume 38, pp. 6477–6485, 2024a.

Zexian Yang, Dayan Wu, Chenming Wu, Zheng Lin, Jingzi Gu, and Weiping Wang. A pedestrian is worth one prompt: Towards language guidance person re-identification. In *CVPR*, pp. 17343–17353, 2024b.

Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven C. H. Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE TPAMI*, 44(6):2872–2893, 2022.

Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *NeurIPS*, 2014.

Guiwei Zhang, Yongfei Zhang, Tianyu Zhang, Bo Li, and Shiliang Pu. Pha: Patch-wise high-frequency augmentation for transformer-based person re-identification. In *CVPR*, pp. 14133–14142, 2023a.

Qian Zhang, Mingxin Zhang, Jinghe Liu, Xuanyu He, Ran Song, and Wei Zhang. Unsupervised maritime vessel re-identification with multi-level contrastive learning. *IEEE Transactions on Intelligent Transportation Systems*, 24(5):5406–5418, 2023b. doi: 10.1109/TITS.2023.3243591.

Huazhong Zhao, Lei Qi, and Xin Geng. Cilp-fgdi: Exploiting vision-language model for generalizable person re-identification. *IEEE TIFS*, 2025.

Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, pp. 1116–1124, 2015.

Zhedong Zheng, Yunchao Wei, and Yi Yang. University-1652: A multi-view multi-source benchmark for drone-based geo-localization. In *ACM MM*, pp. 1395–1403, 2020.

Beier Zhu, Jiequan Cui, and Hanwang Zhang. Robust fine-tuning of zero-shot models via variance reduction. In *NeurIPS*, 2024.

Kuan Zhu, Haiyun Guo, Tianyi Yan, Yousong Zhu, Jinqiao Wang, and Ming Tang. Pass: Part-aware self-supervised pre-training for person re-identification. In *ECCV*, 2022.

A Diagnostic Protocol and Extended Evidence

A.1 Common Trajectory and Uncertainty

Common diagnostic trajectory. All PFR diagnostics follow one baseline source-only adaptation trajectory without ReState. At optimization step t, encoder E_t produces features $z_t(x) = E_t(x)$ for image x, with E_0 as the pretrained encoder. We use the unit-normalized features $\hat{z}_t(x) = z_t(x)/\|z_t(x)\|_2$. Checkpoints are evaluated every 500 steps through the final step $T = 5000$, with $t_e = 500$ as the early reference. The trajectory, source-calibration population, and evaluation manifests are fixed throughout the investigation, connecting early accessibility, subsequent functional change, and historical restoration. Evaluation covers five source datasets and ten target retrieval tasks, together with the identity-disjoint CUTE category-readout task.

The diagnostic trajectory and the GVA component-ablation baseline use separate experimental configurations. Their absolute scores should therefore not be compared directly: each diagnostic intervention is compared with the unmodified representation from its own trajectory, while component gains use the matched baseline in Table 3.

Uncertainty estimates. We report paired 95% percentile intervals from 1,000 identity-cluster bootstrap replicates. Retrieval resamples query identities. For CUTE-10, identities are resampled within each object category, retaining all images of each sampled identity and pooling them into an image-weighted accuracy. Category-level changes are descriptive; intervals are reported for the aggregate probe results. These intervals describe evaluation uncertainty conditional on the fitted representations and fixed gallery; they do not measure variability across training seeds, router initializations, or perturbation histories. Empirical cumulative distribution functions (ECDFs) instead describe variation among individual responses.

A.2 Source Calibration and Subspace Construction

To compare later representations against one fixed early reference, we fit the subspace, alignment, and restoration quantities using source calibration only.

Calibration. Within each source, a fixed half of the training identities supplies calibration images, balanced across identities; the remainder supplies readout training and validation. Calibration images remain fixed across checkpoints. Source evaluation and target images, including all CUTE probe images, are excluded from fitting the shared subspace, coordinate alignment, and restoration scale.

Early discriminative span. We construct a source-balanced Fisher subspace from normalized features at t_e:

$$S_B v_j = \gamma_j (S_W + 10^{-3} I) v_j, \qquad \gamma_1 \geq \cdots \geq \gamma_p, \tag{7}$$

where S_B and S_W are between-identity and within-identity scatter matrices, respectively, computed within each source, normalized by its calibration sample count, and averaged equally across sources. The vector v_j is a generalized eigenvector, and γ_j is its eigenvalue. We retain the smallest rank r accounting for 80% of the nonnegative eigenvalue sum and orthonormalize the retained vectors to obtain Q_A. This defines the fixed projectors in Eq. 1; the rank criterion measures discriminative scatter rather than raw feature variance. The descriptor dimension is $p = 768$.

Coordinate alignment. We apply source-fitted orthogonal alignment:

$$u_t(x) = R_{t \to e} \hat{z}_t(x). \tag{8}$$

where $R_{t \to e} \in \mathbb{R}^{p \times p}$ is the orthogonal transformation fitted on paired source calibration features to align step t with the early reference t_e, and $u_t(x)$ is the resulting aligned feature. Alignment preserves full-feature cosine retrieval while enabling comparison through the fixed projectors.

A.3 Early Source Alignment and Frozen Readouts

We first ask how much identity structure is accessible before encoder adaptation and how its accessibility relates to the rapid early gains. The expanded results distinguish shared from source-specific accessibility.

Table 4: **Pretrained identity accessibility is source-dependent despite rapid early adaptation.** Expanded retrieval scores, shared and source-specific readout gains, readout ratios, and early/late adaptation rates for the five sources.

Dataset	mAP_0	$\mathrm{mAP}^{\mathrm{lin}}$	$\mathrm{mAP}^{\mathrm{lin}}_{\mathrm{src}}$	$\mathrm{mAP}_{\mathrm{early}}$	G^{lin}	$G^{\mathrm{lin}}_{\mathrm{src}}$	ΔG_{src}	G_{early}	η_{read}	η_{src}	v_{early}	v_{late}
Market-1501	1.67	2.90	8.45	61.22	1.23	6.78	5.55	59.55	0.02	0.11	0.1191	0.0043
iPanda-50	18.51	17.82	21.59	28.54	-0.70	3.08	3.78	10.03	-0.07	0.31	0.0201	0.0019
VeRi	4.95	9.73	16.82	52.85	4.78	11.86	7.08	47.90	0.10	0.25	0.0958	0.0040
ATRW	47.76	59.89	63.24	55.34	12.13	15.47	3.35	7.58	1.60	2.04	0.0152	0.0014
VesselReID	26.45	46.22	45.00	51.25	19.77	18.55	-1.22	24.80	0.80	0.75	0.0496	0.0030

Metrics. We measure early retrieval gains and compare the average pace of adaptation before and after the early reference:

$$G_s^{\mathrm{early}} = P_s(t_e) - P_s(0), \tag{9}$$

$$v_s^{\mathrm{early}} = \frac{G_s^{\mathrm{early}}}{t_e}, \qquad v_s^{\mathrm{late}} = \frac{P_s(T) - P_s(t_e)}{T - t_e}. \tag{10}$$

where $P_s(t) = P_{\mathcal{D}_s}(\hat{z}_t)$ is retrieval mAP for source dataset $\mathcal{D}_s$ at step t, s indexes source datasets, and t_e is the early reference step with $0 < t_e < T$. More generally, $P_{\mathcal{D}}(R)$ denotes retrieval mAP on the 0–100 scale for dataset $\mathcal{D}$ using representation R. Here G_s^{early} is the gain over the pretrained state, while v_s^{early} and v_s^{late} are the early and late average gain rates. The rates measure mAP improvement in percentage points per optimization step over the two intervals.

To measure identity structure accessible before encoder adaptation, we fit shared and source-specific linear readouts while keeping E_0 frozen. Their retrieval descriptors are

$$f_{\mathrm{sh}}(x) = \mathrm{normalize}(W_{\mathrm{sh}}\hat{z}_0(x)), \qquad f_s(x) = \mathrm{normalize}(W_s\hat{z}_0(x)). \tag{11}$$

where $W_{\mathrm{sh}} \in \mathbb{R}^{p\times p}$ is the linear map jointly fitted across sources, $W_s \in \mathbb{R}^{p\times p}$ is fitted independently for source s, and f_{sh} and f_s are their respective retrieval representations. We quantify their gains over frozen pretrained retrieval and compare the source-specific gain with early encoder adaptation:

$$G_s^{\mathrm{lin}} = P_{\mathcal{D}_s}(f_{\mathrm{sh}}) - P_s(0), \qquad G_{s,\mathrm{src}}^{\mathrm{lin}} = P_{\mathcal{D}_s}(f_s) - P_s(0), \tag{12}$$

$$\eta_{s,\mathrm{src}} = \frac{G_{s,\mathrm{src}}^{\mathrm{lin}}}{G_s^{\mathrm{early}}}. \tag{13}$$

where G_s^{lin} and $G_{s,\mathrm{src}}^{\mathrm{lin}}$ are the shared and source-specific readout gains, respectively, and $\eta_{s,\mathrm{src}}$ compares the source-specific gain with early encoder adaptation. Table 4 additionally reports $\Delta G_{\mathrm{src}} = G_{s,\mathrm{src}}^{\mathrm{lin}} - G_s^{\mathrm{lin}}$ and $\eta_{\mathrm{read}} = G_s^{\mathrm{lin}}/G_s^{\mathrm{early}}$. Both ratios compare measured gains and can exceed one.

Readout fitting. The readouts in Eq. 11 are identity-initialized linear maps fitted on frozen normalized pretrained features. Training uses a source-balanced metric objective with hard positive and negative pairs selected within each source. Shared and source-specific maps receive the same 1,000-update budget with AdamW. Selection uses source validation mAP, averaged across sources for the shared map and evaluated separately for each source-specific map.

Expanded evidence. Table 4 extends Table 1 with absolute scores that separate shared from source-specific accessibility. Source-specific fitting exposes more identity structure than the shared map on four sources, while VesselReID favors the shared map; both readout gains and their relation to early adaptation vary substantially. Figures 7(a,b) show rapid initial progress followed by slower improvement across all five sources despite their different absolute trajectories. Together, these results characterize source-dependent accessibility within the fitted linear family and the rapid early development along this decreasing-learning-rate trajectory. They establish the pretrained starting point and early reference for the subsequent analysis.

A.4 Complementary Margins and Feature Movement

Figure 2(a) establishes increasing retrieval utility in the fixed complement. We next ask whether this development also appears in the identity relations used by the full descriptor and how feature movement is distributed between the two spans.

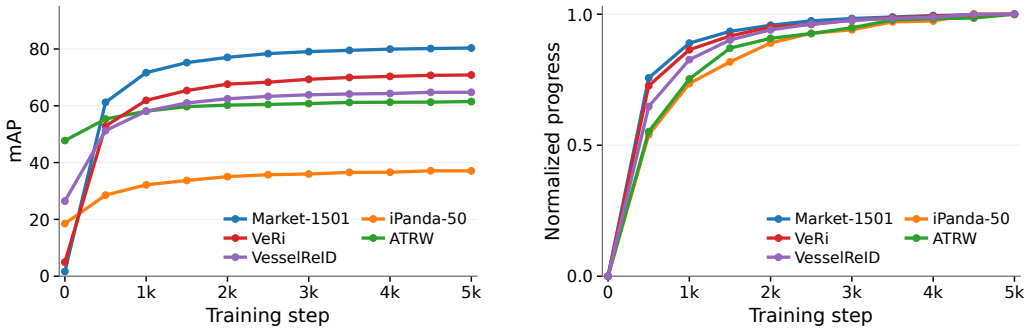

(a) Source retrieval trajectories (b) Normalized source learning progress

Figure 7: **Source retrieval improves rapidly early and more slowly thereafter.** (a) Full-feature mAP trajectories for the five sources. (b) The same trajectories normalized by each source's total improvement from step 0 to 5000.

Projected retrieval. Projected retrieval uses

$$z_t^B(x) = \text{normalize}(P_B u_t(x)), \qquad P_s^B(t) = P_{\mathcal{D}_s}(z_t^B). \tag{14}$$

where $B \in \{A, C\}$ selects a span, P_B is its projector, $z_t^B(x)$ is its separately normalized retrieval descriptor, and $P_s^B(t)$ is the resulting source retrieval mAP. We compare each representation with its own early score and average across sources:

$$\Delta P_s^B(t) = P_s^B(t) - P_s^B(t_e), \qquad \overline{\Delta P}^B(t) = \frac{1}{S}\sum_{s=1}^{S} \Delta P_s^B(t). \tag{15}$$

where $\Delta P_s^B(t)$ is the within-span retrieval gain relative to t_e, S is the number of source datasets, and $\overline{\Delta P}^B(t)$ is their equal-weight macro-average gain. Positive $\Delta P_s^C(t)$ means that the same fixed complementary span becomes more useful for source identity retrieval than it was at the early reference. The Full curve uses u_t without projection and the same own-baseline convention.

Fixed-triplet margins. To connect aggregate projected retrieval with individual identity relations in the full descriptor, we freeze difficult triplets at t_e and track their span-wise margin contributions:

$$m_{q,t}^B = u_t(q)^\top P_B\big(u_t(p_q) - u_t(n_q)\big), \qquad B \in \{A, C\}, \tag{16}$$

where p_q and n_q are the least similar valid positive and most similar different-identity negative for query q at t_e, held fixed thereafter, and $m_{q,t}^B$ is the similarity margin contributed by span B. Without component renormalization, the two margins sum to the full-feature margin. Figure 8 shows that complementary margins improve on Market-1501, iPanda-50, and VeRi, decrease on VesselReID, and have an interval crossing zero on ATRW. These source-dependent responses show that later complementary development affects individual identity relations heterogeneously. They extend the aggregate projected-retrieval finding to the unprojected descriptor, where A and C jointly determine each fixed triplet's margin.

Movement energy. Finally, we ask where subsequent feature movement occurs, independently of whether it improves a fixed triplet. Per-dimension energies account for the different sizes of the two spans:

$$A_{A,s} = \frac{\text{mean}_i \, \|P_A \Delta u_i\|_2^2}{r}, \qquad A_{C,s} = \frac{\text{mean}_i \, \|P_C \Delta u_i\|_2^2}{p - r}, \qquad q_s = \frac{A_{C,s}}{A_{A,s} + 10^{-12}}, \tag{17}$$

where $\Delta u_i = u_{t_b}(x_i) - u_{t_a}(x_i)$ is displacement between adjacent 500-step checkpoints. The mean uses fixed source-s evaluation images; $A_{A,s}$ and $A_{C,s}$ are per-dimension energies.

The ratio q_s compares per-dimension movement in C with that in A. Figures 9(a,b) show decreasing absolute movement in both spans, while the ratio in panel (c) generally rises from its first interval and remains below one. Later adaptation therefore slows overall while movement becomes relatively more prominent beyond the early discriminative span. Combined with the margin analysis, this distinguishes where representation change occurs from how that change affects retrieval.

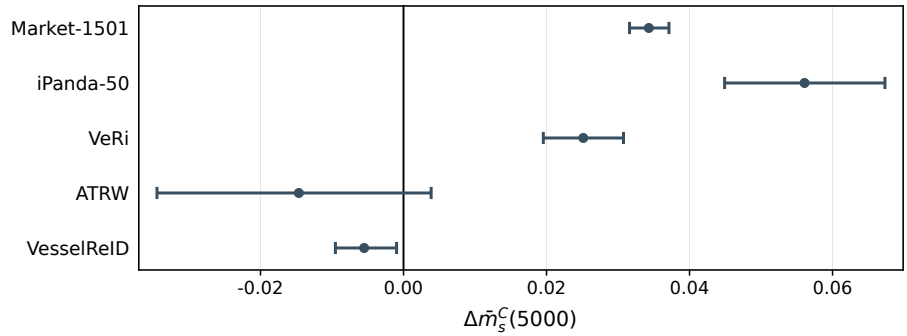

Figure 8: **Later complementary development changes identity margins in a source-dependent manner.** Mean changes from step 500 to 5000 in the complementary-span contribution to fixed-triplet cosine margins; whiskers denote paired 95% identity-cluster bootstrap intervals.

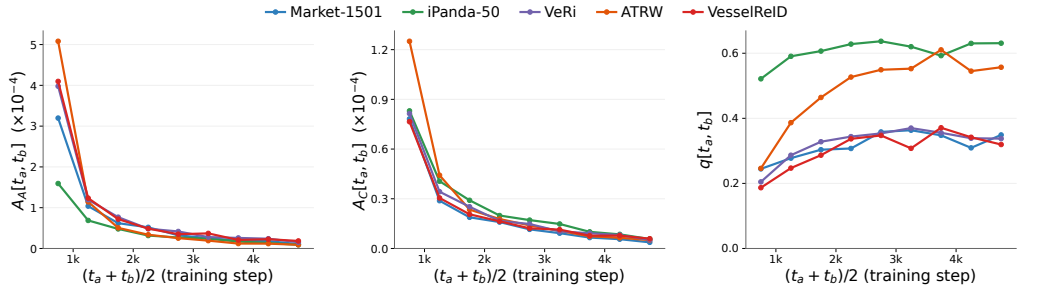

(a) Early-span movement energy (b) Complementary-span movement energy (c) Relative complementary movement

Figure 9: **Later adaptation slows overall while complementary movement becomes relatively more prominent.** Per-dimension movement over adjacent 500-step intervals: (a) early-span energy A_A, (b) complementary-span energy A_C, and (c) relative complementary movement q, comparing the two per-dimension energies. Horizontal positions denote interval midpoints.

A.5 Independent Category Readouts

CUTE-10 uses the ten CUTE categories with ten object identities each, permitting a fixed three-way identity-disjoint split. Each category contributes five identities for classifier fitting, two for validation, and three for testing, yielding 4,904 fitting, 1,968 validation, and 2,932 test images. All images of an identity remain in the same split, and the splits are shared across checkpoints. This is a category-classification diagnostic on a retrieval target dataset, not a separate diagnostic dataset.

All probes use linear classifiers with bias, fitted by regularized cross-entropy using full-batch L-BFGS and validation-selected regularization. Inputs are $P_C u_t$ and u_t, respectively, without post-projection normalization. Probe fitting leaves the encoder and diagnostic subspace unchanged.

Readout trajectories. We compare fixed and refitted category classifiers:

$$\begin{aligned}
\mathcal{A}_C^{\text{fixed}}(t) &= \text{Acc}_{\mathcal{D}_{\text{probe}}}\left(g_{t_e}^C(P_C u_t(x)), c\right), \\
\mathcal{A}_C^{\text{refit}}(t) &= \text{Acc}_{\mathcal{D}_{\text{probe}}}\left(g_t^C(P_C u_t(x)), c\right), \\
\mathcal{A}_{\text{full}}^{\text{refit}}(t) &= \text{Acc}_{\mathcal{D}_{\text{probe}}}\left(g_t^{\text{full}}(u_t(x)), c\right).
\end{aligned} \tag{18}$$

where $g_{t_e}^C$ is the linear category classifier fitted on early complementary features, g_t^C is independently refitted at step t, and g_t^{full} is refitted on the full aligned representation. The operator $\text{Acc}_{\mathcal{D}_{\text{probe}}}$ computes image-weighted classification accuracy on the fixed held-out probe set $\mathcal{D}_{\text{probe}}$, whose identities are disjoint from probe fitting and validation; c is the category label paired with image x. The quantities $\mathcal{A}_C^{\text{fixed}}$, $\mathcal{A}_C^{\text{refit}}$, and $\mathcal{A}_{\text{full}}^{\text{refit}}$ denote the three resulting accuracy trajectories. C-fixed measures whether the early decision boundary remains usable, whereas C-refit measures category accessibility after allowing a new linear readout under the same fitting and validation procedure. Full-refit tests whether the trend extends beyond C.

Table 5: **Refitting mitigates category-readout decline, but a modest loss remains.** Early/final CUTE-10 accuracies and paired changes from step 500 to 5000 for C-fixed, C-refit, and Full-refit. Intervals use category-stratified identity-cluster bootstrap; n and # IDs count test images and identities.

Component	Readout	Early Acc.	Final Acc.	Δ Acc.	95% CI	n	# IDs
C	Fixed	99.73	96.45	-3.27	$[-4.16, -2.39]$	2,932	30
C	Refit	99.73	98.94	-0.78	$[-1.16, -0.44]$	2,932	30
Full	Refit	99.69	99.08	-0.61	$[-0.99, -0.24]$	2,932	30

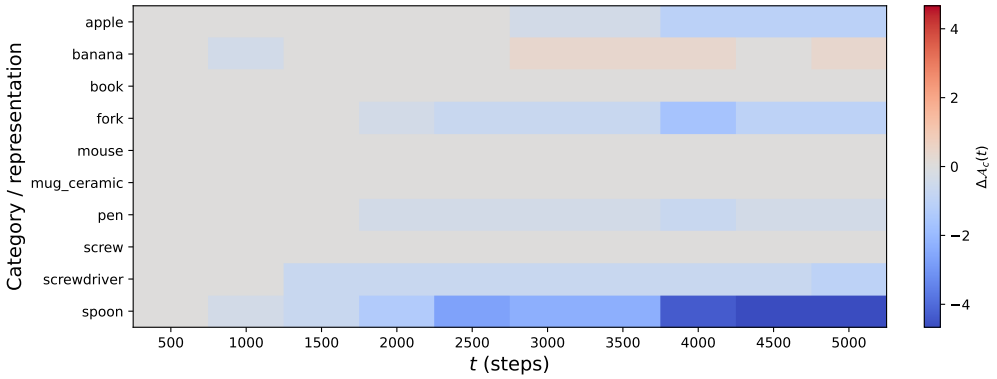

Figure 10: **Category-readout decline is concentrated in a subset of CUTE-10 categories.** C-refit accuracy changes across adaptation relative to each category's step-500 accuracy.

Table 5 summarizes early-to-final changes on the 30 held-out test identities. Refitting reduces the decline relative to retaining the early classifier, but both C-refit and Full-refit still show modest accuracy losses near ceiling, with aggregate paired intervals below zero. Figure 10 shows that these changes are uneven across categories: declines concentrate in a subset, while other categories remain comparatively stable. Category-level changes are descriptive, with three test identities per category.

A.6 Historical Restoration and Query Responses

We test historical retrieval value by first calibrating restoration on source images, then comparing its trajectory and endpoint effects with matched controls, and finally examining the query responses underlying dataset averages.

The shared restoration coefficient matches component energy using source calibration:

$$\mathcal{E}_B(t) = \frac{1}{S}\sum_{s=1}^{S}\frac{1}{n_s}\sum_i \|P_B u_t(x_{si})\|_2^2, \qquad \kappa_B(t) = \sqrt{\frac{\mathcal{E}_B(t)}{\mathcal{E}_B(t_e) + 10^{-12}}}, \tag{19}$$

where x_{si} is calibration image i from source s, n_s is that source's calibration sample count, $\mathcal{E}_B(t)$ is its source-balanced mean component energy, and $\kappa_B(t)$ is the restoration scale for span B. One scale is shared across evaluation datasets. The reverse-A control restores early A while retaining current C. The matched-random control uses a fixed random orthogonal subspace with the same rank and its own source-calibrated scale.

Restoration effects. Relative to current features, the effect is

$$I_{\mathcal{D}}^{C}(t) = P_{\mathcal{D}}(f_t^{\text{restore}\,C}) - P_{\mathcal{D}}(u_t). \tag{20}$$

where $I_{\mathcal{D}}^{C}(t)$ is the restoration-induced mAP change on evaluation dataset $\mathcal{D}$, measured in percentage points using the retrieval functional defined in Appendix A.3. A positive value means that this controlled restoration improves retrieval on $\mathcal{D}$; a negative value favors retaining the current component. The comparison concerns the hybrid representation under this intervention, not the complete early encoder.

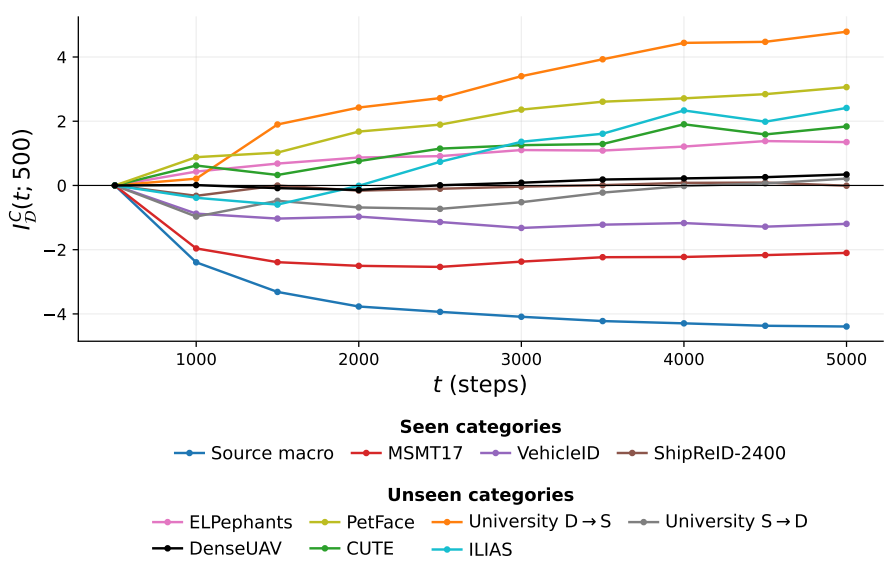

Figure 11: **Historical restoration effects diverge across evaluation settings as adaptation proceeds.** Restoration-induced mAP changes for the source macro and ten evaluation settings across checkpoints, relative to the corresponding unmodified representation.

Figure 11 traces how restoration effects diverge across checkpoints for the source macro and all ten diagnostic targets. Table 6 complements these trajectories with endpoint effects and paired intervals across all evaluated datasets. Every effect compares restored and current representations at the same checkpoint.

Matched controls. Table 6 reports absolute scores, restoration effects $I_{\mathcal{D}}^{C}$, and the localization contrast $L_{\mathcal{D}}^{C} = P_{\mathcal{D}}(\text{restoreC}) - P_{\mathcal{D}}(\text{randomC})$. The latter compares restoring the fitted complement with restoring a matched-random complement. ELPephants and ILIAS benefit from restoration but have negative localization contrasts, while the contrasts for CUTE, PetFace, and DenseUAV have intervals crossing zero. Matched-random restoration recovers comparable or greater utility on these targets, showing that useful historical restoration extends across different subspace choices.

Query-response proportions. To examine variation within each dataset, we measure the restoration-induced change in query-level retrieval:

$$\delta_q^{\text{AP}} = \text{AP}_q^{\text{restore}} - \text{AP}_q^{\text{current}}. \tag{21}$$

where q denotes a valid retrieval query, $\text{AP}_q^{\text{restore}}$ and $\text{AP}_q^{\text{current}}$ are its average precision under the restored and unmodified current representations, respectively, and δ_q^{AP} is their difference. Its sign indicates whether restoration improves, leaves unchanged, or degrades that query's ranking quality. At $t = T$, the fraction of queries improved by restoration is

$$r_{\mathcal{D}}^{+} = \frac{1}{|\mathcal{Q}_{\mathcal{D}}|} \sum_{q \in \mathcal{Q}_{\mathcal{D}}} \mathbf{1}[\delta_q^{\text{AP}} > 0], \tag{22}$$

where $\mathcal{Q}_{\mathcal{D}}$ is the valid query set of dataset $\mathcal{D}$, $|\mathcal{Q}_{\mathcal{D}}|$ is its size, $\mathbf{1}[\cdot]$ is the indicator function, and $r_{\mathcal{D}}^{+}$ is the improved-query fraction. Replacing the condition with $\delta_q^{\text{AP}} < 0$ or $\delta_q^{\text{AP}} = 0$ defines the degraded and unchanged fractions $r_{\mathcal{D}}^{-}$ and $r_{\mathcal{D}}^{0}$, respectively. The three fractions sum to one. Figure 3 displays them as percentages; they measure response prevalence, not effect magnitude.

Query response distributions. The ECDFs in Figure 12 and statistics in Table 7 supplement the response proportions in Figure 3 with effect magnitudes. Distributions vary substantially across datasets, and positive average effects can coexist with degradation for many queries. These comparisons change both query and gallery representations.

Summary. Across the trajectory, source retrieval generally favors current structure, target responses to historical restoration differ in direction and magnitude, and individual queries within one setting

Table 6: **Historical transfer utility extends beyond the fitted complement.** At step 5000, (a) retrieval performance of Current, Restore-C, Reverse-A, and matched Random-C representations; (b) paired restoration and localization effects with 95% query-identity-cluster bootstrap intervals.

(a) Retrieval performance

Dataset	N	Current	Restore C	Reverse A	Random C
Source evaluation					
Market-1501	3,368	80.3487	75.9999	75.2490	66.7990
iPanda-50	1,665	37.0838	33.2711	35.3235	30.6904
VeRi	1,678	70.8462	62.5060	66.7106	56.8971
ATRW	1,063	61.4885	59.9276	59.3893	57.4610
VesselReID	1,794	64.7526	60.8555	61.2140	55.3153
Target evaluation					
MSMT17	11,659	25.0928	22.9924	25.1340	21.4081
VehicleID	17,377	55.9080	54.7126	55.9297	54.3146
ShipReID-2400	706	4.3282	4.3197	4.4735	4.3200
ELPephants	591	11.6481	12.9979	12.2004	13.1829
PetFace	115,708	41.6134	44.6747	42.6161	44.6956
University D→S	37,855	24.4911	29.2778	26.7909	28.8696
University S→D	701	14.7426	14.9558	15.9354	14.8087
DenseUAV	2,331	4.4418	4.7834	4.7538	4.7514
CUTE	17,804	50.0391	51.8760	50.7077	51.9459
ILIAS	1,232	36.0280	38.4408	37.1882	38.7157

(b) Paired restoration effects

Dataset	$I_{\mathcal{D}}^{C}$	95% CI	$L_{\mathcal{D}}^{C}$	95% CI
Source evaluation				
Market-1501	-4.3488	$[-4.6881, -4.0329]$	$+9.2009$	$[+8.6902, +9.7757]$
iPanda-50	-3.8127	$[-5.6925, -1.8701]$	$+2.5807$	$[+1.6784, +3.7079]$
VeRi	-8.3402	$[-9.4008, -7.3406]$	$+5.6089$	$[+4.8857, +6.4105]$
ATRW	-1.5609	$[-2.9895, +0.0003]$	$+2.4666$	$[+0.6411, +4.1305]$
VesselReID	-3.8971	$[-4.2744, -3.5122]$	$+5.5402$	$[+5.0570, +6.0471]$
Target evaluation				
MSMT17	-2.1004	$[-2.2926, -1.9161]$	$+1.5843$	$[+1.4449, +1.7168]$
VehicleID	-1.1954	$[-1.6169, -0.8082]$	$+0.3981$	$[+0.1348, +0.6576]$
ShipReID-2400	-0.0085	$[-0.5167, +0.4391]$	-0.0003	$[-0.2789, +0.2792]$
ELPephants	$+1.3498$	$[+0.7009, +2.0631]$	-0.1849	$[-0.3619, -0.0176]$
PetFace	$+3.0614$	$[+2.9416, +3.1807]$	-0.0209	$[-0.0712, +0.0310]$
University D→S	$+4.7867$	$[+3.7344, +5.8120]$	$+0.4082$	$[+0.0492, +0.7924]$
University S→D	$+0.2132$	$[-0.5236, +0.9618]$	$+0.1472$	$[-0.0821, +0.3995]$
DenseUAV	$+0.3416$	$[+0.0386, +0.6454]$	$+0.0320$	$[-0.0986, +0.1699]$
CUTE	$+1.8368$	$[+1.3444, +2.3467]$	-0.0699	$[-0.1708, +0.0320]$
ILIAS	$+2.4128$	$[+1.7055, +3.1364]$	-0.2749	$[-0.5018, -0.0643]$

can respond oppositely. Matched controls further show that useful restoration is attainable with different subspace choices. Together, these results establish retained retrieval value in earlier feature structure and motivate learning how to reuse historical alternatives.

Table 7: **Historical restoration produces heterogeneous query-level gains and losses within individual datasets.** AP changes are Restore-C minus Current at step 5000, using step 500 as the historical reference. Median and the 10th/90th percentiles (Q_{10}/Q_{90}) are in percentage points; positive, unchanged, and negative response proportions are percentages of queries.

	AP change (pp)			Query proportion (%)		
Dataset	Median	Q_{10}	Q_{90}	Positive	Unchanged	Negative
Source evaluation						
Market-1501	−2.1276	−13.5367	+1.0509	15.7660	14.3409	69.8931
iPanda-50	−3.3000	−11.2870	+2.8838	24.6246	0.0000	75.3754
VeRi	−6.8045	−20.5280	+1.0833	15.5542	2.0858	82.3600
ATRW	−1.2063	−8.7466	+4.8886	36.4064	3.1044	60.4892
VesselReID	−2.5898	−10.8197	+1.4295	19.0635	5.7971	75.1394
Target evaluation						
MSMT17	−0.4242	−11.1518	+4.8170	37.8506	0.1715	61.9779
VehicleID	0.0000	−5.7143	+0.9091	16.2629	53.1795	30.5576
ShipReID-2400	+0.0007	−0.5911	+1.5258	51.5581	1.1331	47.3088
ELPephants	+0.3017	−3.0451	+7.6864	61.4213	0.0000	38.5787
PetFace	0.0000	−1.3354	+13.5016	41.7110	35.5230	22.7659
University D→S	+0.0726	−5.5556	+30.0000	53.3879	16.6530	29.9591
University S→D	−0.0739	−7.5111	+9.2286	41.0842	0.4280	58.4879
DenseUAV	+0.0082	−3.1822	+4.6446	52.1665	0.0000	47.8335
CUTE	+1.0406	−4.0099	+8.7001	62.9184	0.3763	36.7052
ILIAS	+0.1519	−7.4376	+15.8948	56.2500	6.1688	37.5812

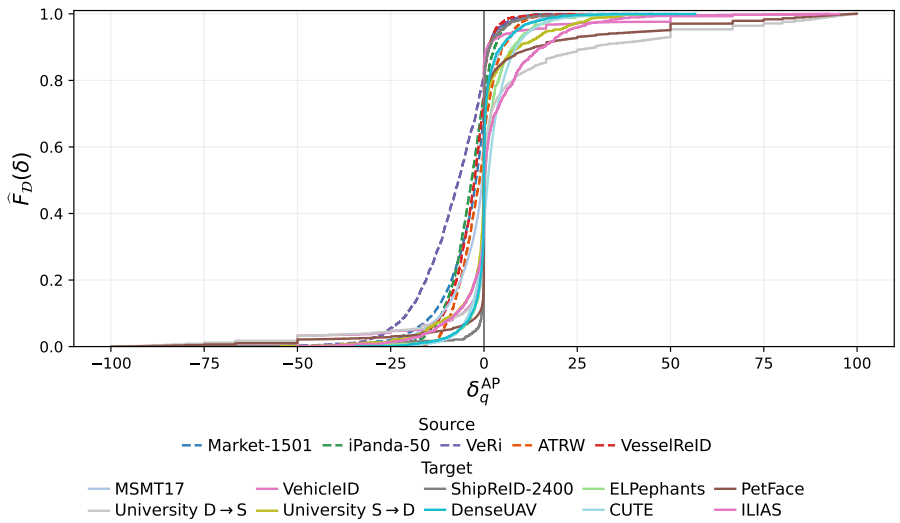

Figure 12: **Restoration effects have broad, dataset-dependent query-level distributions.** ECDFs of restoration-induced AP changes at step 5000, using step 500 as the historical reference, for five source datasets and ten evaluation settings.

B METHOD AND EVALUATION DETAILS

B.1 STATE POOL, ROUTER, AND CHRONOLOGY

Retaining Adaptation Alternatives. The shared encoder takes moving or cached adaptation parameters:

$$E_t^{\mathrm{mov}}(x) = E(x;\theta_t), \qquad E_{T_k}(x) = E(x;\theta_{T_k}), \qquad E_0(x) = E(x;\theta_0), \tag{23}$$

where $E(x;\theta)$ denotes an encoder with a shared frozen pretrained backbone and adaptation parameters θ; θ_t collects all Moving State LoRA factors at optimization step t; and θ_{T_k} is their frozen snapshot cached at step T_k. The Shared Anchor uses the zero-adaptation reference $\theta_0 = 0$. This reference recovers the pretrained function, as does the actual adapter initialization, but need not equal its random factor values. Each historical encoder therefore stores an adapter snapshot rather than a separate backbone.

The active adapted states are indexed by

$$\mathcal{I}_t = \{T_1, \ldots, T_K, \mathrm{mov}\}, \qquad E_s(x) = \begin{cases} E(x;\theta_{T_k}), & s = T_k, \\ E(x;\theta_t), & s = \mathrm{mov}, \end{cases} \tag{24}$$

where K is the number of historical states available at step t, s is a state label, and mov denotes the Moving State. Each cache also copies the moving router head. Historical adapters remain frozen, whereas all router heads remain trainable so that reuse can adjust as the pool grows.

Normalized descriptor and router. Retrieval uses

$$\hat{f}_{\mathrm{TRR}}(x) = \frac{f_{\mathrm{TRR}}(x)}{\|f_{\mathrm{TRR}}(x)\|_2}, \tag{25}$$

where $\hat{f}_{\mathrm{TRR}}$ is the unit-normalized retrieval descriptor. Residualization organizes reuse around adaptation-induced changes without introducing the diagnostic subspaces used to characterize PFR.

Routing consumes raw residuals:

$$r_s(x) = h_s(g(\Delta_s(x))) \in \mathbb{R}^p, \tag{26}$$

where g is a shared $p \to p$ MLP with GELU, h_s is a state-specific linear p-output head, and r_s contains coordinate-wise logits.

The coordinate-wise normalizer is

$$Z_j(x) = 1 + \sum_{s \in \mathcal{I}_t} \exp(r_{s,j}(x)/\tau_r)\,, \tag{27}$$

where $j \in \{1, \ldots, p\}$ indexes coordinates, $r_{s,j}$ is the corresponding state logit, $\tau_r > 0$ is the routing temperature, and Z_j is the coordinate-wise normalizer. The constant 1 is the zero option's exponentiated logit. We obtain

$$w_{s,j}(x) = \frac{\exp(r_{s,j}(x)/\tau_r)}{Z_j(x)}, \qquad w_{0,j}(x) = \frac{1}{Z_j(x)}, \tag{28}$$

where $w_{s,j}$ is the normalized weight of state s and $w_{0,j}$ is the zero-adaptation weight. Since $w_{0,j} + \sum_s w_{s,j} = 1$, Eq. 3 equivalently gives $f_{\mathrm{TRR},j} = w_{0,j}E_{0,j} + \sum_s w_{s,j}E_{s,j}$. Weight assigned to the zero option retains pretrained-anchor content without adding an adaptation residual.

Chronology-aware Regularization. The chronology regularizer uses mean routing mass:

$$\bar{w}_i = \frac{1}{Bp} \sum_{b=1}^{B} \sum_{j=1}^{p} w_{s_i,j}(x_b), \tag{29}$$

where B is mini-batch size, x_b is its bth image, s_i is the ith adapted state in chronological order with the Moving State last, and $\bar{w}_i$ averages its weights over images and coordinates. The regularizer is

$$c_i = \frac{i}{\sum_{k=1}^{N_e} k}, \qquad \mathcal{L}_{\mathrm{chron}} = \sum_{i=1}^{N_e} c_i \bar{w}_i, \tag{30}$$

where $N_e = |\mathcal{I}_t| = K + 1$ is the number of active adapted states, c_i is the normalized chronology cost, and $\mathcal{L}_{\mathrm{chron}}$ penalizes average reliance on later adaptation. The zero option is excluded from

this sum and has zero cost. The prior penalizes later chronological positions while leaving state utilization unconstrained; its cost is determined by order rather than adaptation magnitude.

Pool-size dependence. The largest chronology cost is $2/(N_e + 1)$, so its numerical scale decreases as the pool grows under a fixed loss coefficient. With equal adapted and zero-option logits, the anchor receives $1/(N_e + 1)$ weight. Changing pool size therefore changes both competition and cost scale.

B.2 DTR Updates and Retraction Strength

Local Update Reversal. LUR transforms the current loss-gradient contribution before it enters the stateful optimizer. For the Moving State,

$$g_t = \nabla_{\theta_t} \mathcal{L}_t, \tag{31}$$

where $\mathcal{L}_t$ is the training objective at step t, specified in Eq. 6, and g_t is its gradient with respect to the Moving State factors. Normal and LUR updates use

$$\tilde{g}_t = \begin{cases} g_t, & \text{Normal}, \\ -\gamma_{\text{amp}} g_t, & \text{LUR}, \end{cases} \qquad \theta_{t+1} = \text{OptStep}(\theta_t, \tilde{g}_t; \nu_t), \tag{32}$$

where $\tilde{g}_t$ is the gradient passed to the optimizer, $\gamma_{\text{amp}} \geq 1$ is the DTR amplification factor, ν_t contains stateful optimizer variables, and OptStep denotes one optimizer update. The same gradient transformation applies to router and supervision-head parameters, although they are not part of θ_t.

Initialization Retraction. To regulate accumulated adaptation directly, IR contracts individual Moving State LoRA factors (Hu et al., 2022) toward the zero-adaptation reference:

$$\theta_{t+1}^{(\ell)} = \theta_t^{(\ell)} - \rho_t^{(\ell)} \big(\theta_t^{(\ell)} - \theta_0^{(\ell)}\big), \tag{33}$$

where ℓ indexes a LoRA factor matrix, $\theta_0^{(\ell)} = 0$ is its retraction reference, and $\rho_t^{(\ell)} \in [0, 1]$ is a factor-specific contraction coefficient. The coefficient matches the current gradient scale using the same γ_{amp} as LUR; the exact rule follows below. Thus, the amplification sweep in Fig. 6 varies one common control for both operations. IR replaces LoRA-factor optimizer updates, while other parameters update normally. Normal steps predominate, and LUR and IR are mutually exclusive.

Retraction strength. LoRA uses $W_t = W_0 + B_t A_t$, where W_0 is frozen and A_t, B_t are adaptation factors with zero initial product. IR contracts the factors toward zero, matching its strength to the current gradient scale:

$$\rho_t^{(\ell)} = \min\left(1, \frac{\eta_t^{(\ell)} \gamma_{\text{amp}} \|g_t^{(\ell)}\|_F}{\|\theta_t^{(\ell)}\|_F}\right), \tag{34}$$

where $\eta_t^{(\ell)}$ is the learning rate for factor ℓ, $g_t^{(\ell)}$ is its current loss gradient, and $\|\cdot\|_F$ is the Frobenius norm. Zero factors remain zero. At an affected LoRA projection, the product contracts as $B'A' = (1-\rho_B)(1-\rho_A)BA$, where ρ_A and ρ_B are the two factor-specific coefficients. Thus IR is shrinkage whose timing and strength are gradient-dependent.

Schedule and optimizer scope. Normal, LUR, and IR steps occur with probabilities 0.90, 0.05, and 0.05, respectively, with shared amplification $\gamma_{\text{amp}} = 4$. LUR reverses the incoming gradient contribution while retaining optimizer state, rather than reversing the entire momentum-dependent displacement. IR directly contracts accumulated adapter factors. The two operations thus regulate the current gradient contribution and accumulated adaptation at different temporal scales.

B.3 DSRP Selection Rules

DSRP perturbs two kinds of relations: the classification branch changes sample-to-class assignments, while the embedding branch changes labels used to construct metric neighborhoods. The following rules specify quota-limited selection, category validity, and identity protection.

Quotas and eligibility. The four DSRP quotas are tied to $N_q = 4$. They bound directly modified samples or pairs, not the number of affected mined triplets or the resulting gradient change. Classification quotas count samples in each device batch; embedding quotas count pairs in the combined

batch. Classification selects same-category proposals first and cross-category proposals from the remaining samples, using the probability-gap rules in Section 3.5. Selected target identities must differ from the original labels and have valid category annotations.

Classification selection. Category-aware proposals are

$$j_i^{\text{same}} = \arg \max_{j \neq y_i,\, c(j)=c_i} g_{i,j}^{\text{cls}}, \qquad j_i^{\text{cross}} = \arg \min_{j \neq y_i,\, c(j)\neq c_i} g_{i,j}^{\text{cls}}, \tag{35}$$

where c_i is the image's category, $c(j)$ is the category of identity class j, and j_i^{same} and j_i^{cross} are alternative identity labels. Only classes with valid category annotations are eligible; a sample without a valid proposal is excluded from that selection. We select samples by their proposal scores:

$$\begin{aligned} \mathcal{I}_{\text{same}} &= \text{Top}_{N_q} \left\{ (i, g_{i,j_i^{\text{same}}}^{\text{cls}}) : i \in \mathcal{V}_{\text{same}} \right\}, \\ \mathcal{I}_{\text{cross}} &= \text{Bottom}_{N_q} \left\{ (i, g_{i,j_i^{\text{cross}}}^{\text{cls}}) : i \in \mathcal{V}_{\text{cross}} \setminus \mathcal{I}_{\text{same}} \right\}, \end{aligned} \tag{36}$$

where $\mathcal{V}_{\text{same}}$ and $\mathcal{V}_{\text{cross}}$ contain samples with valid proposals, and Top_{N_q}/Bottom_{N_q} return at most N_q sample indices with the largest/smallest scores. The disjoint selected sets determine

$$\hat{y}_i = \begin{cases} j_i^{\text{same}}, & i \in \mathcal{I}_{\text{same}}, \\ j_i^{\text{cross}}, & i \in \mathcal{I}_{\text{cross}}, \\ y_i, & \text{otherwise}, \end{cases} \tag{37}$$

where $\hat{y}_i$ is the perturbed classification target. Unselected samples retain ordinary identity supervision.

Embedding selection. Eligible pairs are

$$\begin{aligned} \mathcal{P}_{\text{same}} &= \{(i,j) : i < j,\ i,j \in \mathcal{U},\ y_i \neq y_j,\ c_i = c_j\}, \\ \mathcal{P}_{\text{cross}} &= \{(i,j) : i < j,\ i,j \in \mathcal{U},\ y_i \neq y_j,\ c_i \neq c_j\}, \end{aligned} \tag{38}$$

where $\mathcal{U}$ contains unprotected samples after reserving one unchanged image per identity, and $\mathcal{P}_{\text{same}}$/$\mathcal{P}_{\text{cross}}$ denote same-/cross-category different-identity pairs.

Selection combines distance preferences with disjointness and identity-pair diversity. Round-robin pairings generate disjoint candidate sets; within each set, distinct identity pairs take priority and distance ranks fill the two quotas. Among the resulting quota-limited sets, we choose

$$\mathcal{P}^\star = \arg\max_{\mathcal{P} \in \mathfrak{C}_{N_q}}{}_{\text{lex}} \left(|\mathcal{P}|,\ D(\mathcal{P}),\ U(\mathcal{P}) \right), \qquad U(\mathcal{P}) = \sum_{(i,j)\in\mathcal{P}} \big(2\mathbf{1}[c_i = c_j] - 1\big) d_{i,j}, \tag{39}$$

where $\mathfrak{C}_{N_q}$ is the family produced by this candidate construction and $D(\mathcal{P})$ counts distinct unordered identity pairs. Each member of the family has disjoint endpoints and at most N_q pairs of each category relation. The lexicographic maximization prioritizes the number of accepted pairs, then identity-pair diversity, then the signed distance utility U; $\mathbf{1}[\cdot]$ is the indicator function. Selection is performed within this round-robin candidate family.

For every selected pair, we simultaneously exchange the original labels:

$$(\tilde{y}_i, \tilde{y}_j) = (y_j, y_i), \qquad (i,j) \in \mathcal{P}^\star, \tag{40}$$

where $\tilde{y}$ denotes labels used only to construct triplet relations; unselected labels remain unchanged. Since $y_i \neq y_j$, the pair remains different-identity after the exchange. The operation rewires each endpoint's positive/negative neighborhoods. The protection rule preserves one unchanged sample per identity, while batch-hard neighborhoods may still change for other anchors.

The moving and composed paths construct perturbations independently from original labels; category annotations remain unchanged.

B.4 Training Configuration and Inference

GVA and ReState use the backbone, adapter rank, image resolution, and training duration specified in Section 4.2. LoRA modifies the Q/K/V projections in every transformer block while the pretrained backbone weights remain frozen. Training applies horizontal flipping, color jitter, random grayscale,

Table 8: **The benchmark spans category-aligned and category-shifted retrieval across diverse object types.** Statistics of the five training datasets and nine evaluation datasets.

Subset (Train)	# IDs	# Images	# Cat.	Object
Market-1501 (Zheng et al., 2015)	751	12,936	1	Person
iPanda-50 (Wang et al., 2021)	35	5,209	1	Panda
VeRi (Liu et al., 2018)	576	37,715	1	Vehicle
ATRW (Li et al., 2019)	107	1,887	1	Tiger
VesselReID (Zhang et al., 2023b)	623	14,570	1	Vessel
Total (Train)	**2,092**	**72,317**	**5**	

Subset (Test)	# IDs	# Images	# Cat.	Object
MSMT17 (Wei et al., 2018)	3,060	93,820	1	Person
VehicleID (Liu et al., 2016)	2,400	19,777	1	Vehicle
ShipReID-2400 (Liu et al., 2025)	600	4,959	1	Vessel
ELPephants (Körschens & Denzler, 2019)	83	591	1	Elephant
PetFace (Shinoda & Shiohara, 2024)	55,686	115,708	13	Multi
University-1652 (Zheng et al., 2020)	951	52,056	1	Buildings
DenseUAV (Dai et al., 2024)	3,033	20,529	1	Buildings
CUTE (Kotar et al., 2023)	180	18,000	50	Multi
ILIAS (Kordopatis-Zilos et al., 2025)	1,000	5,947	42	Multi
Total (Test)	**66,993**	**331,387**	**111**	

and ImageNet normalization. Batches contain four images per identity, with 256 images per GPU on two GPUs. Optimization uses SGD with momentum 0.9, weight decay 10^{-4}, and a learning rate decreasing linearly from 0.008 to 0.0016. ID classification uses batch-normalized features; batch-hard triplet learning uses normalized descriptors with margin 0.1. Both loss terms have unit weight; the routing temperature and chronology-loss coefficient are also set to 1.

Each branch uses the identity objective

$$\mathcal{L}_{\mathrm{ReID}} = \mathcal{L}_{\mathrm{ID}} + \mathcal{L}_{\mathrm{Tri}}, \tag{41}$$

where $\mathcal{L}_{\mathrm{ID}}$ and $\mathcal{L}_{\mathrm{Tri}}$ denote the identity-classification and triplet metric-learning losses, respectively (Hermans et al., 2017; Luo et al., 2019). Both terms have unit weight in our implementation.

During training, the Moving State learns continuously and supplies periodic snapshots; only its path is supervised before the first cache. Once history is available, TRR composes the retained adaptation residuals, DSRP supplies supervision to both branches, and DTR regulates parameter updates under Eq. 6. The PFR projectors and coordinate alignment are diagnostic tools and are not used by ReState.

At inference, the retained pool includes the terminal cached state and its routing head. TRR combines this pool with the pretrained anchor to produce normalized retrieval descriptors. The independent moving branch and training-only supervision components are omitted; routing weights are normalized over the retained states and the zero-adaptation option. Shared pretrained parameters reduce storage, while inference computes features under each retained state and the anchor. Gallery descriptors can be precomputed; queries retain this multi-state computation.

B.5 Datasets and Evaluation Protocol

We train on Market-1501 (Zheng et al., 2015), iPanda-50 (Wang et al., 2021), VeRi (Liu et al., 2018), ATRW (Li et al., 2019), and VesselReID (Zhang et al., 2023b), covering persons, pandas, vehicles, tigers, and vessels. Evaluation covers 111 dataset-local categories; these counts are summed within datasets rather than globally deduplicated. Seen/unseen refers to ReID training, not encoder pretraining.

Evaluation populations. Category-aligned single-category tasks are MSMT17, VehicleID, and ShipReID-2400. Category-shifted single-category tasks are ELPephants, University D→S and S→D, and DenseUAV. Category-shifted multi-category tasks are PetFace, CUTE, and ILIAS. Together, they cover human, vehicle, vessel, animal, geo-localization, and general-object retrieval. CUTE also supplies an identity-disjoint category probe; its labels train only diagnostic readouts, never the retrieval encoder.

Retrieval metrics. Gallery images are ranked by cosine similarity between normalized descriptors. mAP measures average ranking precision across relevant matches, R1 measures first-match accuracy, and mINP reflects recovery of the lowest-ranked positive.

The retrieval populations preserve dataset-specific gallery and exclusion rules. Market-1501 and VeRi retain their recorded camera exclusions; iPanda-50 and ATRW use the test images as both query and gallery with self-matches excluded. VesselReID retains separate query and gallery sets. VehicleID uses a fixed gallery sample, while PetFace pools its thirteen families into one retrieval population. ILIAS uses its core split of 1,232 queries and 4,715 gallery images, without the YFCC

Table 9: **Historical composition remains effective under multiple routing restrictions.** GVA, terminal-state deployment, and three composition variants are compared across ten tasks. *Terminal Moving Output* uses the final moving encoder without fusion. *Input-conditioned State-wise Routing* predicts one weight per state and input, shared across coordinates. *Static Coordinate-wise Routing* learns an input-independent logit array over states and coordinates. TRR uses input-conditioned coordinate-wise weights.

Method	MSMT17			VehicleID			ShipReID-2400			ELPephants			PetFace		
	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
GVA	25.1	1.5	56.1	55.8	55.8	47.3	4.3	1.8	6.7	11.7	2.3	31.1	41.7	40.1	37.8
Terminal Moving Output	24.9	1.4	56.0	55.6	55.6	47.1	4.1	1.8	6.2	11.4	2.3	29.4	40.6	39.1	36.7
Input-conditioned State-wise Routing	25.4	**1.7**	**56.2**	55.8	55.8	**47.4**	**4.5**	**1.9**	**6.8**	12.3	2.3	**33.5**	**43.3**	**41.7**	**39.4**
Static Coordinate-wise Routing	25.3	**1.7**	55.9	55.8	55.8	47.3	**4.5**	**1.9**	**6.8**	12.3	**2.4**	**33.5**	**43.3**	**41.7**	**39.4**
TRR	**25.5**	**1.7**	55.9	**55.9**	**55.9**	**47.4**	**4.5**	1.8	**6.8**	**12.4**	2.3	**33.5**	**43.3**	**41.7**	**39.4**

Method	University D→S			University S→D			DenseUAV			CUTE			ILIAS		
	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
GVA	23.8	23.8	16.3	15.4	3.1	**31.1**	4.4	1.2	3.9	50.0	**16.2**	90.4	36.0	11.8	61.0
Terminal Moving Output	23.0	23.0	15.6	14.4	2.7	29.2	4.1	1.0	3.9	48.9	15.6	90.0	34.8	10.9	60.3
Input-conditioned State-wise Routing	28.2	28.2	19.9	**15.7**	**3.6**	29.8	**4.8**	**1.3**	4.3	**51.4**	**16.2**	**91.0**	38.1	12.6	**64.0**
Static Coordinate-wise Routing	**28.4**	**28.4**	**20.0**	15.5	**3.6**	29.5	**4.8**	**1.3**	**4.6**	**51.4**	**16.2**	**91.0**	**38.3**	**12.7**	63.9
TRR	27.8	27.8	19.5	**15.7**	**3.6**	29.8	**4.8**	**1.3**	4.5	51.2	16.1	90.9	38.2	**12.7**	63.7

distractor collection; we compute full-gallery retrieval metrics rather than the large-scale benchmark's mAP@1000. DenseUAV evaluates drone queries against the full satellite test gallery, without training images. Queries without valid positive matches are excluded from all retrieval averages.

VICP example protocol. We follow the VICP repository implementation and provide one batch of 128 images as visual in-context examples per evaluation task. Examples are sampled from the dataset's training split when available; otherwise, they are sampled from the pooled query and gallery images. Sampled examples remain in the evaluated query/gallery sets. This pool supplies in-context examples without parameter training. For mixed-category evaluations such as CUTE and PetFace, we likewise provide only one batch for the entire task, rather than separate batches for individual categories, maintaining a common one-batch example budget.

B.6 Parameter-Sweep Configuration

Module context and task aggregation. Each panel in Fig. 6 adds only the corresponding module to GVA: (a) TRR, (b) DTR, and (c) DSRP. The other two modules are disabled. Each default point therefore characterizes the corresponding single-module configuration. The parameter sweeps use six evaluation tasks: category-aligned results are the equal-task mean mAP over MSMT17 and VehicleID; category-shifted results average PetFace, University D→S, CUTE, and ILIAS with equal task weights. The routing visualization uses the complete ten-task suite.

Sweep interpretation. Changing cache interval jointly varies retained-state count, temporal spacing, chronology-cost scale, and zero-option competition (Appendix B.1). DTR varies the shared LUR/IR amplification, while DSRP varies all four operation quotas jointly through $\log_2 N_q$. Defaults are cache interval 500, amplification 4, and quota 4, selected using source-side data only. Figure 6 subsequently characterizes module behavior around these settings on the evaluation suite. The TRR panel thus reflects the complete retention schedule, while the DSRP panel characterizes perturbation extent, complementing the enabled/disabled comparison in the component table.

B.7 Extended Ablation Studies

The following comparisons cover all ten evaluation tasks under the retrieval protocols in Appendix B.5. They examine component interactions, routing flexibility, optimization interventions with and without historical reuse, and structured versus random supervision perturbation. Each table includes GVA, the component-free baseline, as a reference.

Historical composition and routing flexibility. Table 9 shows that all three composition variants improve or match GVA mAP on every task and exceed the terminal moving output, with gains of 4.0–4.6 mAP on University D→S. Input-conditioned state-wise routing and static coordinate-wise routing therefore demonstrate that input adaptivity and fine-grained feature composition can each

Table 10: **Combined trajectory regulation and historical reuse achieve higher mAP on most shifted tasks than either standalone intervention.** GVA, IR alone, LUR alone, and DTR (IR+LUR) with TRR are compared across ten tasks; DSRP is disabled throughout. Checkmarks/crosses mark enabled/disabled components. The first three rows use a single adapted encoder, while the final row uses historical composition.

Components			MSMT17			VehicleID			ShipReID-2400			ELPephants			PetFace		
IR	LUR	TRR	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
×	×	×	25.1	1.5	56.1	55.8	55.8	47.3	4.3	1.8	6.7	11.7	2.3	31.1	41.7	40.1	37.8
✓	×	×	23.2	1.3	53.4	55.3	55.3	47.5	4.4	1.9	6.5	12.2	**2.4**	34.2	43.4	41.8	39.4
×	✓	×	**25.7**	**1.7**	**56.9**	**56.1**	**56.1**	**47.7**	5.1	1.9	**9.1**	12.7	2.3	**35.2**	44.5	42.9	40.6
✓	✓	✓	22.3	1.4	52.0	54.7	54.7	47.0	**5.2**	**2.1**	8.2	**12.8**	2.3	34.9	**44.7**	**43.0**	**40.7**

Components			University D→S			University S→D			DenseUAV			CUTE			ILIAS		
IR	LUR	TRR	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
×	×	×	23.8	23.8	16.3	15.4	3.1	31.1	4.4	1.2	3.9	50.0	16.2	90.4	36.0	11.8	61.0
✓	×	×	27.1	27.1	19.4	17.4	4.4	31.1	5.2	1.4	4.9	52.6	16.8	91.2	39.9	13.6	65.4
×	✓	×	30.9	30.9	21.9	16.9	4.4	31.1	**5.3**	**1.5**	**5.9**	51.3	15.2	**91.6**	40.3	13.8	66.6
✓	✓	✓	**31.5**	**31.5**	**22.8**	**17.5**	**4.5**	**31.7**	5.2	**1.5**	5.4	**53.3**	**17.2**	91.3	**42.7**	**14.8**	**69.3**

Table 11: **Structured perturbation adds transfer gains beyond random relabeling on several shifted tasks.** GVA, random perturbation, and DSRP are compared across ten evaluation tasks. GVA uses original identity supervision. Random selects operation labels randomly instead of using identity-classification logits and triplet-space distances; DSRP uses category- and model-guided selection.

Method	MSMT17			VehicleID			ShipReID-2400			ELPephants			PetFace		
	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
GVA	**25.1**	**1.5**	**56.1**	55.8	55.8	47.3	4.3	1.8	**6.7**	**11.7**	**2.3**	31.1	41.7	40.1	37.8
random	24.8	**1.5**	56.0	55.7	55.7	47.6	**4.4**	**1.9**	**6.7**	11.5	**2.3**	31.1	**42.6**	**41.1**	**38.8**
DSRP	23.5	1.3	54.3	**56.0**	**56.0**	**47.7**	4.3	1.8	6.2	**11.7**	**2.3**	**31.8**	42.4	40.8	38.6

Method	University D→S			University S→D			DenseUAV			CUTE			ILIAS		
	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1	mAP	mINP	R1
GVA	23.8	23.8	16.3	15.4	3.1	31.1	4.4	1.2	3.9	50.0	16.2	90.4	36.0	11.8	61.0
random	25.9	25.9	17.8	16.5	3.2	33.0	**4.9**	**1.3**	**4.9**	51.2	16.7	91.0	36.8	12.0	62.8
DSRP	**26.9**	**26.9**	**18.6**	**17.4**	**3.7**	**33.8**	**4.9**	**1.3**	4.6	**51.3**	**17.0**	**91.1**	**37.6**	**12.6**	**63.1**

make historical states useful. Combining both forms of flexibility in TRR yields similar overall performance rather than a consistent further gain, suggesting that their benefits partly overlap under the current source-trained state pool. Future work could further disentangle these two routing factors and investigate whether richer state formation or more structured routing can better exploit their potential complementarity.

Update reversal, retraction, and reuse. Table 10 shows distinct benefits from local update reversal and accumulated-adaptation retraction. Both standalone interventions improve mAP over GVA on all seven category-shifted tasks. LUR exceeds IR on eight tasks, while IR performs better on University S→D and CUTE, showing that the two temporal scales favor different retrieval settings. Combining both operations with TRR achieves the highest mAP among the tested configurations on six of the seven shifted tasks, bringing together strong performance across these settings. LUR alone remains strongest on MSMT17, VehicleID, and DenseUAV. These results support coupling trajectory regulation with historical reuse for broad shifted-task performance, alongside the distinct utility of each standalone intervention.

Structured versus random supervision perturbation. Table 11 separates useful supervision perturbation from the additional value of structured selection. Random relabeling already improves several shifted tasks over GVA. DSRP further improves both University-1652 directions, CUTE, and ILIAS, including 1.0 mAP on University D→S and 0.8 on ILIAS. Random perturbation remains stronger on MSMT17, ShipReID-2400, and PetFace. Thus, model-guided selection changes the transfer profile of perturbation and adds value on several shifted settings.

C SCOPE AND FUTURE DIRECTIONS

ReState uses fixed-interval retention to keep adaptation alternatives available during learning. A natural extension is to characterize how feature roles and state utility evolve, using these signals to adaptively select, merge, or prune states for preservation and reuse. Historical reuse also introduces additional encoder-state evaluations at inference. Adaptive execution, state compression, or distillation could preserve retrieval benefits while reducing computation.

The PFR diagnostics characterize one source-only ReID adaptation trajectory, and the extended ablations compare the evaluated learning configurations. Extending the analysis to additional pretrained encoders and adaptation objectives would broaden our understanding of how representational utility evolves. Comparing retained histories under a common composition rule could further connect trajectory formation with the utility of the alternatives it produces.