ROW AMPLITUDE AS AN OPTIMIZATION MECHANISM IN TRANSPORT-BASED LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

A transport target determines both a sample's assignment and its weight in learning. We study the latter—its row amplitude—as a downstream optimization variable, separating its effect from changes in conditional assignment. For losses linear in a stopped target, Row-Mass Tempering (RMT) isolates amplitude while preserving assignments, the set of nonzero rows, and total mass. An exact mass–force covariance explains the resulting gradient correction and when scalar rescaling cannot reproduce it. This motivates Row-Amplitude Audit (RAA), which measures the actual optimizer effect and selects a coefficient from initial training observations for prospective validation. Frozen rules meet the declared update-deviation criterion in all three P^{2}OT seeds and all five JUMBOT seeds. Separate fixed-intervention studies establish task value: development-selected RMT improves PROTOCOL/Hdigit clustering accuracy by 5.09 percentage points and adjusted Rand index by 6.56 points over 10 paired seeds; equalizing active row amplitudes improves these metrics by 5.60 and 6.53 points in P^{2}OT/CIFAR100-LT over five paired seeds, with lower class-balanced accuracy and purity. JUMBOT shows the boundary: successful optimizer control yields no observed accuracy gain over the fixed reference. The contribution is a systematic account of an amplitude channel that can be isolated, explained, measured, and controlled, with optimizer controllability and downstream utility evaluated as distinct objectives.

1 INTRODUCTION

Transport-based learning turns a correspondence between samples and targets into a training signal. In clustering, the targets can be prototypes or pseudo-labels; in domain adaptation, they can couple source and target examples. The practical question is usually whether these assignments provide useful supervision. Yet assignment alone does not specify the update: the mass allocated to a sample also determines how strongly it contributes to learning. Separating these two effects is necessary to understand what a transport target changes in the optimizer.

Existing optimal transport (OT) methods already prescribe, normalize, or exploit sample masses, including confidence weighting in partial transport (Caron et al., 2020; Deng et al., 2025; Zhang et al., 2024a; Xue et al., 2025). These uses leave a distinct optimization question: when assignments are held fixed, what changes in the gradient, how does the optimizer transform that change, and can the resulting effect be controlled? A task score alone does not resolve these questions.

Our starting insight is that one target exposes two different quantities. For a nonnegative row with $s_i = \sum_k Q_{ik} > 0$,

$$Q_{i:} = s_i P_i, \qquad P_{ik} = Q_{ik}/s_i. \tag{1}$$

The *conditional assignment* P_i gives relative allocation over targets; the *row amplitude* s_i is the total mass weighting that sample. When the target is *stopped*, meaning held constant during differentiation, and enters the loss linearly, amplitude can change the gradient even with P_i fixed. Because samples share parameters, changing their relative contributions can rotate the aggregate gradient, not merely rescale it. This independence is a property of the downstream loss interface, not arbitrary freedom inside an OT solver.

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.

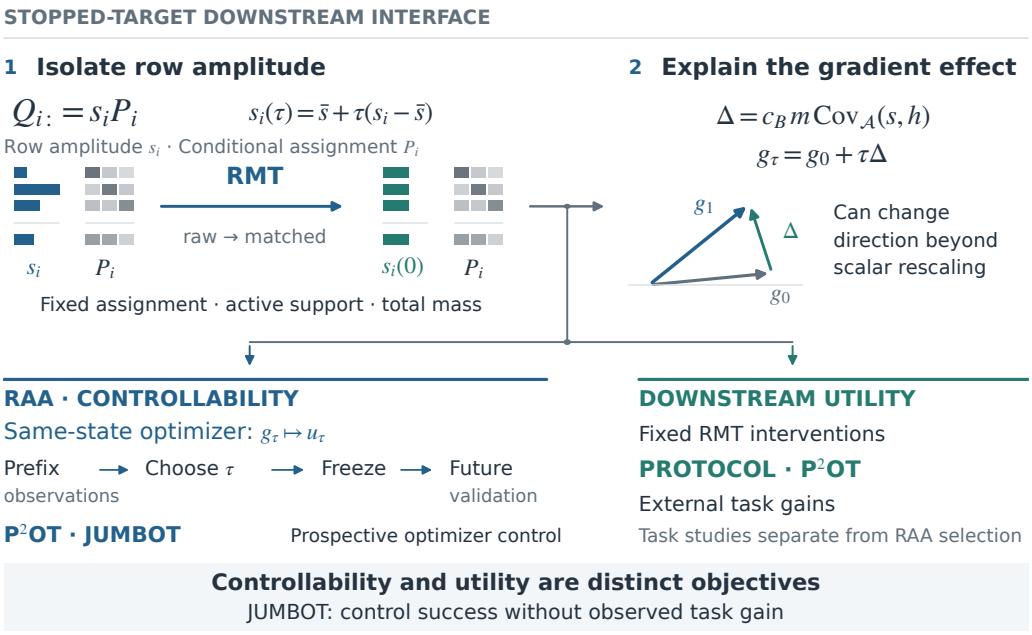

Figure 1: **Row amplitude: mechanism, control, and task utility.** RMT isolates amplitude at the stopped-target interface; the matched endpoint preserves assignments, active support, and total mass. The covariance explains the gradient change (non-collinear example shown schematically). The evidence branches evaluate distinct objectives: RAA tests prospective optimizer control; separate fixed interventions test task utility and do not use RAA-selected coefficients. JUMBOT demonstrates control success without an observed task gain.

Two tools make this insight testable (Fig. 1). To isolate amplitude, *Row-Mass Tempering* (RMT) varies its heterogeneity while preserving conditional assignments, the samples with nonzero mass, and total mass. The resulting gradient correction is exactly a covariance between masses and conditional sample gradients. To measure what that correction does to the update, *Row-Amplitude Audit* (RAA) applies the actual optimizer at the same model and optimizer state. It uses an initial training prefix to select a coefficient, freezes it, and tests the optimizer effect during subsequent training.

This channel has practical value beyond diagnosis. Fixed amplitude interventions improve sample-level clustering metrics in the external PROTOCOL and P^{2}OT frameworks. Separate prospective control studies in P^{2}OT and JUMBOT test the optimizer objective; JUMBOT meets that objective without an observed accuracy gain. Our controlled testbed, TPF, supplies supporting mechanism and dataset-sensitivity evidence.

Contributions.

- **Isolate and explain.** We characterize row amplitude as a downstream optimization channel through RMT, an exact gradient covariance, and the boundary of scalar-rescaling equivalence.
- **Measure and control.** RAA connects that gradient mechanism to same-state optimizer proposals and prospectively tests interventions selected from a training prefix.
- **Establish value and scope.** External fixed-intervention gains establish task value, while JUMBOT and complementary metrics show why controllability and utility require separate evaluation.

2 Related Work

How transport masses enter learning. Mass is already a design choice in transport-based learning. Balanced assignments in SwAV prescribe sample masses (Caron et al., 2020); THESAURUS

row-normalizes couplings (Deng et al., 2025). P^{2}OT uses row mass for progressive confidence weighting (Zhang et al., 2024a), and PROTOCOL combines partial transport with class-rebalanced contrastive learning (Xue et al., 2025). These methods establish useful ways to construct or consume weighted targets. Our question concerns what the final sample amplitudes do to optimization when their conditional assignments are held fixed.

From reweighting to an isolated optimization mechanism. Algebraically, RMT is sample reweighting. Its role here is an intervention that preserves conditional assignments, active support, and total mass, so that the immediate gradient difference isolates amplitude. The contribution lies in studying this channel systematically: the covariance identity explains the gradient change, RAA measures its transformation by the optimizer, and prospective continuations test a declared update objective. Task utility is then evaluated independently. This distinction matters when row masses encode confidence or a curriculum: changing them has an intended weighting interpretation as well as an optimizer effect.

Transport formulation and downstream interface. Unbalanced OT permits variation in transported mass through its marginal penalties (Chizat et al., 2018; Liero et al., 2018). More broadly, transport is used in graphs, semantic assignments, fusion, and view-level weighting across different learning objectives (Yao et al., 2022; Zhang et al., 2024b; Dong et al., 2025; Zhou et al., 2026). JUMBOT weights a differentiable joint feature–label cost by a stopped coupling (Fatras et al., 2021). These settings motivate an interface-level analysis across losses. RMT need not preserve a solver's constraints or optimality. Appendix A.4 treats source-prior realizability as a separate mathematical question.

Local information and later outcomes. Feature-attribution impossibility results distinguish local summaries from counterfactual behavior (Bilodeau et al., 2024), while influence functions use curvature to estimate sensitivity to loss perturbations (Koh & Liang, 2017). Our supporting construction in Appendix E concerns complete local row observations in a smooth stopped-linear class. Its negative result and structural recovery conditions clarify the limits of inferring continuation utility from an optimizer audit.

3 Row Amplitude: Isolation and Gradient Mechanism

To study amplitude independently, the intervention must leave the conditional assignment unchanged at the state being compared. We first specify this loss interface, construct the intervention, and then derive the gradient quantity it isolates.

Scope. Let $Q \in \mathbb{R}_+^{B\times K}$ be the final target actually consumed by

$$\mathcal{L}(\theta; Q) = c_B \sum_{i,k} \mathrm{sg}(Q_{ik})\ell_{ik}(\theta), \qquad c_B > 0, \tag{2}$$

Here B is the number of samples, K the number of targets, and sg the stop-gradient operator; c_B is fixed during the comparison. This includes cost losses weighted by stopped targets and soft-target cross-entropy. Define $\mathcal{A} = \{i : s_i > 0\}$, $m = |\mathcal{A}| > 0$, and $\bar{s} = m^{-1}\sum_{i\in\mathcal{A}} s_i$. Zero rows remain zero; their conditional directions are undefined. If the solver is differentiated, the result describes only the fixed-target direct term. Nonlinear use of the final target is outside this result.

Isolating amplitude with RMT. A comparison of different solver outputs can change assignment and amplitude together. RMT instead acts on one final stopped target: it interpolates each active mass toward the active-row mean. For $\tau \in [0, 1]$, set

$$Q_{i:}^{(\tau)} = \begin{cases} [\bar{s} + \tau(s_i - \bar{s})]P_i, & i \in \mathcal{A}, \\ 0, & i \notin \mathcal{A}. \end{cases} \tag{3}$$

The matched endpoint is $\tau = 0$, and raw is $\tau = 1$ (Fig. 1). RMT preserves active support, conditional directions, and total mass, with $\mathrm{CV}_{\mathcal{A}}[s(\tau)] = \tau\mathrm{CV}_{\mathcal{A}}[s]$. Here CV is the coefficient of variation, $\mathrm{CV}_{\mathcal{A}}[s] = \sigma_{s,\mathcal{A}}/\bar{s}$. It acts on downstream target weights and need not preserve the original transport constraints or optimality. Uniform active masses make the intervention a no-op.

Gradient consequence. Let $g_\tau = \nabla_\phi \mathcal{L}(\theta; Q^{(\tau)})$ in a shared parameter block ϕ, with $Q^{(\tau)}$ held fixed. All vector norms are Euclidean. Each conditional force h_i is the sample gradient under P_i, before multiplication by its row amplitude. These forces give

$$\begin{aligned} h_i &= \sum_k P_{ik} \nabla_\phi \ell_{ik}, \qquad g_0 = c_B \bar{s} \sum_{i \in \mathcal{A}} h_i, \\ \Delta &= g_1 - g_0 = c_B m \, \mathrm{Cov}_{\mathcal{A}}(s, h), \qquad g_\tau = g_0 + \tau \Delta, \\ \|\Delta\|_2 &\le c_B \, m \, \sigma_{s,\mathcal{A}} \sigma_{h,\mathcal{A}}. \end{aligned} \tag{4}$$

Equation (4) identifies the mechanism: amplitude changes the aggregate gradient through its covariance with the conditional forces. Mass variation alone is therefore insufficient to determine the effect. The covariance and standard deviations use population moments over active rows. Batch averaging ($c_B = 1/B$) retains the factor m/B. Nonuniform masses can have zero correction when mass–force covariance cancels. Appendix A defines the moments and proves these identities.

Is amplitude control equivalent to scalar rescaling? For a resolved nonzero g_0, orthogonal projection gives

$$\alpha_{\mathrm{proj}} = \frac{\langle \Delta, g_0 \rangle}{\|g_0\|_2^2}, \quad \Delta_\perp = \Delta - \alpha_{\mathrm{proj}} g_0, \quad \min_{c \in \mathbb{R}} \|g_1 - c g_0\| = \|\Delta_\perp\|. \tag{5}$$

A resolved nonzero transverse component identifies a gradient-direction change that scalar rescaling cannot reproduce. Positive-learning-rate stochastic gradient descent (SGD) equivalence further requires $1 + \alpha_{\mathrm{proj}} > 0$ when both gradients are nonzero. For an adaptive optimizer, RAA evaluates its actual same-state proposal (Section 4).

Solver realizability. For the full positive plan, let s collect its row masses and let P be the row-stochastic matrix with rows P_i. In the entropic unbalanced OT model with Kullback–Leibler (KL) marginal penalties in Eq. (16), a fixed-total amplitude change $s \mapsto s'$ at fixed conditional assignments and cost can be realized by changing only an unrestricted positive source prior exactly when $P^\top s' = P^\top s$. A normalized-prior constraint can eliminate this freedom. Appendix A.4 proves the characterization and identifies which realizable changes remain visible to the gradient. Thus downstream independence and solver realizability are separate properties.

4 Auditing and Controlling the Optimizer Effect

The gradient mechanism does not by itself specify the parameter update: an adaptive optimizer also depends on its incoming state. RAA extends the RMT comparison to this update. With model parameters, optimizer and auxiliary states, and batch randomness held fixed, it evaluates candidate gradients using the final stopped target, the actual loss reduction, and all other loss terms. Its measurements require no task-evaluation labels. The procedure separates three operations: measure the local gradient effect, map it through the optimizer, and test a prefix-selected intervention during continuation.

Measure the optimization effect. The endpoint gradients give their norms and the transverse correction

$$E_\perp = \frac{\|\Delta_\perp\|_2}{\|g_1\|_2 + \|g_0\|_2 + \epsilon_g}, \qquad \epsilon_g > 0. \tag{6}$$

This statistic measures the part of the local gradient change that scalar rescaling cannot reproduce. It uses aggregate gradients; optional per-sample calculations verify the covariance identity. The projection requires a numerically resolved g_0. For several loss terms, their sum determines the total update, so component cancellation matters.

Map gradients to optimizer proposals. A proposal is the parameter displacement the optimizer would make for a candidate gradient at the current state. For a specified optimizer, let u_τ be this displacement. Its deviation from raw is

$$D_\tau = \frac{\|u_\tau - u_1\|_2}{\|u_1\|_2}, \qquad \|u_1\|_2 > 0. \tag{7}$$

Each Adam proposal uses the candidate gradient and identical incoming moments; it is evaluated separately because the optimizer map is nonlinear. A zero raw displacement receives an undefined-reference status. The intended displacement is also checked against the actual stored-parameter change.

Choose and validate an intervention. Specify a candidate grid, a deviation band, and a required fraction of observations in that band. During the initial training prefix, compare candidates at the selected states. Choose the smallest eligible τ, giving the strongest amplitude equalization, and hold it fixed during continuation; abstain when none qualifies. Evaluate D_τ at future states reached under the selected intervention. Section 5.2 tests this procedure with criteria declared before the new runs.

The prospective question is whether prefix eligibility persists at the sampled future states on the selected intervention's own trajectory. This is an empirical control test, not a guarantee for unobserved updates or a rule selected for task improvement. When masses encode confidence, curriculum, or importance, raw remains the reference for that intended weighting; its task value is assessed in a separate comparison. Appendix E gives a supporting smooth-class construction showing why complete local row observations need not identify continuation utility, together with conditions under which recovery is possible.

5 Experiments: Mechanism, Control, Utility, and Boundaries

We organize the evidence by the claim being tested. Mechanism studies ask whether amplitude changes gradient geometry beyond scalar rescaling. Control studies ask whether prefix measurements support a frozen intervention at future probes. Utility studies evaluate separate fixed interventions on task metrics. Boundary studies ask whether success on one objective implies success on the other.

Roles of the implementations. TPF denotes our controlled sample-to-prototype unbalanced OT testbed for multi-view clustering. It lets us vary the amplitude regime and examine dataset sensitivity; it is not proposed as the paper's main method. P^{2}OT supplies cross-entropy with stopped soft targets for mechanism and control tests, and JUMBOT supplies a joint feature–label cost weighted by a stopped coupling for control and boundary tests. Fixed-intervention studies in PROTOCOL and P^{2}OT assess downstream utility.

Metrics and evidence status. Clustering accuracy (ACC) and adjusted Rand index (ARI) measure task outcomes; normalized mutual information (NMI), purity, class-average accuracy (mACC), and cluster occupancy provide complementary criteria. The optimizer-control endpoint is future coverage of the declared deviation band, evaluated per seed. Predictions are frozen before task-metric evaluation, with no outcome-based reselection of seeds or checkpoints. PROTOCOL uses a development-selected coefficient; P^{2}OT gradient diagnostics were specified after outcome access, and its partition analysis is descriptive. RAA confirmation uses new seeds and continuations in already studied settings. The TPF dataset-level prospective studies remain distinct from retrospective and development evidence in Appendix C. JUMBOT uses source supervision and reserves target-test labels for evaluation. Unless stated otherwise, dispersions are population standard deviations. Appendix B gives settings and evaluation rules, and Appendix D specifies the control protocols.

5.1 Mechanism: does amplitude change optimization geometry?

Controlled checks of the amplitude channel. In a frozen synthetic study, mass–force alignment controls the projected correction sign as predicted, with maximum affine error 5.11×10^{-15} over 200 replicates per cell. A 36-run grid in the controlled testbed then varies source-marginal regularization from .125 through .5 to 2 on two datasets with three paired seeds. Increasing ρ_s compresses both epoch-0 mass CV and relative cost-gradient correction in all six dataset–seed units; their mean strong/weak ratios are 0.0830 and 0.0808. Absolute paired ARI sensitivity also contracts between these endpoints in all six units (Appendix C.1). This coordinated contraction is consistent with a weakening amplitude channel as source-marginal regularization increases.

Table 1: **Prospective optimizer control.** Success requires at least 18/20 future probes in band per seed; pooled counts are descriptive. JUMBOT task accuracy is a separate secondary outcome. P^{2}OT confirmation uses imbalance ratio .1, distinct from the .01 task study, and evaluates no task metrics.

Framework	Policy	τ	In band	Seeds passing	Test ACC (%)
P^{2}OT	Shared prefix rule	.875	$58/60$	$3/3$	—
JUMBOT	Per-seed prefix rule	0	$100/100$	$5/5$	97.604
JUMBOT	Fixed reference	.875	$0/100$	$0/5$	97.606

External evidence beyond scalar rescaling. At ten selected P^{2}OT models, RAA probes the selected cosine-head weight block using reconstructed memory banks. These diagnostics were specified after outcome access. At partial-transport budget .5, all 20 assembled probes have $E_\perp \in [.0670, .4720]$, and all 40 solver calls meet their stopping criterion. Thus scalar rescaling cannot reproduce the measured local gradients. At budget 1, the normalized total correction is at most 1.78×10^{-6}. The full gradient measurements and solver termination results appear in Appendix C.2.

What the local result does not establish. A training control scales the raw Adam proposal to the matched proposal's norm at the control's own state, using all trainable parameters and committing raw moments. In P^{2}OT, matched exceeds this implemented control by 5.07 ± 4.73 ARI points at the common final endpoint, but two seeds fail the predefined update-realization criterion. The long-run effect beyond update scale therefore remains unresolved. Appendix C.3 gives the numerical criterion and local implementation checks. JUMBOT provides a separate double-precision (FP64) four-arm comparison, whose task result is reported in Section 5.4.

5.2 Control: do prefix-selected effects persist at future probes?

Here the outcome is optimizer controllability: whether a coefficient chosen before continuation meets a specified update objective on the resulting trajectory. The objective is a proposal deviation $D_\tau \in [.021, .099]$ at prespecified future states. Candidates must have $\tau \leq .875$, reducing active row-mass contrast by at least 12.5%, and enter the band at least nine times in ten prefix probes. The smallest qualifying τ is held fixed during continuation; future success requires at least 18/20 in-band probes along the resulting training trajectory. The band defines a measurable, limited intervention, with a lower bound excluding negligible update changes.

Prospective coverage and coefficient choice. In P^{2}OT, a revised rule is frozen after an initial abstention study and before three new seeds on CIFAR100-LT (imbalance ratio .1). It selects a shared $\tau = .875$ and continues from epoch 10 to 20. Future coverage is 19/20, 20/20, and 19/20: all three seeds meet the success threshold. Recorded updates pass the numerical checks, and the analysis includes finite-iteration solver outputs (Appendix D).

In JUMBOT on the already studied SVHN-to-MNIST task, five fresh seeds each supply 100 raw prefix updates and 400 continuation updates. Every seed selects $\tau = 0$, meeting the criterion at 20/20 future probes. A paired fixed-$\tau = .875$ branch meets the criterion at 0/20 probes in every seed because its update effect falls below the band. Thus the coefficient that met the P^{2}OT objective does not transfer to this setting; the prefix measurements select a coefficient that does (Table 1). Since all five choices are identical, fixed $\tau = 0$ would coincide with RAA here. Coverage is evaluated per seed at the sampled future states.

5.3 Utility: do fixed amplitude interventions improve task metrics?

Whether a controlled optimizer effect improves learning is a separate task question. We test task utility using fixed RMT interventions in two external frameworks. Their representation backbones and transport objectives are retained; neither study evaluates a coefficient selected by RAA.

Task gains in external multi-view clustering. In PROTOCOL (Xue et al., 2025) on Hdigit, RMT tests amplitude intervention within an established multi-view clustering framework. A development

Table 2: Clustering performance on Hdigit under PROTOCOL across 10 paired random seeds. The RMT coefficient $\tau = .5$, selected from the development sweep, is held fixed across the reported runs and compared with the raw plan ($\tau = 1$).

Dataset	Method	ACC (%)	ARI (%)
Hdigit	PROTOCOL (Raw, $\tau = 1$)	84.11	83.19
	PROTOCOL + RMT ($\tau = .5$)	**89.20**	**89.75**
Paired difference (δ, pp)		$\mathbf{+5.09}$	$\mathbf{+6.56}$

Table 3: Task utility and its metric dependence in P^{2}OT on CIFAR100-LT: mean $\pm$ population standard deviation over five paired seeds, in percent (ARI scaled by 100). δ is matched minus raw in percentage points; its dispersion is computed from paired differences before rounding. Counts show positive/negative differences. All ten runs occupy 100 of 100 clusters.

Metric	Matched	Raw	δ	$+/-$
ARI	41.60 ± 3.22	35.08 ± 0.83	$+6.53 \pm 2.78$	$5/0$
ACC	46.96 ± 1.93	41.36 ± 1.09	$+5.60 \pm 2.12$	$5/0$
NMI	68.59 ± 1.17	68.75 ± 0.63	-0.16 ± 1.09	$2/3$
Purity	65.86 ± 1.61	68.59 ± 1.34	-2.73 ± 1.49	$0/5$
mACC	33.02 ± 1.01	36.94 ± 2.37	-3.93 ± 2.60	$0/5$

coefficient sweep identifies $\tau = .5$ as the clearest-performing tested intervention, so this coefficient is held fixed across the reported 10 paired runs rather than selected per seed. Relative to the raw plan ($\tau = 1$), RMT at $\tau = .5$ raises clustering ACC from 84.11% to 89.20% ($+5.09$ pp) and ARI from 83.19% to 89.75% ($+6.56$ pp) (Table 2). The comparison supports task value for a development-selected amplitude intervention.

Task gains and metric dependence in external vision clustering. The P^{2}OT task study uses CIFAR100-LT (imbalance ratio .01), DINO ViT-B/16 (Caron et al., 2021), five paired seeds, and 50 epochs. A training-loss rule selects models and heads; evaluation is fixed after training and before metric access on 10,847 long-tail training samples. Matching improves the primary sample-level outcomes in every seed, raising ARI by 6.53 points and ACC by 5.60 points on average (Table 3). Complementary partition analyses show that the gain is concentrated in the sample-weighted objective, while class-balanced and purity-oriented criteria favor raw. Appendix Fig. 2 characterizes this tradeoff through class-frequency and partition-structure analyses; these measurements are descriptive and are not used to attribute the task effect causally to the audited head-gradient component.

5.4 Boundary: control success need not yield task improvement

The utility studies establish that amplitude intervention can help. JUMBOT tests the converse concern: whether achieving the control objective is itself sufficient evidence of task improvement. We compare the two endpoints within the same prospective branches, then use the controlled testbed to delimit dataset dependence.

Direct separation between optimizer control and task outcome. The same JUMBOT branches that produce the $100/100$ versus $0/100$ coverage contrast have mean target-test accuracies of 97.604% and 97.606% (Table 1). The difference is -0.002 percentage points. The rule achieves the declared optimizer effect without an observed accuracy gain. This descriptive comparison measures two different objectives; it is not an accuracy-equivalence test.

A separate five-seed study compares source-only, raw, matched, and update-scale-control branches after common source training. The modified JUMBOT implementation intervenes on target-image amplitudes, corresponding to columns of the stored source-by-target plan, and includes both the joint cost and source cross-entropy in total-update comparisons. All four branches receive 500 updates, with paired data streams and FP64 model and moment arithmetic. Mean target accuracies are 67.218%, 97.574%, 97.532%, and 97.584%, respectively. Matched minus control is -0.052 points (sample SD 0.116; one of five differences positive). Adaptation is effective in this setup,

while removing amplitude heterogeneity adds no observed accuracy advantage over the scale control. Settings and per-seed results are in Appendices B and C.4.

Dataset dependence. Across nine TPF datasets with ten paired seeds each, matched mean ARI is higher on six and lower on three. The table is descriptive and combines explicitly distinguished evidence tiers (Appendix Table 4). The two dataset-level prospective studies support bounded mean ARI loss under a frozen non-inferiority margin, not superiority. Appendix C retains the complete endpoint table, negative results, and earlier inactive selectors; these delimit generality rather than strengthen the external task-gain claim.

6 Discussion

Implications for transport-based learning. A transport target should be evaluated both as an assignment and as a weighting of sample gradients. The latter can be changed downstream without redesigning the transport formulation. This gives a reusable way to investigate whether a learning effect depends on conditional assignments, amplitude heterogeneity, or their interaction along training. RMT supplies the controlled intervention and RAA measures its optimizer consequence; neither replaces evaluation of the intended task. Their long-term value is this separation of design and evaluation questions across transport-based losses.

From a local mechanism to a training outcome. RMT preserves conditional assignments at each fixed-state comparison. Once interventions produce different updates, later parameters, targets, and assignments can diverge. The covariance identity therefore explains the immediate gradient difference without identifying a particular audited component as the cause of a full-training task gain. The P^2OT comparison beyond update scale remains unresolved because two seeds fail the numerical realization criterion. The descriptive partition analysis explains which metrics favor each trained arm, not how a local gradient change caused the final partition.

Choosing the objective. The measured amplitude effect depends on sample forces and the optimizer state, so mass dispersion alone is insufficient. A declared deviation band makes optimizer control testable, but it does not specify whether confidence weighting, class balance, or sample-weighted accuracy should improve. Raw remains an essential reference when masses encode confidence, curriculum, or importance. The JUMBOT comparison and P^2OT metric tradeoff require task utility to be assessed on its own terms.

Scope and limitations. Downstream independence does not imply arbitrary freedom within an OT solver; Appendix A.4 characterizes a particular source-prior realizability problem. The control results concern measured parameter blocks and sampled future states in the studied settings, without guaranteed coverage of every update or transfer to a new task. TPF assumes aligned complete views and known cluster count. RAA incurs additional gradient and optimizer-proposal evaluations, with no runtime-savings claim. Appendix E gives a general smooth-class distinction between local observations and continuation utility, not a neural-task impossibility theorem. Development selection, retrospective diagnostics, and mixed evidence tiers remain explicitly identified in the experiments and appendices.

7 Conclusion

The same conditional transport assignment can exert different optimization forces because its row amplitude changes how sample gradients are combined. We turn this observation into a systematic study of a downstream optimization channel: RMT isolates it, the mass–force covariance explains it, and RAA measures and prospectively tests control of its optimizer effect. Fixed interventions in PROTOCOL and P^2OT establish task value, while JUMBOT establishes why control success cannot substitute for task evaluation. The resulting perspective makes amplitude an explicit design variable at the stopped-target interface, with controllability and utility treated as separate objectives.

AI USE STATEMENT

Generative AI Assisted literature discovery and Validation, and paper writing and revision. The authors take full responsibility for the validation work and the final content.

REPRODUCIBILITY STATEMENT

The anonymous supplementary material provides selected implementations and supporting records; it does not contain the complete codebase or the PROTOCOL implementation. We will release the full code and experiment configurations in a public GitHub repository after the review decisions are announced on December 16, 2026. Appendix B describes the experimental designs and reported settings; Appendix D specifies prefix selection, continuation, probe locations, and numerical checks. The release will provide framework-specific configurations and run instructions. Appendix B.4 summarizes the reporting scope and planned release.

REFERENCES

Martial Agueh and Guillaume Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011. doi: 10.1137/100805741. URL https://epubs.siam.org/doi/10.1137/100805741.

Blair Bilodeau, Natasha Jaques, Pang Wei Koh, and Been Kim. Impossibility theorems for feature attribution. *Proceedings of the National Academy of Sciences*, 121(2):e2304406120, 2024. doi: 10.1073/pnas.2304406120. URL https://arxiv.org/abs/2212.11870.

Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9912–9924, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/70feb62b69f16e0238f741fab228fec2-Abstract.html.

Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 9650–9660, October 2021. URL https://openaccess.thecvf.com/content/ICCV2021/html/Caron_Emerging_Properties_in_Self-Supervised_Vision_Transformers_ICCV_2021_paper.html.

Wei-Ting Chiu, Pei Wang, and Patrick Shafto. Discrete probabilistic inverse optimal transport. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 3925–3946. PMLR, 2022. URL https://proceedings.mlr.press/v162/chiu22b.html.

Lénaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018. doi: 10.1090/mcom/3303. URL https://www.ams.org/mcom/2018-87-314/S0025-5718-2018-03303-8/.

Marco Cuturi and Arnaud Doucet. Fast computation of Wasserstein barycenters. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 685–693, 2014. URL https://proceedings.mlr.press/v32/cuturi14.html. Part II.

Bowen Deng, Tong Wang, Lele Fu, Sheng Huang, Chuan Chen, and Tao Zhang. THESAURUS: Contrastive graph clustering by swapping fused Gromov–Wasserstein couplings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 16199–16207, 2025. doi: 10.1609/aaai.v39i15.33779. URL https://ojs.aaai.org/index.php/AAAI/article/view/33779. Issue 15.

Zhibin Dong, Meng Liu, Siwei Wang, Ke Liang, Yi Zhang, Suyuan Liu, Jiaqi Jin, Xinwang Liu, and En Zhu. Enhanced then progressive fusion with view graph for multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15518–15527, 2025.

Kilian Fatras, Thibault Sejourne, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3186–3197, 2021. URL https://proceedings.mlr.press/v139/fatras21a.html.

Andreas Fürst, Elisabeth Rumetshofer, Johannes Lehner, Viet T. Tran, Fei Tang, Hubert Ramsauer, David Kreil, Michael Kopp, Günter Klambauer, Angela Bitto, and Sepp Hochreiter. CLOOB: Modern Hopfield networks with InfoLOOB outperform CLIP. In *Advances in Neural Information Processing Systems*, volume 35, pp. 20450–20468, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/8078e76f913e31b8467e85b4c0f0d22b-Abstract-Conference.html.

Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 1885–1894, 2017. URL https://proceedings.mlr.press/v70/koh17a.html.

Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures. *Inventiones Mathematicae*, 211(3):969–1117, 2018. doi: 10.1007/s00222-017-0759-8. URL https://link.springer.com/article/10.1007/s00222-017-0759-8.

Xuqian Xue, Yiming Lei, Qi Cai, Hongming Shan, and Junping Zhang. PROTOCOL: Partial optimal transport-enhanced contrastive learning for imbalanced multi-view clustering. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 70105–70119, 2025. URL https://proceedings.mlr.press/v267/xue25c.html.

Jie Yao, Renjie Lin, Zhenghong Lin, and Shiping Wang. Multi-view clustering with graph regularized optimal transport. *Information Sciences*, 612:563–575, 2022. doi: 10.1016/j.ins.2022.08.117.

Chuyu Zhang, Hui Ren, and Xuming He. P^2OT: Progressive partial optimal transport for deep imbalanced clustering. In *International Conference on Learning Representations*, 2024a. URL https://openreview.net/forum?id=hD3sGVqPsr.

Qian Zhang, Lin Zhang, Ran Song, Runmin Cong, Yonghuai Liu, and Wei Zhang. Learning common semantics via optimal transport for contrastive multi-view clustering. *IEEE Transactions on Image Processing*, 33:4501–4515, 2024b. doi: 10.1109/TIP.2024.3436615.

Taichun Zhou, Zhibin Dong, Hao Tan, Siwei Wang, Xinwang Liu, En Zhu, Di Hu, Tianrui Liu, Chuankun Li, and Kunlun He. Imbalanced view contribution evaluation and refinement for deep incomplete multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 39606–39616, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Zhou_Imbalanced_View_Contribution_Evaluation_and_Refinement_for_Deep_Incomplete_Multi-View_CVPR_2026_paper.html.

A GRADIENT IDENTITIES AND INTERVENTION PROPERTIES

This appendix establishes the isolation properties and exact gradient mechanism used by RMT and RAA. The final subsection separately asks when a downstream amplitude change can also be realized by a specified family of transport problems; it is a boundary on solver-level interpretation, not a prerequisite for the downstream intervention.

All vector norms are Euclidean. The derivations apply to a fixed parameter block ϕ, with the target and its generating context held constant during differentiation.

A.1 SUPPORT, MASS, AND GRADIENT DECOMPOSITION

For a nonnegative Q, let $\mathcal{A} = \{i : s_i > 0\}$, $m = |\mathcal{A}| > 0$, and $\bar{s}_{\mathcal{A}} = m^{-1} \sum_{i \in \mathcal{A}} s_i$. Thus $\bar{s}_{\mathcal{A}} = \bar{s}$ in the main-text notation. Nonnegativity makes every inactive row exactly zero. On active rows, $P_{ik} = Q_{ik}/s_i$ is the unique conditional assignment. Equation (3) gives positive active masses because each is a convex combination of s_i and $\bar{s}_{\mathcal{A}}$, and

$$\sum_{i \in \mathcal{A}} s_i(\tau) = m\bar{s}_{\mathcal{A}}, \qquad s_i(\tau) - \bar{s}_{\mathcal{A}} = \tau(s_i - \bar{s}_{\mathcal{A}}).$$

Thus total mass, active support, and conditional assignments are preserved, while $\mathrm{CV}_{\mathcal{A}}[s(\tau)] = \tau \mathrm{CV}_{\mathcal{A}}[s]$. These are properties of the downstream target; the original transport constraints need not be preserved.

With $h_i = \sum_k P_{ik} \nabla_\phi \ell_{ik}$ and $\bar{h}_{\mathcal{A}} = m^{-1} \sum_{i \in \mathcal{A}} h_i$,

$$\begin{aligned} g_1 &= c_B \sum_{i \in \mathcal{A}} s_i h_i, & g_0 &= c_B \bar{s}_{\mathcal{A}} \sum_{i \in \mathcal{A}} h_i, \\ \Delta &= c_B \sum_{i \in \mathcal{A}} (s_i - \bar{s}_{\mathcal{A}}) h_i = c_B \, m \, \mathrm{Cov}_{\mathcal{A}}(s, h), & g_\tau &= g_0 + \tau\Delta. \end{aligned} \tag{8}$$

The covariance uses population moments:

$$\mathrm{Cov}_{\mathcal{A}}(s, h) = \frac{1}{m} \sum_{i \in \mathcal{A}} (s_i - \bar{s}_{\mathcal{A}})(h_i - \bar{h}_{\mathcal{A}}).$$

Subtracting $\bar{h}_{\mathcal{A}}$ leaves the sum unchanged because the centered masses sum to zero. Define

$$\sigma^2_{s,\mathcal{A}} = \frac{1}{m} \sum_{i \in \mathcal{A}} (s_i - \bar{s}_{\mathcal{A}})^2, \qquad \sigma^2_{h,\mathcal{A}} = \frac{1}{m} \sum_{i \in \mathcal{A}} \|h_i - \bar{h}_{\mathcal{A}}\|_2^2.$$

Vector Cauchy–Schwarz yields

$$\|\Delta\|_2 \le c_B \, m \, \sigma_{s,\mathcal{A}} \sigma_{h,\mathcal{A}}. \tag{9}$$

For batch averaging, $c_B = 1/B$ retains the factor m/B. Nonuniform masses can still have zero covariance with the conditional forces.

A.2 SCALAR EQUIVALENCE AND DERIVATIVE SCOPE

For $g_0 \neq 0$, the decomposition in Eq. (5) gives

$$\|g_1 - c g_0\|_2^2 = (1 + \alpha_{\mathrm{proj}} - c)^2 \|g_0\|_2^2 + \|\Delta_\perp\|_2^2.$$

Minimization over $c \in \mathbb{R}$ proves the residual formula. For two nonzero gradients, positive-step SGD equivalence requires both $\Delta_\perp = 0$ and $1 + \alpha_{\mathrm{proj}} > 0$. At $g_0 = 0$, report the endpoint norms and correction without assigning a projection direction.

If the target is differentiated, the full derivative includes the additional term $c_B \sum_{i,k} \ell_{ik} \nabla_\phi Q_{ik}$. Our identities describe the stopped/direct term. For several target terms plus an unaffected loss, their gradients add; the total candidate gradient is affine in τ at a fixed state. Adam proposals are computed separately using the incoming moments, step count, parameter groups, and all loss terms.

A.3 SUPPLEMENTARY GRADIENT MEASUREMENTS

The normalized total and transverse corrections are

$$E = \frac{\|\Delta\|_2}{\|g_1\|_2 + \|g_0\|_2 + \epsilon_g}, \qquad E_\perp = \frac{\|\Delta_\perp\|_2}{\|g_1\|_2 + \|g_0\|_2 + \epsilon_g}, \quad \epsilon_g > 0.$$

Both use the sum of endpoint norms; $E_\perp$ is not the transverse fraction of Δ. Aggregate endpoint gradients suffice for these statistics. At sparse states, per-sample forces additionally verify the covariance and its bound. A midpoint calculation checks the affine path; numerical resolution and solver termination are recorded independently.

A.4 WHEN AMPLITUDE CHANGES REMAIN TRANSPORT-OPTIMAL

We specialize entropic UOT optimality conditions (Chizat et al., 2018) to Eq. (16), varying only the source prior. This differs from inferring costs from couplings in inverse OT (Chiu et al., 2022).

Fix one view with finite cost matrix $C \in \mathbb{R}^{m\times K}$. Assume $\varepsilon = \varepsilon_{\mathrm{uot}} > 0$, $\rho_s, \rho_t > 0$, and priors $a \in \mathbb{R}^m_{++}, b \in \mathbb{R}^K_{++}$. Write the exact minimizer as $\Pi = \mathrm{diag}(s)P$, with $s \in \mathbb{R}^m_{++}$, $P \in \mathbb{R}^{m\times K}_{++}$, and $P\mathbf{1}_K = \mathbf{1}_m$, and let $r = P^\top s$ and $S = \mathbf{1}_m^\top s$. All m rows are active ($m = B$ for TPF). Here $b = b_t$; all inputs except a, including the entropy reference, remain fixed.

Proposition A.1 (source-prior realizability). For $s' > 0$ with $\mathbf{1}_m^\top s' = S$, the plan $\Pi' = \mathrm{diag}(s')P$ is the minimizer for some positive source prior a' if and only if

$$P^\top s' = P^\top s. \tag{10}$$

When it exists, this prior is unique and equals

$$a'_i = a_i \left(\frac{s'_i}{s_i}\right)^p, \qquad p = 1 + \frac{\varepsilon}{\rho_s} > 1. \tag{11}$$

Thus the realizable amplitudes form $\{s' > 0 : P^\top s' = r\}$, a relatively open convex set of dimension $m - \mathrm{rank}(P)$.

Proof. Strict convexity makes stationarity sufficient and the minimizer unique. At the baseline,

$$C_{ik} + \varepsilon \log \Pi_{ik} + \rho_s \log(s_i/a_i) + \rho_t \log(r_k/b_k) = 0. \tag{12}$$

Subtract these from the equations for Π', and set $d_i = s'_i/s_i$ and $r' = P^\top s'$. Then

$$(\varepsilon + \rho_s) \log d_i - \rho_s \log(a'_i/a_i) + \rho_t \log(r'_k/r_k) = 0 \quad \text{for all } i, k. \tag{13}$$

The row-only terms force r'_k/r_k to be constant in k. Equal total mass makes this constant one, giving Eqs. (10)–(11). Conversely, these identities ensure stationarity. Since $P\mathbf{1}_K = \mathbf{1}_m$, Eq. (10) already preserves total mass; positivity preserves the affine dimension. □

Consequences for RMT and normalized priors. For RMT applied to this full plan, $s(\tau) = s+(1-\tau)(\bar{s}\mathbf{1}_m - s)$, with $\bar{s} = S/m$. Every $\tau \in [0, 1)$ satisfies Eq. (10) exactly when $P^\top(\bar{s}\mathbf{1}_m - s) = 0$; otherwise none does. If the admissible priors must also preserve their total mass, realizability additionally requires

$$\sum_i a_i (s'_i/s_i)^p = \sum_i a_i. \tag{14}$$

This constraint can remove all nontrivial freedom. For example, if $a_i = \alpha s_i$ for some $\alpha > 0$, equal plan mass gives $\sum_i (s_i/S) d_i = 1$. Jensen's inequality for the strictly convex function $x \mapsto x^p$ yields $\sum_i (s_i/S) d_i^p \geq 1$, with equality only when all $d_i = 1$. Thus Eq. (14) forces $s' = s$, even when $\ker(P^\top)$ is nonzero. At this point the first-order normalization condition vanishes on every fixed-mass direction, although the exact feasible set is a singleton. If the prior is held fixed, Eq. (11) also forces $s' = s$.

Corollary A.2 (gradient distinguishability). At a fixed stopped-loss state, let $H = [h_1 \cdots h_m] \in \mathbb{R}^{d_\phi \times m}$ collect the conditional forces in Eq. (8); d_ϕ is the parameter-block dimension. Under the unrestricted positive-prior family of Proposition A.1, the attainable gradient changes have linear span

$$c_B H \ker(P^\top), \qquad \dim\big(H \ker(P^\top)\big) = \operatorname{rank}\begin{bmatrix} P^\top \\ H \end{bmatrix} - \operatorname{rank}(P). \tag{15}$$

The gradient difference is $c_B H(s' - s)$, and every sufficiently small kernel direction preserves positivity. Rank–nullity for H restricted to $\ker(P^\top)$ proves the formula. In particular, all realizable changes are gradient-invisible if and only if $H = LP^\top$ for some $L \in \mathbb{R}^{d_\phi \times K}$. Amplitude freedom inside the solver therefore need not produce a learning signal.

TPF fixes $a = \mathbf{1}_B/B$ and applies RMT downstream, without this inverse construction. P^{2}OT's sliced target and other entropy conventions require separate optimality conditions.

B Implementations and evaluation settings

The settings distinguish three experimental roles: TPF supplies controlled mechanism and dataset-sensitivity studies; external fixed interventions test task utility; and Appendix D specifies prospective optimizer control. Selection rules and evaluation access are reported separately because these roles carry different evidence claims.

B.1 Controlled testbed: objective

TPF is the controlled testbed used to study the amplitude channel, not a competing main method. Its objective below specifies where the stopped transport target enters learning and which other terms remain present in the comparisons.

For aligned views $\{X^{(v)}\}_{v=1}^V$, encoders and projectors produce unit-normalized embeddings $z_i^{(v)} \in \mathbb{R}^d$. We maintain K shared trainable prototypes $\hat{c}_k$ and use cosine cost $D_{ik}^{(v)} = 1 - (z_i^{(v)})^\top \hat{c}_k$. With source prior $a = \mathbf{1}_B/B$ and target prior $b_t \in \Delta^{K-1}$, each view solves

$$\begin{aligned}\Pi_v^\star = \arg\min_{\Pi \geq 0} \langle D^{(v)}, \Pi \rangle + \varepsilon_{\mathrm{uot}} \sum_{i,k} \Pi_{ik}(\log \Pi_{ik} - 1) \\ + \rho_s \,\mathrm{KL}(\Pi \mathbf{1}_K \| a) + \rho_t \,\mathrm{KL}(\Pi^\top \mathbf{1}_B \| b_t).\end{aligned} \tag{16}$$

ρ_s and ρ_t penalize source and target marginals, respectively; the subscript in b_t indexes training time. Generalized KL includes the linear correction terms for nonnegative measures. The iterative solver and updates of b_t are outside backpropagation.

Let $\widetilde{\Pi}_v^{(\tau)}$ be the RMT transform of each view's detached solver output $\Pi_v^\star$, using its active-row mean as in Eq. (3). The transport component is

$$\mathcal{L}_{\mathrm{tr}}(\tau) = \frac{1}{V} \sum_{v=1}^V \left\langle D^{(v)}, \mathrm{sg}\big(\widetilde{\Pi}_v^{(\tau)}\big) \right\rangle. \tag{17}$$

The no-transport-gradient intervention sets its loss coefficient to zero.

Fusion and total loss. In the fusion equation, $p_i^{(v)}$ denotes a view's predicted cluster distribution. The default prediction distribution is the equal-view arithmetic pool

$$\bar{p}_i = \frac{1}{V} \sum_{v=1}^V p_i^{(v)}, \qquad \hat{y}_i = \arg\max_k \bar{p}_i[k]. \tag{18}$$

A fixed-support entropic Wasserstein barycenter is retained only as a fusion-geometry control (Agueh & Carlier, 2011; Cuturi & Doucet, 2014). Symmetric sample-level InfoLOOB (Fürst et al., 2022) aligns views, and a squared-cosine prototype penalty discourages collapse. The objective is

$$\mathcal{L} = \lambda_{\mathrm{info}} \mathcal{L}_{\mathrm{info}} + \lambda_{\mathrm{tr}} \mathcal{L}_{\mathrm{tr}}(\tau) + \lambda_{\mathrm{div}} \mathcal{L}_{\mathrm{div}}. \tag{19}$$

B.2 CONTROLLED TESTBED: EVALUATION DESIGN

The main endpoint is adjusted Rand index (ARI), with Hungarian-matched clustering accuracy (ACC), normalized mutual information (NMI), purity, and occupied-cluster fraction as secondary measurements. TPF assumes aligned complete views and known cluster count. Predictions precede label evaluation; the reported experiments use fixed seeds, endpoints, and selection rules. Dispersions are population standard deviations unless stated otherwise.

The first four datasets combine a three-seed screen with seven independent seeds; labels had informed earlier project development. The next three combine five development seeds and five held-out stochastic seeds. OutdoorScene and WebKB were selected at the dataset level before local feature or label access: seeds 66–75, both endpoints, 50 epochs, a -0.02 ARI margin, and an occupancy guard were fixed in advance. The nine-dataset comparison contains 180 endpoint runs, with 36 additional interior-coefficient runs on the first four datasets. Earlier 200-epoch mechanism studies remain separate.

For each prospective dataset, let $\delta_j = \mathrm{ARI}_{j,\tau=0} - \mathrm{ARI}_{j,\tau=1}$. With $\bar{\delta} = 10^{-1}\sum_{j=1}^{10}\delta_j$, the lower bound is

$$\bar{\delta} - t_{0.975,9}\frac{s_\delta}{\sqrt{10}}, \qquad s_\delta^2 = \frac{1}{9}\sum_{j=1}^{10}(\delta_j - \bar{\delta})^2.$$

$t_{0.975,9}$ is the .975 quantile of Student's t distribution with nine degrees of freedom. The per-dataset one-sided level .025 adjusts for two datasets. Each lower bound must exceed $-.02$, all runs must be numerically valid, and matched mean occupied-cluster fraction must be no more than .05 below raw, with no run below .5.

The paired-t analysis assumes independent seed-level differences and approximate normality; it concerns mean ARI loss, not a guarantee for individual seeds.

The marginal-strength study uses two datasets, three seeds (76–78), $\rho_s \in \{.125, .5, 2\}$, both endpoints, and 50 epochs (36 runs). All 36 runs in this study use detached internal-FP64 solves with tolerance 10^{-6} and a 5,000-iteration ceiling; outputs are cast to model precision. This common setting follows an initial FP32 convergence plateau.

B.3 EXTERNAL IMPLEMENTATIONS

PROTOCOL. The Hdigit task comparison applies downstream RMT in PROTOCOL while retaining its representation backbone and transport objective. A development coefficient sweep selects $\tau = .5$, which is held fixed across the reported 10 paired runs against raw ($\tau = 1$). Table 2 reports mean ACC and ARI and their paired mean differences. This comparison is development-selected task evidence, not a prospective RAA control experiment.

P^{2}OT. The task study uses CIFAR100-LT with imbalance ratio .01, DINO ViT-B/16, two heads, batch size 512, five paired seeds (76–80), and 50 epochs. The actual stopped soft-target cross-entropy consumes the current-batch slice of the memory-bank plan. RMT is applied after slicing. A training-loss rule selects the checkpoint and head. Labels construct the 10,847-sample long-tail subset; task metrics are evaluated on that subset after training and prediction freezing. This is transductive clustering evaluation. Here ρ denotes the transported-mass fraction, distinct from the TPF source-marginal penalty ρ_s.

JUMBOT. Our modified official implementation uses the CNN generator and classifier of JUMBOT (Fatras et al., 2021). SVHN provides labeled source images; MNIST provides unlabeled adaptation images and a separately evaluated labeled test set. The source split uses a fixed class-stratified seed (1910), holding out 7,320 of 73,257 images for source validation. Target labels are not used in training or coefficient selection.

Let f be the feature extractor, h the classifier, and CE cross-entropy. For source index i and target index j, the joint cost is

$$M_{ij} = .1\|f(x_i^s) - f(x_j^t)\|^2 + .1\,\mathrm{CE}(y_i^s, h(f(x_j^t))).$$

The entropic unbalanced solver uses uniform minibatch reference marginals, entropy coefficient .1, marginal relaxation 1, and a recorded stopping threshold 10^{-6}. The plan Π is detached before the network update, and the objective is

$$L = L_{\text{source CE}} + \sum_{ij} \text{sg}(\Pi_{ij}) M_{ij}.$$

The audited sample is a target image: define $s_j = \sum_i \Pi_{ij}$ and $P_{ji} = \Pi_{ij}/s_j$. Thus the stored plan is transposed for the row convention. RMT scales its active columns by $[\bar{s} + \tau(s_j - \bar{s})]/s_j$. It changes the complete joint-cost term; the source CE term is kept and included in total-update comparisons.

Four-arm comparison. Seeds 920–924 each receive ten source epochs (1,310 updates), then 500 updates in each of source-only, raw, matched, and scale-control branches. Batch size is 500. The generator and classifier use separate Adam optimizers at learning rate 2×10^{-4}, reset at the common branch point. Models, moments, costs, and plans use FP64; preprocessing uses FP32. Relevant data streams are paired, and all fixed-endpoint predictions precede target-label evaluation. This modified schedule defines an intervention study, not an original-benchmark reproduction.

Update-scale comparison. For raw and matched proposals u_1 and u_0 at the control's own parameters and incoming optimizer state,

$$u_{\text{control}} = \frac{\|u_0\|_2}{\|u_1\|_2} u_1 \quad (\|u_1\|_2 > 0).$$

A single learning-rate multiplier acts on the concatenated proposal across all trainable parameters and optimizer groups; raw moments are committed. The comparison therefore matches a same-state counterfactual norm. Stored parameter displacement is checked separately from proposal arithmetic.

B.4 Reproducibility and code release

This manuscript reports TPF endpoint summaries, P^{2}OT gradient and partition summaries, and per-seed JUMBOT task results and prospective coverage. These reported summaries are distinct from the implementation and experiment configurations planned for release. Per-sample covariance checks and aggregate-only measurements are implementation options within the same RAA procedure.

Public code and configurations. The anonymous supplement contains selected implementations and supporting records for TPF, P^{2}OT, and JUMBOT, rather than the complete codebase; the PROTOCOL implementation is not included. The full code and experiment configurations will be released in a public GitHub repository after the review decisions are announced on December 16, 2026, with instructions for the reported training and evaluation workflows. The configurations will identify the framework and dataset setting, intervention coefficient or selection rule, seed, training schedule, and endpoint convention.

Reproducing the comparisons. Task comparisons and optimizer-control studies require different workflows. For a task comparison, use the study's paired data streams, coefficient choice, and checkpoint/head rule, then freeze predictions before computing task metrics. For prospective control, preserve the declared prefix/continuation split, incoming optimizer and auxiliary state, candidate grid, eligibility criterion, and future probe schedule in Appendix D. Report coverage per seed; pooled probe counts are descriptive. Check solver termination separately from stored-update realization, and retain the finite-iteration targets included by the original protocol.

Reporting scope. The PROTOCOL/Hdigit table reports paired-run means without a dispersion estimate. Its coefficient is development-selected, so it supports a task comparison under that selection procedure rather than independent prospective confirmation.

C Additional mechanism and task results

This appendix separates supporting checks from the main external evidence. The controlled testbed tests intervention strength and task dependence; P^{2}OT diagnostics measure local geometry and de-

Table 4: Primary ARI at the RMT endpoints, mean ± population std over ten paired seeds. δ = matched minus raw. Paired differences are computed before rounding. Tiers: R+I, retrospective screen plus independent seed extension; D+H, development plus held-out stochastic replication; P-NI, dataset-level prospectively frozen non-inferiority.

Dataset	Tier	Matched $\tau = 0$	Raw $\tau = 1$	Paired δ	$\delta > 0$
MNIST–USPS	R+I	$.8095 \pm .0607$	$.7795 \pm .0657$	$+.0300 \pm .0498$	$7/10$
UCI Multiple	R+I	$.6385 \pm .0612$	$.6352 \pm .0598$	$+.0033 \pm .0299$	$5/10$
Caltech101-20	R+I	$.1073 \pm .0077$	$.1047 \pm .0094$	$+.0027 \pm .0036$	$9/10$
Scene15	R+I	$.1454 \pm .0136$	$.1459 \pm .0133$	$-.0005 \pm .0025$	$4/10$
100Leaves	D+H	$.1706 \pm .0255$	$.1643 \pm .0205$	$+.0063 \pm .0061$	$9/10$
BBCSport	D+H	$.0741 \pm .0253$	$.0708 \pm .0252$	$+.0033 \pm .0052$	$9/10$
MSRC-v5	D+H	$.4623 \pm .0462$	$.4533 \pm .0483$	$+.0090 \pm .0180$	$7/10$
OutdoorScene	P-NI	$.3239 \pm .0265$	$.3274 \pm .0293$	$-.0035 \pm .0097$	$6/10$
WebKB	P-NI	$.0334 \pm .0303$	$.0354 \pm .0331$	$-.0020 \pm .0106$	$5/10$

scribe the resulting partitions; scale controls delimit the long-run interpretation; and the JUMBOT comparison checks for an added task advantage. None of the descriptive analyses upgrades the evidence status of the underlying study.

C.1 Controlled testbed: intervention strength and task dependence

Task effects vary across datasets. Ten paired seeds on nine TPF datasets give higher matched mean ARI on six and lower mean ARI on three (Table 4). The first four rows combine a retrospective screen with independent seeds; the next three combine development and held-out stochastic replication. OutdoorScene and WebKB are dataset-level prospective studies, frozen before local feature or label access. No test on the original four datasets rejects after Holm correction; the nine-dataset table remains descriptive.

On the two prospective datasets, adjusted one-sided paired-t lower bounds of -0.0108 and -0.0100 exceed the frozen -0.02 ARI margin, with full occupancy. Together with the negative mean differences, these bounds support a scoped bounded-mean-loss result rather than superiority (Table 6). Earlier negative evidence is consistent with this separation: transport-gradient removal helps Scene15 but harms BDGP, a stability-based selector does not select an intervention, and a gradient-budget controller is inactive in 410 inspected units. The studies below remain separate from the optimizer-level confirmation in Appendix D.

Transport utility and amplitude utility. The original interventions separate two questions. On Caltech101-20, removing the stopped transport gradient reduces mean ARI from 0.1046 (raw) and 0.1100 (matched) to 0.0331; on Scene15, no transport instead gives the highest ARI (0.1283 versus $0.1168/0.1197$). On prospectively frozen BDGP, raw/matched/no-transport ARI is $0.6875/0.6916/0.0621$, while raw is substantially more seed-stable than matched. Thus transport-gradient utility is distinct from amplitude utility. Arithmetic and Wasserstein fusion also have no consistent winner.

The diagnostics reveal different relationships between gradient and task effects. In all 12 audited states, observed mass–force pairing is more extreme than all 128 row-mass permutations, but UCI has a larger relative gradient correction than MNIST–USPS without a comparable task gain. Endpoint representation gradients remain highly aligned, and a predeclared INFOLOOB-compatibility explanation for the Caltech/Scene difference fails (same-seed joint pass $1/3$). On MNIST–USPS, a modest CKA change accompanies a large cross-seed partition-ARI change (0.7956 raw to 0.9729 matched), consistent with amplification along training trajectories.

Synthetic and continuous interventions. In a frozen synthetic study with 200 replicates per cell, mass–force alignment in $\{-2, -1, 0, 1, 2\}$ and two orthogonal direction-noise levels controls the sign predicted by Eq. (8): every nonzero-alignment cell has the correct mean sign and a correct-sign fraction of 1.0. The maximum affine-identity error is 5.11×10^{-15} and the maximum budget excess is 2.22×10^{-16}. This verifies the predicted gradient relation in the controlled setting.

On four real datasets with seeds 46–48 and $\tau \in \{0, .25, .5, .75, 1\}$, MNIST–USPS ARI decreases from 0.8213 at matched to 0.7352 at raw. UCI peaks weakly at $\tau = .25$ (0.6700 versus $0.6644/.6472$ at the endpoints), while Caltech101-20 and Scene15 are nearly flat. Continuous control therefore exposes strong, weak-interior, and insensitive regimes; no common optimal τ emerges.

Marginal relaxation. As source regularization increases from $.125$ to 2, mass variation and gradient correction shrink in all six dataset–seed units, with mean ratios $.082999$ and $.080787$. Absolute paired ARI sensitivity contracts in all six units and is strictly monotone across the three strengths in five. Signed mean differences change sign: MNIST–USPS gives $+.111111, -.029730, +.001022$; Scene15 gives $-.000394, +.000507, +.000331$. Table 5 reports absolute sensitivity.

Table 5: Mean absolute paired ARI difference $|\mathrm{ARI}_{\tau=0} - \mathrm{ARI}_{\tau=1}|$ over three seeds. This is intervention sensitivity, not improvement.

Dataset	$\rho_s = .125$	$\rho_s = .5$	$\rho_s = 2$
MNIST–USPS	.113715	.029730	.001649
Scene15	.002803	.002026	.000692

Endpoint outcomes and non-inferiority. The Holm-corrected exact tests mentioned above use the seven independent seeds of the original four datasets. Table 6 gives the prospective lower bounds and occupancy checks. Scores on WebKB, BBCSport, and Caltech101-20 are low in absolute terms, further limiting the practical interpretation of bounded mean ARI loss.

Table 6: Frozen prospective paired-ARI non-inferiority (ten seeds/dataset). LB is the Bonferroni-adjusted one-sided lower confidence bound; occupancy reports mean matched/raw fractions.

Dataset	Mean δ	LB	Margin	Occupancy	Gate
OutdoorScene	$-.0035$	$-.0108$	$-.0200$	$1.00/1.00$	pass
WebKB	$-.0020$	$-.0100$	$-.0200$	$1.00/1.00$	pass

Other tested explanations. A cross-seed stability selector abstains on all three development datasets and fails its actionability criterion. An earlier fixed-state gradient-budget rule also remains inactive in all 410 inspected units. It uses $r = \|g_1 - g_0\|_2/(\|g_0\|_2 + \epsilon)$ and $\tau = \min\{1, .5/\max(r, \epsilon)\}$, with $\epsilon > 0$, measured relative to matched; it is distinct from the optimizer deviation measured relative to raw in Section 4. NUS-WIDE occupies all 31 clusters but has ARI $.0608$, and a consensus-distillation stress test combines high view agreement with cluster collapse. Two 100Leaves occupancy incidents were recorded during development; the underlying records are not included in this supplement.

C.2 P^{2}OT: GRADIENT AND PARTITION MEASUREMENTS

The gradient measurements check whether the amplitude mechanism occurs in an external loss. The post-result partition analyses serve a different purpose: they explain which evaluation criteria favor each trained arm, without identifying the local gradient effect as the cause of the task difference.

Interpreting the metric tradeoff. Two post-result identities sharpen this comparison without new predictions. Let n_c be the size of class c and N the total number of samples. Let r_c denote class recall under each arm's original micro-ACC-optimal assignment, and $w_c = Kn_c/N$ its normalized class frequency. Then $\mathrm{ACC} - \mathrm{mACC} = \mathrm{Cov}_c(w_c, r_c)$, where classes receive equal weight in the covariance. Thus $\mathrm{Cov}_c(w_c, r_c^{\mathrm{matched}} - r_c^{\mathrm{raw}})$ is positive in every seed, averaging 0.09527: the sample-weighted gain coexists with a class-average loss. This covariance describes the relation between class frequency and recall change. Also let $G = \mathrm{Purity} - \mathrm{ACC} \geq 0$, the score loss from imposing one-to-one matching instead of independently assigning each cluster its majority class. Mean G falls from 0.27230 to 0.18903, with a decrease in every seed. Consequently $\delta_{\mathrm{ACC}} = \delta_{\mathrm{Purity}} - \delta_G = -0.02725 - (-0.08327) = 0.05602$. The smaller assignment loss offsets lower majority-class purity in this exact score decomposition. Equalizing sample amplitudes also does not

Table 7: P^{2}OT fixed-state RAA, stratified by training arm and diagnostic transport budget. CV averages the 20 component targets per row; E and $E_\perp$ average the 10 two-term probes (five seeds, two probes each), not independent seed replications. $E = \|\Delta\|/(\|g_1\| + \|g_0\| + \epsilon_g)$ and $E_\perp = \|\Delta_\perp\|/(\|g_1\| + \|g_0\| + \epsilon_g)$. The model arm identifies the checkpoint; both gradient endpoints are always measured at that same checkpoint.

Training arm	ρ	Mass CV	E	$E_\perp$
Matched	0.1	.8264	.4753	.4264
Raw	0.1	.3824	.7122	.1715
Matched	0.5	.5099	.3647	.2820
Raw	0.5	.3096	.6800	.1712
Matched	1	1.40×10^{-8}	9.25×10^{-8}	6.16×10^{-8}
Raw	1	7.21×10^{-8}	1.84×10^{-7}	1.71×10^{-7}

equalize class weights: at a frozen target, each class's total matched mass remains proportional to its number of active samples.

Localizing the tradeoff in frozen partitions. We analyze all ten frozen contingency tables, using frequency groups specified before receiving these tables: head ($n_c > 100$, 35 classes), middle ($20 \leq n_c \leq 100$, 35), and tail ($n_c < 20$, 30). Mean class recall under the original mapping changes by $+5.90$, -10.73, and -7.45 percentage points, respectively; each group has the same sign in all five seeds (Fig. 2). The head group contains 81.35% of samples and contributes $+7.50$ points to overall ACC, offsetting middle and tail contributions of -1.71 and -0.20 points. Separately optimizing the macro assignment still gives lower matched accuracy (35.24% versus 39.21%, five negative paired differences), so the aggregate class-average deficit is not solely an artifact of choosing the micro-optimal mapping. This macro-optimized score is reported separately from mACC under the original mapping.

Pair recall increases from 0.2661 to 0.3703 in all five seeds; pair precision decreases on average from 0.5598 to 0.5072, with four decreases and one increase. In every seed, fewer same-class pairs are split, but more different-class pairs share a cluster. Let Y denote class and Z cluster. The sample-weighted conditional entropy $H(Z \mid Y)$ falls from 1.5999 to 1.4517 nats while $H(Y \mid Z)$ rises from 1.0705 to 1.1581, with both directions consistent across seeds. Thus less splitting and more mixing coexist in these empirical partition measures. Head classes account for 96.27% of true same-class pairs and their recovered-pair gain exceeds the total mean gain, while middle and tail gains are negative on average. This locates the pair-recall improvement primarily in the head classes. Moreover, mutual information decreases in all seeds (2.9370 to 2.8494 nats on average), while cluster entropy also falls (4.5369 to 4.3010); nearly unchanged NMI therefore does not mean unchanged unnormalized information.

Gradient measurements. Two fixed 512-sample probes per selected model use reproducible augmentations and reconstructed 5,120-entry memory banks. Each cross-view term uses $c_B = 1/512$. Endpoint and midpoint gradients are evaluated in the selected cosine-head weight block, with per-sample and aggregate calculations checked independently. All 120 component and 60 assembled checks pass; the maximum symmetric-normalized covariance residual is 1.59×10^{-15}. All measured row masses are positive.

At transport budget $\rho = .5$, all 40 solver calls terminate successfully and all 20 assembled measurements have $E_\perp \in [.0670, .4720]$. At $\rho = 1$, $E \leq 1.78 \times 10^{-6}$; only 8/40 targets pass the strict component uniformity test. At $\rho = .1$, 20/40 solver calls miss the stopping criterion and remain included as finite-iteration targets. Across 40 within-state budget sequences, mass CV decreases strictly in all 40, while E and $E_\perp$ decrease in 36 and 27, respectively. Thus mass dispersion alone cannot replace gradient measurement. Component cancellation ratios at $\rho < 1$ range from $.6715$ to $.9966$.

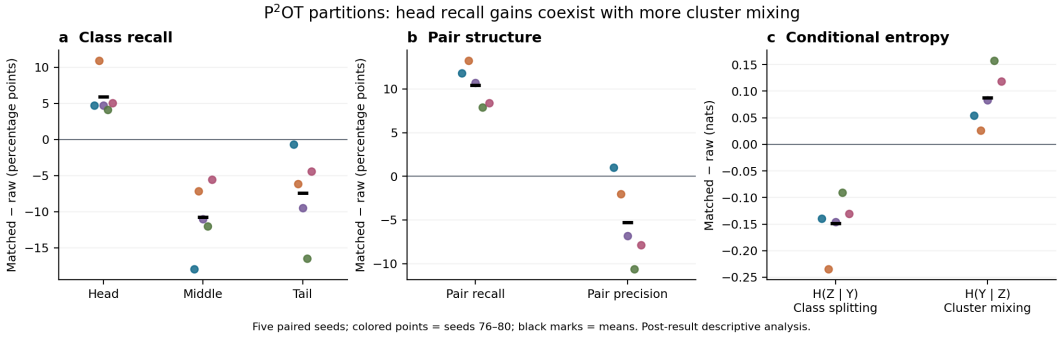

Figure 2: Post-result descriptive analysis of P^{2}OT's frozen partitions. Colored points are the five paired seeds; black marks are means. Matched minus raw changes show higher head recall, lower middle and tail recall, higher pair recall, and lower mean pair precision. Lower $H(Z \mid Y)$ and higher $H(Y \mid Z)$ describe less class splitting and more cluster mixing. These measures characterize the task-metric tradeoff, not its causal connection to a local gradient component.

Interpretation. These analyses characterize the frozen partition tradeoff: stronger sample-weighted performance coexists with lower class-average recall, less class splitting, and more cluster mixing. The group averages and partition measures do not imply uniform class changes or globally nested partitions. Causal linkage to the earlier local gradient effect is outside this analysis.

C.3 Update-scale comparison

The P^{2}OT scale comparison adds five 50-epoch runs to the ten endpoint runs on the same split. All 7,241,472 trainable parameters enter the proposal norm. The primary comparison fixes epoch 50 and head 0; loss-selected evaluation is secondary. All 30 predictions are frozen before task evaluation.

Table 8: Scale-control primary endpoint: means across five seeds, multiplied by 100. Differences are matched minus control in points. Dispersion is the population standard deviation of paired differences; counts give positive/negative differences. These are descriptive outcomes under a numerically limited control.

Metric	Raw	Scale control	Matched	Paired difference	$+/-$
ARI	35.35	35.41	40.48	$+5.07 \pm 4.73$	4/1
ACC	41.84	41.56	46.05	$+4.49 \pm 3.59$	4/1
NMI	68.79	68.66	68.30	-0.36 ± 1.15	3/2
Purity	68.51	68.25	65.55	-2.70 ± 1.35	0/5
mACC	36.68	36.93	32.48	-4.44 ± 2.91	1/4

The primary matched-minus-control ARI differences are $+9.9082$, $+9.1551$, $+6.5019$, $+2.7136$, and -2.9527 points. Control minus raw averages $+.0577$ points. At the secondary loss-selected endpoint, matched minus control averages $+5.8372$ and control minus raw $+.6920$ points, both positive in all seeds. Checkpoint and head selection change together between these endpoint rules.

The predeclared numerical criterion requires at least 99% of each seed's 1,050 updates to realize the intended vector within 1% relative error. Seeds 76 and 80 fail (Table 9); all 41 exceedances occur in the final epoch. The full-training effect beyond update scale therefore remains unresolved.

A separate five-seed, three-arm, 21-step study uses FP32 frozen-backbone computation and FP64 trainable parameters and moments. All 15 cells pass their numerical criterion, but ARI, Purity, and mACC remain unchanged at the tiny final learning rate. This verifies local update realization at the tested states.

Table 9: Original scale-control realization criterion. Passing fraction is evaluated separately for every seed.

Seed	Passing steps	Passing fraction	Exceedances	Criterion
76	1030/1050	98.0952%	20	Fail
77	1050/1050	100%	0	Pass
78	1050/1050	100%	0	Pass
79	1050/1050	100%	0	Pass
80	1029/1050	98.0000%	21	Fail

Table 10: JUMBOT four-arm fixed-endpoint target-test accuracy (percent). Paired differences are in percentage points. These are modified-schedule intervention results, not original-paper benchmark reproduction.

Seed	Source-only	Raw	Matched	Control	Matched–control
920	68.73	97.78	97.78	97.83	$-.05$
921	67.79	97.68	97.68	97.69	$-.01$
922	69.09	97.55	97.51	97.59	$-.08$
923	65.34	97.37	97.11	97.33	$-.22$
924	65.14	97.49	97.58	97.48	$+.10$
Mean	67.218	97.574	97.532	97.584	$-.052$

C.4 JUMBOT: FIXED-ENDPOINT TASK COMPARISON

The matched-minus-control accuracy difference averages $-.052$ points with sample standard deviation .116 (denominator $n - 1$); one of five differences is positive. The large gain over source-only shows effective adaptation in this setting, while the amplitude intervention adds no observed accuracy advantage over the update-scale control.

D PROSPECTIVE OPTIMIZER CONTROL

These studies test future optimizer-effect coverage after prefix selection. Their primary endpoint is the declared deviation-band criterion. The P^{2}OT confirmation evaluates no task metrics, while JUMBOT reports task accuracy as a separate secondary outcome. Earlier rule development is retained below to make the prospective boundary explicit.

D.1 SELECTION AND EVALUATION RULE

Each candidate proposal uses the same current parameters, complete optimizer state, and stopped target. All trainable parameters enter D_τ in Eq. (7). The grid is $\{0, .25, .5, .75, .875, 1\}$, with selectable candidates $\tau \leq .875$. The nominal band $.02 \leq D_\tau \leq .1$ has a .001 guard, giving $[.021, .099]$. At least nine of ten prefix probes must be in band; the smallest eligible τ is selected. If no candidate qualifies, the rule abstains. Future success requires at least 18/20 in-band probes during training with the selected coefficient. These thresholds define the study's update objective.

D.2 P^{2}OT

The control study uses CIFAR100-LT at imbalance ratio .1 (19,573 samples, 38 batches per epoch). The learning-rate and transport schedules retain their 50-epoch horizons. Three new seeds (86–88) supply ten-epoch prefixes, with probes at epochs 5 and 10 and batch positions 0, 9, 18, 27, 37. One shared candidate must meet the prefix criterion in every seed.

Rule selection. An earlier five-seed study (81–85) used the grid without .875 and only an upper deviation bound. Nontrivial candidate counts were $0, 0, 0, 6$ out of 50 for $\tau = 0, .25, .5, .75$; the rule abstained and stopped after the prefixes. The revised grid and two-sided band were fixed before

seeds 86–88. The confirmation therefore tests new seeds and continuations in an already studied setting.

Continuation. Each epoch-10 state branches into raw and selected-action training through epoch 20, restoring model, Adam moments, criterion state including memory banks and solver warm starts, and random states. Epoch-seeded nonpersistent workers pair sample order and augmentations. Future probes occur at epochs 12, 15, 18, 20 and the same five batch positions. Candidate deviations are recomputed at each branch's own state. No task metrics are evaluated in this study.

Table 11: Future control: in-band counts for the selected $\tau = .875$. Counts summarize repeated probes within each seed.

Stage	Seed 86	Seed 87	Seed 88	Total
Prefix selection	10/10	10/10	10/10	30/30
Action own trajectory (primary)	19/20	20/20	19/20	58/60
Raw counterfactual (secondary)	19/20	18/20	19/20	56/60

Only $\tau = .875$ qualifies at the new prefixes, with 30/30 observations in band. The action's two exceeding probes are seed 86, epoch 12, batch 18 ($D_\tau = .103105$), and seed 88, epoch 12, batch 27 ($D_\tau = .110328$); all solver calls at both probes meet their stopping criterion. Full matching at the same action states gives deviations .388394–.846117, with 0/60 in band; this is a counterfactual proposal comparison, not a separately trained branch.

Table 12: Update realization and solver convergence. "All calls" requires every solver call within a probe to meet its stopping criterion. Primary coverage uses all probes; the last column summarizes the converged subset.

Stage	Native updates pass	Converged calls	All calls	Subset in band
Prefix	30/30	47/120	11/30	11/11
Raw	60/60	234/240	57/60	53/57
Action	60/60	193/240	46/60	44/46

All 150 recorded native-update checks pass, with maximum relative vector error 6.4012×10^{-5}. The numerical criterion is at least 99% of recorded updates within 1% error per seed, which requires all ten or twenty recorded updates to pass. Coverage uses all finite-iteration targets. Solver convergence is less frequent on the selected branch than raw, with each branch measured at its own evolving states.

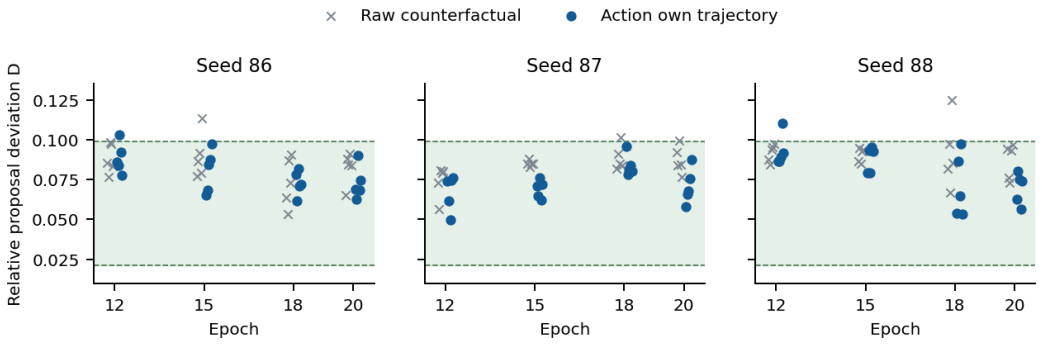

Figure 3: Frozen $\tau = .875$ evaluated at prescribed future P^{2}OT probes. Each point is one batch. The shaded band is $[.021, .099]$; horizontal offsets separate observations. Blue points measure the selected action on its own trajectory (primary); crosses measure counterfactual proposals along raw training (secondary). Per-seed action coverage is 19/20, 20/20, and 19/20. Repeated probes are not independent seed replications.

D.3 JUMBOT

Seeds 940–944 use the same source-training setup, then 100 raw prefix updates. The prefix grid is $\{0, .25, .5, .75, .875, 1\}$, with probes at offsets $0, 10, \ldots, 90$. Eligible actions have $\tau \leq .875$ and at least nine proposals in $[.021, .099]$. Each seed chooses the smallest eligible τ and freezes it; otherwise it retains raw. Unlike the earlier P^{2}OT experiment's shared choice, selection here is per seed. All five choose $\tau = 0$, each with nine passing prefix probes.

From each prefix endpoint, raw, fixed .875, and RAA branches continue for 400 updates with incoming optimizer state preserved and paired data streams. Twenty probes occur at offsets $0, 20, \ldots, 380$. Primary success requires a qualifying coefficient and at least 18 in-band probes during continuation. Target accuracy at the fixed final endpoint is secondary. The task was already known from development; the seeds and continuations are new.

Table 13: JUMBOT frozen-rule outcomes. Accuracies are percentages; coverage counts are per seed. All five seeds select the matched endpoint.

Seed	RAA coverage	Fixed .875 coverage	Raw ACC	RAA ACC	Fixed ACC
940	20/20	0/20	97.82	97.87	97.77
941	20/20	0/20	97.47	97.38	97.39
942	20/20	0/20	97.62	97.55	97.59
943	20/20	0/20	97.35	97.29	97.30
944	20/20	0/20	97.87	97.93	97.98
Mean			97.626	97.604	97.606

The fixed-$\tau = .875$ deviations lie below the lower band boundary. RAA's selections coincide with a fixed-zero policy and yield no observed accuracy gain. Five seeds are the evaluation units; the 100 probes form a descriptive total.

E LOCAL OBSERVATIONS AND CONTINUATION UTILITY

This supporting result clarifies why optimizer measurement alone is insufficient to infer continuation utility in the stated smooth class. It does not explain the external task gains or establish an impossibility result for the tested neural architectures. The recovery results identify additional structure that can make continuation utility identifiable.

The construction below preserves indistinguishability of complete local row losses while changing a specified continuation utility. It uses a standard smooth extension argument, complementing the diagnostic limitations studied for feature attribution (Bilodeau et al., 2024).

Observation model. Let $z = (\theta, a, \omega, \xi, t)$ contain the parameters, auxiliary state (such as teacher, memory bank and buffers), incoming optimizer state, batch/randomness, and schedule position. At a query state, a paired first-order row oracle returns the consumed targets Q, row losses ℓ_{ik} and their frozen-context parameter derivatives, together with the value and gradient of any other loss term and declared reduction/support metadata. All parameter blocks whose updates are to be explained are included. The optimizer implementation is known. Raw and tempered candidates share this context; the target is not re-solved under each coefficient.

Aggregate RAA statistics follow from the endpoint gradients; covariance verification additionally uses individual conditional forces. Combining all gradient terms with ω determines each candidate's next optimizer proposal. A finite audit is a transcript $\mathcal{T}_N = ((z_j, \mathcal{O}(z_j)))_{j=1}^N$; adaptive queries may depend on earlier replies. A quantity is identified by this transcript if it is constant over the declared model class compatible with the replies. The oracle excludes source-code inspection and uncounted training or outcome queries. Any training feedback used by a selector must enter the transcript.

The relevant continuation quantity is specified separately: for a known utility V and horizon T, let $J_T = V(z_T^1) - V(z_T^0)$. Recomputed-and-detached targets can produce nonconservative fields. The fixed-target subclass below suffices for the negative result, even with known V.

Lemma E.1 (joint realization). Given smooth functions $F, \Psi : \mathbb{R}^p \to \mathbb{R}$, choose $c > 0$ and $s > d > 0$. Two fixed active rows with masses $s + d, s - d$, common conditional target e_1, and losses

$$\ell_1 = \frac{F}{2cs} + \frac{\Psi}{2cd}, \qquad \ell_2 = \frac{F}{2cs} - \frac{\Psi}{2cd}$$

realize the full RMT family

$$L_\tau = c[(s + \tau d)\ell_1 + (s - \tau d)\ell_2] = F + \tau\Psi.$$

Thus both the matched objective and its amplitude correction are prescribed. The proof is direct expansion. If F and Ψ agree between two worlds on a region, their complete row functions agree there as well.

Positive row losses are not guaranteed for arbitrary F, Ψ, but the following construction supplies them globally. A strictly positive smooth loss ℓ_i also has an exact two-class cross-entropy representation: use probabilities $(e^{-\ell_i}, 1 - e^{-\ell_i})$ and target e_1. This is a realization in a general smooth model class. It neither asserts realizability by a fixed CNN nor that the target solves a specified OT optimization problem.

Proposition E.2 (local observations do not identify continuation utility). There exist two smooth fixed-target, positive two-row instances with the same matched objective F, initialization, GD step size, and known utility $V = -F$, for which every row function agrees on an open neighborhood of initialization, but J_2 has opposite signs. At initialization the matched gradient is nonzero and the amplitude correction is nonzero and orthogonal to it.

Proof. Write $x = (x_1, x_2)$ and take

$$F(x) = \tfrac{1}{2}\|x\|^2 + 10, \quad x_{\text{init}} = (1, 0), \quad \eta = \tfrac{1}{2}, \quad r = (\tfrac{1}{2}, -\tfrac{1}{2}).$$

Let χ be a smooth cutoff in $[0, 1]$, equal to one near r, supported in $B(r, 1/10)$. Set

$$q_+ = (\tfrac{1}{2}, -\tfrac{3}{2}), \quad q_- = (-\tfrac{3}{2}, -\tfrac{3}{2}), \qquad \Psi_\pm(x) = x_2 + \chi(x) q_\pm^\top (x - r).$$

On $O = B(x_{\text{init}}, 1/10)$ both perturbations equal x_2, so all row functions and all their derivatives agree there. Initially $\nabla F = e_1$ and $\nabla\Psi_\pm = e_2$. The matched two-step endpoint is $(1/4, 0)$. Both raw trajectories first reach r, where

$$\nabla(F + \Psi_+)(r) = (1, -1), \qquad \nabla(F + \Psi_-)(r) = (-1, -1).$$

Their second endpoints are respectively $(0, 0)$ and $(1, 0)$, giving

$$J_2^+ = 1/32, \qquad J_2^- = -15/32.$$

Use Lemma E.1 with $c = s = 1, d = 1/2$. Then $\ell_{1,2} = F/2 \pm \Psi_\pm$ are strictly positive, since $F/2 \pm x_2 \geq 4$ and the bump magnitude is at most $\|q_\pm\|/10 < 1$. Each $F + \tau\Psi_\pm$ is coercive for $\tau \in [0, 1]$. This completes the construction. □

All adaptive queries confined to O receive identical replies in the two worlds, even if the oracle returns the complete local function germ. With equal priors, a randomized rule forced to predict the sign errs with probability $1/2$; couple its random tape across worlds. Abstention requires a separate coverage/loss convention. This statement has no common prescribed smoothness constant and no restriction that the chosen step size satisfy a global descent bound. It is not an Adam-specific lower bound or a semantic-accuracy impossibility result.

Nor does it assert that any finite number of unrestricted queries is insufficient: with known V, directly simulating two finite trajectories resolves their finite-horizon comparison. The observation restriction and prediction target are essential.

Proposition E.3 (recovery under a declared affine structure). Suppose the full gradient fields are known a priori to have the form

$$G_0(x) = A_0 x + b_0, \qquad \Delta(x) = A_\Delta x + b_\Delta,$$

with no unknown auxiliary-state evolution. Under the same paired first-order oracle, $p + 1$ affinely independent state queries identify both fields. Hence a specified deterministic optimizer, initial

optimizer state, and schedule determine every finite continuation; its utility is recoverable if V is also supplied.

Proof. For queries $x^{(0)}, \dots, x^{(p)}$, let $X = [x^{(1)} - x^{(0)}, \dots, x^{(p)} - x^{(0)}]$. It is invertible. Multiplying the corresponding gradient-difference matrix by X^{-1} recovers A_0; the initial value recovers b_0. The same calculation recovers A_Δ, b_Δ. All subsequent updates can then be computed. This includes finite simulation of a specified Adam map, without asserting its convergence. □

The structural assumption is prior information, not something certified by fitting a finite sample. If it holds only on a region, recovery applies only while the trajectory stays in that region. This result is a sufficient recovery scheme, not a claim that $p + 1$ queries are necessary for a particular utility, and not a low-cost procedure when p is large.

Quadratic special case. For $F(x) = \frac{1}{2}x^\top Hx - b^\top x + c$, $H \succ 0$, and $\Delta = d$ constant,

$$x_*^\tau = H^{-1}(b - \tau d), \qquad F(x_*^\tau) - F(x_*^0) = \tfrac{1}{2}\tau^2 d^\top H^{-1} d \geq 0.$$

The displacement is the familiar inverse-Hessian sensitivity structure (Koh & Liang, 2017). It must not be confused with the change in the minimum of the *perturbed* objective. For GD with $B = I - \eta H$, write $B^{[T]}$ for its Tth power. Then

$$x_T^\tau = x_*^\tau + B^{[T]}(x_{\text{init}} - x_*^\tau), \qquad e_T = x_T^\tau - x_T^0 = -\tau H^{-1}(I - B^{[T]})d,$$

$$F(x_T^\tau) - F(x_T^0) = (x_T^0 - x_*^0)^\top H e_T + \tfrac{1}{2} e_T^\top H e_T.$$

The finite-horizon sign can differ from the nonnegative limiting loss change. The limit follows under $0 < \eta < 2/\lambda_{\max}(H)$. These calculations illustrate structural recovery; they are not assumptions imposed on the neural experiments.