1. Supplementary Material

1.1. Full protocol and state definitions

The online contract is batch size one, one adaptation attempt per sample, and post-action evaluation after the permitted parameter/state action. The stream order is fixed before execution. Online adaptation receives the image or frame and the text vocabulary, but no target mask, target class label, oracle shift label, video-boundary label, or future metric. A reset-style arm restores the declared initial condition at its matched boundary. A persistent arm carries the mutable adaptation state to the next sample.

The carried state is decomposed as follows:

The source anchor is the protected pre-stream copy of the optimizer-managed visual LayerNorm affine weights and biases. The clean Adam state is the construction-time state loaded into the existing optimizer object. During F2C Freeze, the inherited Controller continues to evolve and can change the parameter-group learning rate; its trusted-checkpoint rollback remains independently executable. That rollback is part of the canonical persistent baseline lineage and remains enabled in both the baseline and Freeze-to-Confirm runs. The method-specific operation is conditional source-active-state recovery after confirmation.

1.2. Per-seed matched results

Table S3 reports the direct values from the canonical three-seed registry. Values are foreground macro mIoU percentages and are rounded to four decimal places only for presentation.

The paired gains are $+30.37 \pm 0.58$ on VOC21, $+7.89 \pm 0.56$ on COCO, and $+15.79 \pm 0.47$ on YTVIS. All nine paired seed differences are positive. The standard deviations in these aggregate statements are sample standard deviations across the three paired differences.

1.3. Complete cross-method scope

The following tables expose the seed scope of the validated comparator rows. These rows are not pooled with the matched Persistent MLMP Base–Freeze-to-Confirm comparison. PCPM and SG-only/SG2 are historical or factorized comparator configurations; none is a component of Freeze-to-Confirm.

Class-sparse variants are reported separately because they change the protocol scope. COCO PCPM is 30.1062, 29.7677, and 31.4459 for seeds

42/43/44; COCO SG2 is 30.4073, 31.8504, and 32.8129. YTVIS PCPM is 40.1840, 39.4483, and 43.3803; YTVIS SG2 is 54.6716, 53.7299, and 52.1195. These values are not pooled into the fixed mixed-protocol results. Historical seed42-only rows in the registry are likewise kept outside the three-seed summaries.

1.4. Complete matched intervention matrix

The matched replay uses a fixed registry of 24 distinct selected events: eight events from each of VOC21, COCO, and YTVIS. For a job–event replay, Q_H is the mean foreground macro mIoU over the metric values in the first H order-matched replay positions; unevaluable entries remain NA and are excluded rather than replaced. H16 denotes Q_{16} and does not denote 16 events. Dataset-balanced quantities first average the eight events within each dataset and then average the three dataset means.

The audited replay design contains 36 logical policy–dataset jobs. Three dynamic memory jobs were held and excluded, leaving 33 runnable jobs. Each runnable job covers the eight registered events for its dataset, yielding 264 job–event replay records (33×8), each with 32 replay samples from which the reported horizons are computed. Thus, 264 is a record count across actions and datasets, not a count of distinct selected events.

The Temporary Freeze support +0.0440 is the dataset-balanced prevention contrast between Temporary-Freeze-then-Resume and Continue. The Adam-only value −0.0115 is the balanced action contrast

$$Q_{16}(\text{Reset Adam}) - Q_{16}(\text{Continue}).$$

The Controller-related support +0.0084 is the matched incremental contrast between the active-reset control and the corresponding confirmed-policy condition, in which the persistent branch changes from retained Controller state to active reset. The Active-state Restore support +7.0844 is the selected-regime intervention support assigned to restoring the active visual-LN parameters in the final evidence synthesis. It is an intervention-support quantity rather than the raw pooled Adam or Controller contrast.

Immediate Restore and F2C policy gains are dataset-balanced $Q_{16}(\text{policy}) - Q_{16}(\text{Continue})$ differences. A false-intervention regret is $Q_{16}(\text{Continue}) - Q_{16}(\text{policy})$ on an initially triggered event where that loss exceeds the fixed 0.1-point replay tolerance; the reported worst regret is the maximum over such events. ConfirmationValue is the dataset-balanced difference between

the confirmed F2C policy and the corresponding always-Restore-after-Freeze policy. EscalationValue is the dataset-balanced difference between the confirmed F2C policy and the Temporary-Freeze-then-Resume policy. All of these are measured in foreground-mIoU points, but the contrasts are not interchangeable.

The Active-state Restore support and the complete F2C $K = 4$ gain are different quantities. The first is intervention support for an active-state action; the second is complete policy gain relative to Continue. Their values match to four decimal places in this replay, but the labels do not denote one quantity.

The finite interval comparison is:

Leave-one-event-out analysis for the retained interval has a minimum gain of 4.3398 and median gain of 7.1029. These results support the evaluated confirmation trade-off, but they do not establish a universal optimum.

1.5. Method pseudocode and state semantics

The public controller can be expressed as the following state machine. The source anchor and construction-time Adam state are prepared before the stream.

> **Initialize:** protect θ_{src} and the construction-time Adam state o_{clean}; initialize the inherited Controller and set the F2C phase to ADAPT.
> **Pre-sample boundary:** before x_t, when m_{t-1} is available, compute $r_t^{\text{pre}} = (-m_{t-1} - \mu_{\text{C1}})/\sigma_{\text{C1}}$. If the phase is ADAPT and $r_t^{\text{pre}} > 0$, open an Alarm before x_t and freeze x_t through x_{t+3}.
> **Process a sample:** run the inherited Controller and ordinary MLMP forward/loss/backward path. In F2C Freeze, suppress the normal Adam optimizer step but retain gradient clearing, prediction/evaluation, Controller evolution, EMA observation, and trusted-checkpoint rollback machinery. After sample x_j, the Controller observes ℓ_j and updates m_j.
> **Confirm:** before x_{t+4}, compute $r_{t+4}^{\text{conf}} = (-m_{t+3} - \mu_{\text{C1}})/\sigma_{\text{C1}}$. The decision uses only this boundary score; intermediate scores are not averaged and need not remain continuously positive.
> **Resume:** if $r_{t+4}^{\text{conf}} \leq 0$, retain the active state, clear the F2C Freeze gate, and process x_{t+4} normally.
> **Restore:** if $r_{t+4}^{\text{conf}} > 0$, set $\theta \leftarrow \theta_{\text{src}}$, load o_{clean} into the existing Adam optimizer, retain C and H, clear the F2C Freeze gate, and process x_{t+4} normally.

The inherited Controller updates the EMA as $m_j = \ell_j$ for its first available observation and $m_j = (1-\alpha)m_{j-1} + \alpha\ell_j$ otherwise, with $\alpha = 0.1$ and no bias correction. F2C Restore does not reset the Controller or discard EMA/history. The inherited trusted-checkpoint rollback also remains active and is distinct: it restores a trusted adapted visual-LN/Adam checkpoint, whereas F2C Restore targets the pre-stream source visual-LN anchor and construction-time Adam state. The latter is an active-state rather than full-model reset and does not require a second full-model copy.

1.6. Efficiency details

Timing uses a fixed COCO 1100-order prefix, three repeats per method, the same A30 GPU, and identical host execution settings. Data loading and one-time initialization are excluded. Host timing does not explicitly synchronize CUDA. All 24 strict repeat checks passed. Static state values are structural accounting; sampled GPU used memory is not an allocator peak.

The reported near-base compute statement is limited to this controlled benchmark. Allocator peak and resident-set absolute values were unavailable from the measurement closure, so small latency differences are not treated as exact algorithmic overhead.

1.7. Backbone variation

The ViT-B/16 variation is reported separately from the main ViT-L/14 tables as a scope check and is not pooled with the main comparison. This ViT-B/16 table does not evaluate Freeze-to-Confirm and therefore does not constitute a backbone-robustness test of F2C.

The ViT-B/16 table does not alter the main method identity or pool with the ViT-L/14 matched Base/F2C gains.

1.8. Telemetry and qualitative provenance

Across the nine full-stream Freeze-to-Confirm runs, the controller telemetry contains 143 alarms, 74 Resume outcomes, and 69 Restore outcomes. These counts are behavior evidence for the executed controller and are not precision or recall against a ground-truth harmful-shift detector.

The qualitative examples use four preselected VOC21 cases, in this order: `VOC21_B16`/order 05986, `VOC21_A14`/order 10636, `VOC21_A13`/order 04889, and `VOC21_D20`/order 12085. For each case, the input, Persistent MLMP Base prediction, Freeze-to-Confirm prediction, and GT are taken from the corresponding fixed exported case assets. Raw prediction masks are unchanged.

GT is used only offline for the semantic overlays and for the Error change panel: green marks Base-incorrect/Freeze-to-Confirm-correct regions, red marks Base-correct/Freeze-to-Confirm-incorrect regions, gray marks regions incorrect for both, and no overlay marks regions correct for both. The four examples are illustrative and do not provide aggregate evidence; the quantitative conclusions do not depend on them.

1.9. Reproducibility

Protocol identities, matched-result scopes, seed assignments, and replay definitions were fixed before manuscript reporting and were checked against the corresponding experiment and analysis records. The main and supplementary tables preserve these fixed result identities without pooling incompatible protocol variants or seed scopes. Implementation, protocol definitions, and reproducibility materials are planned for public release with the paper according to the stated data and code availability policy.

1.10. T0 and K derivation note

Let ℓ_j be the ordinary MLMP adaptation objective computed before the optional optimizer step. The inherited Controller sets $m_j = \ell_j$ for the first available observation and otherwise updates $m_j = (1-\alpha)m_{j-1} + \alpha\ell_j$, with $\alpha = 0.1$ and no bias correction. Before x_t, F2C evaluates

$$\begin{aligned} r_t^{\mathrm{pre}} &= \frac{-m_{t-1} - \mu_{\mathrm{C1}}}{\sigma_{\mathrm{C1}}}, \\ \mu_{\mathrm{C1}} &= -1.2564001647222331, \\ \sigma_{\mathrm{C1}} &= 0.8664717902121067. \end{aligned} \tag{1}$$

The mean and population standard deviation are fixed C1 statistics, are not updated online, and use no normalization epsilon. The negative sign is applied before normalization. $T_0 = 0$ is the semantic zero crossing of this score, and an Alarm opens only under the strict inequality $r_t^{\mathrm{pre}} > 0$. No ground-truth label, future sample, future metric, or separately tuned persistence threshold is used. The evidence supports describing this as a fixed controller boundary, not as an optimum threshold or a result of a dense threshold sweep.

The interval length was retained from the audited finite candidate set $K \in \{1, 4, 8\}$. In the deployed state machine, $K = 4$ means that an Alarm before x_t freezes exactly x_t through x_{t+3} and confirmation evaluates the current fixed-C1 score from m_{t+3} before x_{t+4}. It does not require the intermediate

scores to remain positive. F2C has no additional threshold, patience, averaging rule, or cooldown; inherited Controller patience and rollback settings are separate baseline behavior. $K = 4$ has confirmation value +2.0881, escalation value +7.0871, and complete-policy gain +7.0844; $K = 8$ provides no audited safety improvement, while $K = 1$ changes the decision pattern. The leave-one-event-out audit has minimum gain 4.3398 and median gain 7.1029. Thus the supported statement is that $K = 4$ was selected from the audited confirmation trade-off and frozen for evaluation. The evidence does not support calling it a universal or exhaustive sensitivity optimum.

Table S1: Evaluated protocol identities. Persistent and reset-gap rows are separate result identities.

Dataset	Persistent protocol	Backbone	Metric	Stream semantics
VOC21	VOC21 mixed persistent	ViT-L/14	foreground macro mIoU	mixed stream, 16-category cycle
COCO-FG80	COCO mixed persistent	ViT-L/14	foreground macro mIoU	mixed foreground stream
YTVIS-FG40	YTVIS natural-video persistent	ViT-L/14	natural-video stream; ineligible frames do not enter the denominator	
COCO reset-gap	COCO:batch-reset vs persistent	ViT-L/14	foreground macro mIoU	separate MLMP reset-gap, TENT, and CLIPArTT comparison
YTVIS reset-gap	YTVIS:batch-reset vs persistent	ViT-L/14	foreground macro mIoU	separate reset-gap comparison with natural video order

Table S2: State semantics used by the persistent controller.

State component	Role in the online execution	Normal, Freeze, and Restore behavior
Active visual-LN θ	Adaptable visual affine parameters	Update when permitted in Normal; no normal optimizer update in F2C Freeze; source-anchor Restore
Adam state	Optimizer moments compatible with θ	Normal step when permitted; no normal step/moment update in F2C Freeze; construction-state Restore
Controller C	Risk and action state, including inherited rollback behavior	Continue evolving in Normal and Freeze; keep on Restore
EMA/history H	Online risk trajectory and historical context	Continue evolving in Normal and Freeze; keep on Restore

Table S3: Per-seed matched persistent results for Persistent MLMP Base and Freeze-to-Confirm.

Dataset	Seed	Persistent MLMP Base	Freeze-to-Confirm
VOC21	42	43.0476	72.7577
VOC21	43	41.7309	72.5195
VOC21	44	43.1769	73.7808
COCO	42	24.2132	31.6055
COCO	43	25.1263	32.9239
COCO	44	25.0447	33.5365
YTVIS	42	39.7805	55.4270
YTVIS	43	37.0582	53.3729
YTVIS	44	37.0081	52.4126

Table S4: Validated persistent comparator rows, foreground macro mIoU by seed. All rows use ViT-L/14 and the persistent protocol named in the first column. PCPM, SG-only, and SG2 are historical/factorized comparator configurations; none is a component of Freeze-to-Confirm.

Dataset / protocol	Method	Seed 42	Seed 43	Seed 44
VOC21:mixed:persistent	PCPM†	72.6656	73.8715	71.2973
VOC21:mixed:persistent	SG-only†	70.3590	69.9992	72.0137
VOC21:mixed:persistent	SG2†	70.3281	69.9645	71.9819
COCO:mixed:persistent	PCPM†	26.2522	30.1179	30.2292
COCO:mixed:persistent	SG-only†	30.5324	31.8789	32.7368
COCO:mixed:persistent	SG2†	30.4973	31.8463	32.7143
COCO:mixed:persistent	SAR	31.8768	33.0642	33.8707
COCO:mixed:persistent	CoTTA	31.8982	31.9688	32.6862
YTVIS:natural:persistent	PCPM†	49.4096	48.4272	47.7017
YTVIS:natural:persistent	SG-only†	54.7078	53.6700	52.6231
YTVIS:natural:persistent	SG2†	54.7006	53.6629	52.2321
YTVIS:natural:persistent	SAR	54.7279	52.3732	53.8567
YTVIS:natural:persistent	CoTTA	30.2591	26.7705	26.1390

† Historical/factorized comparator configuration. PCPM's VOC21 accumulator and the YTVIS timing and logging qualifiers remain attached to the fixed reported identity.

Table S5: Separate reset-gap scope. Each cell gives reset and persistent foreground macro mIoU for seeds 42/43/44; the two protocols are not merged with the persistent rows in Table S4.

Dataset	Method	Seed 42	Seed 43	Seed 44
COCO	MLMP reset-gap	30.4405 / 30.1516	31.7807 / 31.1423	32.6488 / 32.2490
COCO	TENT	30.5666 / 28.0699	31.9854 / 30.0089	32.3638 / 30.4523
COCO	CLIPArTT	29.4020 / 26.1546	31.3300 / 28.0202	31.4734 / 28.3326
YTVIS	MLMP reset-gap	51.5461 / 46.1306	50.4894 / 45.2581	50.2560 / 44.6841
YTVIS	TENT	52.3717 / 53.4919	51.7683 / 53.2011	50.3059 / 51.2303
YTVIS	CLIPArTT	50.5260 / 50.3128	49.0466 / 48.2368	48.5083 / 47.9958

Each entry is reset / persistent. This table is the reset-gap configuration and is not the Persistent MLMP Base identity.

Table S6: Matched intervention and recovery-timing matrix. Every numeric evidence entry is a difference of H16 foreground-mIoU utilities over the balanced 24-event scope. The labels denote different intervention, policy, confirmation, escalation, or regret contrasts and therefore should not be treated as the same estimand. Values are reported at the precision needed to distinguish their definitions.

Intervention	Action definition	Balanced H16 evidence	Interpretation
Continue	No intervention	Reference	Baseline replay policy
Temporary Freeze	Normal Adam step suppressed; execution/observation continue	Support +0.0440	Weak prevention at the observed boundary
Adam-only intervention	Modify Adam state without active-state restoration	Mean action replay −0.0115	Adam alone is near null in the pooled action replay
Controller-related control	Matched Controller-related control	Support +0.0084	Weak incremental Controller-related effect in this matched control
Active-state Restore	Restore active visual-LN θ to source anchor	Intervention support +7.0844	Recoverable active-state degradation in selected regimes
Immediate Restore	Restore at the initial alarm	Gain +4.4967; worst regret 77.9774	Helpful in some regimes but vulnerable to early false intervention
F2C, $K = 1$	Freeze one sample, then confirm or restore	Gain +4.2136; confirmation 0.0000; regret 77.9861	Short interval changes the action pattern
F2C, $K = 4$	Freeze four samples, then same-signal confirmation	Gain +7.0844; confirmation +2.0881; escalation +7.0871; regret 10.2810	Fixed deployed timing choice
F2C, $K = 8$	Freeze eight samples, then same-signal confirmation	Gain +4.6674; confirmation +1.9234; escalation +4.6645; regret 10.2823	Longer interval adds no audited safety improvement
F2C active-reset control	F2C with active reset control	Gain +7.0928; regret 10.2759	Secondary control, not the deployed identity
Archive preservation	Preserve archive/history state	Support +0.0929	Historical-state contribution in the matched replay

Table S7: Audited interval comparison used to retain $K = 4$. All columns are dataset-balanced H16 foreground-mIoU-point contrasts. The table is a finite candidate comparison, not a dense hyperparameter sweep.

K	Prevention	Delay	Confirmation	Escalation
1	+0.0440	−0.2831	0.0000	+4.1697
4	−0.0027	+0.4995	+2.0881	+7.0871
8	+0.0029	−1.7528	+1.9234	+4.6645

Table S8: Controlled timing and static state details. Timing values are mean milliseconds per order with repeat standard deviation. PCPM, SG-only, and SG2 are retained here as historical/factorized comparator configurations for state-overhead context; none is a component of Freeze-to-Confirm.

Method	ms/order	static state (MiB)	full-model copies
Persistent MLMP Base	158.28 ± 10.10	1.17	1
Freeze-to-Confirm	158.94 ± 11.95	1.56	1
PCPM	266.51 ± 3.72	37.52	1
SG-only	154.89 ± 4.31	1.17	1
SG2	193.59 ± 6.36	37.52	1
SAR	216.57 ± 6.28	1.17	1
CoTTA	203.96 ± 6.96	3264.38	3

Table S9: ViT-B/16 variation. Values are mean ± sample standard deviation over seeds 42/43/44; the exact seed vectors are shown in the last column.

Dataset	Method	Mean	Std.	Seed vector (42, 43, 44)
COCO	NoAdapt	25.35	0.33	24.9829, 25.4748, 25.6024
COCO	Persistent MLMP Base	19.09	1.47	18.6238, 17.9185, 20.7411
COCO	PCPM	25.95	0.20	25.8430, 26.1887, 25.8288
COCO	SG2	25.67	0.30	25.3591, 25.6915, 25.9509
YTVIS	NoAdapt	48.33	0.83	48.9674, 48.6385, 47.3905
YTVIS	Persistent MLMP Base	38.06	2.31	40.6191, 37.4121, 36.1448
YTVIS	PCPM	50.60	1.06	51.8149, 49.8646, 50.1174
YTVIS	SG2	52.79	1.13	53.1356, 53.7061, 51.5312