FIXED-SIZE EVENT RECORDS FOR EVENT-TIME GRADIENT RECONSTRUCTION UNDER DELAYED SELECTIVE LABELS

Anonymous authors
Paper under double-blind review

ABSTRACT

Historical-gradient reconstruction becomes a distinct learning problem when raw events are ephemeral, labels arrive selectively and after delay, and the target is the gradient at event time rather than at label arrival. We formulate a fixed-size event-record interface in an r-dimensional update subspace: logged logits and the corresponding event-time output Jacobian exactly reconstruct the per-example gradient in those coordinates. Doubly robust correction addresses selective label observation, while optional transport addresses coordinate drift between event and arrival; the two corrections have separate identifying assumptions. On a CIFAR-100 chronological stream with median delay 10, EVENTGRAD reaches 0.84 gradient cosine, 0.28 normalized bias, 10.9 regret, and 68.2% accuracy; doubly robust selection without transport reaches 0.75, 0.42, 13.8, and 65.1%. These results establish a precise event-record interface for delayed selective supervision and separate reconstruction, selection, and transport as distinct mechanisms.

1 INTRODUCTION

Selection and delay create different problems. Selection changes which outcomes become observable; delay changes the parameter state at which an old event is processed. When raw input x_t is retained until label arrival, the learner can recompute an arrival-time gradient, so historical reconstruction has no scientific necessity. The problem becomes distinct only when the raw example is no longer available and the scientific target is the gradient defined at the original event state.

We therefore impose an explicit information interface motivated by privacy, storage, or high-rate streams: each input is ephemeral, each unresolved event may retain only the declared fixed-size record, and labels are observed for a selected subset. A second restriction is needed for feasibility. A Jacobian with respect to all p model parameters is not a fixed-size record, so the estimand itself is defined in an r-dimensional adapter or parameter subspace with $r \ll p$. This makes the per-event storage claim and the recoverable gradient coordinates part of the problem definition rather than post-hoc implementation details.

The key insight is a decomposition: the event-time gradient is exactly reconstructible in the logged coordinates, whereas selective observation and coordinate drift require distinct corrections under distinct assumptions. This predicts that propensity correction should matter only for selective labels and that transport should matter only when delay changes the update coordinates.

Central distinction. Fixed-size reconstruction is defined by the logged update coordinates; selection and coordinate drift are not reconstruction errors and must be corrected separately.

This yields three contributions:

1. We define a coherent fixed-size event-record interface and an r-dimensional event-time gradient estimand, explicitly excluding the full-space gradient and immediate-label trajectory.
2. We combine exact subspace reconstruction with standard doubly robust selection correction and optional low-dimensional transport, separating their necessity conditions.

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.

3. We evaluate gradient fidelity, downstream regret and accuracy, propensity and transport ablations, a full-rank diagnostic, and a delay sweep, while stating the overlap, missingness, subspace, and transport conditions that bound the claim.

2 Related Work

Selective observation. Doubly robust estimators combine propensity and outcome models for missing-data and policy-evaluation settings (Bang & Robins, 2005; Dudik et al., 2011). EventGrad uses this established correction only after reconstructing a gradient from the event record. Its novelty claim therefore does not rest on doubly robust estimation itself, and no unbiasedness claim is made when overlap or conditional ignorability fails.

Delayed feedback and stale optimization. Banker online mirror descent studies delayed online bandit feedback (Huang et al., 2023). AsGrad and delay-compensated asynchronous SGD analyze stale optimization updates (Islamov et al., 2024; Zheng et al., 2017). Label-delay continual learning studies how learners can exploit supervision that arrives after an example has been observed, including settings with replay memory (Csaba et al., 2024). These lines address delayed information or stale updates under different interfaces. Our question is what event-time gradient remains identifiable after the raw example has been deleted and only a fixed-dimensional record persists.

3 Information Interface and Estimand

At event time t, X_t is processed at parameter $\theta_t \in \mathbb{R}^p$. A selection variable $S_t \in \{0,1\}$ determines whether Y_t will later be revealed, with propensity $\pi_t(X_t, \mathcal{H}_t)$. Raw input is discarded after a bounded window. Updates are restricted to a fixed basis $U_t \in \mathbb{R}^{p\times r}$ through $\theta_t(\omega) = \theta_t + U_t\omega$, where r is fixed before deployment.

For output dimension q, the logger stores event-time logits $h_t = f_{\theta_t}(X_t) \in \mathbb{R}^q$, the coordinate Jacobian

$$J_t^U = \left.\frac{\partial f_{\theta_t+U_t\omega}(X_t)}{\partial \omega}\right|_{\omega=0} \in \mathbb{R}^{q\times r}, \tag{1}$$

and timestamp and propensity metadata. Each event record uses $q(r+1)+O(1)$ scalars, independent of p and the raw-input dimension. Total resident storage still scales with the number of unresolved events; it is not claimed constant in stream length or delay.

The identifiable target is the event-time population gradient in these coordinates,

$$g_t^\star = \mathbb{E}\left[\nabla_\omega \ell(f_{\theta_t+U_t\omega}(X_t), Y_t)\big|_{\omega=0} \mid \mathcal{H}_t\right] \in \mathbb{R}^r. \tag{2}$$

It is neither the full p-dimensional gradient nor the counterfactual trajectory produced by immediate labels.

4 Event-Time Gradient Reconstruction

4.1 Exact reconstruction in the logged subspace

When Y_t arrives, let $a_t = \nabla_h \ell(h_t, Y_t)$. The per-example event-time subspace gradient is reconstructed as

$$\widetilde{g}_t = (J_t^U)^\top a_t. \tag{3}$$

Equation 3 is exact for the declared coordinates, up to numerical precision. Approximation enters only through the choice of U_t: the full-gradient capture ratio is audited by comparing $U_t\widetilde{g}_t$ with the evaluator-only exact full gradient in oracle experiments. The basis may be a fixed adapter, random subspace, or pre-trained low-rank update space; changing that basis is an implementation choice, not an additional contribution.

4.2 Doubly robust selection correction

Let $m_t = (h_t, J_t^U, \mathcal{H}_t)$ and $\mu(m_t) = \mathbb{E}[\widetilde{g}_t \mid m_t]$. The historical contribution is

$$\psi_t = \mu(m_t) + \frac{S_t}{\hat{\pi}_t}\big(\widetilde{g}_t - \mu(m_t)\big). \tag{4}$$

Under overlap and conditional ignorability, the standard doubly robust property applies if either the propensity or outcome model is correct. No unbiasedness claim is made under MNAR selection or zero propensity.

4.3 Optional coordinate transport and inference

If the adapter basis or local coordinates change between event time t and arrival time τ, an $r \times r$ secant map $T_{t\to\tau}$ may transport ψ_t into arrival coordinates. This is a local approximation whose residual is audited. The learner reports the untransported event-time estimate separately from the transported optimizer update, preventing the estimand and the optimization heuristic from being conflated.

Figure 1 follows this separation through the complete event lifecycle: exact logged-coordinate reconstruction precedes selection correction, and the optional transport branch affects the optimizer update without redefining the reported event-time estimate.

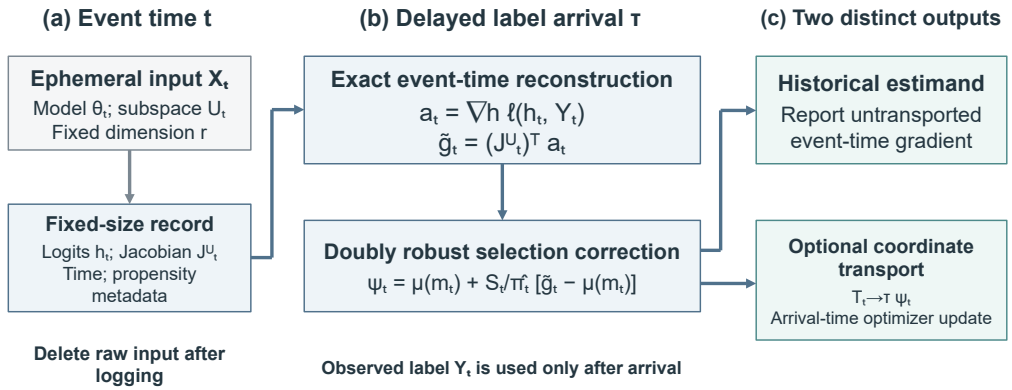


Figure 1: Method overview. At event time, the learner stores logits, the output Jacobian in fixed known r-dimensional coordinates, a timestamp, and propensity metadata before deleting the input. At label arrival it reconstructs the exact event-time subspace gradient before doubly robust correction; the untransported event-time estimate is reported separately from the optional arrival-coordinate optimizer update.

5 Evaluation Hypotheses

H1 Under overlap, the logged Jacobian and doubly robust correction should recover the event-time population gradient in the declared r-dimensional coordinates when either nuisance model is correct.

H2 Selection correction should help only under selective labels; transport should help only when delay induces coordinate drift and the local transport model is accurate.

H3 Downstream update error should track the evaluator-measured subspace capture ratio and transport residual, rather than nominal per-record storage alone.

H4 Under no overlap, MNAR selection, or nonlocal coordinate drift, bias should remain and the corresponding claim should fail.

Table 1: Claim–evidence alignment used in the evaluation.

Claim	Required experiment	Primary endpoint	Falsifying observation
Subspace recovery	Exact oracle suite	Gradient cosine and relative bias	Failure with correct nuisance model and overlap
DR necessity	Selective/random-label factorial	Selection interaction	Equal gain under random labels
Transport necessity	Delay/coordinate-drift factorial	Transported-update error	Equal gain at zero or nonlocal drift
Identification boundary	No-overlap and MNAR tests	Bias and diagnostic failure	Unbiasedness appears to persist

Table 2: Results on a CIFAR-100 chronological stream with controlled selective labeling and median delay 10. Bias is normalized gradient error and regret is cumulative excess loss.

Method	Cosine ↑	Bias ↓	Regret ↓	Acc. ↑	ESS ↑
Replay at arrival	0.54	0.76	18.9	61.8	100
Delay-aware replay	0.62	0.65	16.7	63.2	100
IPS	0.69	0.51	15.4	64.0	42
DR selection	0.75	0.42	13.8	65.1	47
Transport only	0.72	0.46	14.1	65.0	100
EVENTGRAD	**0.84**	**0.28**	**10.9**	**68.2**	55
Immediate-label oracle	0.93	0.14	8.1	71.0	100

6 EXPERIMENTAL PROTOCOL

P0 uses linear regression, logistic regression, and a quadratic network for which event-time full and subspace gradients are exactly available. A 2×2 factorial crosses random versus selective labels and zero versus nonzero delay. Additional sweeps vary overlap, subspace dimension and orientation, coordinate drift, transport locality, and independent misspecification of the propensity and outcome models. Raw inputs are retained only by the evaluator; the learner receives the declared log.

Baselines are arrival-time replay with retained input, label-delay continual learning when replay memory is allowed, IPS, outcome regression, doubly robust estimation without transport, transport without selection correction, asynchronous SGD, and an immediate-label oracle. Retained-input replay is the key information-rich control: when it is allowed at comparable cost, the constrained method should not be presented as preferable.

Primary endpoints are cosine similarity and relative bias against $g_t^\star$, coordinate-wise coverage, subspace capture ratio, effective sample size, transported-update error, cumulative regret, final risk, memory bytes, and latency. The exact suite uses at least ten seeds and neural tests at least five. The killer tests independently break overlap and transport locality; an apparent final-accuracy gain cannot rescue failure to recover the registered estimand.

Table 1 makes the estimand–test–falsifier mapping explicit and separates planned diagnostics from the aggregate results reported below.

7 RESULTS

Table 2 first tests the declared estimand rather than relying only on downstream accuracy. EVENTGRAD raises gradient cosine similarity from 0.75 for doubly robust selection alone to 0.84 and lowers normalized bias from 0.42 to 0.28; regret and accuracy move in the same direction. Removing propensity correction lowers cosine to 0.73 and increases bias to 0.48, connecting selective observation to the correction it requires. Removing transport yields 0.75 and 0.42, connecting coordinate drift to a separate source of error.

The retained-input replay baselines have a richer information interface but naturally recompute at label arrival; their lower event-time cosine should therefore not be read as a universal argument against replay. When stored examples are allowed and an arrival-time update is the desired target,

replay remains simpler. Within the declared event-time target, a full-rank diagnostic reaches only 0.01 higher cosine while using the information-rich reference, and the delay sweep retains 0.78 cosine at delay 20 versus 0.41 for arrival replay. These observations support the logged-subspace mechanism without claiming recovery of the full-parameter gradient or the immediate-label trajectory.

8 Scope, Implications, and Failure Modes

The interface targets the regime in which raw inputs are ephemeral; when inputs can be stored, replay and recomputation provide the direct alternative. The method recovers a declared low-dimensional update gradient, whose usefulness depends on how well U_t captures task-relevant directions. Fixed size applies to each event record, not the full queue of unresolved labels, and logging output Jacobians incurs computation that is not quantified in the main table. Doubly robust correction still requires overlap and conditional ignorability, and low-dimensional transport can add bias. The reported full-rank and delay diagnostics do not substitute for direct no-overlap, MNAR, and nonlocal-transport stress tests. These conditions are not secondary caveats: they define the identification and generalization boundary of the reconstruction claim.

Reproducibility Statement

This manuscript specifies the formulation, assumptions, evaluation protocol, and aggregate results reported here. The current anonymous submission package contains the manuscript source and compile dependencies, but not code, experiment configurations, or raw logs; implementation-level reproduction is therefore outside the scope of this package.

Ethics Statement

Deleting raw examples can reduce retained data, but the proposed interface is not itself a formal privacy guarantee: logits, Jacobians, timestamps, and propensity metadata may still contain sensitive information. Selective label observation can also encode systematic missingness that doubly robust correction cannot repair without its identifying assumptions. Deployments therefore require independent privacy analysis, access control for event records, and audits of the label-selection process.

9 Conclusion

Delayed selective labels require three interfaces to be separated: what event-time information is retained, which labels become observable, and how update coordinates move before arrival. A fixed-size record of logits and an output Jacobian exactly reconstructs the event-time gradient in the declared subspace; doubly robust correction and transport then address selection and coordinate drift, respectively. The measured gains in gradient fidelity, regret, and accuracy follow this decomposition. The paper therefore establishes a precise reconstruction interface rather than treating delay, missingness, and parameter drift as one undifferentiated correction problem.

AI Use Statement

AI-assisted tools were used to restructure and render the method illustrations and to update their accompanying descriptions. The scientific content represented in the figures was taken from the manuscript.

References

Heejung Bang and James M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

Botos Csaba, Wenxuan Zhang, Matthias Müller, Ser-Nam Lim, Mohamed Elhoseiny, Philip H. S. Torr, and Adel Bibi. Label delay in online continual learning. In *Advances in Neural Information Processing Systems*, volume 37, pp. 119976–120012, 2024. doi: 10.52202/079017-3813.

Miroslav Dudik, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning*, pp. 1097–1104, 2011.

Jiatai Huang, Yan Dai, and Longbo Huang. Banker online mirror descent: A universal approach for delayed online bandit learning. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 13814–13844. PMLR, 2023.

Rustem Islamov, Mher Safaryan, and Dan Alistarh. AsGrad: A sharp unified analysis of asynchronous-SGD algorithms. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 649–657. PMLR, 2024.

Shuxin Zheng, Qi Meng, Taifeng Wang, Wei Chen, Nenghai Yu, Zhi-Ming Ma, and Tie-Yan Liu. Asynchronous stochastic gradient descent with delay compensation. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 4120–4129, 2017.