

# Jacobian-Velocity Bounds for Deployment Risk Under Covariate Drift

Jonathan R. Landers  
Independent Researcher  
New York, NY, USA  
jonathan.robert.landiers@gmail.com

**Abstract**—We study long-horizon deployment of a frozen predictor under dynamic covariate shift. A time-domain Poincaré inequality reduces temporal risk volatility to derivative energy, and a Jacobian-velocity theorem identifies directional tangent energy along the deployment path as the governing quantity under explicit along-path regularity and domination assumptions. Under low-rank drift, that quantity reduces to directional Jacobian energy in the drift subspace, motivating drift-aligned tangent regularization (DTR) and a matched monitoring proxy. Rather than smoothing the network isotropically, DTR penalizes sensitivity only along estimated drift directions. We validate the theorem-to-method pipeline in four experiments: a synthetic benchmark for the time-domain inequality, a controlled synthetic comparison against isotropic Jacobian regularization, and two frozen-deployment studies on the UCI Air Quality and Tetouan power-consumption datasets. DTR reduces risk volatility and directional gain in the controlled low-rank regime, beats isotropic smoothing there, and gives validation-selected deployment gains on both real datasets when the Air Quality drift subspace is estimated from target-orthogonal sensor motion. Moderate drift-subspace misspecification is tolerable while orthogonal misspecification largely removes the benefit.

**Index Terms**—covariate drift, deployment risk, Jacobian regularization, temporal robustness, distribution shift

## I. INTRODUCTION

Distribution shift is often discussed statically, either through source-target generalization gaps in domain adaptation [1] or through dataset-shift taxonomies that separate covariate, label, and concept shift [2]. Long-horizon deployment is different. A model may be frozen while the environment moves continuously, so the relevant object is not a single out-of-domain error but a risk trajectory. Temporal benchmarks and deployment studies make the same point from different angles. Even when labels and architecture are fixed, performance can move materially over time [3]–[6]. The practical question is therefore not only whether shift occurs, but how violently performance can move once deployment begins.

This paper develops a geometric answer. If the feature stream follows a path  $X_t$  and the predictor has local tangent map  $J_f(X_t)$ , then instability is governed not by shift magnitude alone but by the directional interaction between the data velocity  $\dot{X}_t$  and the model tangent geometry. The dangerous quantity is  $J_f(X_t)\dot{X}_t$ . The intuition is simple. The same deployment path can be almost harmless for one model and damaging for another, because only motion through steep tangent directions accumulates into large score variation. If the stream moves

through a flat direction, drift can be present without being operationally dangerous. If it moves through a steep direction, even gentle but persistent motion can compound into long-horizon degradation. The paper’s central bound formalizes this. Under explicit along-path regularity and directional domination assumptions,

$$\text{Var}_U(r(U)) \leq \frac{\beta^2 T}{\pi^2} \int_0^T \mathbb{E} \|J_f(X_t)\dot{X}_t\|^2 dt.$$

First, a time-domain Poincaré inequality converts temporal volatility into derivative energy. Second, a Jacobian-velocity theorem bounds that energy under explicit along-path assumptions on the deployment performance field. Third, a low-rank drift specialization naturally motivates a drift-aligned regularizer and a matched monitoring proxy. Fourth, we test the resulting story with a synthetic time-domain sanity check, a compact isotropic-versus-directional comparison, and two real frozen-deployment studies on Air Quality and Tetouan. The empirical sequence mirrors the logic of the paper: verify the temporal inequality, compare competing regularizers in the idealized low-rank regime, then move to real deployments where the drift subspace must be estimated. Concretely, the paper identifies directional tangent energy as the quantity governing frozen-model deployment volatility, uses the resulting low-rank reduction to motivate drift-aligned tangent regularization (DTR) and a matched hazard proxy, and gives empirical evidence that the same directional geometry remains visible in both controlled synthetic experiments and two real field deployments.

The paper is deliberately scoped to frozen-model deployment under covariate drift. It does not address general distribution shift or test-time adaptation, but asks specifically what local geometry governs frozen-model degradation and whether that geometry remains visible on real drifting data. Code, scripts, and supplementary materials are available in a public repository [7]. A companion project page with figures and summaries is available at [8].

## II. RELATED WORK

The paper sits closest to temporal distribution shift and dynamic deployment evaluation. Classical shift analyses study source-target error under changed input distributions [1], dataset-shift taxonomies separate covariate, label, and concept shift [2], and concept-drift work emphasizes sequential adaptation under evolving streams [3]. More recent benchmark-driven

work has made temporal structure explicit, both in naturally shifted datasets and in stylized train-deploy protocols [4]–[6]. Related theory also studies heterogeneous distribution shift without an explicit temporal deployment path [9]. Here the focus is the volatility of one frozen model along a deployment path rather than a worst-case shift gap.

The paper is also adjacent to test-time adaptation [10]–[12]. Those methods update the model online to track changing data. Here the predictor remains frozen. That restriction is deliberate. It isolates the geometric mechanism of degradation itself and yields a quantity that can be analyzed, regularized, and monitored without mixing in the dynamics of online parameter updates.

Finally, the work draws on two technical lines. The regularization side is related to Jacobian-based control of robustness [13], [14] and to empirical evidence linking sensitivity to generalization [15], but the object here is directional rather than isotropic. Only tangent energy in likely drift directions enters the bound. The monitoring side is related to shift detection, uncertainty under shift, and sequential change detection [16]–[22]. Our hazard score is not presented as a sequentially optimal detector. It is a model-aware proxy matched to the same drift geometry used in training.

### III. SETUP AND ASSUMPTIONS

Let  $t \in [0, T]$  index deployment time. Let  $X_t \in \mathbb{R}^d$  denote the covariate process, and let  $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$  be a frozen scalar score with row Jacobian  $J_f(x) := \nabla f_\theta(x)^\top$ . We study a scalar deployment performance field  $g_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$  and the associated trajectory

$$r(t) := \mathbb{E}[g_\theta(X_t)]. \quad (1)$$

For classification, a natural example is

$$g_\theta(x) = \mathbb{E}_{Y \sim P_0(\cdot|x)}[\ell(f_\theta(x), Y)],$$

but the theorem below does not rely on this representation alone. Instead, we state the regularity needed to make the risk-trajectory argument defensible. The point of this setup is to separate two roles that are often conflated in shift analyses. The model is frozen, while the covariate stream is allowed to move.

To measure temporal instability, draw  $U \sim \text{Unif}[0, T]$  and define

$$\text{Var}_U(r(U)) = \mathbb{E}_U[(r(U) - \mathbb{E}_U r(U))^2]. \quad (2)$$

The quantity in (2) is the notion of deployment volatility studied throughout the paper.

*Assumption A1 (trajectory regularity).* The sample paths  $t \mapsto X_t$  are almost surely absolutely continuous and

$$\mathbb{E} \int_0^T \|\dot{X}_t\|^2 dt < \infty.$$

*Assumption A2 (chain rule under the expectation).* The field  $g_\theta$  is weakly differentiable,  $t \mapsto g_\theta(X_t)$  is almost surely absolutely continuous, and

$$g_\theta(X_t) - g_\theta(X_s) = \int_s^t \nabla g_\theta(X_u)^\top \dot{X}_u du,$$

for all  $0 \leq s < t \leq T$  on almost every sample path, with

$$\mathbb{E} \int_0^T |\nabla g_\theta(X_t)^\top \dot{X}_t| dt < \infty.$$

*Assumption A3 (score Jacobian well posed and directional domination along deployment).* There exists a measurable matrix field  $J_f$  such that  $J_f(X_t)$  is well defined almost surely for almost every  $t \in [0, T]$ ,

$$\mathbb{E} \int_0^T \|J_f(X_t) \dot{X}_t\|^2 dt < \infty,$$

and there exists  $\beta > 0$  such that

$$|\nabla g_\theta(X_t)^\top \dot{X}_t| \leq \beta \|J_f(X_t) \dot{X}_t\|$$

for almost every  $t \in [0, T]$  on almost every sample path.

Assumption A3 is the paper’s main compatibility condition. It should not be read as a consequence of dynamic covariate shift alone. Rather, it asserts that along the realized deployment path, the local change in the chosen performance field is dominated by the local change in the score, and that the score Jacobian is well posed there.

**Remark 1** (What A3 means, when it holds, and when it fails). *A3 is deliberately along-path rather than pointwise in every ambient direction. The proof only needs control of  $\nabla g_\theta(X_t)^\top \dot{X}_t$  by  $J_f(X_t) \dot{X}_t$  along the realized velocity.*

*A3 holds in the natural composition case. If  $g_\theta(x) = h(f_\theta(x))$ , then*

$$\nabla g_\theta(X_t)^\top \dot{X}_t = h'(f_\theta(X_t)) J_f(X_t) \dot{X}_t,$$

*so A3 follows with  $\beta = \sup |h'|$  whenever  $J_f(X_t)$  is well defined almost surely for almost every  $t$ . For Bernoulli cross-entropy against a fixed soft target,  $|h'| \leq 1$ , giving  $\beta = 1$ . The theorem is therefore not ReLU-specific. Smooth predictors fit directly, and piecewise-linear predictors such as ReLU networks fit as soon as the deployment path avoids their nondifferentiability set almost surely for almost every deployment time.*

*A3 does not follow from pure covariate shift alone. The general risk  $g_\theta(x) = \mathbb{E}_{Y \sim P_0(\cdot|x)}[\ell(f_\theta(x), Y)]$  carries an extra  $\nabla \eta(x)$  term from the label conditional  $\eta(x)$ , not necessarily dominated by the score-Jacobian term even when  $\eta$  is time-invariant. A3 also fails under simultaneous concept shift, and at score saturation wherever  $\nabla f_\theta(x) = 0$  but  $\nabla g_\theta(x) \neq 0$ .*

### IV. MAIN THEOREM PACKAGE

The first step is purely temporal.

**Lemma 1** (Time-domain derivative-energy control). *If  $r$  is absolutely continuous on  $[0, T]$ , then*

$$\text{Var}_U(r(U)) \leq \frac{T}{\pi^2} \int_0^T (r'(t))^2 dt. \quad (3)$$

*Proof.* Let  $\bar{r} = T^{-1} \int_0^T r(t) dt$ . The Wirtinger inequality yields

$$\int_0^T (r(t) - \bar{r})^2 dt \leq \frac{T^2}{\pi^2} \int_0^T (r'(t))^2 dt.$$

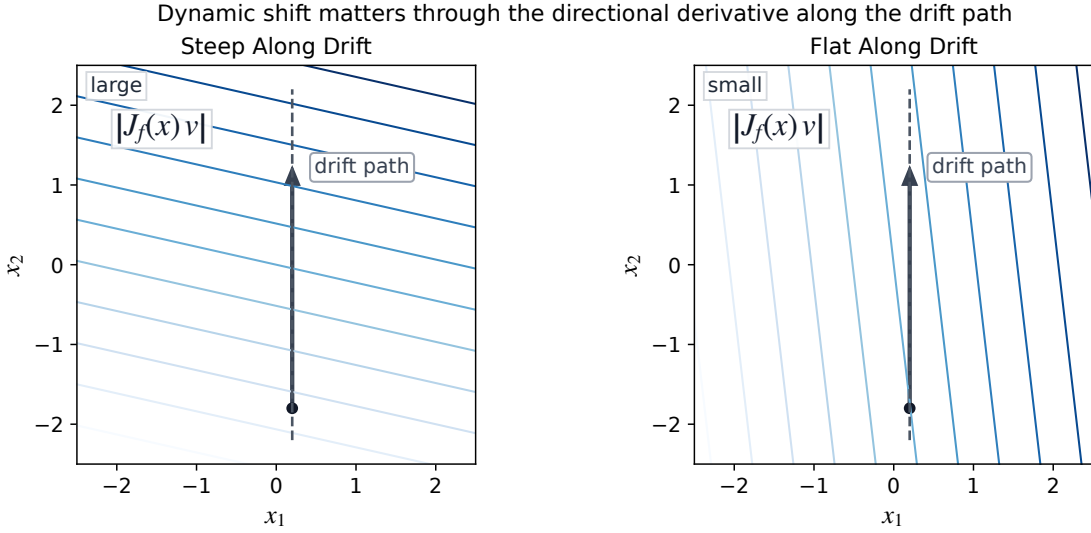


Fig. 1. Geometric intuition for drift-aligned instability. The deployment path is the same in both panels, but the local tangent geometry is different. When the drift direction aligns with a steep tangent direction, the score changes rapidly, whereas when the model is flat along that same direction, the score remains stable. This figure visualizes the directional quantity that appears in Theorem 1 and (4).

Divide by  $T$ .

Lemma 1 says volatility is impossible without derivative energy; see (3). The next result identifies the geometric driver of that energy.

**Theorem 1** (Jacobian-velocity control of risk volatility). *Assume A1–A3. Then  $r$  is absolutely continuous and*

$$\text{Var}_U(r(U)) \leq \frac{\beta^2 T}{\pi^2} \int_0^T \mathbb{E} \|J_f(X_t) \dot{X}_t\|^2 dt. \quad (4)$$

*Proof.* Fix  $0 \leq s < t \leq T$ . By Assumption A2 and Fubini,

$$\begin{aligned} r(t) - r(s) &= \mathbb{E} \left[ \int_s^t \nabla g_\theta(X_u)^\top \dot{X}_u du \right] \\ &= \int_s^t \mathbb{E} [\nabla g_\theta(X_u)^\top \dot{X}_u] du. \end{aligned}$$

Hence  $r$  is absolutely continuous and, for almost every  $t$ ,

$$r'(t) = \mathbb{E} [\nabla g_\theta(X_t)^\top \dot{X}_t].$$

Jensen's inequality and Assumption A3 then give

$$\begin{aligned} (r'(t))^2 &\leq \mathbb{E} [(\nabla g_\theta(X_t)^\top \dot{X}_t)^2] \\ &\leq \beta^2 \mathbb{E} \|J_f(X_t) \dot{X}_t\|^2. \end{aligned}$$

Apply Lemma 1 and (3).  $\square$

Theorem 1 isolates a single instability mechanism for the trajectory in (1). Deployment volatility is driven by accumulated tangent amplification along the realized motion of the data stream, as captured by (4). Figure 1 is a geometric rendering of that statement. In that sense, the theorem turns Jacobian-based sensitivity analysis and regularization [13]–[15] into a statement about temporal deployment risk rather than static local robustness. The bound also has a useful operational reading. Volatility requires two ingredients at once, data motion

$\square$  and tangent gain. If either one is small, long-horizon fluctuation remains small.

The low-rank regime gives the algorithmic specialization.

**Corollary 1** (Drift-subspace bound). *Assume the hypotheses of Theorem 1. If*

$$\dot{X}_t = V a_t + \rho_t, \quad V \in \mathbb{R}^{d \times k},$$

*with  $a_t \in \mathbb{R}^k$ ,  $\rho_t \in \mathbb{R}^d$ ,  $V^\top V = I_k$ , and  $V^\top \rho_t = 0$ , then*

$$\text{Var}_U(r(U)) \leq \frac{2\beta^2 T}{\pi^2} (B_V + B_\rho), \quad (5)$$

*where*

$$B_V := \int_0^T \mathbb{E} \|J_f(X_t) V\|_F^2 \|a_t\|^2 dt,$$

$$B_\rho := \int_0^T \mathbb{E} \|J_f(X_t) \rho_t\|^2 dt.$$

*Proof.* By Theorem 1 and (4),

$$\begin{aligned} \|J_f(X_t) \dot{X}_t\|^2 &= \|J_f(X_t) V a_t + J_f(X_t) \rho_t\|^2 \\ &\leq 2 \|J_f(X_t) V a_t\|^2 + 2 \|J_f(X_t) \rho_t\|^2. \end{aligned}$$

Since  $V$  has orthonormal columns,

$$\|J_f(X_t) V a_t\|^2 \leq \|J_f(X_t) V\|_F^2 \|a_t\|^2.$$

Integrate over  $t$ .  $\square$

Corollary 1 is the paper's main reduction. When  $B_\rho$  is small, the leading stability term in (5) is governed by the directional Jacobian energy in the drift subspace rather than by isotropic smoothness. This is precisely the regime in which low-dimensional temporal drift can be a useful approximation [5] and low-rank change models become algorithmically plausible [20].

## V. METHOD

The theorem package suggests a simple design principle. If instability is created by motion through steep directions, then regularization should flatten the model along directions deployment is expected to traverse, not everywhere at once.

### A. Drift-aligned tangent regularization

Let  $V$  estimate the dominant drift subspace from unlabeled deployment covariates. Two minimal estimators are sufficient for the present paper. For window mean  $\mu_t$  and lag  $\Delta$ , one is a mean-difference direction

$$\Delta\mu_t = \mu_t - \mu_{t-\Delta}, \quad v_t = \frac{\Delta\mu_t}{\|\Delta\mu_t\|},$$

or the top- $k$  principal directions of the difference cloud  $\{X_t - X_{t-\Delta}\}$  over a rolling window.

With a fixed subspace estimate  $V$ , we train using

$$\mathcal{L}_{\text{DTR}}(\theta) = \mathbb{E}_{(X,Y)} \ell(f_\theta(X), Y) + \lambda \mathbb{E}_X \|J_f(X)V\|_F^2. \quad (6)$$

The objective in (6) is not global Jacobian smoothing. It is a directional sensitivity constraint. The model is flattened only along directions expected to dominate future drift. Corollary 1 explains why this is the quantity worth shrinking, and it distinguishes DTR from isotropic Jacobian penalties that aim at broader smoothness or margin control [13], [14]. It is also conceptually adjacent to empirical sensitivity analyses that track related Jacobian quantities without imposing directional structure [15]. In other words, DTR concentrates regularization budget on directions the future trajectory is expected to traverse, rather than spreading it uniformly over directions the deployment path may never visit.

### B. Monitoring score

The same geometry yields a theorem-motivated proxy diagnostic. With current drift-subspace estimate  $V_t$ , define speed  $s_t$ , gain  $g_t$ , and product  $h_t$  by

$$\begin{aligned} s_t &:= \|\Delta\mu_t\|/\Delta, & g_t &:= \mathbb{E}\|J_f(X_t)V_t\|_F^2, \\ h_t &:= s_t^2 g_t. \end{aligned} \quad (7)$$

In the experiments we also report short rolling averages of  $h_t$ , which better match the accumulated-energy form of the volatility bound.

**Proposition 1** (Rank-1 bookkeeping for the hazard score). *Consider the rank-1 monitoring proxy used below, so that  $V_t = v_t \in \mathbb{R}^d$  is a unit vector. Suppose the block mean difference satisfies*

$$\frac{\Delta\mu_t}{\Delta} = v\bar{a}_t + \bar{\rho}_t, \quad \|v\| = 1, \quad v^\top \bar{\rho}_t = 0,$$

where  $\bar{a}_t$  is the block-averaged aligned drift coefficient and  $\bar{\rho}_t$  is the block-averaged residual drift. Write

$$v_t = \cos \theta_t v + \sin \theta_t u_t, \quad u_t^\top v = 0, \quad \|u_t\| = 1,$$

and define

$$\begin{aligned} G_{\parallel,t} &:= \mathbb{E}\|J_f(X_t)v\|^2, \\ G_{\perp,t} &:= \mathbb{E}\|J_f(X_t)u_t\|^2, \\ C_t &:= \mathbb{E}\langle J_f(X_t)v, J_f(X_t)u_t \rangle. \end{aligned}$$

Then

$$s_t^2 = |\bar{a}_t|^2 + \|\bar{\rho}_t\|^2$$

and

$$g_t = \cos^2 \theta_t G_{\parallel,t} + \sin^2 \theta_t G_{\perp,t} + 2 \sin \theta_t \cos \theta_t C_t.$$

Consequently,

$$\begin{aligned} h_t &= (|\bar{a}_t|^2 + \|\bar{\rho}_t\|^2) \\ &\quad \times \left( \cos^2 \theta_t G_{\parallel,t} + \sin^2 \theta_t G_{\perp,t} + 2 \sin \theta_t \cos \theta_t C_t \right). \end{aligned}$$

*Proof.* The identities are immediate from orthogonal decompositions. The first uses  $(\Delta\mu_t)/\Delta = v\bar{a}_t + \bar{\rho}_t$  with  $v^\top \bar{\rho}_t = 0$ . The second expands  $J_f(X_t)v_t = \cos \theta_t J_f(X_t)v + \sin \theta_t J_f(X_t)u_t$  and takes expectations.  $\square$

Proposition 1 makes the proxy gap explicit in the rank-1 monitoring setting for the score in (7). The finite-window drift statistic contributes the residual term  $\|\bar{\rho}_t\|^2$ , while angular error mixes aligned, orthogonal, and overlap energies through  $\theta_t$ ,  $G_{\perp,t}$ , and  $C_t$ . In that precise bookkeeping sense,  $h_t$  matches the leading low-rank quantity from Corollary 1 up to block averaging, residual drift, and subspace-estimation error rather than reproducing the corollary integrand exactly. It becomes large only when the stream is moving quickly *and* the predictor is steep along that motion. Relative to dataset-shift diagnostics [16], uncertainty diagnostics under shift [17], and sequential detection procedures [18], [19], [22], the point of  $h_t$  is alignment rather than optimality. It monitors the same directional mechanism singled out by the theory without claiming to be a sequentially optimal detector. This symmetry between training and monitoring is deliberate. The same geometry used to suppress future instability is reused to measure when future instability is likely.

## VI. EXPERIMENTS

*Evaluation protocol.* Hyperparameters use only training/validation windows; deployment is held out for final evaluation.  $\lambda$  is selected by validation loss, with validation directional gain only a secondary tie-breaker. Metrics are post-selection matched-seed means: volatility is (2), derivative energy the empirical integral in (3), directional gain drift-direction Jacobian energy, and terminal risk final-block loss.

### A. Synthetic time-domain sanity check

The synthetic study uses the smallest setup that exposes the paper's mechanism. There is one stable signal coordinate, one drifting nuisance coordinate, and a frozen classifier whose tangent geometry can either amplify or suppress that drift. The drift direction is deliberately known so the theorem-to-method link can be read cleanly [5], [6]. Labels are sampled as




Fig. 2. Synthetic time-domain sanity check. Each point is one trained model evaluated over the drifting deployment path. The horizontal axis is the derivative-energy quantity from Lemma 1 and (3), and the vertical axis is empirical volatility  $\text{Var}_U(r(U))$ . DTR shifts models toward lower energy and lower volatility.

$Y \sim \text{Bernoulli}(1/2)$  and we write  $S = 2Y - 1 \in \{-1, +1\}$ . Conditional on  $S$ , features are generated as

$$x_1 \sim \mathcal{N}(1.05S, 0.90^2), \quad x_2 \sim \mathcal{N}(1.25S + \delta(t), 0.55^2).$$

The first coordinate is a stable signal. The second is a nuisance coordinate that is predictive at training time but drifts during deployment. The label mechanism is held fixed throughout. Only  $\delta(t)$  changes. The benchmark is intentionally simple. Its job is to make the paper’s directional mechanism observable with as little confounding structure as possible.

The drift schedule is monotone and smooth, with speed

$$\dot{\delta}(t) = 0.55 + 1.05e^{-((t-0.33)/0.10)^2} + 0.75e^{-((t-0.76)/0.08)^2}, \\ \delta(0) = 0,$$

so the stream contains two periods of rapid motion. We train a small ReLU classifier on samples drawn at  $t = 0$  and evaluate it on fresh samples along the deployment path. The sweep uses  $\lambda \in \{0, 0.01, 0.03, 0.08\}$  over 20 random seeds. DTR uses the true drift direction  $V = e_2$  so that this experiment isolates the theorem-to-method link rather than subspace-estimation error.

Figure 2 plays the lemma-linked role in the paper’s narrative arc. Each point is one trained model, plotted by empirical risk volatility in (2) versus the derivative-energy quantity from Lemma 1 and (3). The points sit below the diagonal, as they should, and DTR moves the cloud toward the lower-left corner. The saved summaries also respect the full Jacobian-velocity upper bound from Theorem 1, but that bound is numerically looser, so the figure emphasizes the sharper temporal step that anchors the rest of the pipeline. Averaging over 20 seeds, mean volatility falls from  $3.25 \times 10^{-3}$  for standard training to  $2.39 \times 10^{-4}$  across the DTR sweep, while mean directional gain falls from 41.5 to 1.85.

TABLE I  
DIRECTIONAL COMPARISON SHOWING THE MEAN OVER THE NONZERO- $\lambda$  SWEEP. LOWER IS BETTER ON ALL METRICS.

Method	Deriv. energy	Volatility	Dir. gain	Term. risk
Standard	$5.15 \times 10^{-3}$	$3.25 \times 10^{-3}$	41.5	0.189
Isotropic	$5.85 \times 10^{-4}$	$4.32 \times 10^{-4}$	2.17	0.172
DTR	$2.90 \times 10^{-4}$	$2.39 \times 10^{-4}$	1.85	0.136

### B. Directional versus isotropic smoothing

The main algorithmic question is not whether Jacobian regularization helps in general, but whether directional smoothing beats isotropic smoothing when drift is concentrated in one known direction. We answer that question in the same synthetic environment by replacing the DTR penalty with an isotropic Jacobian penalty

$$\lambda \mathbb{E} \|\nabla f_\theta(X)\|^2,$$

while keeping the directional penalty

$$\lambda \mathbb{E} \|J_f(X)V\|_F^2$$

for DTR, with the same architecture and 20 matched seeds over  $\lambda \in \{0.01, 0.03, 0.08\}$ . Under rank-1 drift, isotropic smoothing spends budget in directions the deployment path does not substantially traverse, whereas DTR spends it only along the realized motion. Table I summarizes the mean values across the nonzero- $\lambda$  sweep. DTR improves on both baselines on every metric. Figure 3 (left) shows the matched  $\lambda = 0.03$  comparison.

We then test drift-subspace misspecification by rotating the one-dimensional subspace used by DTR. Writing the estimated direction as

$$\hat{v}_\alpha = (\sin \alpha, \cos \alpha),$$

its alignment with the true drift direction  $e_2$  is  $\cos \alpha$ . With the correct subspace, the directional advantage remains strongest. With a mild  $20^\circ$  rotation, volatility rises by a factor of 1.39 and terminal risk by a factor of 1.18 relative to aligned DTR, but both remain far below the standard model. The same cosine-sine bookkeeping used in Proposition 1 helps explain that pattern: small angular error mainly rescales the aligned contribution, whereas large misalignment transfers weight to orthogonal energy. With an orthogonal subspace, the effect largely disappears. Derivative energy rises by a factor of 39.8, volatility by 29.6, and directional gain by 18.9. Figure 3 (right) therefore sharpens the paper’s claim in the intended way. DTR is more targeted than isotropic Jacobian smoothing when the drift geometry is right, and its gains depend on keeping the estimated drift subspace reasonably aligned with realized motion.

### C. Real field deployments

We next ask whether the same directional story survives on real deployments where the drift subspace must be estimated from unlabeled covariates. We use two complementary regression settings: Air Quality as the direct sensor-drift case,

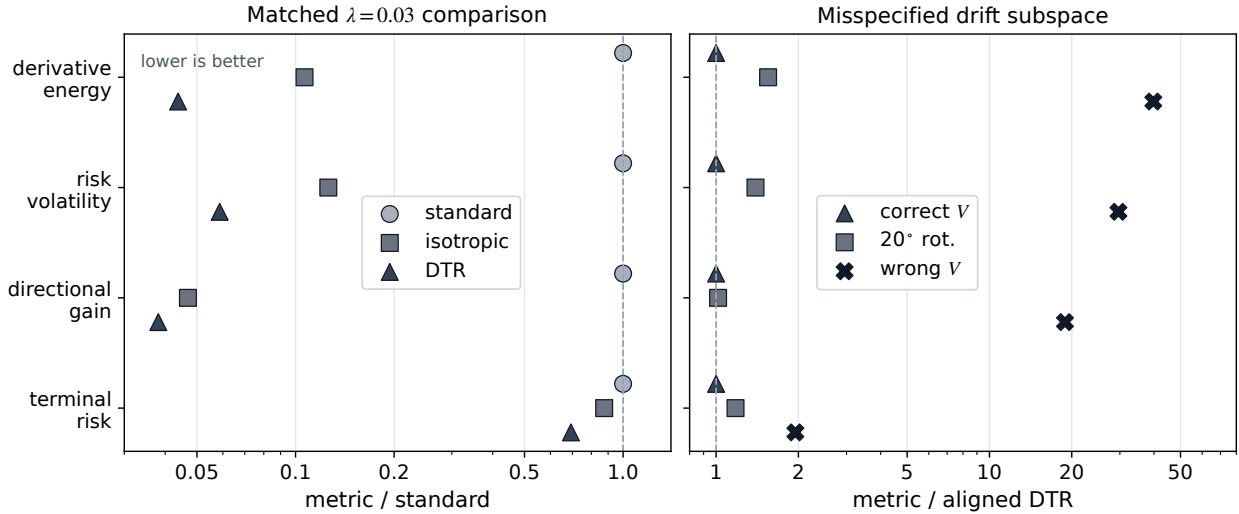


Fig. 3. Synthetic directional comparisons. Left shows matched  $\lambda = 0.03$  ratios for derivative energy, risk volatility, directional gain, and terminal risk, normalized by the standard model. DTR is lower than isotropic Jacobian regularization on all four metrics. Right shows the same metrics normalized by aligned DTR for the correct subspace, a mild  $20^\circ$  rotation, and a wrong subspace. Mild misspecification is tolerable, but wrong subspaces are not.

and Tetouan as a weather-driven load-prediction task with fixed label semantics and explicit exogenous covariates. In the real-data experiments, the drift subspace is estimated retrospectively from unlabeled deployment covariates; a fully online deployment would replace this with rolling estimates of  $V_t$ .

The UCI Air Quality dataset records hourly outputs of a five-sensor gas array together with temperature and humidity in an Italian city, and the associated field-study paper explicitly discusses sensor drift and seasonal influences over long deployments [23]. We predict the reference CO concentration  $\text{CO}(\text{GT})$  from the five sensor channels PT08.S1–PT08.S5 together with temperature,  $RH$ , and  $AH$ . Rows with the dataset’s  $-200$  missing-value marker are removed, leaving 1573 training points from the first 12 weeks, 580 validation points from the next 4 weeks, and 5191 deployment points grouped into 20 biweekly blocks. The predictor is frozen after training, and deployment risk is blockwise mean-squared error. Because this dataset is explicitly a sensor-drift setting, the primary Air Quality DTR subspace is estimated from the sensor channels only. We first remove the supervised linear target direction learned on the training window, then take the top two singular vectors of consecutive biweekly mean-shift vectors in that target-orthogonal sensor space and embed them back into the full feature space. This keeps weather covariates available to the predictor while avoiding a DTR penalty on the main target-predictive sensor direction. We compare standard training, isotropic Jacobian regularization, and DTR over 10 matched random seeds. For reporting the deployment monitor, we use the blockwise rank-1 diagnostic: if  $\mu_t$  is the standardized covariate mean in block  $t$  and  $\Delta = 1$  block, then

$$s_t = \|\mu_t - \mu_{t-1}\|, \quad v_t = \frac{\mu_t - \mu_{t-1}}{\|\mu_t - \mu_{t-1}\|},$$

$$g_t = \mathbb{E}[(\nabla f_\theta(X_t)^\top v_t)^2], \quad h_t = s_t^2 g_t.$$

We then test a second real benchmark on the UCI Tetouan City power-consumption dataset [24]. It is a ten-minute 2017 weather-and-load series with three zone-level targets; we predict Zone 1 power consumption from temperature, humidity, wind speed, and diffuse-flow channels. Using the same frozen-model protocol, we train on January–April, validate on May–June, and deploy on July–December in six monthly blocks, yielding 17280/8784/26352 training/validation/deployment points.

Table II reports validation-MSE-selected results. Air Quality chooses  $\lambda = 0.08$  for isotropic smoothing and  $\lambda = 0.003$  for DTR, while Tetouan chooses  $\lambda = 3 \times 10^{-4}$  for isotropic smoothing and  $\lambda = 10^{-2}$  for DTR. The Air Quality result is now favorable but appropriately qualified: DTR improves mean deployment MSE, volatility, and terminal risk relative to standard training, while isotropic smoothing has the lowest mean MSE. The Air Quality gain column is the full block-direction gain used by the diagnostic; the DTR penalty itself acts on the target-orthogonal sensor subspace. Tetouan gives the stronger all-around real-data directional result, with DTR improving both deployment MSE and volatility relative to both baselines (Figure 4). Paired seed comparisons avoid per-block significance tests. On Air Quality, DTR improves deployment MSE and volatility in 9/10 seeds against standard training, with bootstrap 95% CIs for the paired mean differences below zero. On Tetouan, DTR improves volatility in 8/10 seeds against standard training and 8/10 seeds against isotropic smoothing, again with bootstrap 95% CIs below zero.

Table III shows why the Air Quality subspace definition matters. If  $V$  is estimated from all covariates, validation-MSE selection chooses a large DTR penalty that suppresses block-direction gain but worsens deployment risk. The target-orthogonal sensor subspace avoids that failure: validation-MSE selection chooses  $\lambda = 0.003$ , and DTR improves MSE and volatility in 9/10 matched seeds against standard training. A weather-residualized sensor subspace gives an even lower mean

TABLE II  
VALIDATION-SELECTED REAL DEPLOYMENT RESULTS. ENTRIES ARE MEAN  $\pm$  STANDARD DEVIATION OVER 10 MATCHED RANDOM SEEDS. LOWER IS BETTER FOR ALL METRICS.

Dataset	Method	Deploy MSE	Volatility	Dir. gain	Terminal risk
Air Quality	Standard	0.449 $\pm$ 0.069	0.073 $\pm$ 0.023	0.079 $\pm$ 0.008	0.165 $\pm$ 0.034
Air Quality	Iso. Jacobian	0.415 $\pm$ 0.030	0.077 $\pm$ 0.008	0.053 $\pm$ 0.005	0.177 $\pm$ 0.045
Air Quality	DTR	0.432 $\pm$ 0.058	0.069 $\pm$ 0.020	0.079 $\pm$ 0.008	0.161 $\pm$ 0.028
Tetouan	Standard	$(1.08 \pm 1.11) \times 10^8$	$(1.07 \pm 1.66) \times 10^{16}$	0.024 $\pm$ 0.024	$(1.57 \pm 0.41) \times 10^7$
Tetouan	Iso. Jacobian	$(1.01 \pm 0.81) \times 10^8$	$(7.20 \pm 10.24) \times 10^{15}$	0.023 $\pm$ 0.017	$(1.61 \pm 0.41) \times 10^7$
Tetouan	DTR	$(6.82 \pm 4.87) \times 10^7$	$(3.07 \pm 4.32) \times 10^{15}$	0.013 $\pm$ 0.008	$(1.45 \pm 0.45) \times 10^7$

TABLE III  
AIR QUALITY DTR SUBSPACE ABLATION. ENTRIES ARE MEANS OVER 10 MATCHED SEEDS. THE FIRST TWO DTR ROWS ARE VALIDATION-MSE SELECTED; THE WEATHER-RESIDUALIZED ROW IS A MODERATE-PENALTY SENSITIVITY POINT. WINS ARE PAIRED DTR-VS-STANDARD MSE/VOLATILITY WINS.

Subspace	$\lambda$	MSE	Vol.	Gain	Wins
All covariates	0.080	0.485	0.112	0.044	4/10, 2/10
Target-orth. sensors	0.003	0.432	0.069	0.079	9/10, 9/10
Weather-resid. sensors	0.030	0.411	0.061	0.064	9/10, 6/10

TABLE IV  
MONITORING-SCORE ABLATION ON SELECTED DTR DEPLOYMENTS. ENTRIES ARE SPEARMAN CORRELATIONS WITH NEXT-BLOCK SQUARED RISK CHANGE  $(r_{t+1} - r_t)^2$ . HIGHER IS BETTER; ROLLING COLUMNS USE FEWER VALID TRANSITIONS.

Dataset	Drift $s_t^2$	Gain $g_t$	Product $h_t$	Roll-2 $h_t$	Roll-3 $h_t$
Air Quality	0.169	-0.099	0.044	0.307	0.353
Tetouan	-0.339	0.328	0.218	0.305	0.335

MSE and volatility at a moderate  $\lambda = 0.03$ , but validation MSE selects a larger penalty for that variant, so it is reported only as a subspace sensitivity check. The practical lesson is the same as the misspecification experiment: DTR is useful when the estimated subspace represents nuisance drift rather than broad covariate motion or target signal.

Finally, we evaluated the matched hazard score against block-to-block risk movement rather than raw risk level. This is the more natural target for the theory, since the bound controls temporal volatility. Table IV reports Spearman correlations with next-block squared risk change. The one-block product  $h_t$  is noisy, but a short rolling average of the matched score is positively associated with future risk movement on both real deployments. On Air Quality, the roll-2 score exceeds drift-only with a bootstrap 95% confidence interval separated from the drift-only estimate. On Tetouan, the point estimate is also positive, but uncertainty is larger because there are only six deployment blocks per seed. We therefore treat the score as an informative theorem-matched volatility diagnostic rather than as a validated sequential detector.

## VII. DISCUSSION AND LIMITATIONS

The crucial caution is Assumption A3. Dynamic covariate shift by itself does not imply that the conditional-risk field is controlled by the score Jacobian along the realized path. We

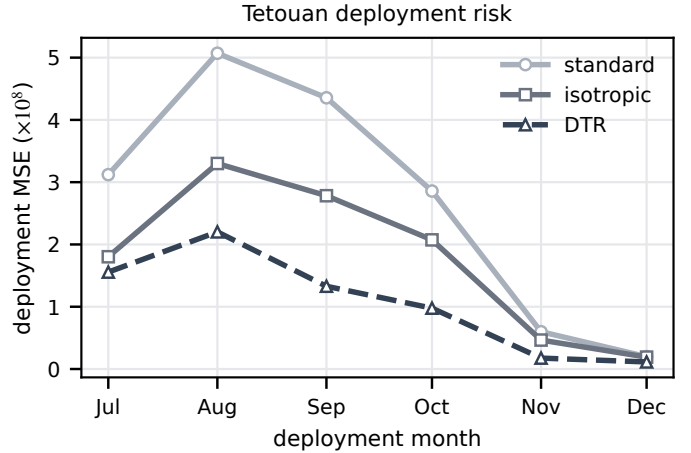


Fig. 4. Tetouan frozen deployment over six monthly blocks. DTR stays below both standard and isotropic training throughout the deployment horizon, with the largest gap in the late-summer high-risk regime.

impose that domination explicitly because it is exactly what the theorem uses. The result is also not tied to ReLU networks in particular. ReLU models are one admissible class when deployment avoids kink sets almost surely for almost every deployment time, but smooth models fit the same theorem just as well.

The four empirical pieces play different roles. The time-domain sanity check asks whether the Poincaré step is visible in trained models at all. The isotropic-versus-directional comparison asks whether the low-rank specialization has real design content. The Air Quality and Tetouan studies ask whether the same directional quantities still matter on real frozen deployments once the drift subspace must be estimated. That progression matches the claim: isolate the temporal bound, test the low-rank design implication, then check whether the same directional quantities remain visible on noisier field data. The Air Quality subspace ablation is informative precisely because it separates broad covariate motion from nuisance sensor drift. An all-covariate subspace over-regularizes, while a target-orthogonal sensor subspace turns the same DTR objective into a validation-selected deployment improvement.

The low-rank story also depends on estimating a drift subspace. The synthetic misspecification study suggests that moderate angular error is tolerable but large misalignment is not. In the real-data studies we estimate a fixed subspace

from unlabeled future covariates, which is more realistic than the aligned synthetic setting but still retrospective. A real online deployment would require rolling estimation of  $V_t$ , and the quality of the monitoring score would then depend on residual drift energy and subspace-estimation error. Finally, the Jacobian-velocity bound is an upper bound: useful because it is actionable, not because it is tight in every regime.

### VIII. CONCLUSION

Long-horizon robustness under dynamic shift is best viewed as a tangent-control problem. The data stream contributes a velocity, the model contributes a local linear map, and deployment volatility is governed by their interaction. The central message is simple. Deployment becomes unstable when environmental motion repeatedly passes through directions to which the model is locally sensitive. The theorem package makes this precise, DTR acts on it, and the experiments show where that directional geometry is strongest: controlled low-rank drift, target-orthogonal sensor drift on Air Quality, and the Tetouan frozen deployment. Air Quality adds the complementary lesson that estimating  $V$  is part of the method: drift-aligned regularization helps when the subspace captures nuisance motion, but an overly broad subspace can flatten useful predictive signal.

Relative to temporal-shift benchmarks and static shift analyses [1], [4]–[6], the paper contributes a theorem-backed mechanism: the same environmental motion can encounter different directional tangent gains across frozen models. Relative to test-time adaptation [10]–[12], it studies the frozen-model regime directly. Relative to Jacobian smoothing and shift detection [13], [14], [16], [22], it ties training-time regularization and monitoring to the same directional drift sensitivity. If future motion is low-dimensional, robustness should be enforced, measured, and monitored in those directions rather than spread uniformly over the ambient space.

### ACKNOWLEDGMENT

The author used Python packages and AI-assisted coding tools, including Codex and Claude Code, for code development, algebraic checks, proof-verification scaffolding, repository hygiene, and editorial refinement. The author conceptualized, reviewed, and finalized all mathematical statements, experiments, code, interpretations, and final text and is solely responsible for the content. The author received no external funding and reports no conflicts of interest.

### REFERENCES

- [1] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. Wortman Vaughan, "A theory of learning from different domains," *Machine Learning*, vol. 79, no. 1–2, pp. 151–175, 2010.
- [2] J. G. Moreno-Torres, T. Raeder, R. Alaiz-Rodriguez, N. V. Chawla, and F. Herrera, "A unifying view on dataset shift in classification," *Pattern Recognition*, vol. 45, no. 1, pp. 521–530, 2012.
- [3] J. Gama, I. Zliobaite, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 44:1–44:37, 2014.
- [4] P. W. Koh, S. Sagawa, H. Marklund, S. M. Xie, M. Zhang, A. Balsubramani, W. Hu, M. Yasunaga, R. L. Phillips, I. Gao, T. Lee, E. David, I. Stavness, W. Guo, B. A. Earnshaw, I. Haque, S. Beery, J. Leskovec, A. Kundaje, E. Pierson, S. Levine, C. Finn, and P. Liang, "Wilds: A benchmark of in-the-wild distribution shifts," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 139, 2021, pp. 5637–5664.
- [5] H. Yao, C. Choi, B. Cao, Y. Lee, P. W. Koh, and C. Finn, "Wild-time: A benchmark of in-the-wild distribution shift over time," in *Advances in Neural Information Processing Systems*, 2022.
- [6] E. Han, C. Huang, and K. Wang, "Model assessment and selection under temporal distribution shift," in *Proceedings of the 41st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 235, 2024, pp. 17 374–17 392.
- [7] J. R. Landers, "jacobian-velocity-bounds," GitHub repository, 2026, [Online]. Available: <https://github.com/jonland82/jacobian-velocity-bounds>. Accessed: May 6, 2026.
- [8] —, "Jacobian-velocity bounds project page," Project website, 2026, [Online]. Available: <https://jonland82.github.io/jacobian-velocity-bounds/>. Accessed: May 6, 2026.
- [9] M. Simchowitz, A. Ajay, P. Agrawal, and A. Krishnamurthy, "Statistical learning under heterogeneous distribution shift," in *Proceedings of the 40th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 202, 2023, pp. 31 800–31 851.
- [10] Y. Sun, X. Wang, Z. Liu, J. Miller, A. A. Efros, and M. Hardt, "Test-time training with self-supervision for generalization under distribution shifts," in *Proceedings of the 37th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 119, 2020, pp. 9229–9248.
- [11] D. Wang, E. Shelhamer, S. Liu, B. Olshausen, and T. Darrell, "Tent: Fully test-time adaptation by entropy minimization," in *International Conference on Learning Representations*, 2021.
- [12] S. Gui, X. Li, and S. Ji, "Active test-time adaptation: Theoretical analyses and an algorithm," in *International Conference on Learning Representations*, 2024.
- [13] J. Sokolic, R. Giryes, G. Sapiro, and M. R. D. Rodrigues, "Robust large margin deep neural networks," *IEEE Transactions on Signal Processing*, vol. 65, no. 16, pp. 4265–4280, 2017.
- [14] T. Kim and H. Yang, "An infinite-width analysis on the jacobian-regularised training of a neural network," in *Proceedings of the 41st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 235, 2024, pp. 24 584–24 657.
- [15] R. Novak, Y. Bahri, D. A. Abolafia, J. Pennington, and J. Sohl-Dickstein, "Sensitivity and generalization in neural networks: An empirical study," *International Conference on Learning Representations*, 2018.
- [16] S. Rabanser, S. Günnemann, and Z. C. Lipton, "Failing loudly: An empirical study of methods for detecting dataset shift," in *Advances in Neural Information Processing Systems*, 2019.
- [17] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, and J. Snoek, "Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift," in *Advances in Neural Information Processing Systems*, 2019.
- [18] S. Wu, E. Diao, J. Ding, T. Banerjee, and V. Tarokh, "Robust quickest change detection for unnormalized models," in *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, ser. Proceedings of Machine Learning Research, vol. 216, 2023, pp. 2314–2323.
- [19] Y. Cao and Y. Xie, "Multi-sensor gradual change detection," in *Proceedings of the 53rd Annual Allerton Conference on Communication, Control, and Computing*, 2015, pp. 827–834.
- [20] Y. Xie and L. M. Seversky, "Sequential low-rank change detection," in *Proceedings of the 54th Annual Allerton Conference on Communication, Control, and Computing*, 2016, pp. 128–133.
- [21] V. Krishnamurthy and L. Snow, "Quickest change detection using time inconsistent anticipatory and quantum decision modeling," in *Proceedings of the 58th Annual Allerton Conference on Communication, Control, and Computing*, 2022, pp. 1–8.
- [22] A. Cooper and S. P. Meyn, "Quickest change detection using mismatched cusum extended abstract," in *Proceedings of the 60th Annual Allerton Conference on Communication, Control, and Computing*, 2024, pp. 1–7.
- [23] S. De Vito, E. Massera, M. Piga, L. Martinotto, and G. Francia, "On field calibration of an electronic nose for benzene estimation in an urban pollution monitoring scenario," *Sensors and Actuators B: Chemical*, vol. 129, no. 2, pp. 750–757, 2008.
- [24] A. Salam and A. El Hibaoui, "Power consumption of tetouan city," UCI Machine Learning Repository, 2018, [Dataset].