

Environment-Aware Learning of Smooth GNSS Covariance Dynamics for Autonomous Racing

Y. Deemo Chen^{*,†}, Arion Zimmermann^{*,†}, Thomas A. Berrueta[†], and Soon-Jo Chung[†]

Abstract—Ensuring accurate and stable state estimation is a challenging task crucial to safety-critical domains such as high-speed autonomous racing, where measurement uncertainty must be both adaptive to the environment and temporally smooth for control. In this work, we develop a learning-based framework, LACE, capable of directly modeling the temporal dynamics of GNSS measurement covariance. We model the covariance evolution as an exponentially stable dynamical system where a deep neural network (DNN) learns to predict the system’s process noise from environmental features through an attention mechanism. By using contraction-based stability and systematically imposing spectral constraints, we formally provide guarantees of exponential stability and smoothness for the resulting covariance dynamics. We validate our approach on an AV-24 autonomous racecar, demonstrating improved localization performance and smoother covariance estimates in challenging, GNSS-degraded environments. Our results highlight the promise of dynamically modeling the perceived uncertainty in state estimation problems that are tightly coupled with control sensitivity.

I. INTRODUCTION

Accurate and reliable state estimation is a key component of safety-critical autonomous systems, where centimeter-level precision is often essential for safe navigation and operation [1]. Autonomous racing in competition circuits such as the Indy Autonomous Challenge (IAC) represents an extreme instance of this requirement, demanding not only high-accuracy localization but also exceptionally smooth state estimates to ensure stable vehicle control at speeds exceeding 150 mph [2]. However, maintaining this performance is complicated by environmental factors that degrade sensor measurements. When a vehicle passes through Global Navigation Satellite System (GNSS)-degraded areas, such as under bridges or near buildings, multipath interference and poor satellite visibility can cause measurement noise to fluctuate dramatically, increasing from less than a meter to over 50 meters in less than a second.

In high-speed autonomous systems, standard estimation methods, such as the Extended Kalman Filter (EKF), face a dual challenge: the measurement noise covariance must be both adaptive with respect to the environment and temporally smooth and stable. The need for smoothness is driven by the downstream control system, where abrupt changes in perceived measurement uncertainty can destabilize the vehicle by triggering aggressive maneuvers. Simultaneously, the covariance must adapt to rapid changes in GNSS qual-

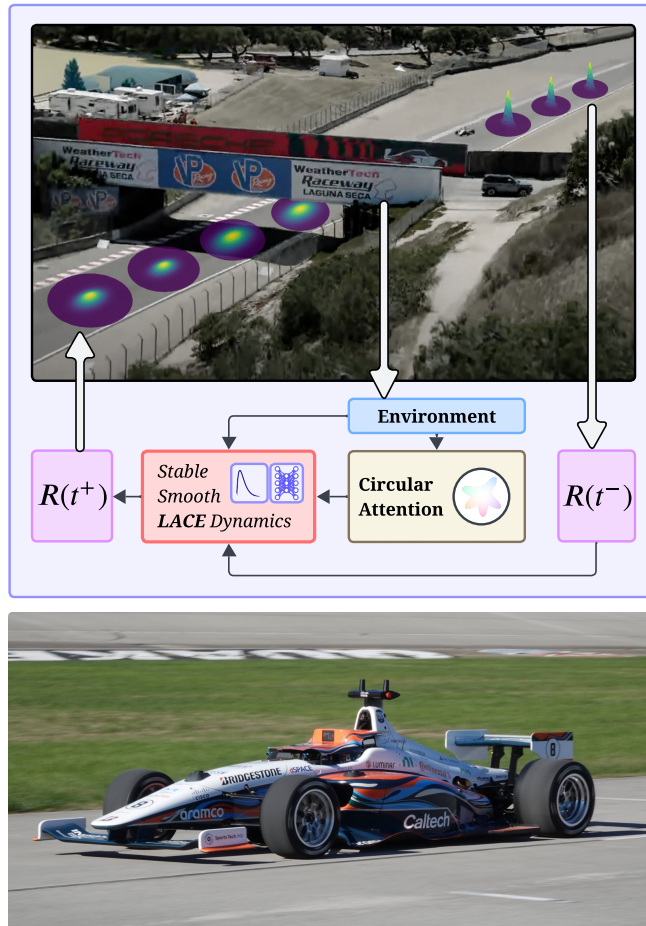


Fig. 1. LACE applied to IAC racecars. **Top:** Illustration of the learned GNSS covariance evolution as the racecar passes under a concrete bridge at Laguna Seca Raceway. The structure degrades GNSS measurement quality, causing the covariance to broaden. For clarity, the covariance distribution is visualized in 2D. **Bottom:** Caltech IAC AV-24 racecar.

ity to prevent the filter from becoming dangerously overconfident in noisy data or overly cautious with clean data. Satisfying this dual requirement is particularly challenging for conventional covariance models. A constant covariance, while inherently smooth, is not adaptive and leads to poor performance at environmental extremes [3]. Moreover, models based on receiver-provided metrics, such as Dilution of Precision (DOP), adapt to satellite geometry but can change erratically and, more importantly, cannot account for local effects like signal multipath [4]. As a result, standard approaches fail to consider the coupled dynamics between state estimation and control.

Recent methods address environmental adaptation and

*Equal contribution.

[†]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA.

Video: <https://youtu.be/fkOKFMRtwyk>

This project was in part funded by Beyond Limits.

smoothness as isolated objectives without modeling covariance evolution as a dynamical system. For environmental awareness, DNNs have proven effective at learning an end-to-end mapping from raw sensor data to uncertainty, such as predicting instantaneous measurement covariance directly from sensor observables like images without needing hand-coded features [5]. This has been extended to simultaneously estimate both process and measurement noise from temporal data streams using multitask learning architectures [6]. Other approaches use intermediate representations such as crowdsourced GNSS signal quality maps to predict measurement noise in urban areas [7]. Some non-parametric framework predicts covariance by weighing empirical errors from nearby samples in a feature space [8]. This approach has since been applied to 3D point cloud registration [9] and has been extended with expectation-maximization to enable learning without ground truth state information [10]. While these methods capture complex sensor-to-uncertainty mappings, they output static covariance estimates without temporal dynamics. In other words, the learned models can produce discrete jumps in uncertainty that may destabilize downstream systems.

In contrast, techniques that prioritize temporal smoothness are often reactive and less capable of adapting to abrupt environmental shifts. Classical adaptive filtering, for example, adjusts noise parameters based on recent performance by analyzing its innovation sequence [11]. Classical techniques rooted in innovation-based adaptive estimation operate on the principle that an optimal filter produces a white-noise innovation sequence and make use of this property to iteratively identify the process and measurement noise covariances [12]. Alternatively, methods based on multiple-model adaptive estimation run banks of parallel filters with different noise model assumptions and compute final state estimates as their weighted averages [13]. Smoothing techniques offer another path to stability by processing data in both forward and backward passes to produce optimal state trajectories, which is particularly effective during sensor outages [14]. However, these methods are either reactive or constrained to near-real-time post-processing, and cannot preemptively adapt to the rapid, environment-driven changes captured by predictive models.

To bridge the gap between purely predictive and reactive methods, some approaches explicitly model covariance as a regression problem. These methods construct a functional mapping from a feature space to a covariance matrix, using either parametric forms [15], [16] or non-parametric techniques like Gaussian Processes [17]. While effective, these regression methods model the output of the uncertainty process but not the process itself. By framing each prediction independently, they neglect the inherent temporal correlation in how state uncertainty evolves. For example, sequential architectures may process temporal features [6], but they still generate a series of static covariance estimates without explicitly modeling the transition between them. This can lead to discontinuities that are problematic for control, a challenge that smoothing techniques address but only in

post-processing [14]. In order to simultaneously achieve environmental adaptation and tunable smoothness there is a need for frameworks capable of learning the underlying covariance temporal dynamics.

In this work, we address this gap by introducing LACE (*Learning Adaptive Covariance Evolution*), a learning-based framework that directly models the temporal dynamics of GNSS measurement covariance in an environmentally-aware fashion. We model the covariance evolution as an exponentially stable linearized time-varying dynamical system. On the one hand, a DNN learns to predict the system's process noise matrix, $Q(t)$, from environmental features. In particular, an attention mechanism produces spatial embeddings of the robot's state in order to capture sharp and highly localized environmental features, such as bridges or buildings. On the other hand, A is constructed to be globally exponentially stable with spectral constraints, ensuring that the covariance evolution is both stable and smooth by design. Our contributions are then the following: 1) we propose a novel architecture that unifies environmental adaptation and tunable smoothness; 2) we prove the exponential stability and convergence of the covariance dynamics using contraction analysis [18], [19]; and 3) we validate our approach on an AV-24 autonomous racecar at the Laguna Seca Raceway, demonstrating improved performance in GNSS-degraded areas.

II. PROBLEM DEFINITION

Consider a robot with state $x(t) \in \mathbb{R}^m$, which includes the position vector $p(t) \in \mathbb{R}^3$. We define the GNSS measurement residual as $\epsilon(t) = p(t) - y(t)$, where $y(t)$ is the received GNSS position. Assuming a common coordinate frame, we model these residuals as realizations of a heteroscedastic Gaussian distribution,

$$\epsilon(t) \sim \mathcal{N}(0, R(t)), \quad (1)$$

where the time-varying covariance matrix $R(t)$ is the quantity we aim to estimate from the robot's estimated state, $\hat{x}(t)$, and GNSS quality metrics, $g(t)$, such as DOP measures and satellite count.

While modeling the residuals with a heteroscedastic Gaussian may seem restrictive, it is a consistent assumption within the Multi-state Constraint Kalman Filter (MSCKF [20], [21]), which makes Gaussian measurement noise assumptions.

In addition, temporal smoothness of the predicted covariance matrix is essential for stable state estimation and control. Rapid collapses of the covariance ellipsoid can cause abrupt changes in position estimates, leading to controller instability, particularly in high-performance autonomous racing. We quantify smoothness through the evolution of the log-determinant of $R(t)$ and impose the bound:

$$\frac{d}{dt} \log \det R(t) \geq -r_{\max}, \quad (2)$$

where $r_{\max} > 0$ is the maximum contraction rate of the ellipsoid. Even though this smoothness metric does not

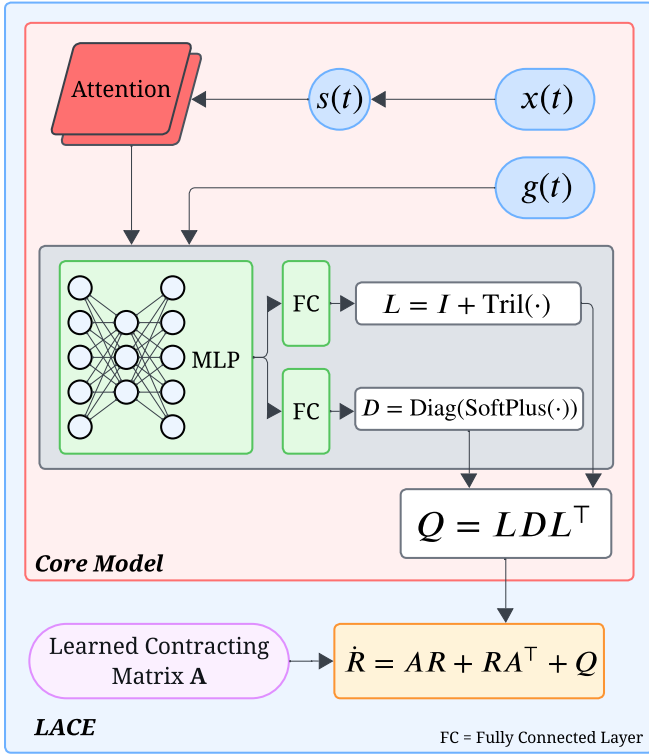


Fig. 2. LACE architecture overview. The core model produces a symmetric positive-definite matrix, while during both training and inference, LACE propagates this output through its dynamical model to guarantee stability and enforce temporal smoothness.

bound the contraction rate along all eigendirections, as the spectral norm would, it is a fast, practical and differentiable metric.

III. METHOD

As part of our method, we first infer at each timestep t a symmetric positive-definite (SPD) matrix Q from our core model using the robot's estimated state $\hat{x}(t)$ and the GNSS quality factors $g(t)$. This matrix can then be utilized in two distinct ways: (i) directly as the desired GNSS measurement covariance, where smoothness in the dynamics is controlled by means of a loss term from Section III-F; or (ii) as a reference covariance within a Lyapunov dynamics formulation, as outlined in Section III-D, for which smoothness is guaranteed *by design*. In this section, we discuss the model architecture of LACE, summarized in Fig. 2, the employed loss function that guarantees smoothness and finally prove the stability of our method.

A DNN is used to infer Q from the estimated state of the robot $\hat{x}(t)$ and from the GNSS quality factors $g(t)$ that consist of various DOP metrics (GDOP, HDOP, VDOP, etc.) and the number of satellites used by the GNSS receiver.

A. Preprocessing

To reduce state dimensionality, $\hat{x}(t)$ is projected onto arc-length coordinates, namely the longitudinal progress $s(t)$ along the reference trajectory and its corresponding velocity $\dot{s}(t)$. This reduction is justified by the fact that the racecar

repeatedly follows the same time-optimal trajectory and is later verified in Section IV by evaluating the model on multiple datasets following slightly different trajectories.

All inputs are normalized prior to being provided to the DNN. Specifically, DOP values are transformed using a logarithmic scale since they can vary over several orders of magnitude and often exhibit extreme values. The number of satellites is normalized per sequence. The longitudinal progress $s(t)$ is normalized to be between 0 and 1. Finally, the velocities $\dot{s}(t)$ are normalized, so that $\int_0^T \dot{s}(t) dt = T$, where T is the sequence duration. This preprocessing facilitates stable optimization during training and improves the generalization performance of the network.

B. Spatial Embedding and Attention Mechanism

A global single-query attention mechanism [22] is applied to the arc-length position $s(t)$ in order to encode spatial dependencies that may influence the output matrix. The inclusion of this mechanism is motivated by the physical environment of the racetrack, where surrounding structures such as bridges and buildings are present, and is aligned with previous work on attention-based map encoders that emphasize localized and sparse environmental features for control [23].

To this end, the query position is encoded into a unit-length complex query $Q = e^{i2\pi s(t)} \in \mathbb{C}$ and similarity with the key vector $K \in \mathbb{T}^a$ is given by:

$$\text{Sim}(Q, K) := \Re(Q^* K), \quad (3)$$

which measures periodic alignment along the track with each key. Based on this similarity function, the attention output is computed as:

$$\tilde{s} = W \left(V^\top \text{Softmax} \left(\frac{1}{\tau} \Re(Q^* K) \right) \right), \quad (4)$$

where $V \in \mathbb{R}^{a \times b}$ associates a value embedding for each key location, $\tau > 0$ is the temperature parameter controlling the sharpness of the attention weights, and $W \in \mathbb{R}^{c \times b}$ is the learnable output projection. With a low temperature τ , this formulation enables the attention mechanism to focus on highly localized spatial dependencies, as depicted in Figure 3. This attention mechanism is essential for learning the locations of sparse and localized environmental features.

C. Core Model

A multi-layer perceptron encodes the spatial embedding $\tilde{s}(t)$, the longitudinal velocity $\dot{s}(t)$, and the GNSS quality factors $g(t)$ into a higher dimensional embedding $\phi \in \mathbb{R}^p$.

The embedding $\phi = \phi(\tilde{s}(t), \dot{s}(t), g(t))$ is finally used to produce an output positive-definite matrix Q . To ensure positive definiteness, we parameterize the process noise matrix using the LDL decomposition:

$$Q(\phi) = L(\phi)D(\phi)L(\phi)^\top, \quad (5)$$

where $L(\phi) \in \mathbb{R}^{n \times n}$ is unit lower triangular and $D(\phi) \in \mathbb{R}^{n \times n}$ is diagonal with positive entries. This decomposition guarantees $Q(\phi) \succ 0$ by construction and provides

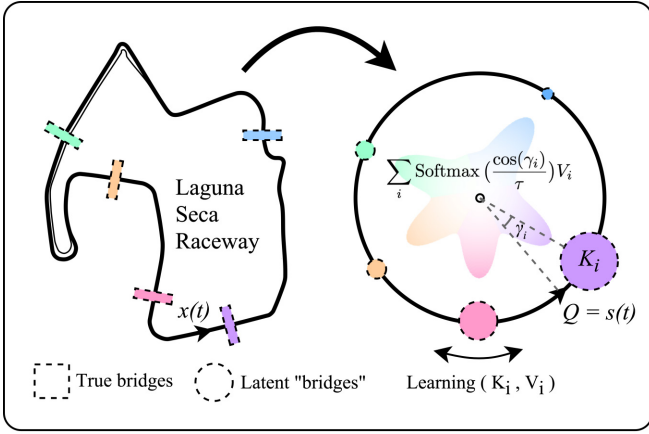


Fig. 3. Illustration of the robot's position $x(t)$, reparameterized into arc-length coordinates $s(t)$ and processed by a low-temperature attention mechanism to capture localized environmental features.

a unique [24] representation for positive definite matrices. Note that as exponential stability of $R(t)$ is maintained (Section III-G) we intentionally do *not* employ spectral normalization to bound the Lipschitz constant of the DNN. This is because it is desirable for $Q(t)$ to be sensitive to $x(t)$ and $g(t)$.

We parametrize the LDL factorization similar to [15], with $D_{ii} = \ln(1 + \exp(a_i^\top \phi))$ ensuring strict positivity and $L_{ij} = b_{ij}^\top \phi$ for $i > j$, where L is unit lower triangular and $a_i, b_{ij} \in \mathbb{R}^p$. This parameterization efficiently represents all positive definite matrices.

D. Smoothness through LACE

In the LACE framework, we model the dynamics of the measurement noise covariance using the Lyapunov differential equation:

$$\dot{R}(t) = A(t)R(t) + R(t)A(t)^\top + Q(t), \quad (6)$$

where $A(t) \in \mathbb{R}^{n \times n}$ is a transition matrix governing the covariance dynamics (e.g., $A(t) = \frac{\partial f(x,t)}{\partial x}$), $R(t) \in \mathbb{R}^{n \times n}$ represents the *measurement* noise covariance matrix that we wish to estimate, rather than the state estimation covariance. In this formulation, $Q(t) \in \mathbb{R}^{n \times n}$, serves as the process noise matrix that drives the evolution of $R(t)$ and is obtained from the core model introduced in Section III-C. In practice, $A(t)$ can be designed to encode a desired noise dissipation process or learned as a data-driven initialization for deployment tuning (Section IV-E).

Assumption 1 (Positive definiteness): By definition, the covariance matrix must be positive semi-definite. In this work, we additionally constrain $R(t) \in \mathbb{S}_{++}^n$ for all $t \geq 0$ to allow for the well-posedness of (2). This property can be guaranteed by ensuring $Q(t) \succ 0$ and $R(0) \succ 0$.

We show that imposing a spectral constraint on $A(t)$ enforces smoothness as defined in (2).

Theorem 1: Let $A(t) \in \mathbb{R}^{n \times n}$ be time-varying with eigenvalues $\{\lambda_i(t) \in \mathbb{C}\}_{i=1}^n$ and let $R(t) \in \mathbb{S}_{++}^n$ be differentiable. If the eigenvalues of $A(t)$ are constrained such that for all

$t \geq 0$, there is a $\lambda_{\min} < 0$, such that:

$$\lambda_{\min} \leq \Re(\lambda_i(t)) < 0, \quad \forall i \in \{1, \dots, n\}, \quad (7)$$

then choosing $\lambda_{\min} \geq -r_{\max}/(2n)$ satisfies the smoothness constraint from (2).

Proof: Using the definition of the smoothness constraint in (2) and Assumption 1, we write

$$\frac{d}{dt} \log \det R(t) = \text{tr}(R(t)^{-1} \dot{R}(t)). \quad (8)$$

Substituting $\dot{R}(t)$ in (8) with the dynamics in (6) yields:

$$\frac{d}{dt} \log \det R(t) = 2 \text{tr}(A(t)) + \text{tr}(R(t)^{-1} Q(t)).$$

Because of the positive definiteness of R and Q , the term $\text{tr}(R^{-1}Q)$ is positive, yielding the bound:

$$\frac{d}{dt} \log \det R(t) \geq 2 \text{tr}(A(t)) = 2 \sum_{i=1}^n \lambda_i(t) \geq 2n\lambda_{\min}.$$

Substituting $\lambda_{\min} = -r_{\max}/(2n)$ gives the final claim. ■

Theorem 1 therefore allows us to bound the contraction rate of the covariance ellipsoid by imposing a spectral constraint on the eigenvalues of $A(t)$.

E. Data-efficient Implementation

In practice, (6) must be discretized at the GNSS sampling period Δ_t , yielding the discrete-time Lyapunov recursion:

$$R_{k+1} = A_{d,k} R_k A_{d,k}^\top + Q_{d,k}, \quad (9)$$

where $(A_{d,k}, Q_{d,k})$ denote the discretized system and process noise matrices.

Assumption 2 (Time-invariant diagonalizable dynamics): We assume that A is time-invariant and diagonalizable as $A = U\Lambda U^{-1}$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ and $V \in \mathbb{C}^{n \times n}$ is an invertible matrix. The eigenvalues $\{\lambda_i\}_{i=1}^n$ are learned parameters and satisfy the constraint specified in (7).

Unrolling the discrete recurrence of (9) yields:

$$R_t = A_d^t R_0 (A_d^t)^\top + \sum_{j=0}^{t-1} A_d^j Q_{d,t-1-j} (A_d^j)^\top. \quad (10)$$

Let $\tilde{R}_t = U^{-1} R_t U^{-\top}$ and $\tilde{Q}_t = U^{-1} Q_t U^{-\top}$ and ignoring the initial condition R_0 that contracts exponentially as shown in Section III-G, we can rewrite (10) as:

$$\tilde{R}_t = \sum_{j=0}^{t-1} \Lambda^j \tilde{Q}_{t-1-j} (\Lambda^j)^\top, \quad (11)$$

since Λ is diagonal, each entry of the covariance matrix evolves independently as:

$$[\tilde{R}_t]_{ij} = \sum_{k=0}^{t-1} (\lambda_i \lambda_j)^k [\tilde{Q}_{t-1-k}]_{ij}. \quad (12)$$

This convolutional structure eliminates the sequential dependency in the recurrence in (9), and is well-suited for GPU hardware since it is parallelizable using the Fast Fourier Transform (FFT) with zero-padding, similar to the S4 state-space-model architecture [25], where a recurrent state update is reformulated as a parallelizable convolution.

F. Loss Function

Given residuals $\epsilon(t)$ and inferred covariance $R(t)$, we can define the loss function as the Negative Log-Likelihood of the assumed heteroscedastic Gaussian model from (1) and incorporate an additional loss term to control the smoothness of the inferred measurement covariance, so that:

$$\mathcal{L}(t) = \mathcal{L}_{\text{NLL}}(t) + \lambda \mathcal{L}_{\text{smooth}}(t), \quad (13)$$

where $\lambda > 0$ and

$$\mathcal{L}_{\text{NLL}}(t) = \sum_{k=1}^T \log |R(t)| + \epsilon(t)^\top R(t)^{-1} \epsilon(t). \quad (14)$$

In addition, we use an additional loss term that penalizes violations of (2) through a squared hinge loss:

$$\mathcal{L}_{\text{smooth}}(t) = \min\left(0, r_{\max} + \frac{d}{dt} \log \det R(t)\right)^2. \quad (15)$$

Note that this penalty is inactive in the full LACE architecture since we have shown in Theorem 1 that $\frac{d}{dt} \log \det R(t) \geq -r_{\max}$ by design. This penalty is nonetheless useful when comparing LACE against other baseline models for which smoothness is not guaranteed.

G. Contraction Stability Analysis

Consider the covariance propagation dynamics in (6), where $R \in \mathbb{S}_{++}^n$ is the positive definite covariance matrix, the time-varying $A(t) \in \mathbb{R}^{n \times n}$ is exponentially stable (contracting) with $\mu = -\frac{1}{2} \lambda_{\max}(A + A^\top) > 0$, and $Q(t) \in \mathbb{S}_{++}^n$ is the time-varying positive definite process noise matrix learned from environmental features. We can also assume the learned process noise Q_t is uniformly bounded: $\|Q_t\|_F \leq Q_{\max}$ for all $t \geq 0$.

Theorem 2: The covariance dynamics (6) is incrementally exponentially stable. Any two solutions $R_1(t)$ and $R_2(t)$ with different initial conditions satisfy:

$$\|R_1(t) - R_2(t)\|_F \leq \|R_1(0) - R_2(0)\|_F e^{-2\mu t}, \quad (16)$$

where $\mu = -\frac{1}{2} \lambda_{\max}(A + A^\top) > 0$.

Proof: Let $R_1(t)$ and $R_2(t)$ be two solutions of (6) with different initial conditions. Define the error $\Delta R = R_1 - R_2$. The error dynamics are:

$$\frac{d\Delta R}{dt} = A\Delta R + \Delta R A^\top. \quad (17)$$

Note that Q_t cancels as it is identical for both trajectories at time t . Also, $\Delta R^\top = \Delta R$ (as the difference of two symmetric matrices). Consider the squared Frobenius norm $V = \|\Delta R\|_F^2 = \text{tr}(\Delta R^\top \Delta R) = \text{tr}(\Delta R^2)$. Taking the time derivative:

$$\dot{V} = 2\text{tr}((A + A^\top)\Delta R^2) \leq -4\mu V, \quad (18)$$

since A is contracting with $A + A^\top$ has all negative eigenvalues.



Fig. 4. Dense LiDAR map of the Laguna Seca Raceway from running offline FGO based SLAM framework with color coded by altitude.

Using the comparison lemma yields $V(t) \leq V(0)e^{-4\mu t}$. With a properly scaled bounded numerical error $\sup_t d(t) = \bar{d}$, we find an exponentially converging error bound:

$$\|\Delta R(t)\|_F \leq \|\Delta R(0)\|_F e^{-2\mu t} + \frac{\bar{d}(1 - e^{-2\mu t})}{2\mu}, \quad (19)$$

which indicates the size of the error decreases with larger μ properly selected for fast convergence of $R(t)$ [19]. ■

IV. RESULTS

We implemented LACE on an AV-24 autonomous racecar and conducted multiple experiments at the Laguna Seca Raceway [2]. Our evaluation strategy leverages a high-fidelity, offline pipeline to generate ground-truth state estimates from data collected during on-track experiments. We use a LiDAR-based simultaneous localization and mapping (SLAM) offline framework to produce accurate trajectories. This pseudo-ground-truth enables a rigorous assessment of our lightweight, online GNSS covariance learning model. The following sections present our baselines, evaluation metrics, and the performance of our method integrated into the vehicle's localization stack by inferring a $R(t) \in \mathbb{R}^{3 \times 3}$ GNSS covariance matrix.

A. Datasets

We collected sensor data from multiple practice sessions at Laguna Seca Raceway using an industrial-grade sensor suite. This setup included Luminar Iris LiDARs (20 Hz), NovAtel OEM7 GNSS receivers with RTK corrections (20 Hz), and a VectorNav VN310 IMU (200 Hz). Then, to establish a high-fidelity ground truth we employed an offline state estimation pipeline using GLIM [26], a factor-graph optimization (FGO) based SLAM framework that allows GPU-based dense LiDAR scan registration and optimal fusion with IMU and GNSS sensor measurements. To ensure robustness against signal degradation, GNSS measurements were only incorporated as factors in areas known to be free of interference with full RTK availability to reduce global drift. In open-sky segments, the SLAM trajectory agreed with RTK-fixed GNSS solutions to within 5 cm, supporting its use as a reference trajectory. The resulting dense map from this process is shown in Fig. 4. By running the framework on all datasets, we obtained the pseudo-ground truth state $x(t)$ for a given run.




Fig. 5. Train and evaluation loss of our method (LACE) and MLP, yellow stars indicate the lowest part of the evaluation loss. LACE shows a much faster training convergence rate.




Fig. 6. Guaranteed exponential convergence of covariance evolution for LACE, demonstrated by different initial covariances $R(0)$.

B. Training

To evaluate the effectiveness of our proposed dynamical framework, we compare LACE against a one-shot covariance prediction baseline using a multi-layer perceptron (MLP) that directly predicts the covariance matrix $R(t)$ at each timestep without explicitly modeling its temporal evolution. Its architecture is illustrated in Fig. 2 as the *Core Model* with Q being the estimated covariance matrix. To ensure a fair comparison that isolates the impact of our dynamic modeling, the MLP baseline is constructed with the same architecture and number of learnable parameters as the covariance head in our LACE model. Both models were trained for a fixed number of epochs to minimize the NLL objective. As illustrated in Fig. 5, our method not only achieves a better evaluation loss, but also converges significantly faster, demonstrating the benefits of explicitly modeling the covariance dynamics.

C. Exponential Convergence

Our theoretical analysis in Section III-G guarantees that the covariance dynamics are incrementally exponentially stable, making our LACE framework robust to initial conditions. This property ensures that for any arbitrary positive definite initial covariance $R(0)$, the system will converge to the same covariance trajectory dictated by the learned dynamics. We validate this theoretical result experimentally, as shown in Fig. 6. By initializing the system with a diverse set of covariance matrices, we demonstrate that all trajectories indeed converge exponentially to the same unique evolution, confirming the practical stability and predictability of our approach.

D. Track Prediction

Over the course of a lap at the Laguna Seca Raceway, our autonomous racecar encounters multiple pedestrian and vehicle bridges that severely degrade GNSS quality. To analyze the practical performance of LACE on the race track, we compared its predictions against two common empirical models in addition to the MLP baseline. The first empirical approach uses a *Constant* covariance, parameterized by a scalar such that $R_{\text{constant}} = cI$ with $c \in \mathbb{R}$ and $I \in \mathbb{R}^{3 \times 3}$ is the identity matrix. The second, a heuristic *Bubble* model, attempts to account for known GNSS degradation zones by linearly scaling c by the distance to the preset bridge locations with some additional fixed padding.

As illustrated in Fig. 7, our analysis reveals the limitations of the baseline methods relative to our solution. The *Constant* model, by design, cannot capture any environmental variations in GNSS quality. While the *Bubble* model provides a coarse approximation by increasing uncertainty near bridges (B0, B1, B2, and B3), its hand-tuned nature prevents it from capturing the true, time-varying signal quality. As a result, it fails to recognize the minimal degradation at another bridge (B2) during this particular lap, thereby sacrificing estimation performance. The purely adaptive MLP baseline, while responsive, produces covariance estimates that can be highly irregular subject to rapid changes. In contrast, our LACE model not only provides a better representation of the underlying residual distribution, especially in areas with erratic GDOP values, but does so while maintaining a significantly smoother covariance trajectory as a direct benefit of our explicit dynamic modeling.

E. Tunable Dynamic Matrix

A key advantage of LACE is the ability to not only learn an optimal system matrix A but also to manually tune its dynamics for specific operational needs. While the framework learns a set of eigenvalues, λ_i^* , that optimally minimize the loss function (Eq. 14), these values can be manually adjusted to alter the model's sensitivity. As demonstrated in Fig. 8, this provides a powerful mechanism for balancing statistical optimality with practical robustness. Although the learned eigenvalues λ^* yield the lowest average loss (Table I), they can occasionally underestimate uncertainty during moments of rapid signal degradation. By selecting a less negative

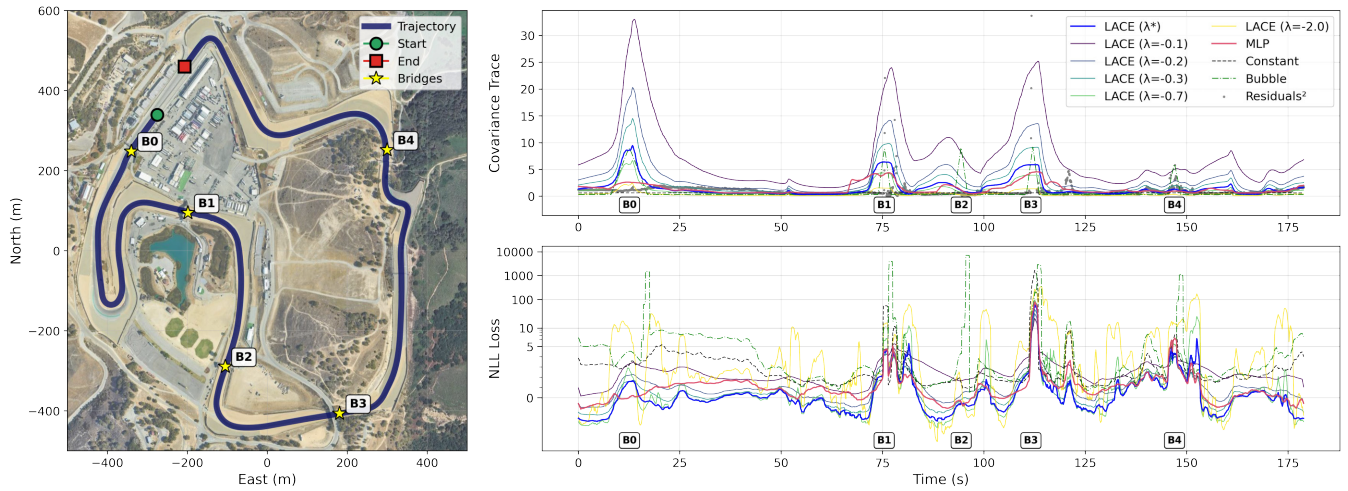


Fig. 7. Covariance model comparison. **Left:** satellite view of a practice lap at Laguna Seca Raceway, with bridges that potentially degrade GNSS measurements labeled. **Top right:** traces of the covariance matrices from different models. Our method successfully encloses the residuals while maintaining smoothness. **Bottom right:** negative log likelihood of evaluated models, LACE overall has the lowest loss suggesting a better estimation of the GNSS uncertainty.




Fig. 8. Effects of different choice of eigenvalues on the predicted covariance. λ^* denotes the learned value. Lower λ biases the learned process noise matrix $Q(t)$ more resulting overall more conservative $R(t)$.

eigenvalue (e.g., $\lambda = -0.1$), we can make the covariance evolution respond more aggressively to the learned process noise $Q(t)$. This results in a more conservative estimate that more consistently bounds the measurement residuals. This inherent flexibility is a significant practical benefit, allowing the system to be tuned for varying levels of environmental adversity and providing a balance between adaptivity and smoothness during deployment.

F. Estimator in the Loop

The Multisensor-aided Inertial Navigation System (MINS [27]) is a MSCKF based state estimation framework that is lightweight, robust, and supports diverse sensor configurations. We adopted a modified version of MINS as part of our localization stack. To further improve real-time performance, we converted our DNN model to TensorRT (TRT), enabling efficient inference on the onboard GPU of our racecar. In this section, we present the results of our

TABLE I. Average loss across multiple laps. While λ^* achieves the lowest average loss, a smaller λ (such as $\lambda = -0.1$) can help robustify the model performance by biasing towards the degradation signals picked up by $Q(t)$.

Model	Avg Loss	Std
LACE (λ^*)	2.9694	5.3995
LACE ($\lambda = -0.1$)	5.3642	0.5913
LACE ($\lambda = -0.2$)	4.4280	1.1541
LACE ($\lambda = -0.3$)	3.9281	1.7011
LACE ($\lambda = -0.7$)	3.1205	3.8769
LACE ($\lambda = -2.0$)	3.2145	11.8982
MLP	3.4167	3.2236
Constant	5.8491	33.0986
Bubble	4.2820	14.8187

method integrated within this state-of-the-art sensor fusion framework.

We evaluated our method across different areas of the track. Figure 9 illustrates estimator performance at Bridge 3 (B3 in Fig. 7), which introduces the strongest GNSS disturbances due to its size and material. LACE adapts to the highly degraded measurements and provides reliable uncertainty estimates to the EKF. In contrast, the other methods allow overconfident updates, leading to sharp deviations in state estimation which will eventually trigger unrecoverable behavior in the high-fidelity race controller. These results highlight the advantage of LACE to the downstream systems, particularly under extreme conditions.

V. CONCLUSION

We presented LACE, a method for estimating and adapting GNSS measurement covariance for high-speed autonomous systems. In particular, our method achieves superior environmental adaptivity while guaranteeing temporal smoothness for the downstream feedback controller. By modeling covariance evolution as a provably stable dynamical system with spectral constraints, our LACE framework provides a novel method for learning environment-aware uncertainty.

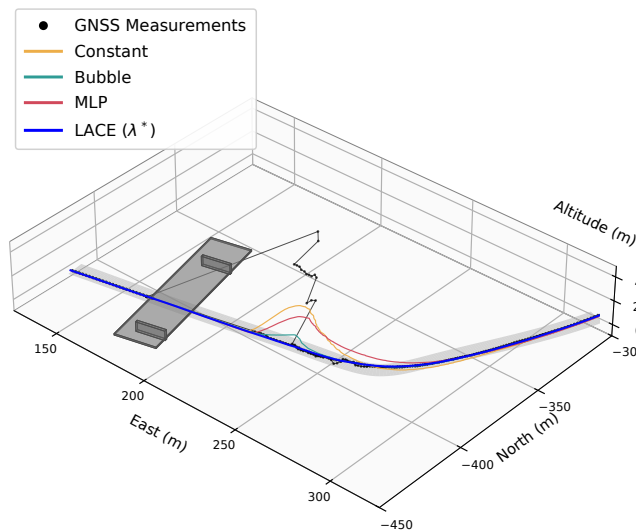


Fig. 9. Trajectories estimated from MINS with different covariance models. As the racecar passes under the bridge (illustrated by the gray slabs), the GNSS quality becomes strongly disturbed, as indicated by the spread of measurements. Our method, with an optimal λ , successfully captures the rapid changes in noise distribution, whereas other methods produce overconfident covariance estimates that mislead the estimator.

We also derived formal guarantees of exponential stability and smoothness bound for the covariance dynamics, ensuring that the resulting estimates are smooth by design. Through validation with data from an autonomous racecar in a real-world racing environment, we demonstrated the effectiveness of our framework in producing smoother and more accurate covariance predictions in challenging, GNSS-degraded conditions.

For future work, the extension of our dynamical modeling approach to other sensing modalities, such as LiDAR and vision, could provide a unified framework for multi-sensor covariance estimation. Furthermore, moving beyond Gaussian noise assumptions to handle arbitrary distributions (such as using flow models) would further broaden its capability to diverse sensing configurations with higher granularity. Finally, exploring the explicit co-design of the covariance model and the downstream vehicle controller could lead to new paradigms in integrated state estimation and control, further enhancing the stability and performance of safety-critical autonomous systems.

VI. ACKNOWLEDGMENT

The authors thank R. Blanchard, M. Anderson, X. Zuo, P. Spieler, J. Lathrop, J. Cho, N. Ranganathan, and the other members of the Caltech Racer team.

REFERENCES

- [1] T. D. Barfoot, *State Estimation for Robotics*, 2nd ed. Cambridge University Press, 2024.
- [2] Indy Autonomous Challenge, “AV-24 Racecar.” [Online]. Available: <https://www.indyautonomouschallenge.com/racecar>
- [3] S. Akhlaghi, N. Zhou, and Z. Huang, “Adaptive adjustment of noise covariance in Kalman filter for dynamic state estimation,” in *2017 IEEE Power Energy Soc. Gen. Meeting*, July 2017, pp. 1–5.

- [4] E. S. Lohan and K. Borre, “Accuracy limits in multi-GNSS,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 52, no. 5, pp. 2477–2494, Oct. 2016.
- [5] K. Liu, K. Ok, W. Vega-Brown, and N. Roy, “Deep Inference for Covariance Estimation: Learning Gaussian Noise Models for State Estimation,” in *2018 IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 1436–1443.
- [6] F. Wu, H. Luo, H. Jia, F. Zhao, Y. Xiao, and X. Gao, “Predicting the Noise Covariance With a Multitask Learning Model for Kalman Filter-Based GNSS/INS Integrated Navigation,” *IEEE Trans. Instrum. Meas.*, vol. 70, pp. 1–13, 2021.
- [7] I. Smolyakov, M. Rezaee, and R. B. Langley, “Resilient multipath prediction and detection architecture for low-cost navigation in challenging urban areas,” *Navigation*, vol. 67, no. 2, pp. 397–409, June 2020.
- [8] W. Vega-Brown, A. Bachrach, A. Bry, J. Kelly, and N. Roy, “CELLO: A fast algorithm for Covariance Estimation,” in *2013 IEEE Int. Conf. Robot. Autom. (ICRA)*, Karlsruhe, Germany, May 2013, pp. 3160–3167.
- [9] D. Landry, F. Pomerleau, and P. Giguère, “CELLO-3D: Estimating the Covariance of ICP in the Real World,” in *2019 Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 8190–8196.
- [10] W. Vega-Brown and N. Roy, “CELLO-EM: Adaptive sensor models without ground truth,” in *2013 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Nov. 2013, pp. 1907–1914.
- [11] A. H. Mohamed and K. P. Schwarz, “Adaptive Kalman Filtering for INS/GPS,” *J. Geod.*, vol. 73, no. 4, pp. 193–203, May 1999.
- [12] R. Mehra, “On the identification of variances and adaptive Kalman filtering,” *IEEE Trans. Automat. Contr.*, vol. 15, no. 2, pp. 175–184, Apr. 1970.
- [13] D. Magill, “Optimal adaptive estimation of sampled stochastic processes,” *IEEE Trans. Automat. Contr.*, vol. 10, no. 4, pp. 434–439, Oct. 1965.
- [14] K.-W. Chiang, T. T. Duong, J.-K. Liao, Y.-C. Lai, C.-C. Chang, J.-M. Cai, and S.-C. Huang, “On-line smoothing for an integrated navigation system with low-cost mems inertial sensors,” *Sensors*, vol. 12, no. 12, pp. 17 372–17 389, 2012.
- [15] H. Hu and G. Kantor, “Parametric covariance prediction for heteroscedastic noise,” in *2015 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Hamburg, Germany, Sept. 2015, pp. 3052–3057.
- [16] P. D. Hoff and X. Niu, “A Covariance Regression Model,” *Statistica Sinica*, vol. 22, no. 2, Apr. 2012.
- [17] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard, “Most likely heteroscedastic Gaussian process regression,” in *Proc. 24th Int. Conf. Mach. Learn. (ICML)*, June 2007, pp. 393–400.
- [18] W. Lohmiller and J.-J. E. Slotine, “On contraction analysis for nonlinear systems,” *Automatica*, vol. 34, no. 6, pp. 683–696, June 1998.
- [19] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, “Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview,” *Annu. Rev. Control*, vol. 52, pp. 135–169, 2021.
- [20] A. I. Mourikis and S. I. Roumeliotis, “A Multi-State Constraint Kalman Filter for Vision-aided Inertial Navigation,” in *2007 IEEE Int. Conf. Robot. Autom. (ICRA)*, Apr. 2007, pp. 3565–3572.
- [21] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, “OpenVINS: A Research Platform for Visual-Inertial Estimation,” in *2020 IEEE Int. Conf. Robot. Autom. (ICRA)*, Paris, France, May 2020, pp. 4666–4672.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 30, 2017.
- [23] J. He, C. Zhang, F. Jenelten, R. Grandia, M. Bächer, and M. Hutter, “Attention-based map encoding for learning generalized legged locomotion,” *Sci. Robot.*, vol. 10, no. 105, p. eadv3604, 2025.
- [24] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed. Soc. Ind. Appl. Math., Jan. 2002.
- [25] A. Gu, K. Goel, and C. Ré, “Efficiently modeling long sequences with structured state spaces,” in *Int. Conf. Learn. Represent. (ICLR)*, 2022.
- [26] K. Koide, M. Yokozuka, S. Oishi, and A. Banno, “GLIM: 3D range-inertial localization and mapping with GPU-accelerated scan matching factors,” *Robot. Auton. Syst.*, vol. 179, p. 104750, Sept. 2024.
- [27] W. Lee, P. Geneva, C. Chen, and G. Huang, “Mins: Efficient and robust multisensor-aided inertial navigation system,” *J. Field Robot.*, vol. 42, no. 7, pp. 3252–3284, 2025.