ICRA 2026 · VIENNA · IEEE INT'L CONFERENCE ON ROBOTICS AND AUTOMATION

Environment-aware learning of smooth GNSS covariance dynamics for autonomous racing

Y. Deemo Chen† · Arion Zimmermann† · Thomas A. Berrueta · Soon-Jo Chung

California Institute of Technology — Department of Computing and Mathematical Sciences · † Equal contribution

01 *The dual challenge*

Autonomous racing at **>150 mph** demands centimeter-level localization. But as the racecar passes under concrete bridges or near buildings, GNSS noise jumps from **<1 m** to **>50 m** in less than a second.

Adaptive
React fast to environmental disturbances so the filter is never over-confident in degraded GNSS.

Smooth
Slow recovery in covariance so the downstream controller is never destabilized by uncertainty jumps.

The gap. Predictive DNNs are adaptive but discontinuous. Adaptive filters are smooth but reactive. Both treat covariance as an output, instead of a dynamical system.

02 *Architecture*

A DNN ingests the estimated arc-length state $\hat{s}(t)$ and GNSS quality factors $g(t)$ (DOP, satellite count) and outputs a symmetric positive-definite process noise matrix Q via an LDL decomposition.

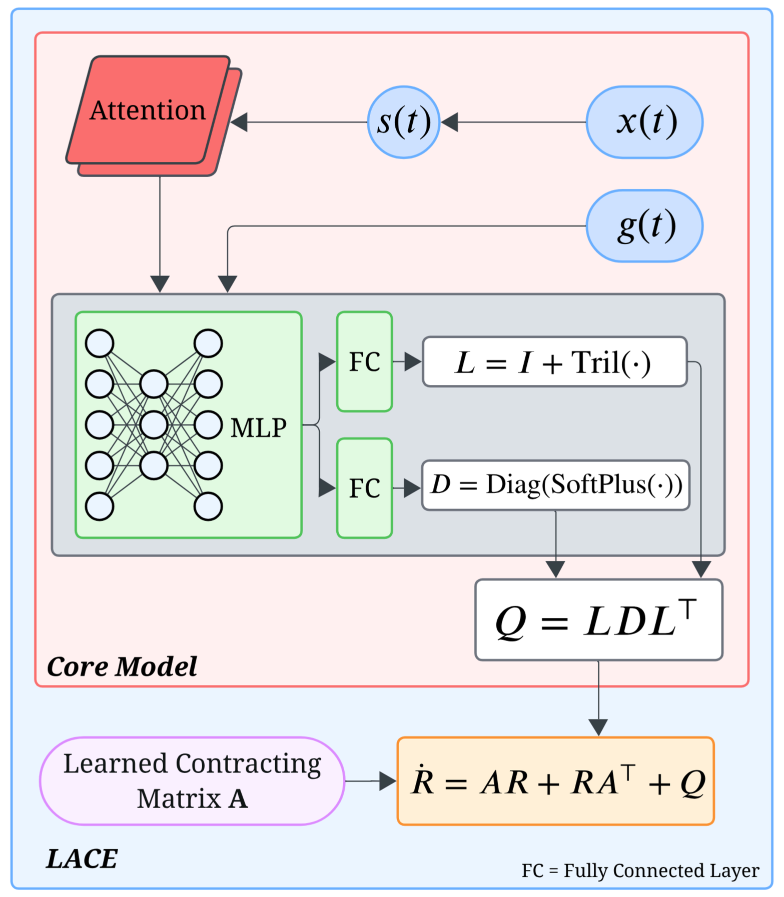

Fig. 2. Attention encodes the arc-length position; an MLP predicts L and D so that $Q = LDL^\top \succ 0$ by construction. LACE wraps Q inside a contracting Lyapunov system.

03 *Contributions*

1. A new architecture. LACE unifies environmental adaptation with tunable conservativity by treating covariance as a learned dynamical system.
2. ***Formal guarantees.*** Contraction analysis yields *exponential stability* and a spectral-constraint smoothness bound on $\frac{d}{dt}\log\det R(t)$.
3. ***Real-world validation.*** Deployed on an *AV-24 IAC racecar* at Laguna Seca; smoother and more accurate covariance under GNSS-degraded passes.

04 *Circular attention*

Position is reparameterized into arc-length $s(t)$ and encoded as a unit-length complex query, capturing localized features along the track.

$$\tilde{s} = W\left(V^\top \mathrm{Softmax}\left(\tfrac{1}{\tau}\Re(Q^* K)\right)\right)$$

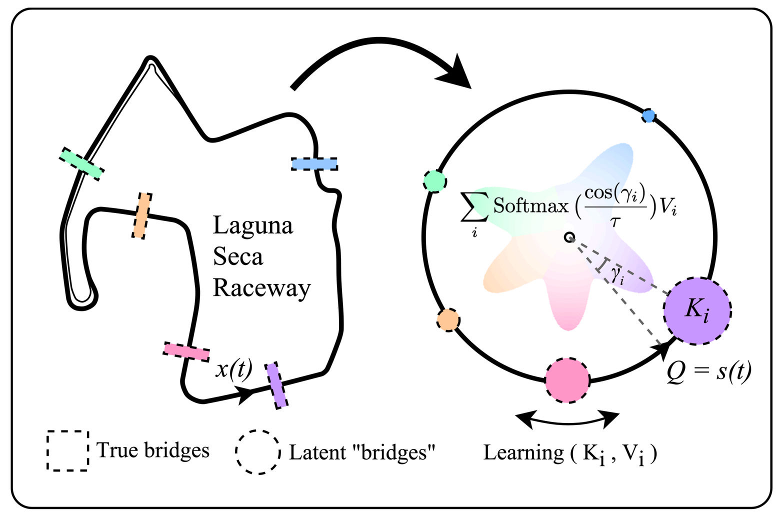

Fig. 3. Low temperature τ sharpens attention onto sparse environmental features (bridges, structures).

05 *Offline optimized GT*

Fig. 4. Dense LiDAR map of Laguna Seca through FGO-based offline SLAM pipeline.

LACE · *Learning Adaptive Covariance Evolution*

Learning Environment-Adaptive *Uncertainty Dynamics.*

A learned Lyapunov system that is *smooth by design* and *environment-aware* — producing uncertainty the controller can actually trust at racing speed.

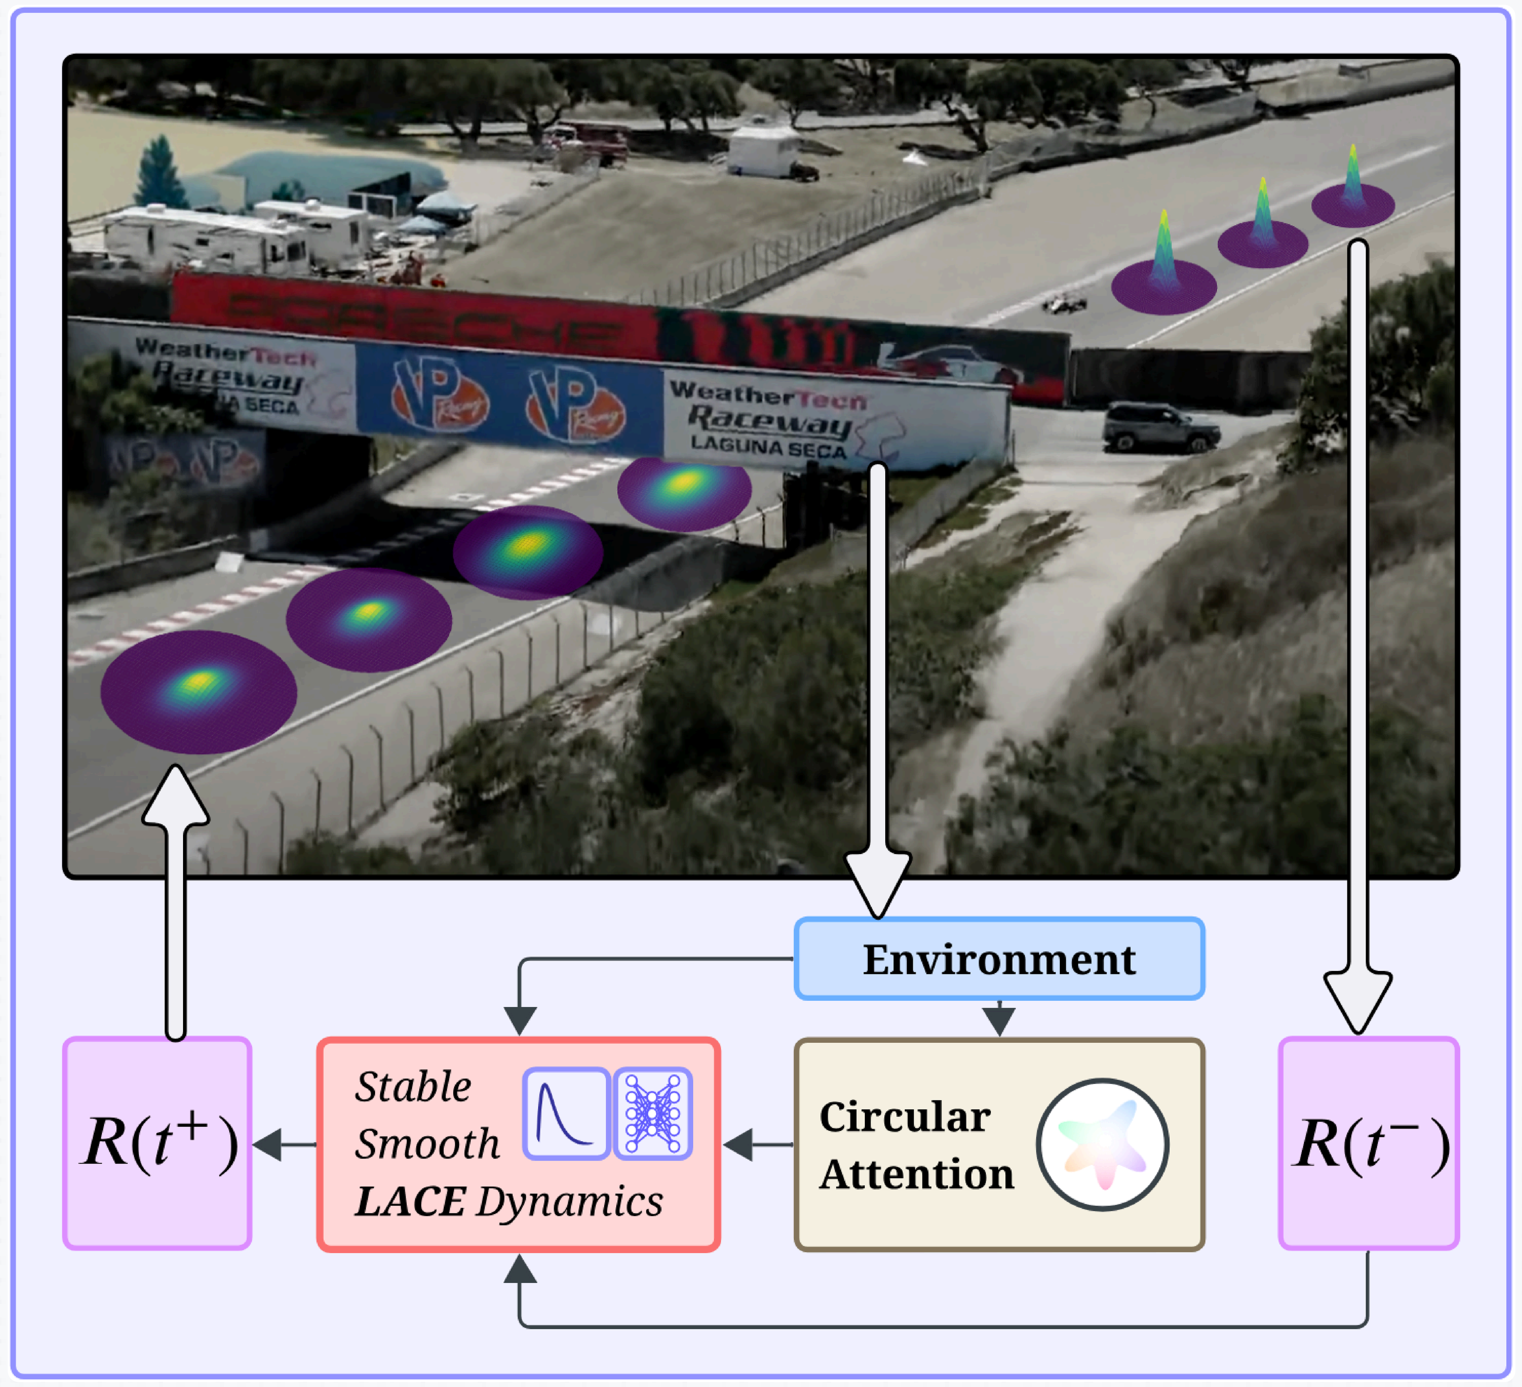

Fig. 1. GNSS covariance evolves smoothly as the racecar approaches and clears a concrete bridge: environmental cues (x(t),g(t)) drive a stable, contracting dynamical model of R(t).

Platform Caltech AV-24 Indy Autonomous Challenge racecar.

06 *Theory*

Model the measurement covariance as a Lyapunov differential equation driven by the learned process noise $Q(t)$:

Covariance dynamics

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

A is learned and spectrally constrained; $Q(t)$ is environmentally driven.

THEOREM 1 · SMOOTHNESS BY DESIGN

Guaranteed smoothness

If $\Re\lambda_i(A) \in [\lambda_{\min}, 0)$ for all i and

$$\lambda_{\min} \geq -\tfrac{r_{\max}}{2n},$$

then the log-volume of the covariance ellipsoid contracts no faster than $r_{\max}$:

$$\tfrac{d}{dt}\log\det R(t) \geq -r_{\max}.$$

Why it matters. Smoothness is enforced by a single spectral constraint on A, instead of a soft penalty.

THEOREM 2 · STABILITY BY CONTRACTION

Exponential convergence

With contraction rate $\mu = -\frac{1}{2}\lambda_{\max}(A + A^\top) > 0$, any two solutions satisfy

$$\|\Delta R(t)\|_F \leq \|\Delta R(0)\|_F\, e^{-2\mu t}.$$

Why it matters. Disturbances and initial-condition variances decay exponentially.

07 *Tunable conservativity*

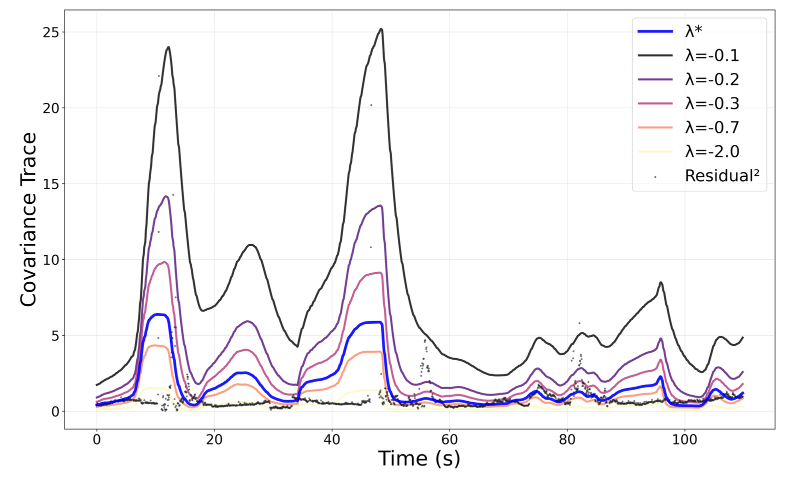

Fig. 8. The learned λ^* minimizes loss; making λ less negative biases LACE to be more conservative.

$\lambda \longrightarrow$ Conservativity Inference time spectral knob. Adjust on-track during deployment.

08 *Results*

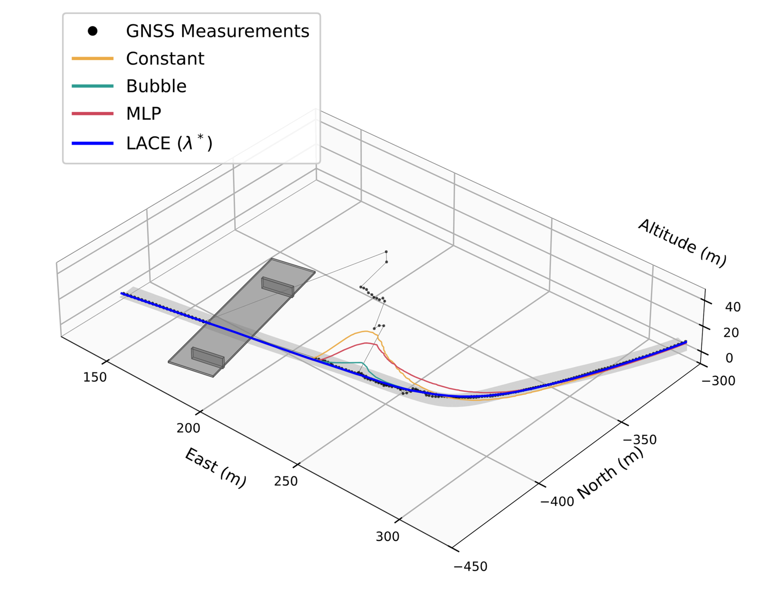

Fig. 9. 3D estimator trajectories under Bridge B3. LACE absorbs disturbances while baselines drift.

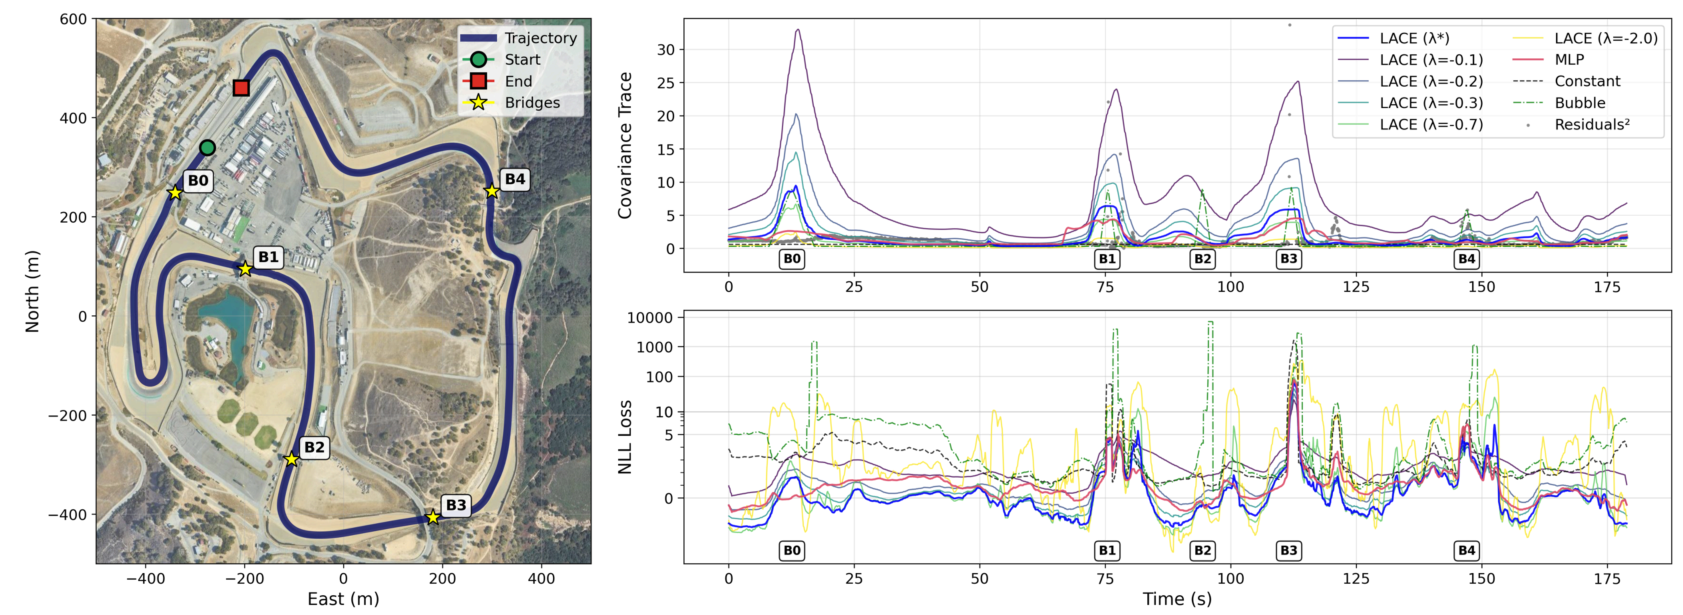

Fig. 7. Covariance traces (top) and NLL loss (bottom) across a lap at Laguna Seca.

ACKNOWLEDGMENTS

The authors thank R. Blanchard, M. Anderson, P. Spieler, J. Lathrop, J. Cho, N. Ranganathan, and the Caltech IAC team.
This project was in part funded by *Beyond Limits*.

arXiv paper
arxiv.org/abs/2602.21366

Video
youtu.be/fkOKFMRtwyk