



Riemannian Geometry-Based Structural Stability Prediction in Dynamic Vehicle Systems

Sk. Rukshana Begum ¹ , Mesfin Dagne ²

¹ Sushila Devi Bansal College of Engineering, Indore, Madhya Pradesh , India-453331; brukshanaece@gmail.com

² Department of CSE, Adama Science and Technology University, Adama, Ethiopia, mesfindagne@gmail.com

* Corresponding Author: Sk. Rukshana Begum ; brukshanaece@gmail.com

Abstract: Vehicles with modern high performance run with complex multi-sensor data-streams under extreme aerodynamic and structural loading where there is a lack of accurate modelling in conventional linear state-space. This paper presents an extended Riemannian geometry framework for structural stability prediction in multi-sensor dynamic vehicle systems using an embedding of multi-sensor measurements onto a curved Riemannian manifold which better retains the true nature of the dynamic system behavior than the Euclidean distance assumption. The framework combines the synchronized acquisition of the sensors, projection of symmetric positive definite (SPD) covariance matrices for the manifolds, calculation of the geodesic distance, calculation of curvature using Christoffel symbols, calculation of a composite instability index (Ψ), which is based on the curvature, geodesic divergence and velocity-weighted contributions. The system is run on an instrumented Formula type car at 1000 Hz on all the IMUs and 240 fps in all speed zones ranging from 80 to 340 km/h. Experimental comparison to the linear, Kalman filter and LSTM baselines shows that the proposed Riemannian framework achieves the highest structural state prediction accuracy of 96%, 3% higher than the LSTM baseline, the shortest prediction lead time (182 ms), and the lowest variance ($\sigma = 0.03$) among the three approaches. The findings pave the way towards a mathematically robust, yet practically applicable basis for next generation structural health monitoring for safety-critical applications in vehicle structures.

Keywords: Structural Stability, Differential Geometry, Sensor Fusion, SPD Manifold, Instability Index.

1. Introduction

In today's high performance vehicles, which have extreme structural and aerodynamic loads, the data streams originating from one or more sensors are multi-sensor, nonlinear, and complex and are difficult to accurately model with traditional linear state-space techniques. This paper proposes a Riemannian geometry framework for the prediction structural stability on a dynamic vehicle system based on embedding multi-sensor measurements into a curved Riemannian manifold where the geodesic relationships between sensors will better capture the actual system dynamics than Euclidean-distance based assumptions. The framework is synchronized sensor acquisition, manifold projection of symmetric positive definite (SPD) covariance matrices, computation of geodesic distance, analysis of curvature by means of Christoffel symbols, a composite instability index Ψ which combines, via weighted contributions, the curvature data and geodesic divergence and the velocity data.

The system is tested on the simulated Formula-version car equipped with IMUs sampling up to 1000 Hz and cameras sampling at 240 fps over speeds ranging from 80 to 340 km/h.

When evaluated against a linear baseline, Kalman filter and LSTM baselines, the proposed Riemannian framework not only provides the highest structural state prediction accuracy (96%) among all approaches, but also yields the lowest variance ($\sigma = 0.03$) and the prediction lead time of 182 ms, which is the smallest value among all approaches. These results provide a firmer mathematical basis for and a pragmatic framework for next-generation structural health monitoring in safety-critical vehicle applications based on Riemannian geometry. The current system has one drawback that is it depends on supervised training labels that are generated by human domain experts, and this leads to label noise and limits the system to types of faults that are considered when constructing the training set. The following deployment scenarios might encounter new fault modes not present in the training distribution, which might cause the classifier to give inaccurate predictions. One way to reduce this risk would be to incorporate uncertainty quantification techniques, specifically Monte Carlo dropout or deep ensembles, which would highlight out-of-distribution inputs that would need human review to correct, rather than giving more negative effects of over confident misclassifications. The integrated cross-



entropy classification loss with the differentiable latency penalty is a principled method for the accuracy-latency co-optimisation. The method makes it easier for the model to learn network architectures that provide a better Pareto front between the two competing objectives of latency and the accuracy of its fitness function, as compared to random architecture search.

2. Literature Survey

The structural health monitoring for ground vehicles and aerospace structures has been decades-long, and is one of the applications that has been integral to vibration analysis and modal identification. Doebling et al. [1] showed that frequency shifts and changes in mode shapes of structural vibrations can confidently detect damage accumulation in composite and metallic structures. These frequency domain techniques, however, rely on the assumption of time-invariance of the structural characteristics of interest under monitoring and cannot be applied to the analysis of structural response to a change in operating conditions, which happens many times during a competitive motorsport event. The incorporation of probabilistic state estimation and sensor noise modeling in the SHM pipeline was made possible by the introduction of the fusion architectures based on Kalman filters, which significantly improved the results. Kalman's pioneering research on optimal linear estimation, [3], set the stage for a class of fault detection systems that are able to combine multiple, possibly different, sensor data streams and propagate uncertainty estimates.

The Extended Kalman filter (EKF) and the Unscented Kalman filter (UKF) expanded the ability of the Kalman filter to situations with mildly-nonlinear dynamics; however, underlying assumptions of using linear state transitions and Gaussian noise models become increasingly restrictive in highly-nonlinear high speed systems. Today, Zarchan and Musoff present a thorough treatment of these drawbacks [9]. The use of machine learning methods greatly improved the predictability other than through explicit physical models. Learning-sequential sensor data, especially the dependencies between sequential data, is done with great success by recurrent neural networks, such as long short-term memory (LSTM) networks of Hochreiter and Schmidhuber [2]. Later, Bahdanau et al. showed how attention can more recently be used to model sequences by selectively focusing on the steps that are relevant for their diagnosability in the task [10].

The latter, however, are all in Euclidean feature space and so still suffer from the "geometric misrepresentation" mentioned above. Alongside this research direction lie graph-based methods, and kernel methods, both of which have some significance when it comes to inter-sensor dependency modeling. In reference [8] sparse Gaussian graphical models (SGM) were used to learn the conditional independence structure of the multi modal correlations of the telemetry data of a vehicle, thus obtaining compact models. These approaches would build the graph structure statically at training time, not suitable for such changes in graph topology that can result from structural damage or changes in configuration, while the Riemannian manifold gracefully handles the changes of Graph covariance geometry without the need for redesigning the Graph. Given the non-Euclidean nature of structured data, differential geometry and manifold learning have proven themselves to be valuable tools. Lee gives the basic tools for computing geodesics, curvature and

parallel transport on differentiable manifolds of any dimension in his book on smooth manifolds [4]. The space of symmetric positive definite matrices (SPD manifold), with the affine-invariant metric, has been used to great advantage in brain-computer interface (BCI) classification [5], computer vision [6] and medical imaging [7]. In a test case involving simulated data with strong nonlinear covariance structure, Pennec et al. clearly showed that Euclidean techniques provide less consistent estimates than Riemannian techniques. The present work is the first one to implement the complete Riemannian manifold analysis in ground vehicles for structural health monitoring applications operating in high-speed conditions.

3. Research Gap and Existing System

Based on a structured literature search three gaps are identified that inspire the present work. Review of previous literature reveals three gaps, which inspired the present study. Firstly, although it is well known that vehicle structures have nonlinear behaviours, the vast majority of vehicles currently using SHM still employ linear state space models or Euclidean-distance based anomaly detectors. This establishes a systematic deviation between the geometry of data and the geometry of the processing algorithm, influencing in a negative way the accuracy exactly in the high-speed high-load operating regimes where it is most safety-critical. Secondly, although manifold learning methods like Isomap, locally linear embedding (LLE) and diffusion maps have been used for standard dimensionality reduction problems, they have not yet been used in the uncertainly quantification framework of safety-critical real-time monitoring applications. Riemannian statistics on the SPD manifold, on the other hand, offer closed-form expression for Fréchet means and the SPD covariance operators, which are completely compatible with the probabilistic inference. Third, currently used vehicle telemetry fusion systems operate each sensor modality independently or perform the linear combination of all the modalities with static fusion weights for them. The proposed weighted fusion in this paper is geodesic, meaning that the relative importance of multiple modalities can be dynamically adjusted based upon the structural state of the vehicle, without fixed weights; this provides a more flexible, physically-motivated model than the fixed-weight models. The present work fills in all three gaps simultaneously in an integrated Riemannian framework.

4. Proposed System

The proposed system are in the process of designing a seven-stage pipeline from the raw information of the sensors which lead to real-time stability classification.




Fig. 1 Proposed Riemannian geometry-based structural stability prediction architecture.

Fig. 1 provides an illustration of the architecture, with the modalities of available sensors as entry point and the operator dashboard as exit. The architecture does not return to the operator, but takes the data through the pre-processing and the feature extraction stages, shown as a black box in between. Manifold projection is then performed, and the data is passed to a geodesic analysis, which in turn produces a stability prediction.

4.1. Mathematical Foundation

The primary mathematical object is the Riemannian manifold (M, g) which is M regarded as the space of $n \times n$ symmetric positive definite matrices, equipped with the affine-invariant metric tensor, g . The inner product of two vectors u, v at P in the tangent space $T_P M$ given the metric tensor is:

$$g_i^N = \langle \partial\phi/\partial x_i, \partial\phi/\partial x^N \rangle$$

The embedding map, ϕ from the parameter space to the ambient space. Geodesic equation for shortest path curves on M is:

$$d^2x^R/dt^2 + \Gamma_{iN}^R (dx_i/dt)(dx^N/dt) = 0$$

Γ^{RO} are the “Christoffel symbols of 2nd kind”, which tell us how the coordinate basis vectors change as we move around the manifold. These symbols are dependent on the metric tensor and its first derivatives and the non-vanishing ones define the intrinsic curvature of the manifold. The Christoffel symbols of the second kind can be calculated explicitly from the partial derivatives of the components of the metric tensor, and are given by:

$$\Gamma_{iN}^R = \frac{1}{2} g^{Rl} (\partial g^l/\partial x^N + \partial g^{lN}/\partial x_i - \partial g^{iN}/\partial x^l)$$

where g is the inverse metric tensor, and hence g^{Rl} is its components. On the SPD manifold with the affine-invariant metric these symbols have closed form expressions that can be numerically evaluated efficiently without resorting to finite differences. Let $\text{Exp}_P: T_P M \rightarrow M$ be the exponential map, and let $\text{Log}_P: M \rightarrow T_P M$ be the inverse of Exp_P . Let $\text{Exp}_P: T_P M \rightarrow M$ be the exponential map, and let $\text{Log}_P: M \rightarrow T_P M$ be the inverse of Exp_P . For the SPD manifold, these maps are given by

$$\text{Exp}_P(v) = P^{1/2} \exp(P^{-1/2} v P^{-1/2}) P^{1/2}$$

and

$$\text{Log}_P(Q) = P^{1/2} \log(P^{-1/2} Q P^{-1/2}) P^{1/2},$$

respectively, where $\exp$ and $\log$ denote the matrix exponential and logarithm.

4.2. Composite Instability Index

The instability index Ψ provides a scalar summary of the vehicle's current structural risk level, combining contributions from three independent Riemannian descriptors: the curvature scalar C derived from the Riemann curvature tensor, the geodesic divergence G measuring separation rate between adjacent state trajectories, and the velocity-weighted vibrational power V . The composite index is:

$$\Psi = (\omega_1 C + \omega_2 G + \omega_3 V) / 3$$

4.3. Parallel Transport and Cross-Window State Comparison

One of the most basic problems in sequential state analysis of structures is the comparison of instantaneous rates of change of the states (tangent vectors) at different points on the manifold. In the rest of Euclidean space, vectors at different points are already

in the same space, so that this is quite easy: in the same manifold, such a vector subtraction is not geometrically correct since vectors at points can belong to different tangent spaces. The Principled solution is Riemannian parallel transport. Let P and Q be two points in M , and let $\gamma(t)$ be a curve from P to Q . A Riemannian parallel transport of a tangent vector v at P along γ brings v to a vector v' at Q with the same length and same angle relative to the metric of M . The parallel transport equation is:

$$D/dt (\tau(t)) = \nabla_{\gamma'(t)} \tau(t) = 0$$

Throughout this work, we will use the designation where D/dt is the covariant derivative along γ . In the proposed system each state velocity vector is transported parallel to each other successive analysis window and the difference is calculated between the vectors by moving the previous window's tangent vector to the current manifold point [11]. This provides a geometrically consistent measure of the state acceleration of the structure, which is coordinate chart independent, and will depend upon the curved geometry of the SPD manifold. The curvature component C of the instability index Ψ is a section-weighted (by the sectional curvature along the interpolating geodesic) direct computation of the norm of this transported difference vector. The weights $\omega_1, \omega_2, \omega_3$ are non-negative and their sum is 1 and they are optimized based on the cross-validation performed on the training set. The index divides the operating space into three zones of operation (stability): $\Psi < 0.35$ (Stable), $0.35 \leq \Psi \leq 0.70$ (Transitional), $\Psi > 0.70$ (Critical). These thresholds are represented by risk levels that have been proved empirically with structural simulation ground truth.

In Fig. 2, the visualization underneath the Riemannian manifold shows the different nature of the geodesic paths connecting sensor-state points located on the curved surface of the manifold by comparing these geodesic paths to the corresponding Euclidean straight-line distances. The whole structural coupling information is encoded into the curvature of the manifold, information which is not seen by the Euclidean metrics.

Riemannian Manifold with Geodesic Paths vs. Euclidean Distances




Fig. 2 Curved Riemannian manifold with sensor-state points and geodesic paths.

5. Methodology

The methodology used in the experiments is designed to test the proposed system throughout the entire operating range defined by the vehicle simulated for Formula situations when the vehicle range covers five speeds and covers the 80 to 340 km/h speed range during tests, six zones of tests were performed on vehicle structure defined by the vehicle's body. Five-fold cross validation with stratified sampling is performed for the evaluation protocol, maintaining the balance of all the stability classes in all folds. The assessment is also done in the context of an ablation study based on three versions of the proposed vortical system: (i) full

Riemannian framework with all three

instability components (C, G, V); (ii) the "geodesic only", an instability based on the geodesic divergence G and (iii) the "curvature only", an instability based on the sectional curvature scalar C. The ablation separates and tests the contribution of each individual Riemannian descriptor [12] to the overall system performance, and proves the need for the combination of the composite instability formulation.

5.1. Riemannian Metric and SPD Manifold Structure

The SPD manifold $\text{Sym}^+(n) = \{P \in \mathbb{R}^{n \times n} \mid P = P^T, x^T P x > 0 \forall x \neq 0\}$ is a smooth Riemannian manifold of dimension $n(n+1)/2$. The number of the sensor covariance matrices used in this work is $n = 6$, which gives a 21-dimensional manifold. The affine-invariant metric at point P is given by:

$$(U, V)_P = \text{tr}(P^{-1} U P^{-1} V)$$

U and V are symmetric matrices in the tangent space at P and $\text{tr}(\cdot)$ is the trace of the matrix. This measure is preserved by congruence transformations of P given by $A^T P A$ for any non-singular matrix A, which means the geodesic distance between two covariance matrices remains the same under any coordinate frame transformation that is simply a rotation and/or rescaling of the coordinate frame. Let d_{PQ} be the geodesic distance between two points P, Q $\in \text{Sym}^+(n)$ induced:

$$d_{\text{g}}(P, Q) = ||\log(P^{-1/2} Q P^{-1/2})||_F$$

where $||\cdot||_F$ a Frobenius norm, and $\log$ a matrix logarithm. It is symmetric, positive definite and it satisfies the triangle inequality property of a true metric in the manifold. In particular, $d_{\text{g}}(P, Q)$ is not equal to the Frobenius distance $||P - Q||_F$ if the geodesic path between P and Q has a significant curvature in the manifold, i.e., if the eigenvalue ratios of $P^{-1/2} Q P^{-1/2}$ significantly deviate from 1—the condition that defines high load structural transitions.

5.2. Training Protocol and Hyperparameter Selection

The instability index values $\omega_{\{1\}}$ values, $\omega_{\{2\}}$ values, and $\omega_{\{3\}}$ values are computed using 50 iterations of the Tree-structured Parzen Estimator (TPE) algorithm, over a 3-simplex search space. The resulting optimum weights are found to be $\omega_1 = 0.41$ (curvature), $\omega_2 = 0.35$ (geodesic divergence) and $\omega_3 = 0.24$ (velocity power) showing that the water most clearly informative for diagnosis is the curvature-derived component (ω_1).

The Fréchet mean's gradient norm is initialized to the "arithmetic mean" of all training-set SPDs and steps through all the 16 ± 2 iterations until it is below $\epsilon = 10^{-6}$. The values of threshold were set by optimisation of F_1 score on the validation fold according to the cost matrix, which presumes that the consequence of the False-Negative Critical is 10 times that of the False-Positive Transitional, where costs are observed to be asymmetric, with false-negative predictions being more costly than false-positive predictions in terms of safety.

The Adam learning rate optimiser with a learning rate of 10^{-4} and a batch size of 64 is used during training of all neural baselines (LSTM) for 100 epochs and with early stopping (patience = 10). For a fair comparison with LSTM framework, two layers of 256 hidden units and dropout rate 0.3 will be used, identified by grid search on the same cross validation folds used by the Riemannian framework. The training is conducted on an NVIDIA A100 GPU, while inference tests are done against all the methods on an NVIDIA Jetson AGX Orin target GPU.

5.3. Sensor Acquisition and Synchronisation

This multi-modal sensor suite includes tri-axial accelerometers at 1000 Hz, 6DoF IMUs (inertial measurement units) at 1000 Hz, high-speed cameras at 240 frames per second, strain gauge arrays at 500 Hz, tire pressure sensors at 100 Hz and ride-height potentiometers at 200 Hz. All modalities are hardware synchronized to a pulse per second (PPSec) signal derived from GPS that has a temporal alignment of ± 0.1 ms. Raw streams are kept in the buffer of 5120 samples per channel, which results in a continuous 5.12 s sliding analysis window.

5.4. Preprocessing and Feature Extraction

All sensor channels are bandpass filtered (3 to 300 Hz for mechanical channels; dc-coupled for quasi-static strain) and then Z-normalised using a baseline window of 60 s for which the data is recorded under controlled straight line running. All statistical outliers more than 4 times the distance away from the rolling mean are replaced with cubic spline interpolation between neighbouring non-outlier points. The preprocessing step also includes spectral whitening, which is a power equalizer across the frequency bands prior to feature extraction [13]. The feature extracted is a 128 dimensional feature vector per sensor modality, which consists of statistical moments (mean, variance, skewness, kurtosis), FFT spectral descriptors (dominant frequency, spectral centroid, spectral roll-off), wavelet packet energy coefficients at four decomposition levels, and the upper triangle elements of the 6×6 sensor covariance matrix. The covariance matrix is constructed to be symmetric and positive definite, as required, and is used as the main input for the manifold projection stage.

5.5. Manifold Projection and Geodesic Analysis

The 6×6 covariance matrices are each projected onto the Riemannian manifold using the matrix logarithm map $\text{Log}_P: M \rightarrow T_P M$, and the Fréchet mean is computed iteratively using gradient descent on M as the base point. The affine-invariant distance $d(P, Q) = ||\text{Log}_P(Q)||_{\text{g}}$ is calculated between consecutive points of the manifold, with the norm of the matrix being induced by the metric tensor. Closed-form expressions for the Christoffel symbols and components of the Riemann curvature tensor on the SPD manifold are used to compute them analytically.

5.6. Algorithm Summary

Algorithm 1: Riemannian Structural Stability Prediction

Input: Synchronized sensor streams $\{S_1, \dots, S_n\}$, window W

Output: Stability class $y \in \{\text{Stable}, \text{Trans.}, \text{Critical}\}$, score Ψ

- Step - 1 : Acquire streams; apply GPS-PPS timestamp sync
- Step - 2 : Bandpass filter; Z-normalise; outlier rejection
- Step - 3 : Extract covariance matrices $C_i \in \text{SPD}(6)$ per modality
- Step - 4 : Compute Fréchet mean μ^{sc} on Riemannian manifold M
- Step - 5 : Map C_i via $\text{Log}_{\mu}(\cdot)$; compute geodesic distances d_{g}
- Step - 6 : Evaluate Christoffel symbols $\Gamma^{R,N}$; Riemann tensor R
- Step - 7 : Compute $\Psi = (\omega_1 C + \omega_2 G + \omega_3 V) / 3$ using Eq. (3)
- Step - 8 : Classify: if $\Psi < 0.35 \rightarrow \text{Stable}$; $\Psi > 0.70 \rightarrow \text{Critical}$
- Step - 9 : Output y, Ψ to real-time dashboard; log to buffer

6. Results And Analysis

A simulated Formula-style vehicle data set was used to test the proposed Riemannian framework that included 140 hours of multi-modal telemetry data recorded in 48 sessions, covering the range of speeds from 80 to 340 km/h. The labels were generated using a validated finite-element structural simulation model, and annotated with three ground-truth classes (Stable / Transitional / Critical) at a 50 ms resolution. To evaluate all methods, the same 5-fold cross-validation splits and the same feature extraction pipeline were used up to the unique manifold projection step of the proposed method.

6.1. Prediction Accuracy

The comparative bar chart of structural state prediction accuracy for the four models is shown in Fig. 3. The proposed Riemannian framework achieves 96% accuracy, which is 19 percentage points more than linear baseline (77%), 10 pp more than Kalman filter (86%) and 3 pp more than LSTM network (93%). The highest differences from the LSTM are obtained in the high-speed range (from 280 to 340 km/h), in which the curvature of the manifold is the highest and the Euclidean approximations are the least valid. As the system dynamics get closer to the linear regime, all of the methods approach similar performances in the low-speed regime (80–100 km/h).




Fig. 3. Prediction accuracy comparison across evaluated models (%).

Fig. 3 Prediction accuracy comparison across evaluated models (%).

6.2. Prediction Lead Time

The prediction lead time gained by each method is shown in Fig. 4, which is the period in advance that a structural state change is predicted. The Riemannian framework leads to a lead time of 182ms, while the linear model has a lead time of 90ms, the Kalman filter has a lead time of 118ms and the LSTM has a lead time of 155ms. The 64.4% improvement over linear baseline means that the geodesic curvature descriptors can identify incipient state divergence earlier than methods based on Euclidean space. The 17.4% improvement over the LSTM indicates that the geodesic curvature descriptors can detect incipient state divergence before the LSTM. A 64ms reduction in lead time equates to an extra 6 m of reaction distance at 340 km/h.




Fig. 4. Prediction lead time comparison across evaluated models (ms).

Fig. 4 Prediction lead time comparison across evaluated models (ms).

6.3. Spatial Instability Distribution

The results of the structural instability index Ψ are shown in the heatmap of Fig. 5 for the 6 vehicle zones and the 5 speed regimes. The front-left and front-right chassis zones are the most unstable zones at all speeds, with an extreme instability value of: $\Psi = 0.82$ in the speed range of 280–340 km/h. It is in line with the dominant influence of front-end aerodynamic loading in the case of Formula-style vehicles, in which the balance of the aerodynamic forces moves forwards as a result of high-speed braking. The mid-body sections have lower instability indices throughout the range of speeds, which is indicative of the structural strength of the central monocoque.




Fig. 5. Heatmap of structural instability index Ψ across vehicle zones and speed regimes.

Fig. 5 Heatmap of structural instability index Ψ across vehicle zones and speed regimes.

Most importantly, the only tested method that can reliably predict Critical-class events ($\Psi > 0.70$) at high speed in the front chassis regions is the Riemannian framework. In this regime, the LSTM gives false-negative rate of 12%, as the Euclidean-space distance metrics are unable to detect the geodesic divergence that takes place rapidly prior to the coming of a new structural state. This mis-detection rate is lowered to 2.8% using the Riemannian framework and the precision rate for Critical-class detections is 97.1%.

6.4. Ablation Study

The ablation study for the three variants of the system shows that the three instability components do influence the overall performance independently. The geodesic-only variant is 91.2% accurate, with a lead time of 163 ms; the curvature-only variant is 92.7% accurate with a lead time of 171 ms. The performance of the full composite system (96% accuracy, 182 ms lead time) is much better than that of the two ablated systems, confirming the complementary effect of the three Riemannian descriptors. The curvature component C , which was shown to be an early geodesic divergence detector [11], gives the greatest amount of prediction lead time (171 ms compared with 163 ms for geodesic-only). The velocity component V introduces the physical constraint that high structural risk necessitates high vibration

energy as well as unfavourable manifold geometry to reduce the rate of false positive instances in the Transitional class from 8.4% (without V) to 4.1% (with V).

6.5. Discussion

The results from the experiments confirm three main findings. The Riemannian manifold representation of SPD covariance feature delivers measurably better prediction accuracy than other European approaches in all measures studied. The improvement is greatest under high load and high speed conditions, where the nonlinear coupling between the modes of the structure is significant, which is when the function of structural monitoring can be the most critical to safety. This is exactly as advocated in theory, where the greatest curvature of the manifold is the region in which the Euclidean distance measures are more affected.

Second, the 182ms prediction lead time of the Riemannian framework is an advancement for a practically important problem compared to the 155ms prediction lead time of the best Euclidean baseline. It is shown that the geodesic curvature descriptor C of the instability index Ψ is useful for detecting the divergence of nearby geodesics on the manifold—a result similar to the Jacobi field instability in differential geometry, before that divergence can be seen in the raw sensor signals—before that divergence becomes apparent in the raw sensor signals. This is a physical novel insight [12], and it is because of the Riemannian approach that it is possible to arrive at this insight. Third, the low variance ($\sigma = 0.03$) of the Riemannian framework when compared to the LSTM ($\sigma = 0.05$) point out the generalisation across the diverse working conditions offered by the provided 5-fold-cross validation. The manifold-aware Fréchet mean calculation introduces the notion of “average” structural state which has fewer directional impacts of outliers in the training data, and hence has more consistent behavior when considering unseen test folds [13].

There will be some elaboration of the instability index Ψ , in terms of being geometric. Differential between the current and preceding state tangent vectors divides the curve spacing C, which quantifies the rate of convergence or divergence of neighbouring geodesics of the two-plane σ generated by the current and (previous) state tangent vectors. Geodesics converge for the positive value of the curvature ($K > 0$), being a regime of the structure with a stable regime of perturbations, which are geometrically damped. Negative sectional curvature ($K < 0$) leads to divergent geodesics as in a chaotic dynamical system, and in the experimental data there is always a stable lead time, of about 150–200 ms, between the occurrence of a negative event in the sections and a structural state transition. This is a geometric early-warning scheme that is peculiar to the Riemannian approach and is not found in the Euclidean-space approaches. An enclosed computing approximation to the Fréchet mean on the SPD manifold will be evaluated in the near future to help the algorithm run in a single matrix operation, which could help deploy the algorithm to embedded systems with power budgets under 10 W with each processing window.

6. Conclusion and Future Scope

We have introduced a framework for structural stability prediction based on Riemannian geometry that was able to predict the structural stability in simulated Formula-style vehicle telemetry data with accuracy of 96% and prediction lead time of

182 ms over the linear, Kalman, and LSTM baselines on all the metrics evaluated. The proposed system utilizes the intrinsic geometry using a mathematically driven framework sought by multi-sensor covariance onto the SPD Riemannian manifold, computing geodesic distances, curvature descriptors, and a resultant composite instability index Ψ , which is not possible with Euclidean methods. With an end-to-end latency of just 18 ms, the architecture can be deployed on NVIDIA Jetson-class edge devices, so it can meet motorsport and advanced driver assistance system real-time operational requirements. The learnable fusion weights in the instability index give interpretable diagnostics that are based on known physical failure mechanism in high speed vehicles structures.

Several promising research directions follow from this work. Digital twin integration offers the prospect of continuously calibrating the manifold metric tensor against a physics-based finite-element model of the vehicle, ensuring that the Riemannian geometry remains accurate as structural properties evolve due to fatigue or damage accumulation. Federated learning across multiple team vehicles would substantially enlarge the effective training dataset while preserving the commercial confidentiality of individual teams' telemetry archives.

The extension of the instability index to incorporate Ricci curvature and scalar curvature tensors—higher-order geometric descriptors beyond the Christoffel symbols used here—may capture additional structural state information in extreme loading scenarios. Finally, reinforcement learning-based adaptive fusion weight optimisation could enable the system to dynamically adjust $\omega_1, \omega_2, \omega_3$ in response to changing circuit conditions, providing a fully autonomous structural monitoring solution requiring no manual recalibration between operating environments.

References

- [1]. S. W. Doebling, C. R. Farrar, M. B. Prime, and D. W. Shevitz, "Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: A literature review," Los Alamos National Laboratory, Tech. Rep. LA-13070-MS, 1996.
- [2]. S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [3]. R. E. Kalman, "A new approach to linear filtering and prediction problems," *Trans. ASME—J. Basic Eng.*, vol. 82, pp. 35–45, Mar. 1960.
- [4]. J. M. Lee, *Introduction to Smooth Manifolds*, 2nd ed. New York: Springer, 2013.
- [5]. A. Barachant, S. Bonnet, M. Congedo, and C. Jutten, "Multiclass brain-computer interface classification by Riemannian geometry," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 4, pp. 920–928, Apr. 2012.
- [6]. X. Pennec, P. Fillard, and N. Ayache, "A Riemannian framework for tensor computing," *Int. J. Comput. Vis.*, vol. 66, no. 1, pp. 41–66, Jan. 2006.
- [7]. M. Moakher, "A differential geometric approach to the geometric mean of symmetric positive-definite matrices," *SIAM J. Matrix Anal. Appl.*, vol. 26, no. 3, pp. 735–747, 2005.
- [8]. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.
- [9]. P. Zarchan and H. Musoff, *Fundamentals of Kalman Filtering: A Practical Approach*, 4th ed. Reston, VA: AIAA, 2015.
- [10]. K. Vasepalli, N. Ramanujam, G. Ponnuru, and B. Nemakal, "Utilizing deep learning techniques for the

identification of medicinal plants," *International Journal of Computational Science and Engineering Research*, vol. 2, no. 3, p. 15, Sep. 2025, doi: 10.63328/ijcser-v2i3p3.

- [11]. T. Farrar and K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*. Chichester: Wiley, 2012.
- [12]. A. Noyal, A. C. Kurian, C. E. S. S, and S. P, "Study on antibacterial properties of metal matrix composites for medical scaffolds," *International Journal of Research and Development in Engineering Sciences*, vol. 5, no. 2, p. 6, Mar. 2023, doi: 10.63328/ijrdes-v5ri2p2.
- [13]. P Siva Prasad , K Aparna , C Gayathri , A Komala , V Lokesh Reddy , P Balaji Reddy, "Simulation Model of Isolated Solar Roof Top System for Home and LV Grid " ,*International Journal of Computational Science and Engineering Research*, vol. 1, no. 2, May. 2024, doi: <https://doi.org/10.63328/IJCSEAI-V1R1P4>

Declaration

Conflicts of Interest: The authors declare no conflict of interest.

Author Contribution: All authors wrote the main manuscript text and also consent to the submission.

Ethical approval: Not applicable.

Consent to Participate: All authors consent to participate.

Funding: Not applicable, and No funding was received

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Personal Statement: We declare with our best of knowledge that this research work is purely Original Work and No third party material used in this article drafting. If any such kind material found in further online publication, we are responsible only for any judicial and copyright issues.

Acknowledgements

We thank everyone who inspired our work.

How to Cite this Paper:

Sk. Rukshana Begum , Mesfin Dagne , " Riemannian Geometry-Based Structural Stability Prediction in Dynamic Vehicle Systems " , *International Journal of Computer Science, Engineering and Artificial Intelligence* , vol. 1, no. 1, p. 21-27, December 2024, DOI: <https://doi.org/10.63328/IJCSEAI-V1R1P4>