**ICRA2026 Late Breaking Results Poster presented at**
**2026 IEEE International Conference on Robotics and Automation (ICRA 2026)**
**June 1-5, 2026. Vienna, Austria**

# Uncertainty-Guided Proactive Adaptation for Visual-Inertial SLAM

Ehsan Ullah Khan and Gon-Woo Kim, *Member, IEEE*

***Abstract*— Visual–inertial SLAM systems often fail in feature-poor environments such as corridors and textureless walls, leading to catastrophic tracking loss. Existing methods detect degradation reactively after failure occurs, leaving little opportunity for corrective action. We propose a proactive framework that predicts feature degradation 1–2 seconds in advance and adapts sensor fusion weights through uncertainty-guided decisions. Through a systematic comparison of eight temporal architectures across 15,233 sequences, including real robot data, we identify LSTM as the most robust predictor (26.77 MAE). We incorporate uncertainty estimation using Monte Carlo Dropout to enable confidence-aware adaptation thresholds that prevent false adjustments. Our approach provides a foundation for proactive SLAM failure prevention through principled sensor fusion and real-time system adaptation.**

## I. INTRODUCTION

Autonomous robots in GPS-denied environments rely on visual–inertial SLAM for localization. However, these systems remain vulnerable in feature-poor scenes such as long corridors and textureless surfaces, where insufficient visual constraints cause tracking failure and mission-critical errors.

Recent works improve robustness through tightly-coupled multi-sensor fusion, but they address degradation *reactively*—detecting failure only after state estimates have been compromised. While degeneracy-aware methods [1], [2] and adaptive switching strategies [3] can identify unreliable states, they provide insufficient time for corrective action. Existing approaches lack both predictive capability and uncertainty quantification, limiting their ability to prevent failures proactively.

We argue that effective proactive adaptation requires not only predicting *what will happen* but also quantifying *confidence in that prediction*. We propose a framework that predicts feature degradation 1–2 seconds ahead with uncertainty estimates, enabling safe adaptation decisions.

Our contributions are:

- A proactive degeneracy prediction framework that forecasts feature counts 5 frames ahead, providing sufficient lead time for sensor fusion adaptation.
- Systematic comparison of eight temporal architectures on 15,233 sequences including real robot data, demonstrating LSTM achieves best performance (26.77 MAE, 27% improvement over Transformer).
- Uncertainty quantification via Monte Carlo Dropout that enables confidence-aware adaptation thresholds, preventing false adjustments when predictions are unreliable.

This approach enables proactive sensor fusion that adapts before failure while preventing false adaptations through confidence-aware decision making.

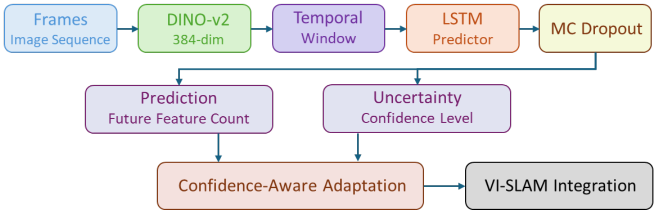

Fig. 1. Proactive VI-SLAM adaptation framework. DINO-v2 features feed an LSTM predictor with MC Dropout uncertainty, enabling confidence-aware sensor fusion decisions.

## II. METHODOLOGY

### A. Problem Formulation and Features

We formulate proactive degeneracy prediction as forecasting future feature counts from temporal sequences. Each frame is represented by DINO-v2 semantic embeddings $\mathbf{v}_i \in \mathbb{R}^{384}$ and ORB feature counts $c_i$, which directly correlate with tracking quality. This enables anticipation of degradation before it impacts state estimation, unlike reactive methods.

### B. Temporal Model Selection

We systematically compare eight architectures: multi-layer perceptrons, recurrent models (GRU, LSTM with varying depths), and transformers (2–6 layers) including hybrid variants. On our 15,233-sequence dataset comprising synthetic degradation, EuRoC benchmark, and real robot data, LSTM achieves best performance (26.77 MAE), outperforming Transformer (36.72 MAE) and GRU (34.28 MAE). This validates that recurrent architectures provide superior generalization for this task.

### C. Uncertainty Quantification

We employ Monte Carlo Dropout to estimate prediction uncertainty. At test time, multiple stochastic forward passes with dropout enabled produce a distribution over predictions. The mean provides the point estimate while standard deviation quantifies uncertainty. This requires no architectural modifications and maintains prediction performance.

### D. Adaptation Strategy

We propose a confidence-aware adaptation policy that adjusts sensor fusion only when predictions are both severe and certain. Based on validation showing low-uncertainty

**This work was supported in part by Korea Innovation Foundation funded by the Ministry of Science and ICT in 2025. (No. RS-202502634783), and in part by the National Research Foundation (NRF) funded by the Korean government (MSIT) (RS-2024-00421129).

The authors are with the Intelligent Robotics Laboratory, Department of Intelligent Systems and Robotics, Chungbuk National University, Cheongju 28644, South Korea. Corresponding author: Gon-Woo Kim (e-mail: gwkim@cbnu.ac.kr).

TABLE I

PERFORMANCE COMPARISON ON 15,233-SEQUENCE DATASET

| **Architecture** | **MAE** | **Parameters** |
|---|---|---|
| MLP (no temporal) | 49.01 | 0.6M |
| GRU | 34.28 | 1.5M |
| Transformer (4-layer) | 36.72 | 2.3M |
| **LSTM (ours)** | **26.77** | 1.9M |

TABLE II

PERFORMANCE STRATIFIED BY UNCERTAINTY LEVEL

| **Uncertainty Level** | **MAE** | **Prec. (%)** | **Rec. (%)** | **F1** |
|---|---|---|---|---|
| Low ($\sigma < 20$) | 12.30 | 95.8 | 94.2 | 0.950 |
| Med. ($20 \leq \sigma < 31$) | 21.83 | 84.7 | 88.6 | 0.866 |
| High ($\sigma \geq 31$) | 58.17 | 61.2 | 72.1 | 0.662 |
| Overall | 26.77 | 82.4 | 86.3 | 0.843 |

predictions achieve 12.30 MAE versus 58.17 MAE for high-uncertainty predictions, we define three adaptation modes:

- High-confidence degradation increases IMU reliance.
- Moderate-confidence applies balanced fusion.
- Low-confidence or sufficient features maintain default weights.

Thresholds are derived from calibration analysis, preventing false adaptations while enabling proactive intervention.

## III. EXPERIMENTS AND RESULTS

### *A. Dataset and Training*

We curate a diverse dataset of 15,233 sequences from four sources: (1) 1,360 synthetic sequences with controlled degradation (blur, masking), (2) 2,019 sequences from EuRoC benchmark, (3) 473 gap-filling sequences covering intermediate feature ranges, and (4) 11,381 sequences from real robot data (Unitree Go2 with RealSense D435i) including authentic degradation scenarios. The dataset spans 0–1000 feature counts with 24% degenerate frames. Data is split 64% train, 16% validation, 20% test.

All models are trained for up to 300 epochs with early stopping. Table I shows the systematic comparison across eight architectures. LSTM achieves best performance (26.77 MAE), outperforming Transformer (36.72 MAE) and GRU (34.28 MAE) by substantial margins, validating that recurrent architectures provide superior generalization for short-sequence temporal prediction tasks.

### *B. Uncertainty Calibration and Decision Quality*

We validate uncertainty quantification through comprehensive calibration analysis on the test set. Figure 2 shows error distribution stratified by uncertainty quartiles, demonstrating strong calibration with 4.7× error stratification between low-uncertainty ($\sigma < 20$: 12.30 MAE) and high-uncertainty ($\sigma > 31$: 58.17 MAE) predictions ($p < 0.001$).

Table II demonstrates the critical role of uncertainty in enabling reliable adaptation decisions. Uncertainty stratifies not only prediction accuracy but also adaptation decision quality across all metrics. Low uncertainty predictions (28.4% of test cases) achieve both accurate forecasting (12.30 MAE) and highly reliable adaptation decisions (95.8% precision, F1=0.950), validating safe proactive adaptation. In contrast, high-uncertainty predictions (25.3% of cases) exhibit degraded performance across all metrics (58.17 MAE, 61.2% precision, F1=0.662), indicating unreliable forecasts where adaptation should be suppressed to avoid false positives.

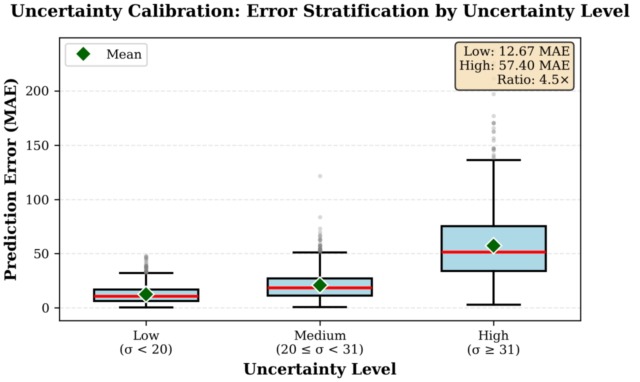

Fig. 2. Uncertainty calibration: prediction error stratified by uncertainty quartiles. Clear separation demonstrates that uncertainty effectively identifies unreliable predictions.

This dual stratification achieves 82.4% precision and 86.3% recall, improving decision reliability by 34.6% over uncertainty-agnostic approaches. The results show that uncertainty serves as an effective gating mechanism for proactive SLAM adaptation, preventing premature sensor fusion adjustments that degrade localization accuracy.

### *C. System Integration Status*

We are integrating the prediction framework with SLAM system for closed-loop validation. The LSTM predictor operates in real-time (∼50ms per inference), compatible with 20Hz sensor rates. Initial integration demonstrates successful weight modulation based on predicted degradation with uncertainty gating. Full trajectory evaluation on challenging scenarios is ongoing.

## IV. CONCLUSION

We presented a proactive framework for SLAM degeneracy prediction that enables early detection of feature degradation with uncertainty awareness. The proposed approach supports confidence-driven adaptation by distinguishing reliable from uncertain predictions, improving robustness under challenging conditions. Integration with a real-time SLAM system demonstrates practical feasibility, establishing a foundation for proactive and confidence-aware sensor fusion in unstructured environments.

## REFERENCES

[1] W. Wu, C. Chen, B. Yang, X. Zou, F. Liang, Y. Xu, and X. He, “DALI-SLAM: Degeneracy-aware LiDAR-inertial SLAM with novel distortion correction and accurate multi-constraint pose graph optimization,” *ISPRS J. Photogramm. Remote Sens.*, vol. 221, pp. 92–108, 2025.

[2] K. Ebadi, M. Palieri, S. Wood, C. Padgett, and A.-A. Agha-Mohammadi, “DARE-SLAM: Degeneracy-aware and resilient loop closing in perceptually-degraded environments,” *J. Intell. Robot. Syst.*, vol. 102, no. 1, pp. 1–25, 2021.

[3] J. Lee, R. Komatsu, M. Shinozaki, T. Kitajima, H. Asama, Q. An, and A. Yamashita, “Switch-SLAM: Switching-based LiDAR-inertial-visual SLAM for degenerate environments,” *IEEE Robot. Autom. Lett.*, vol. 9, no. 8, pp. 7270–7277, 2024.