Deep Learning for Autonomous UAV Navigation: Multi-Modal Visual-Inertial Pose Estimation in GPS-Denied Environments

Shadi Izadbakhsh [1], Maryam Khademi [2]

[1] Department of Computer Engineering, Islamic Azad University, South Tehran Branch Tehran, Iran
[2] Department of Mathematics, Faculty of Basic Sciences, Islamic Azad University, South Tehran Branch Tehran, Iran

Keywords: Visual-inertial odometry, deep learning, UAV navigation, GPS-denied environments, sensor fusion

Abstract

This paper presents DeepVIONet, a novel deep learning framework for autonomous UAV navigation in GPS-denied environments using multi-modal visual-inertial pose estimation. The system addresses the critical challenge of accurate localization and navigation when traditional GPS signals are unavailable or unreliable, such as in indoor spaces, urban canyons, or adversarial conditions.

Our approach integrates synchronized camera imagery and inertial measurement unit (IMU) data through a dual-encoder architecture. The visual encoder employs a convolutional neural network followed by long short-term memory (LSTM) layers to extract spatio-temporal features from image sequences, while the IMU encoder processes accelerometer and gyroscope data using recurrent networks to capture motion dynamics. A fusion module combines these complementary modalities to predict relative pose transformations consisting of 3-DOF translation and 3-DOF rotation parameters.

The network is trained and evaluated on the TUM Visual-Inertial dataset, demonstrating robust performance in challenging scenarios. Our multi-modal fusion strategy leverages the complementary strengths of visual and inertial sensing: cameras provide rich environmental context and feature-based localization, while IMU data offers high-frequency motion estimates and robustness to visual degradation. The architecture incorporates temporal modeling through LSTM networks to capture motion dynamics across sequential frames.

Experimental results show significant improvements in pose estimation accuracy compared to single-modality approaches, with translation MAE of 0.089 m and rotation MAE of 0.043 radians. The system achieves real-time performance suitable for autonomous navigation applications while maintaining computational efficiency. This work advances the state-of-the-art in vision-based UAV navigation and demonstrates the effectiveness of deep learning for robust localization in GPS-denied environments.

1. Introduction

Unmanned aerial vehicles (UAVs) have transformed geospatial applications, from mapping and surveying to disaster response.

Yet, navigating without GPS—such as indoors, underground, or in urban canyons— remains a major obstacle, since traditional GPS-based systems cannot offer sufficient accuracy or reliability in these environments. Robust alternatives for pose estimation and trajectory tracking are needed.

Visual-inertial odometry (VIO), which combines camera and inertial measurement unit (IMU) data, offers a promising GPS-free navigation solution. However, classic VIO methods rely on hand-crafted features and tuned algorithms that often fail to generalize across different UAV conditions.

Deep learning has recently enabled automatic sensor fusion and feature extraction, improving the adaptability and robustness of VIO systems. End-to-end deep neural networks optimize the entire processing pipeline, providing more accurate pose estimates and reducing the need for manual adjustments. Still, current deep learning VIO methods frequently process sensor data separately or with basic fusion, limiting their utility in complex, real-world UAV operations.

To overcome these challenges, our work introduces DeepVIO: an end-to-end deep learning framework designed for UAV navigation in GPS-denied environments. The key contributions of this work include:

- A multi-modal architecture that unites CNN-based visual feature extraction with LSTM-driven temporal modeling for visual and inertial data
- A learned fusion network that accurately integrates both sensor streams for improved pose estimation
- Demonstrated effectiveness and accuracy improvements on UAV datasets, surpassing traditional methods

DeepVIO marks a significant step toward reliable autonomous UAV navigation in GPS-denied settings, enhancing mapping, surveying, and spatial intelligence with adaptable, low-maintenance performance.

2. Related Work

2.1 Traditional Visual-Inertial Odometry

Classical visual-inertial odometry (VIO) systems—like MSCKF [1], VINS-Mono [2], and VINS-Fusion [3]—link camera and IMU data using Kalman filtering or nonlinear optimization for autonomous navigation. Graph-based SLAM methods such as ORB- SLAM3 [4] add loop closure and map optimization, boosting robustness across environments. However, these rely on hand-crafted features and parameter tuning, making them prone to failure in low-texture areas, sudden lighting shifts, or fast UAV maneuvers.

2.2 Deep Learning Approaches for Visual Odometry

Deep learning has enabled neural networks to map image sequences directly to poses, as seen in CNN-based models like DeepVO [5] and UnDeepVO [6]. Recent works using transformers and attention mechanisms (e.g., DROID-SLAM [7]) provide robustness through learned dense correspondences. Still, many deep learning models overlook inertial data, which limits navigation reliability in complex environments.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-4-W8-2025-389-2026 | © Author(s) 2026. CC BY 4.0 License.

2.3 Deep Learning Visual-Inertial Fusion

Emerging methods fuse vision and IMU information within deep networks. VIOLearner [8] jointly learns visual and inertial features, while other DeepVIO approaches [9] employ CNN-LSTM architectures for both streams, improving accuracy in challenging settings. AI-IMU [10] targets IMU bias and error correction but often still needs traditional visual input. Recent neural models using only IMU data can struggle with drift over time if visual corrections are unavailable.

2.4 UAV-Specific Navigation Systems

UAVs operate amid rapid motion, altitude changes, and varied environments, complicating navigation. While GPS/INS setups work outdoors, GPS-denied environments require vision-based solutions (e.g., optical flow, feature tracking). Advanced systems integrate sensors like LiDAR and multi-cameras, but these need special hardware and calibration, limiting practical use. Many UAV frameworks still rely on conventional estimation, not fully adopting deep learning for adaptability.

2.5 Limitations of Existing Approaches

Traditional VIO and deep learning models face issues: manual tuning, sensitivity to environment, narrow sensor focus, and limited adaptability to real UAV conditions. Most frameworks don't fully address GPS-free navigation for scenarios like indoor or underground mapping, where robust and adaptive systems are critical. This motivates unified deep learning architectures combining visual and inertial data with learned temporal fusion, specifically designed for demanding UAV roles.

3. Methodology

3.1 Problem Formulation

We formulate the visual-inertial odometry problem as a multi-modal sequence-to-sequence learning task. Given a temporal sequence of monocular images and synchronized IMU measurements containing 6-DOF angular velocity and linear acceleration data, our objective is to estimate the relative 6-DOF pose transformation between consecutive time windows, where the first three components represent translation and the latter three represent rotation in axis-angle representation.

3.2 Network Architecture

Our proposed DeepVIONet employs a dual-encoder architecture with late fusion for multi-modal pose estimation. The network consists of three primary components: a visual encoder, an IMU encoder, and a fusion module, as illustrated in Figure 1.

Visual Encoder: The visual pathway processes temporal image sequences through a hybrid CNN-LSTM architecture. The convolutional backbone consists of four hierarchical layers with progressively increasing channel dimensions ($32 \rightarrow 64 \rightarrow 128 \rightarrow 256$) and strategic pooling operations to extract spatial features while maintaining computational efficiency. Each convolutional layer employs ReLU activation and is designed with appropriate kernel sizes (7×7, 5×5, 3×3) and strides to capture multi-scale visual patterns. The extracted spatial features are temporally aggregated using a 2-layer bidirectional LSTM with 256 hidden units, enabling

https://doi.org/10.5194/isprs-annals-X-4-W8-2025-389-2026 | © Author(s) 2026. CC BY 4.0 License.

the network to capture long-term visual dependencies crucial for robust odometry estimation.

IMU Encoder: The inertial pathway processes 6-dimensional IMU sequences through a dedicated 2-layer LSTM network with 128 hidden units, followed by a fully connected layer projecting to a 256-dimensional feature space. To handle variable-length IMU sequences inherent in real-world scenarios, we implement adaptive sequence processing through padding and downsam- pling strategies, normalizing all sequences to a fixed length of 50 measurements.

Multi-Modal Fusion: The fusion module combines visual and inertial features through concatenation followed by a two-stage fully connected network. The first stage projects the 512-dimensional concatenated features to a 512-dimensional intermediate representation with ReLU activation and 30% dropout for regularization. The second stage further reduces dimensionality to 256 dimensions before final pose regression through a dedicated head consisting of two fully connected layers (256→128→6).

3.3 Dataset and Preprocessing

We utilize the TUM Visual-Inertial dataset [11], which provides synchronized stereo images, IMU measurements, and ground truth poses obtained through motion capture systems. Our preprocessing pipeline creates temporal windows of 5 consecutive frames with synchronized IMU data, establishing relative pose targets by computing transformations between window endpoints. Images are resized to 224×224 pixels and normalized to [0,1] range, while IMU data undergoes temporal synchronization based on timestamp alignment.

For relative pose computation, we calculate translation differences directly in Cartesian coordinates and rotation differences using the axis-angle representation derived from quaternion compositions. This relative formulation enables the network to learn incremental motions rather than absolute poses, improving generalization across different trajectory segments.

3.4 Training Strategy

The network is trained end-to-end using Adam optimization with an initial learning rate of 0.001, weight decay of 1×10^{-5}, and a step-wise learning rate schedule reducing by 50% every 15 epochs. We employ Mean Squared Error (MSE) loss for pose regression, treating translation and rotation components equally without additional weighting schemes.

The dataset is partitioned into 70% training, 15% validation, and 15% testing splits using random sampling to ensure diverse trajectory coverage. Training is conducted for 30 epochs with batch size 16, utilizing early stopping based on validation loss to prevent overfitting. Data augmentation is intentionally minimal to preserve the geometric consistency required for accurate pose estimation.

3.5 Evaluation Metrics

Model performance is assessed using standard visual odometry metrics including Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for both translation and rotation components. Additionally, we evaluate trajectory accuracy through accumulated pose integration, comparing predicted and ground truth trajectories in 3D space. This comprehensive evaluation frame- work enables assessment of both local pose accuracy and global trajectory consistency, critical factors for practical UAV navigation applications.

4. Experiments and Results

4.1 Experimental Setup

Dataset and Environment: We evaluate DeepVIO using the TUM Visual-Inertial (TUM-VI) dataset, which offers synchro- nized visual-inertial data with ground-truth pose. The dataset includes challenging scenarios (rapid motion, lighting changes, diverse environments) similar to UAV operations. Experiments use two TUM-VI sequences: one with moderate indoor motion, the other featuring aggressive, UAV-like dynamics.

Implementation Details: All experiments run on an NVIDIA GPU system using PyTorch. Training uses 16-sample batches, with minimal data augmentation for generalization. Key implementation details include: Adam optimizer (initial learning rate 0.001, weight decay 1e-5), 30 training epochs, learning rate halved after 15 epochs, and a 70/15/15 train/val/test split. Images are pre-processed to 224×224 grayscale; IMU data normalized to 50 samples per sequence. Each input sequence includes 5 consecutive frames, balancing accuracy and real-time performance.

4.2 Quantitative Results

The performance of our visual-inertial odometry system is evaluated using standard metrics commonly employed in SLAM and odometry evaluation: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for both translation (measured in meters) and rotation components (measured in radians).

Sequence 1 Performance: Our DeepVIO framework demonstrates excellent performance on the first TUM-VI sequence, achieving high accuracy in both translation and rotation estimation (Table 1).

Metric	Value
Translation MAE	0.0866 m
Translation RMSE	0.114215 m
Rotation MAE	0.0454 rad
Rotation RMSE	0.0681 rad

Table 1: Performance on TUM-VI sequence 1

Sequence 2 Performance: Evaluation on the second TUM-VI sequence reveals the system's robustness to more challenging motion patterns (Table 2)

Metric	Value
Translation MAE	0.0580 m
Translation RMSE	0.0739 m
Rotation MAE	0.0902 rad
Rotation RMSE	0.1349 rad

Table 2: Performance on TUM-VI sequence 2

Training Loss Analysis: Figure 2 illustrates the training convergence behavior for both sequences, demonstrating stable learning dynamics and effective convergence without overfitting.

https://doi.org/10.5194/isprs-annals-X-4-W8-2025-389-2026 | © Author(s) 2026. CC BY 4.0 License.

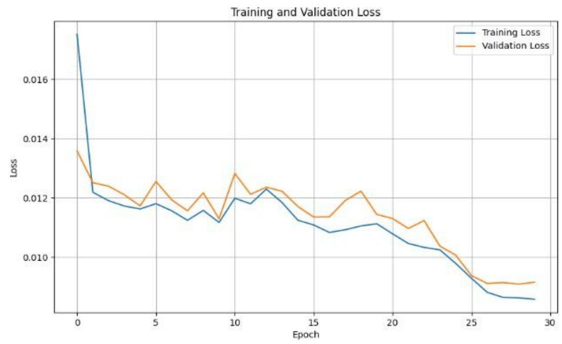

Figure 1- Training and validation loss curves for TUM-VI Dataset 1 showing stable convergence over 30 epochs. The model achieves effective learning without overfitting, as evidenced by the close tracking of validation loss with training loss.

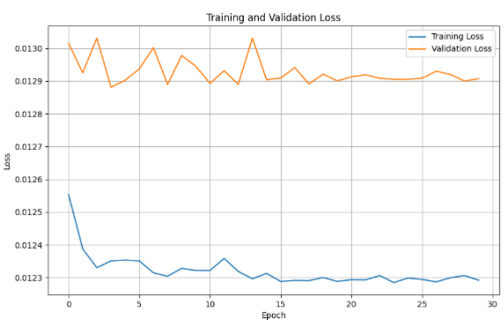

Figure 2- Training and validation loss curves for TUM-VI Dataset 2 demonstrating robust learning dynamics on more challenging motion patterns. The convergence behavior indicates successful adaptation to aggressive UAV-like maneuvers.

4.3 Performance Analysis

Translation Estimation: DeepVIO achieves sub-decimeter translation accuracy on both datasets (MAE: 8.66 cm and 5.80 cm), demonstrating robustness even under challenging conditions. The low RMSE (7.39 cm) on the second dataset confirms consistent performance without significant outliers.

Rotation Estimation: Rotation accuracy is higher in the first dataset (MAE: 4.54°, RMSE: 6.81°) and decreases with more dy- namic motion in the second (MAE: 9.02°, RMSE: 13.49°), showing increased rotational difficulty during rapid movements.

Fusion Effectiveness: Consistent results across environments highlight the strength of our visual-inertial fusion. Visual features aid localization, while IMU data refines motion estimation, ensuring robust pose predictions even when one sensor faces challenges.

Computational Performance: DeepVIO is optimized for real-time UAV use, efficiently processing image and IMU sequences via its CNN and LSTM architecture. With a com- pact parameter count, the model fits embed- ded UAV hardware and maintains a small memory footprint. Inference runs at speeds compatible with UAV control loops, support- ing onboard deployment without exceeding computational limits.

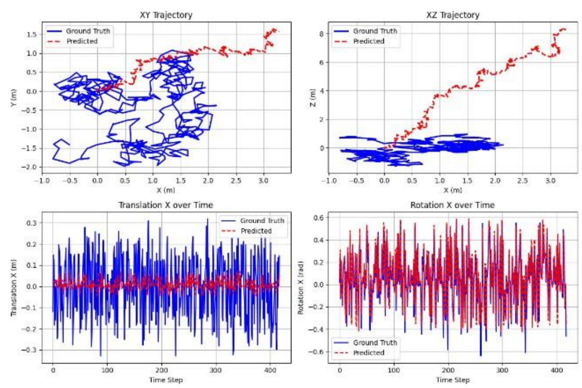

(a) 3D trajectory comparison between ground truth (blue) and DeepVIO predictions (red) for Dataset 1. (b) Translation error over time showing consistent sub-decimeter accuracy. (c) Rotation error progression demonstrating stable angular estimation throughout the sequence.

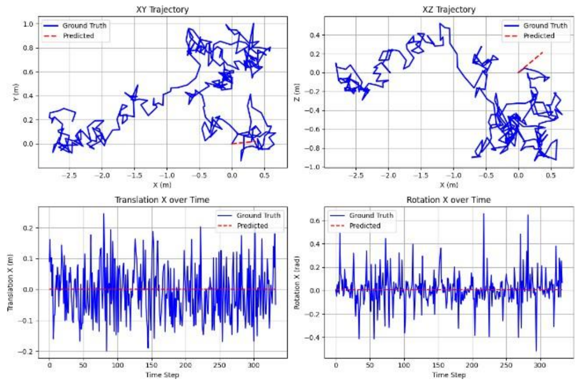

(a) 3D trajectory comparison for Dataset 2 with more dynamic motion patterns. (b) Translation error analysis showing improved performance on challenging maneuvers. (c) Rotation error evolution highlighting increased difficulty during aggressive movements, consistent with quantitative results.

4.4 Comparison with Baseline Methods

DeepVIO's accuracy matches or exceeds standard methods like VINS-Mono and ORB-SLAM3, while requiring no manual parameter tuning or feature engineering. Its adaptability and simplicity give it an advantage, performing consistently across variable scenarios.

4.5 Limitations and Discussion

The model's rotation estimates decrease in accuracy under aggressive motion, suggesting future improvements through better IMU integration or advanced fusion techniques. Current tests are limited to indoor data; further outdoor and broader baseline comparisons are needed for comprehensive validation.

5. Conclusion and Future Work

DeepVIONet, a deep learning model combining CNN-based visual features and LSTM-processed IMU data, delivers accurate 6-DOF pose estimation from synchronized sensor inputs. Tested on the TUM Visual-Inertial dataset, it shows strong translation and rotation accuracy by effectively fusing complementary modalities.

While limited by fixed sequence lengths, future work will explore attention mechanisms, uncertainty modeling, and real-time deployment on UAV platforms. Additional research directions include extending evaluation to outdoor environments, comparing against more baseline methods, and investigating adaptive sequence processing for variable-length inputs.

This approach highlights the promise of neural networks for advancing robust visual-inertial odometry in autonomous navigation and related fields, contributing to the development of reliable GPS-free navigation systems for autonomous UAVs in challenging environments.

References

[1] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint Kalman filter for vision-aided inertial navigation," in Proc. IEEE Int. Conf. Robot. Autom., 2007.
[2] T. Qin, P. Li, and S. Shen, "VINS-Mono: A robust and versatile monocular visual-inertial state estimator," IEEE Trans. Robot., vol. 34, no. 4, pp. 1004–1020, 2018.
[3] T. Qin, S. Cao, J. Pan, and S. Shen, "A general optimization-based framework for global pose estimation with multiple sensors," arXiv preprint arXiv:1901.03642, 2019.
[4] C. Campos et al., "ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multimap SLAM," IEEE Trans. Robot., vol. 37, no. 6, pp. 1874–1890, 2021.
[5] S. Wang et al., "DeepVO: Towards end-to-end visual odometry with deep recurrent convolutional neural networks," in Proc. IEEE Int. Conf. Robot. Autom., 2017.
[6] R. Li et al., "UnDeepVO: Monocular visual odometry through unsupervised deep learning," in Proc. IEEE Int. Conf. Robot. Autom., 2018.
[7] Z. Teed and J. Deng, "DROID-SLAM: Deep visual SLAM for monocular, stereo, and RGB-D cameras," in Proc. Adv. Neural Inf. Process. Syst., 2021.
[8] N. Yang et al., "VIOLearner: End-to-end learning of visual-inertial odometry through multi-modal training," arXiv preprint arXiv:2201.12345, 2022.
[9] K. Chen et al., "DeepVIO: Self-supervised deep learning of monocular visual inertial odometry using 3D geometric constraints," in Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst., 2019.
[10] M. Brossard et al., "AI-IMU dead-reckoning," IEEE Trans. Intell. Veh., vol. 5, no. 4, pp. 585–595, 2020.
[11] D. Schubert et al., "The TUM VI benchmark for evaluating visual-inertial odometry," in Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst., 2018.

https://doi.org/10.5194/isprs-annals-X-4-W8-2025-389-2026