

Adaptive Neural Dynamics for Robust Geometric LiDAR-Inertial State Estimation on UAVs

Sitian Peng¹  and Rui Wang^{2,1,3} *

¹ Shanghai Institute of Intelligent Science and Technology, Tongji University,
Shanghai 201804, China
2510239@tongji.edu.cn

² College of Electronic and Information Engineering, Tongji University, No. 4800
Cao'an Highway, Jiading District, Shanghai 201804, China
ruiwang@tongji.edu.cn

³ National College of Elite Engineers, Tongji University, Shanghai 201804, China

Abstract. This paper presents an enhanced LiDAR-inertial odometry system that addresses the challenge of spatiotemporal drift through an adaptive neural stochastic differential equation (SDE) framework. Traditional IMU bias models based on random walk assumptions often fail to capture the complex, high-order non-linearities inherent in low-cost inertial sensors during high-dynamic UAV maneuvers. To bridge the gap between classical geometric constraints and deep learning, our approach introduces three key innovations: an adaptive hybrid bias model that dynamically integrates learned neural dynamics with linear motion priors; a neural SDE framework that explicitly models stochastic bias evolution and heteroscedastic uncertainties; and a differentiable covariance propagation mechanism that ensures consistency between learned representations and geometric state estimation within an invariant extended Kalman filter. Extensive evaluations on self-collected UAV datasets, validated by a Vicon motion capture system, demonstrate that our method consistently outperforms the baseline FAST-LIO2, providing superior accuracy in both position and orientation estimation. Furthermore, our approach offers reliable uncertainty estimates crucial for complex spatial perception and robust autonomous flight in challenging environments.

Keywords: LiDAR-inertial odometry · IMU bias estimation · neural stochastic differential equations · adaptive filtering · uncertainty quantification

1 Introduction

The rapid evolution of Unmanned Aerial Vehicles (UAVs) has revolutionized tasks such as search and rescue, infrastructure inspection, and obstacle-aware autonomous navigation [30]. A fundamental prerequisite is robust state estimation in GPS-denied environments [2], where GNSS performance deteriorates

* Corresponding author.

due to signal attenuation or multi-path effects [24]. In such scenarios, LiDAR-inertial odometry (LIO) [4, 11, 14, 21] has emerged as a predominant solution, leveraging the complementary strengths of LiDAR geometric constraints and high-frequency IMU motion estimates [33, 45]. State-of-the-art frameworks like FAST-LIO2 [44] typically utilize an Invariant Extended Kalman Filter (IEKF) for efficient point-to-map registration in structured environments.

However, during high-dynamic maneuvers, these systems are constrained by the accuracy of IMU bias estimation. Low-cost MEMS IMUs exhibit complex, time-varying biases influenced by motor vibrations and mechanical stresses [17]. Traditional random walk [8] or first-order Gauss-Markov models [19] often fail to capture the high-order non-linearities inherent in aggressive flight dynamics [41], leading to suboptimal filter convergence and significant drift during rapid orientation changes.

To address this, researchers are exploring learning-based perspectives ranging from sequence-learning [31] to Neural Ordinary Differential Equations (NODEs) [43]. While promising, these neural representations often lack explicit mechanisms for capturing heteroscedastic vibrations and fail to provide the rigorous uncertainty quantification required for safety-critical filtering. Consequently, integrating these black-box models into a mathematically consistent IEKF framework remains a significant challenge.

In this work, we propose Neural Stochastic Differential Equation-enhanced LiDAR-Inertial Odometry (**NeuroSDE-LIO**), an adaptive neural Stochastic Differential Equation (SDE) framework for robust IMU bias representation. Our approach extends the neural ODE paradigm to incorporate explicit stochastic modeling, capturing both deterministic bias evolution and motor-induced heteroscedastic uncertainties. To ensure mathematical rigor, we derive a differentiable covariance propagation scheme that maps learned dynamics’ uncertainties into IEKF state prediction while maintaining geometric consistency. Furthermore, an adaptive hybrid representation dynamically balances neural dynamics with classical linear priors, ensuring robustness across diverse flight regimes from hovering to aggressive manual piloting.

Extensive evaluations on UAV datasets, validated by Vicon ground truth, demonstrate that NeuroSDE-LIO consistently outperforms the baseline FAST-LIO2 [44]. Our contributions are summarized as follows:

- We present a Neural SDE framework that captures non-linear bias evolution and heteroscedastic motor noise within a unified stochastic representation.
- We derive a differentiable uncertainty propagation mechanism enabling seamless integration of learned bias dynamics into IEKF covariance prediction.
- We propose an adaptive mixing mechanism to synergize neural dynamics with classical linear priors, enhancing stability across varying motion complexities.
- We provide empirical validation on high-dynamic flight data, demonstrating significant improvements in localization accuracy and robustness.

2 Related Work

LiDAR-Inertial Odometry (LIO) has progressed from loosely-coupled designs [34, 46] toward tightly-coupled frameworks that emphasize continuous-time formulations [26] and multi-sensor consistency [51]. The FAST-LIO family [44, 45] introduced direct point-to-map registration and Invariant Extended Kalman Filter (IEKF) based mapping to achieve high computational efficiency. However, despite recent progress in observability analysis and adaptive techniques [3, 39, 47], most state-of-the-art LIO systems still rely on simplified stochastic models. These traditional models often fail to adequately capture the complex, non-linear temporal dynamics of IMU biases, especially under the significant vibration and mechanical stress typical of high-dynamic UAV flight.

To overcome the limitations of analytical bias modeling, researchers have increasingly turned to learning-based IMU representations. Sequence-learning models, including Transformers [31, 36] and CNNs [35, 38, 42], have demonstrated superior denoising capabilities by capturing long-range temporal dependencies. These models are typically integrated into classical estimators [18, 49] or unified optimization pipelines for online calibration [22, 23]. More recently, Neural Ordinary Differential Equations (NODEs) have been employed to model bias evolution as continuous vector fields, offering an elegant way to accommodate irregular sampling without requiring explicit bias ground truth [10, 20, 29]. Nevertheless, existing NODE frameworks predominantly treat bias evolution as a deterministic process, largely neglecting the inherent heteroscedastic stochasticity present in real-world IMU measurements.

Stochastic Differential Equations (SDEs) provide a mathematically rigorous framework for modeling such random disturbances within filtering formulations [5, 6, 28]. While SDEs have been extensively used for observability evaluation [12] and stability proofs in invariant filtering [15, 27], the emergence of Neural SDEs [7, 7] has opened new avenues for learning non-linear and non-stationary dynamics. Although Neural SDEs have seen applications in trajectory prediction [1] and general sensor fusion [25, 37], their specific integration into tightly-coupled LIO frameworks remains a significant research gap.

Our work, NeuroSDE-LIO, builds upon the IEKF backbone of FAST-LIO2 by replacing simplified random-walk models with learned Neural SDE representations. By addressing limitations of deterministic learning models and bridging the gap between deep representational power and rigorous filtering, our approach makes the learned bias dynamics expressive while remaining consistent with the uncertainty propagation requirements of safety-critical navigation filters.

3 Methodology

This section presents the theoretical foundation and architectural design of the proposed NeuroSDE-LIO system, structured around four key innovations: adaptive hybrid bias modeling, a neural SDE framework, differentiable covariance propagation, and incremental online learning.

3.1 System Overview and Architectural Innovation




Fig. 1: Architecture of the NeuroSDE-LIO system, integrating adaptive neural SDE bias learning within the LiDAR-inertial framework.

NeuroSDE-LIO provides an adaptive neural SDE formulation for IMU bias estimation within a tightly-coupled LIO pipeline. As shown in Fig. 1, the system maintains a feedback loop: (1) raw IMU data is corrected via the adaptive hybrid model; (2) the neural SDE predicts bias dynamics based on motion context; (3) differentiable covariance propagation maps learned stochasticity into the IEKF; and (4) lightweight online adaptation updates low-dimensional calibration parameters under consistency checks. This architecture leverages deep representational power while preserving the theoretical rigor of classical filtering.

3.2 Adaptive Hybrid Bias Modeling

Motion Complexity Assessment and Theoretical Motivation The fundamental insight driving our adaptive approach is that the effectiveness of learning-based bias estimation varies significantly with motion characteristics. During periods of limited motion excitation, simple linear models often outperform complex neural networks due to their stability and lower variance. Conversely, complex maneuvers with rich dynamic content provide the excitation necessary for sophisticated learned models to demonstrate their superior representational capacity.

To operationalize this insight, we introduce a novel real-time motion complexity metric $C(t)$ that quantifies the diversity and intensity of vehicle motion through a sliding window analysis:

$$C(t) = \frac{1}{\tau} \int_{t-\tau}^t \left(\|\tilde{\omega}_s\|^2 + \|\tilde{\mathbf{a}}_s - \hat{\mathbf{R}}_s^T \mathbf{g}\|^2 \right) ds \quad (1)$$

where τ represents the temporal assessment window (typically 2–5 seconds), $\tilde{\omega}_s, \tilde{\mathbf{a}}_s \in \mathbb{R}^3$ denote raw gyroscope and accelerometer measurements at time s , and $\hat{\mathbf{R}}_s \in SO(3)$ is the current orientation estimate. This formulation captures both rotational and translational excitation while compensating for gravitational effects in the accelerometer measurements via the estimated attitude. The squared norms ensure that the metric is sensitive to both the magnitude and variability of motion, providing a comprehensive characterization of motion complexity.

Adaptive Mixing Mechanism and Stability Analysis The core innovation of our hybrid approach lies in the dynamic weighting mechanism that seamlessly transitions between learned and linear bias models based on real-time motion complexity. The adaptive mixing weights $\alpha_i(t)$ are computed through a smooth, differentiable function:

$$\alpha_i(t) = \sigma \left(\beta_i \left(\frac{C(t) - C_{\min}}{C_{\max} - C_{\min}} - \gamma_i \right) \right), \quad i \in \{g, a\} \quad (2)$$

where $\sigma(\cdot)$ represents the sigmoid function, β_i control the transition sharpness, and γ_i are bias parameters that determine the switching threshold. This formulation ensures that during low-complexity motion ($C(t) \rightarrow C_{\min}$), the system predominantly relies on the stable linear model ($\alpha \rightarrow 0$), while during high-complexity maneuvers ($C(t) \rightarrow C_{\max}$), it leverages the expressive power of the learned model ($\alpha \rightarrow 1$). The thresholds $C_{\min}$ and $C_{\max}$ are not universal constants: in our experiments they are obtained from a short static-and-aggressive calibration segment for each UAV/IMU setup, while the neural SDE weights remain unchanged.

The stability properties of this adaptive system can be analyzed through Lyapunov theory. Consider the hybrid bias estimation error $\tilde{\mathbf{b}} = \mathbf{b} - \hat{\mathbf{b}}$. Under mild assumptions of bounded motion complexity and Lipschitz continuous bias dynamics, the adaptive system exhibits exponential convergence in mean square:

$$\mathbb{E}[\|\tilde{\mathbf{b}}_t\|^2] \leq K e^{-\lambda t} \mathbb{E}[\|\tilde{\mathbf{b}}_0\|^2] + \frac{\text{tr}(\boldsymbol{\Sigma})}{\lambda_{\min}(\mathbf{Q})} \quad (3)$$

where $K, \lambda > 0$ are constants dependent on the adaptive mixing parameters, and $\lambda_{\min}(\mathbf{Q})$ represents the minimum eigenvalue of the process noise covariance matrix.

Hybrid Bias Dynamics Formulation The complete bias dynamics integrate both learned and linear components through the adaptive mixing mechanism:

$$db_t^g = [\alpha_g(t)f_g(b_t^g, \tilde{\mathbf{u}}_t; \boldsymbol{\theta}_g) + (1 - \alpha_g(t))\mathbf{L}_g\tilde{\mathbf{u}}_t] dt + \boldsymbol{\Sigma}_g^{1/2}d\mathbf{W}_t^g \quad (4)$$

$$db_t^a = [\alpha_a(t)f_a(b_t^a, \tilde{\mathbf{u}}_t; \boldsymbol{\theta}_a) + (1 - \alpha_a(t))\mathbf{L}_a\tilde{\mathbf{u}}_t] dt + \boldsymbol{\Sigma}_a^{1/2}d\mathbf{W}_t^a \quad (5)$$

where $\tilde{\mathbf{u}}_t = [\tilde{\boldsymbol{\omega}}_t^T, \tilde{\mathbf{a}}_t^T]^T \in \mathbb{R}^6$ is the IMU input vector, $\mathbf{L}_g, \mathbf{L}_a \in \mathbb{R}^{3 \times 6}$ represent linear model parameters, $\mathbf{W}_t^g, \mathbf{W}_t^a$ are independent Wiener processes, and $\boldsymbol{\Sigma}_g, \boldsymbol{\Sigma}_a \in \mathbb{R}^{3 \times 3}$ are diffusion matrices. This formulation provides: (1) robustness across diverse motion patterns, (2) graceful degradation, and (3) explicit modeling of both deterministic and stochastic bias components.

3.3 Neural Stochastic Differential Equation Framework

Extended State Space Formulation To seamlessly integrate learning within the filtering framework, we extend the conventional LiDAR-inertial state vector to incorporate bias learning parameters. The state vector is defined on the manifold $\mathcal{M} = SO(3) \times \mathbb{R}^{12} \times \mathbb{R}^p$:

$$\mathbf{X} = \left[\underbrace{\mathbf{R}, \mathbf{p}, \mathbf{v}}_{\text{kinematic states}}, \underbrace{\mathbf{b}^g, \mathbf{b}^a}_{\text{bias states}}, \underbrace{\boldsymbol{\theta}_g, \boldsymbol{\theta}_a}_{\text{network params}}, \underbrace{\mathbf{L}_g, \mathbf{L}_a}_{\text{linear models}}, \underbrace{\alpha_g, \alpha_a}_{\text{weights}} \right]^T \quad (6)$$

This extended formulation enables joint estimation of both navigation states and learning parameters. For real-time tractability, the covariance of $\boldsymbol{\theta}$ and $\mathbf{L}$ is often treated in a decoupled optimization thread, while $\mathbf{b}^i$ remain in the core filter.

Complete System Dynamics and Theoretical Properties The continuous-time system dynamics capture the deterministic evolution and stochastic disturbances on the Lie Group:

$$d\mathbf{R}_t = \mathbf{R}_t(\tilde{\boldsymbol{\omega}}_t - \mathbf{b}_t^g - \mathbf{n}_t^g)^\wedge dt \quad (7)$$

$$d\mathbf{v}_t = [\mathbf{R}_t(\tilde{\mathbf{a}}_t - \mathbf{b}_t^a - \mathbf{n}_t^a) + \mathbf{g}] dt \quad (8)$$

$$d\mathbf{p}_t = \mathbf{v}_t dt \quad (9)$$

$$db_t^g = [\alpha_g f_g + (1 - \alpha_g)\mathbf{L}_g\tilde{\mathbf{u}}_t] dt + \boldsymbol{\Sigma}_g^{1/2}d\mathbf{W}_t^g \quad (10)$$

$$db_t^a = [\alpha_a f_a + (1 - \alpha_a)\mathbf{L}_a\tilde{\mathbf{u}}_t] dt + \boldsymbol{\Sigma}_a^{1/2}d\mathbf{W}_t^a \quad (11)$$

where $\mathbf{n}_t^g, \mathbf{n}_t^a$ represent high-frequency white noise and $(\cdot)^\wedge : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ denotes the skew-symmetric matrix operator. The system preserves the following properties: Markov Property for efficient recursive estimation; Stochastic Stability under bounded drift; and Geometric Consistency on $SO(3)$.

Neural Network Architecture and Representational Analysis The networks f_g and f_a employ an enhanced architecture:

$$f_i(\mathbf{b}^i, \tilde{\mathbf{u}}_t; \boldsymbol{\theta}_i) = \text{MLP}_i \left([\mathbf{b}^i; \tilde{\mathbf{u}}_t; \dot{\tilde{\mathbf{u}}}_t; C(t)]; \boldsymbol{\theta}_i \right) + \text{ResBlock}_i([\mathbf{b}^i; \tilde{\mathbf{u}}_t]; \boldsymbol{\theta}_i^{\text{res}}) \quad (12)$$

This architecture incorporates: Multi-Scale Inputs including measurement derivatives $\dot{\tilde{\mathbf{u}}}_t$ for high-frequency variations; Motion Context $C(t)$ to guide the regime switching; Residual Learning for incremental corrections; and Efficiency with only 14,243 parameters per network.

3.4 Differentiable Covariance Propagation

Continuous-Time Uncertainty Dynamics The covariance matrix $\mathbf{P}_t$ evolves according to the differential Lyapunov equation:

$$\frac{d\mathbf{P}_t}{dt} = \mathbf{F}_t \mathbf{P}_t + \mathbf{P}_t \mathbf{F}_t^T + \mathbf{G}_t \mathbf{Q}_t \mathbf{G}_t^T \quad (13)$$

where $\mathbf{F}_t = \frac{\partial \delta \dot{\mathbf{x}}}{\partial \delta \mathbf{x}}$ is the error-state transition Jacobian, and $\mathbf{G}_t, \mathbf{Q}_t$ define the process-noise input. This extends classical EKF propagation to the learned manifold.

Neural Network Jacobian Computation We derive analytical expressions for the bias blocks of the error-state transition matrix $\mathbf{F}_t$:

$$\mathbf{F}_{\mathbf{b}^i \mathbf{b}^i} = \alpha_i \frac{\partial f_i}{\partial \mathbf{b}^i} + (1 - \alpha_i) \frac{\partial (\mathbf{L}_i \tilde{\mathbf{u}}_t)}{\partial \mathbf{b}^i} + \frac{\partial \alpha_i}{\partial C} \mathbf{J}_C (f_i - \mathbf{L}_i \tilde{\mathbf{u}}_t) \quad (14)$$

where $\mathbf{J}_C = \frac{\partial C}{\partial \delta \phi}$ denotes the Jacobian of motion complexity with respect to the orientation error state, as detailed in the supplementary material. This term explicitly couples the adaptive weighting mechanism with the system’s geometric uncertainty, ensuring that orientation errors correctly propagate into the bias estimation process.

While the linear components are computed analytically, automatic differentiation is leveraged to obtain the gradients of the neural network f_i with respect to its inputs and parameters. This hybrid approach enables efficient, real-time uncertainty propagation through the learned dynamics, ensuring the scalability and robustness of the NeuroSDE framework on embedded platforms like the NVIDIA Orin NX.

Discrete-Time Implementation We employ a second-order approximation for discrete propagation:

$$\mathbf{P}_{k+1|k} = \boldsymbol{\Phi}_k \mathbf{P}_{k|k} \boldsymbol{\Phi}_k^T + \mathbf{Q}_d, \quad \boldsymbol{\Phi}_k = \exp(\mathbf{F}_k \Delta t) \approx \mathbf{I} + \mathbf{F}_k \Delta t + \frac{1}{2} \mathbf{F}_k^2 \Delta t^2 \quad (15)$$

The sparse block-diagonal structure of $\boldsymbol{\Phi}_k$ ensures $\mathcal{O}(n)$ complexity for real-time operation.

3.5 Incremental Online Learning Framework

Sliding Window Optimization Formulation The Neural SDE weights are initialized by offline pre-training on held-out flights using Vicon-aided Rauch–Tung–Striebel smoothing labels for IMU bias residuals. They are not re-trained from scratch for each flight. During deployment, the reported system keeps the pre-trained neural weights fixed and only updates the low-dimensional linear branch and gating parameters, denoted by $\boldsymbol{\eta} = \{\beta_i, \gamma_i\}_{i \in \{g, a\}}$, in a short 2–5s sliding window. Updates are gradient-clipped and skipped when residual or normalized innovation squared (NIS) consistency checks fail, so neural gradients never overwrite the IEKF nominal state. We solve:

$$\min_{\mathbf{L}, \boldsymbol{\eta}} \sum_{t=k-W}^k \|\mathbf{z}_t - \mathbf{h}(\mathbf{X}_t)\|_{\Sigma_t^{-1}}^2 + \lambda_\theta \|\boldsymbol{\theta} - \boldsymbol{\theta}_{\text{pre}}\|^2 + \lambda_L \|\mathbf{L} - \mathbf{L}_{\text{prev}}\|_F^2 + \mathcal{L}_\alpha \quad (16)$$

where $\mathbf{z}_t$ are LiDAR measurements, $\mathbf{h}(\cdot)$ is the measurement model, $\boldsymbol{\theta}_{\text{pre}}$ is the offline pre-trained network, and regularization prevents drift away from the validated prior. This separation addresses the low-excitation case: the neural branch has sufficient capacity to fit simple motion, but pose residuals alone do not make IMU bias fully observable, so the linear branch acts as an observability prior and α_i is a confidence gate rather than a hard switch.

Multi-Objective Loss Function Design The training objective balances pose accuracy, parameter reliability, and filter stability through a multi-objective loss function. The total loss $\mathcal{L}$ is formulated as a weighted sum of four synergistic components:

$$\mathcal{L} = \mathcal{L}_{\text{pose}} + \lambda_1 \mathcal{L}_{\text{reg}} + \lambda_2 \mathcal{L}_{\text{uncertainty}} + \lambda_3 \mathcal{L}_{\text{smoothness}} \quad (17)$$

$$\mathcal{L}_{\text{uncertainty}} = \frac{1}{N} \sum_{i=1}^N [\text{tr}(\mathbf{P}_i) - \log \det(\mathbf{P}_i + \epsilon \mathbf{I})] \quad (18)$$

$$\mathcal{L}_{\text{smoothness}} = \frac{1}{N} \sum_{i=1}^N \left(\|\dot{\mathbf{b}}_i^g\|^2 + \|\dot{\mathbf{b}}_i^a\|^2 \right) \quad (19)$$

Each component addresses a specific requirement of the learning-enhanced estimation process:

- Pose Accuracy Loss ($\mathcal{L}_{\text{pose}}$): This term minimizes the residual between the estimated navigation states and the reference constraints within the sliding window. It ensures that the neural SDE-derived biases lead to physically consistent trajectories:

$$\mathcal{L}_{\text{pose}} = \frac{1}{N} \sum_{i=1}^N \left(\|\log(\hat{\mathbf{R}}_i \mathbf{R}_i^T)\|^2 + \|\hat{\mathbf{p}}_i - \mathbf{p}_i\|^2 \right) \quad (20)$$

where $\hat{\mathbf{R}}_i \in SO(3)$ and $\hat{\mathbf{p}}_i \in \mathbb{R}^3$ represent the estimated orientation and position, respectively.

- Parameter Regularization ($\mathcal{L}_{\text{reg}}$): To prevent overfitting to transient sensor noise and ensure numerical stability during online optimization, we apply regularization to the network weights $\boldsymbol{\theta}$ and linear parameters $\mathbf{L}$:

$$\mathcal{L}_{\text{reg}} = \|\boldsymbol{\theta}_g\|^2 + \|\boldsymbol{\theta}_a\|^2 + \|\mathbf{L}_g\|_F^2 + \|\mathbf{L}_a\|_F^2 \quad (21)$$

where $\|\cdot\|_F$ denotes the Frobenius norm for the linear transition matrices.

- Uncertainty Calibration ($\mathcal{L}_{\text{uncertainty}}$): This term calibrates the filter’s confidence by optimizing the structure of the error-state covariance matrix $\mathbf{P}_i$. The trace term $\text{tr}(\mathbf{P}_i)$ minimizes the total estimation uncertainty, while the negative log det term prevents the covariance from collapsing into a singular matrix, maintaining a well-conditioned filter.
- Temporal Smoothness ($\mathcal{L}_{\text{smoothness}}$): By penalizing the squared norm of the bias derivatives $\dot{\mathbf{b}}^i$, this component ensures that the learned bias estimates follow a continuous and plausible evolution, effectively filtering out high-frequency estimation jitter.

The weighting factors $\lambda_1, \lambda_2, \lambda_3$ are adaptively tuned to balance these competing objectives, facilitating a robust integration of deep representational power and classical estimation theory.

4 Experimental Results

4.1 Experimental Setup

Experimental Platform To validate the proposed NeuroSDE-LIO system in high-dynamic scenarios, we developed a customized quadrotor UAV platform, as shown in **Fig. 2(a)**. The hardware suite comprises a Livox Mid-360 LiDAR for 3D spatial sensing and a high-frequency BMI088 IMU. The onboard computation is handled by an NVIDIA Orin NX embedded computer, which executes the real-time EKF propagation and neural SDE inference.

For ground truth validation, a 32-camera Vicon Vantage motion capture system is employed within a $50 \times 30 \times 8$ m³ flight arena (**Fig. 2(b)**). The UAV is equipped with a specific constellation of retro-reflective markers, allowing the Vicon cameras (**Fig. 2(c)**) to track its 6-DOF pose with millimeter-level accuracy by capturing reflected infrared signals. This setup ensures highly reliable reference data for assessing the performance of our adaptive SDE framework during aggressive maneuvers.

Datasets and Evaluation Metrics We conduct comprehensive evaluations on three diverse datasets to assess performance under various operating conditions. The UrbanLoco dataset [40] provides complex urban driving scenarios with dense structures and GPS-denied environments. The KITTI dataset [13] serves as a




Fig. 2: Experimental setup for high-dynamic flight validation. (a) The customized quadrotor UAV platform integrated with a Livox Mid-360 LiDAR and an NVIDIA Orin NX module. (b) The flight arena equipped with a Vicon Vantage motion capture system consisting of multiple cameras. (c) Close-up view of a Vicon Vantage camera mounted on the ceiling for high-precision infrared tracking.

public structured-road sanity check, rather than the primary evidence for high-dynamic UAV flight. Crucially, we employ a self-collected UAV dataset, recorded using a specialized aerial platform in a laboratory environment. This dataset is characterized by aggressive flight dynamics, severe motor-induced vibrations, and high-speed maneuvers, specifically designed to test the limits of inertial modeling under extreme conditions. For this dataset, a Vicon motion capture system is utilized to provide millimeter-level ground truth for validation. The flights used for testing are held out from offline neural pre-training and hyper-parameter selection.

For quantitative evaluation, we employ established metrics: Absolute Pose Error (APE) measures the global consistency of the trajectory, and Relative Pose Error (RPE) evaluates local tracking accuracy over short segments. We also assess computational efficiency including processing time, memory footprint, and CPU utilization to ensure practical feasibility for onboard embedded systems. In the supplementary videos, the estimator is further coupled with the FUEL exploration planner [52] to demonstrate closed-loop autonomous navigation.

Compared Methods We compare **NeuroSDE-LIO** against several representative state-of-the-art approaches: **ORB-SLAM3** [9] is run in stereo-inertial mode on the synchronized camera-IMU stream; **LIO-SAM** [33] represents graph-based optimization with loop closure capabilities; and **FAST-LIO2** [44] serves as our direct IEKF baseline. Additionally, we evaluate a controlled neural ODE variant, **FAST-LIO2-NODE**: it uses the same FAST-LIO2 backend, training loss, and bias-input interface as our method, but replaces the stochastic Neural SDE with a deterministic NODE drift model and disables the learned diffusion head. To address recent high-dynamic and motion-degradation baselines, we also evaluate SUPER [32], SR-LIO [48], Adaptive-LIO [50], and iG-LIO [11] on the UAV-Lab sequences. NVIDIA cuVSLAM [16] is included only as a hardware-

Table 1: Overall localization accuracy comparison (APE in meters)

Method	Urban-LoCo	KITTI	UAV-Lab	Average
ORB-SLAM3	1.87 _(↑0.98)	1.08 _(↑0.62)	2.98 _(↑1.64)	1.98 _(↑1.08)
LIO-SAM	2.13 _(↑1.24)	1.27 _(↑0.81)	3.45 _(↑2.11)	2.28 _(↑1.38)
FAST-LIO2	1.42 _(↑0.53)	0.73 _(↑0.27)	2.13 _(↑0.79)	1.43 _(↑0.53)
FAST-LIO2-NODE	1.15 _(↑0.26)	0.61 _(↑0.15)	1.78 _(↑0.44)	1.18 _(↑0.28)
NeuroSDE-LIO (Ours)	0.89	0.46	1.34	0.90

Table 2: Additional recent baselines on UAV-Lab. APE is in meters; RPE is in deg/m. cuVSLAM is a VIO runtime reference and is not ranked as an LIO method.

Method	Type	APE	RPE	Hz
FAST-LIO2	LIO	2.13 _(↑0.79)	1.87 _(↑0.75)	22.1 _(↑2.9)
SUPER	LIO	1.62 _(↑0.28)	1.41 _(↑0.29)	16.8 _(↓2.4)
SR-LIO	LIO	1.74 _(↑0.40)	1.53 _(↑0.41)	13.5 _(↓5.7)
Adaptive-LIO	LIO	1.56 _(↑0.22)	1.35 _(↑0.23)	18.1 _(↓1.1)
iG-LIO	LIO	1.48 _(↑0.14)	1.29 _(↑0.17)	11.9 _(↓7.3)
NeuroSDE-LIO	LIO	1.34	1.12	19.2
cuVSLAM	VIO ref.	1.91 _(↑0.57)	1.68 _(↑0.56)	28.5 _(↑9.3)

aware VIO runtime reference on the Jetson Orin NX, not as a sensor-equivalent LIO competitor.

4.2 Quantitative Results

Overall Localization Accuracy Table 1 compares APE across all datasets, with the largest gain appearing on the high-dynamic UAV-Lab sequences.

Our method improves average APE by 37.4% over FAST-LIO2 and 23.7% over FAST-LIO2-NODE. On UAV-Lab, it reduces FAST-LIO2 error by 37.1%, showing the benefit of stochastic bias modeling under aggressive dynamics.

Table 2 adds recent UAV-Lab baselines under the same LiDAR-IMU logs. These stronger motion-aware LIO methods improve over FAST-LIO2, while NeuroSDE-LIO remains the most accurate at real-time throughput. Cerberus-style legged VIO and imperative-learning SLAM are discussed qualitatively because their sensors and optimization interface are not drop-in comparable.

Fig. 3 further shows four independent manual flights with Vicon ground truth. NeuroSDE-LIO remains closest to the reference, especially during rapid heading changes where FAST-LIO2 visibly deviates.

Local Tracking Precision Analysis In addition to global APE, Table 3 reports RPE for local tracking accuracy.

NeuroSDE-LIO reduces average RPE by 45.2% over FAST-LIO2, mainly by modeling high-frequency gyroscope bias dynamics and heteroscedastic motor noise.




Fig. 3: Qualitative trajectory comparison across four representative manual flight missions (a-d) conducted in our laboratory environment. Each plot displays five trajectories: the **Groundtruth** provided by a **Vicon motion capture system** (dashed black), three state-of-the-art baselines (**ORB-SLAM3**, **LIO-SAM**, and **FAST-LIO2**), and our **NeuroSDE-LIO** (red).

Table 3: Relative Pose Error (RPE) comparison (deg/m)

Method	Urban-LoCo	KITTI	UAV-Lab	Average
ORB-SLAM3	1.65 ($\uparrow 0.92$)	0.94 ($\uparrow 0.56$)	2.45 ($\uparrow 1.33$)	1.68 ($\uparrow 0.94$)
LIO-SAM	1.87 ($\uparrow 1.14$)	1.12 ($\uparrow 0.74$)	2.89 ($\uparrow 1.77$)	1.96 ($\uparrow 1.22$)
FAST-LIO2	1.28 ($\uparrow 0.55$)	0.68 ($\uparrow 0.30$)	1.87 ($\uparrow 0.75$)	1.28 ($\uparrow 0.54$)
FAST-LIO2-NODE	1.02 ($\uparrow 0.29$)	0.55 ($\uparrow 0.17$)	1.52 ($\uparrow 0.40$)	1.03 ($\uparrow 0.29$)
NeuroSDE-LIO (Ours)	0.73	0.38	1.12	0.74

Uncertainty Quantification Performance Table 4 evaluates uncertainty calibration with normalized innovation squared (NIS); an ideal filter places about 95% of NIS values within the 95% confidence interval.

NeuroSDE-LIO reaches 93.0% average consistency, closest to the ideal 95%. ORB-SLAM3, LIO-SAM, and FAST-LIO2-NODE remain more over-confident, especially on UAV-Lab.

Table 4: Uncertainty estimation consistency: Percentage of NIS values within the 95% confidence interval (Higher is better).

Method	KITTI (%)	UAV-Lab (%)	Average (%)
ORB-SLAM3	78.5 _(↓15.6)	65.8 _(↓26.0)	72.2 _(↓20.8)
LIO-SAM	72.4 _(↓21.7)	61.2 _(↓30.6)	66.8 _(↓26.2)
FAST-LIO2	84.2 _(↓9.9)	70.5 _(↓21.3)	77.4 _(↓15.6)
FAST-LIO2-NODE	88.9 _(↓5.2)	79.3 _(↓12.5)	84.1 _(↓8.9)
NeuroSDE-LIO (Ours)	94.1	91.8	93.0

Table 5: Leave-one-out and requested NODE ablation on UAV-Lab.

Variant	Base	Hybrid	NODE	SDE	APE	RPE
FAST-LIO2	yes	no	no	no	2.13 _(↑0.79)	1.87 _(↑0.75)
Base+Hybrid	yes	yes	no	no	1.76 _(↑0.42)	1.55 _(↑0.43)
Base+Hybrid+NODE	yes	yes	yes	no	1.52 _(↑0.18)	1.34 _(↑0.22)
Full w/o Hybrid	yes	no	yes	yes	1.63 _(↑0.29)	1.42 _(↑0.30)
Full NeuroSDE-LIO	yes	yes	yes	yes	1.34	1.12

Table 6: Computational performance comparison

Method	Processing Time (ms)	Memory (MB)	CPU Usage (%)
FAST-LIO2	45.2 _(↓6.9)	312 _(↓73)	62.3 _(↓8.9)
FAST-LIO2-NODE	58.7 _(↑6.6)	428 _(↑43)	78.5 _(↑7.3)
Full w/o mixing	51.6 _(↓0.5)	381 _(↓4)	70.4 _(↓0.8)
NeuroSDE-LIO (Ours)	52.1	385	71.2

4.3 Ablation Studies

Table 5 gives the requested leave-one-out/NODE ablation. “Base+Hybrid” has no neural dynamics, “Base+Hybrid+NODE” disables SDE diffusion, and “Full w/o Hybrid” removes the adaptive gate.

Both neural dynamics and the hybrid gate are necessary: NODE improves over the non-neural hybrid baseline, SDE further improves under heteroscedastic noise, and removing the gate degrades APE from 1.34 m to 1.63 m.

4.4 Computational Performance

Table 6 reports embedded computational cost.

NeuroSDE-LIO runs at 52.1 ms per LiDAR scan (19.2 Hz), with only 15.3% latency increase over FAST-LIO2 and less overhead than FAST-LIO2-NODE. The w/o-mixing row shows that the adaptive gate is nearly free; IMU pre-integration and state propagation remain at 200 Hz, and the 385 MB footprint fits the NVIDIA Orin NX.




Fig. 4: GT vs. estimated IMU bias on one held-out UAV-Lab flight. Ground-truth bias is obtained by Vicon-aided offline smoothing; the plotted sequence is not used for pre-training or hyper-parameter selection.

4.5 Analysis and Interpretations

The results show positional and orientation gains of 37.4% and 45.2%, respectively. Fig. 4 directly validates bias estimation on a held-out flight: mean 6D bias error decreases from 0.05329 to 0.02115; gyro error from 0.00914 to 0.00363 rad/s; and accelerometer error from 0.05214 to 0.02070 m/s².

Together with the NIS results, this confirms that differentiable covariance propagation integrates learned bias stochasticity into geometric estimation while preserving 200 Hz IMU propagation on the Orin NX.

5 Conclusion

This paper presents **NeuroSDE-LIO**, an adaptive Neural SDE framework for robust LiDAR-inertial state estimation under high-dynamic UAV motion. By combining linear priors, stochastic learned bias dynamics, and differentiable covariance propagation within an IEKF, the method improves localization accuracy, uncertainty reliability, and direct IMU-bias agreement on Vicon-validated flight data. The results indicate that learned stochastic bias models can be integrated into geometric LIO while preserving real-time onboard operation.

Limitations and Future Work. Neural SDE inference adds overhead compared with vanilla EKFs, and extreme sensor saturation remains challenging when inertial measurements exceed the observable correction range. Future work will extend the framework to multi-IMU and cross-sensor transfer settings, improve robustness under rare impact-like events, and study certifiable safety guarantees for learning-enhanced EKFs.

Acknowledgments and Disclosure of Funding. This work was supported in part by the Fundamental Research Funds for the Central Universities under Grant 22120260207 and in part by the National Natural Science Foundation of China under Grant 62271352.

References

1. Alam, M.M., Soares, A., Rodrigues-Jr, J.F., Spadon, G.: Physics-informed neural networks for vessel trajectory prediction: learning time-discretized kinematic dynamics via finite differences. In: Proceedings of the 19th International Symposium on Spatial and Temporal Data. pp. 55–65 (2025)
2. Alghamdi, S., Alahmari, S., Yonbawi, S., Alsaleem, K., Ateeq, F., Almushir, F.: Autonomous navigation systems in gps-denied environments: a review of techniques and applications. In: 2025 11th International Conference on Automation, Robotics, and Applications (ICARA). pp. 290–299. IEEE (2025)
3. Altawaitan, A., Stanley, J., Ghosal, S., Duong, T., Atanasov, N.: Learned imu bias prediction for invariant visual inertial odometry. *IEEE Robotics and Automation Letters* **10**(11), 11872–11879 (2025)
4. Bai, C., Xiao, T., Chen, Y., Wang, H., Zhang, F., Gao, X.: Faster-lio: Lightweight tightly coupled lidar-inertial odometry using parallel sparse incremental voxels. *IEEE Robotics and Automation Letters* **7**(2), 4861–4868 (2022)
5. Barrau, A., Bonnabel, S.: The invariant extended kalman filter as a stable observer. *IEEE Transactions on Automatic Control* **62**(4), 1797–1812 (2017)
6. Bedrossian, J., Blumenthal, A., Punshon-Smith, S.: A regularity method for lower bounds on the lyapunov exponent for stochastic differential equations. *Inventiones mathematicae* **227**(2), 429–516 (2022)
7. Boral, A., Wan, Z.Y., Zepeda-Núñez, L., Lottes, J., Wang, Q., Chen, Y.f., Anderson, J., Sha, F.: Neural ideal large eddy simulation: Modeling turbulence with neural stochastic differential equations. *Advances in neural information processing systems* **36**, 69270–69283 (2023)
8. Buchanan, R., Agrawal, V., Camurri, M., Dellaert, F., Fallon, M.: Deep imu bias inference for robust visual-inertial odometry with factor graphs. *IEEE Robotics and Automation Letters* **8**(1), 41–48 (2022)
9. Campos, C., Elvira, R., Rodríguez, J.J.G., M. Montiel, J.M., D. Tardós, J.: Orbslam3: An accurate open-source library for visual, visual-inertial, and multimap slam. *IEEE Transactions on Robotics* **37**(6), 1874–1890 (2021)
10. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. *Advances in neural information processing systems* **31** (2018)
11. Chen, Z., Xu, Y., Yuan, S., Xie, L.: ig-lio: An incremental gicp-based tightly-coupled lidar-inertial odometry. *IEEE Robotics and Automation Letters* **9**(2), 1883–1890 (2024)
12. Dong, T., Wang, D., Fu, F., Li, M., Hou, B., Xu, C.: Analytical nonlinear observability analysis for spacecraft autonomous relative navigation. *IEEE Transactions on Aerospace and Electronic Systems* **58**(6), 5875–5893 (2022)
13. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The international journal of robotics research* **32**(11), 1231–1237 (2013)
14. Huang, K., Zhao, J., Zhu, Z., Ye, C., Feng, T.: Log-lio: A lidar-inertial odometry with efficient local geometric information estimation. *IEEE Robotics and Automation Letters* **9**(1), 459–466 (2023)
15. Kim, J.W., Mehta, P.G.: Variance decay property for filter stability. *IEEE Transactions on Automatic Control* **69**(12), 8140–8155 (2024)
16. Korovko, A., Slepichev, D., Eftorov, A., Dzhumamuratova, A., Kuznetsov, V., Rabeti, H., Biswas, J., Pouya, S.: cuvslam: Cuda accelerated visual odometry and mapping. *arXiv preprint arXiv:2506.04359* (2025)

17. Laidig, D., Seel, T.: Vqf: Highly accurate imu orientation estimation with bias estimation and magnetic disturbance rejection. *Information Fusion* **91**, 187–204 (2023)
18. Lee, S., Kim, H.B., Kim, K.S.: Legged robot state estimation using invariant neural-augmented kalman filter with a neural compensator. In: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 15445–15452 (2025)
19. Lethander, K.A., Taylor, C.N.: Conservative estimation of inertial sensor errors using allan variance data. *NAVIGATION: Journal of the Institute of Navigation* **70**(3) (2023)
20. Li, F., Tang, Y.: Multi-sensor fusion boolean bayesian filtering for stochastic boolean networks. *IEEE Transactions on Neural Networks and Learning Systems* **34**(10), 7114–7124 (2023)
21. Li, K., Li, M., Hanebeck, U.D.: Towards high-performance solid-state-lidar-inertial odometry and mapping. *IEEE Robotics and Automation Letters* **6**(3), 5167–5174 (2021)
22. Li, S., Li, X., Chen, S., Zhou, Y., Wang, S.: Targetless spatiotemporal calibration of multi-lidar multi-imu system based on continuous-time optimization. *IEEE Transactions on Industrial Informatics* **20**(5), 7565–7575 (2024)
23. Li, S., Li, X., Chen, S., Zhou, Y., Wang, S.: Two-step lidar/camera/imu spatial and temporal calibration based on continuous-time trajectory estimation. *IEEE Transactions on Industrial Electronics* **71**(3), 3182–3191 (2024)
24. Liu, W., Ren, Y., Guo, R., Kong, V.W., Hung, A.S., Zhu, F., Cai, Y., Wu, H., Zou, Y., Zhang, F.: Slope inspection under dense vegetation using lidar-based quadrotors. *Nature Communications* **16**(1), 7411 (2025)
25. Lou, J., Zhang, R.: Neural ode-based continuous imu debiasing for robust gnss/ins integration. In: 2025 IEEE SENSORS. pp. 1–5. IEEE (2025)
26. Lv, J., Lang, X., Xu, J., Wang, M., Liu, Y., Zuo, X.: Continuous-time fixed-lag smoothing for lidar-inertial-camera slam. *IEEE/ASME Transactions on Mechatronics* **28**(4), 2259–2270 (2023)
27. McDonald, C., Yüksel, S.: Stochastic observability and filter stability under several criteria. *IEEE Transactions on Automatic Control* **69**(5), 2931–2946 (2023)
28. Nakka, Y.K., Chung, S.J.: Trajectory optimization of chance-constrained nonlinear stochastic systems for motion planning under uncertainty. *IEEE Transactions on Robotics* **39**(1), 203–222 (2022)
29. Park, S., Kim, K., Lee, J., Choo, J., Lee, J., Kim, S., Choi, E.: Vid-ode: Continuous-time video generation with neural ordinary differential equation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2412–2422 (2021)
30. Peng, S., Wang, R., Liang, Y., Wang, L.: Mbd-planner: A real-time obstacle avoidance framework for uavs via feature-domain motion blur decoupling. *IEEE Robotics and Automation Letters* **11**(5), 6226–6233 (2026)
31. Rao, B., Kazemi, E., Ding, Y., Shila, D.M., Tucker, F.M., Wang, L.: Ctin: Robust contextual transformer network for inertial navigation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 5413–5421 (2022)
32. Ren, Y., Zhu, F., Lu, G., Cai, Y., Yin, L., Kong, F., Lin, J., Chen, N., Zhang, F.: Safety-assured high-speed navigation for mavs. *Science Robotics* **10**(98), eado6187 (2025)
33. Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., Rus, D.: Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5135–5142 (2020)

34. Shan, T., Englot, B., Ratti, C., Rus, D.: Lvi-sam: Tightly-coupled lidar-visual-inertial odometry via smoothing and mapping. In: 2021 IEEE international conference on robotics and automation (ICRA). pp. 5692–5698. IEEE (2021)
35. Shao, X., Zhang, L., Zhang, T., Shen, Y., Li, H., Zhou, Y.: A tightly-coupled semantic slam system with visual, inertial and surround-view sensors for autonomous indoor parking. In: Proceedings of the 28th ACM international conference on multimedia. pp. 2691–2699 (2020)
36. Singh, A.: Transformer-based sensor fusion for autonomous driving: A survey. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3312–3317 (2023)
37. Singh, S., Ramirez, F.M., Varley, J., Zeng, A., Sindhvani, V.: Multiscale sensor fusion and continuous control with neural cdes. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 10897–10904 (2022)
38. Wang, S., Clark, R., Wen, H., Trigoni, N.: End-to-end, sequence-to-sequence probabilistic visual odometry through deep neural networks. The International journal of robotics research **37**(4-5), 513–542 (2018)
39. Wang, Y., Song, W., Lou, Y., Zhang, Y., Huang, F., Tu, Z., Liang, Q.: Rail vehicle localization and mapping with lidar-vision-inertial-gnss fusion. IEEE Robotics and Automation Letters **7**(4), 9818–9825 (2022)
40. Wen, W., Zhou, Y., Zhang, G., Fahandezh-Saadi, S., Bai, X., Zhan, W., Tomizuka, M., Hsu, L.T.: Urbanloco: A full sensor suite dataset for mapping and localization in urban scenes. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 2310–2316. IEEE (2020)
41. Wu, Y., Yin, L., Guo, S., Qin, Y., et al.: Accurate and steady inertial pose estimation through sequence structure learning and modulation. Advances in Neural Information Processing Systems **37**, 42468–42493 (2024)
42. Xiao, J., Liu, Y., Wei, X.: Region-aware sequence-to-sequence learning for hyperspectral denoising. In: European Conference on Computer Vision. pp. 218–235. Springer (2024)
43. Xiong, D., Lu, Y., Han, W., Jiang, Z., Wang, X., Huang, Y., Li, T.: Neural ode network-based dynamic object trajectory prediction. In: 2025 9th International Conference on Robotics and Automation Sciences (ICRAS). pp. 160–166. IEEE (2025)
44. Xu, W., Cai, Y., He, D., Lin, J., Zhang, F.: Fast-lio2: Fast direct lidar-inertial odometry. IEEE Transactions on Robotics **38**(4), 2053–2073 (2022)
45. Xu, W., Zhang, F.: Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. IEEE Robotics and Automation Letters **6**(2), 3317–3324 (2021)
46. Ye, H., Chen, Y., Liu, M.: Tightly coupled 3d lidar inertial odometry and mapping. In: 2019 international conference on robotics and automation (ICRA). pp. 3144–3150. IEEE (2019)
47. Yu, Y., Fan, Y., Lai, G., Shi, C., Sun, F., Meng, X.: Tight fusion of odometry, kinematic constraints, and uwb ranging systems for state estimation of integrated aerial platforms. IEEE Transactions on Automation Science and Engineering **22**, 18297–18313 (2025)
48. Yuan, Z., Lang, F., Xu, T., Yang, X.: Sr-lio: Lidar-inertial odometry with sweep reconstruction. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7862–7869 (2024)
49. Zhang, Z., Li, Y., Wang, J., Liu, Z., Jiang, G., Guo, H., Zhu, W.: A hybrid data-driven and learning-based method for denoising low-cost imu to enhance ship navigation reliability. Ocean Engineering **299**, 117280 (2024)

- 50. Zhao, C., Hu, K., Xu, J., Zhao, L., Han, B., Wu, K., Tian, M., Yuan, S.: Adaptive-lio: Enhancing robustness and precision through environmental adaptation in lidar inertial odometry. *IEEE Internet of Things Journal* **12**(9), 12123–12136 (2025)
- 51. Zhong, X., Li, Y., Zhu, S., Chen, W., Li, X., Gu, J.: Lvio-sam: A multi-sensor fusion odometry via smoothing and mapping. In: 2021 IEEE International Conference on Robotics and Biomimetics (ROBIO). pp. 440–445. IEEE (2021)
- 52. Zhou, B., Zhang, Y., Chen, X., Shen, S.: Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning. *IEEE Robotics and Automation Letters* **6**(2), 779–786 (2021)