International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Toward Agentic Driver Intent Prediction from Smartphone Inertial Sensors: A Conceptual Framework for Future Study

Gaurav Pokharkar
Independent Researcher
gpokhark@asu.edu

Abstract

Predicting driver intent, such as braking, lane changes, and turns, before execution remains a critical challenge for advanced driver assistance systems (ADAS). Existing approaches rely on expensive dedicated hardware (cameras, LiDAR, CAN bus) and generic models that fail to capture individual driving signatures. This paper presents an agentic framework for driver intent prediction using only smartphone inertial measurement unit (IMU) sensors, specifically accelerometer and gyroscope data. We introduce a three-tier agent architecture consisting of a Signal Processing Agent for extracting pre-motion micro-patterns, a Driver Profiling Agent that maintains personalized driving behavior models through continual learning, and a Prediction Agent that fuses contextual features with driver-specific signatures to anticipate maneuvers before execution. To address the fundamental signal-to-noise ratio challenge inherent in weak pre-motion signals, we propose a wavelet-enhanced temporal attention network (WE-TAN) that selectively amplifies discriminative micro-patterns while suppressing sensor noise. No new data collection, model training, or quantitative testing is reported in the present paper. Instead, the contribution is a conceptual research framework and a detailed future validation pipeline covering data collection, labeling, model training, baseline comparison, robustness analysis, and simulation-in-the-loop evaluation. The goal is to provide a credible roadmap for how smartphone-only personalized intent prediction can be studied rigorously in future work.

Keywords: Driver intent prediction, smartphone sensors, agentic AI, personalized driving model, lane change prediction, braking prediction

1 Introduction

Anticipating driver maneuvers before they occur is a foundational capability for next-generation intelligent transportation systems. Accurate prediction of braking events, lane changes, and turning maneuvers with lead times of 2–4 seconds can enable proactive collision avoidance, adaptive cruise control refinement, and real-time risk assessment [1], [2]. However, the majority of existing driver intent prediction systems depend on vehicle-mounted sensors, such as forward-facing cameras, LiDAR arrays, radar units, or direct CAN bus telemetry, that are expensive, vehicle-specific, and unavailable in the vast majority of vehicles on the road today [3].

Smartphones, by contrast, are nearly ubiquitous. Modern devices embed high-frequency inertial measurement units (IMUs) comprising tri-axial accelerometers and gyroscopes capable of sampling at 100–400 Hz [4]. These sensors capture the kinematic signature of vehicle motion with sufficient fidelity to detect driving events [5]. Yet, leveraging smartphone IMU data for predictive (rather than reactive) intent recognition introduces a fundamental challenge: the pre-motion signals that precede a maneuver, subtle micro-decelerations before braking, minor lateral drift before a lane change, characteristic speed modulation before a turn, are extremely weak relative to road noise, device vibration, and mounting variability [6].

This signal-to-noise ratio (SNR) problem is compounded by inter-driver variability. Two drivers approaching the same intersection may exhibit vastly different preparatory behaviors: one may decelerate gradually over 4 seconds while another maintains speed until 1.5 seconds before the turn [7]. Population-level models trained on aggregate data inevitably learn a blurred average of these diverse signatures, leading to high false positive rates and poor generalization to individual drivers [8].

We argue that addressing these challenges requires a fundamentally different architectural paradigm, one inspired by agentic AI principles [9], [10]. Rather than deploying a monolithic prediction model, we propose a multi-agent system in which specialized agents collaboratively handle signal processing, driver profiling, and intent prediction. Each agent maintains its own state, adapts to incoming data, and communicates through a structured message-passing protocol. This decomposition offers three key advantages:

1) Specialization: The Signal Processing Agent applies wavelet-domain denoising tailored to the specific noise characteristics of each device-vehicle combination, extracting clean micro-pattern features from raw IMU streams.
2) Personalization: The Driver Profiling Agent builds and continuously updates a per-driver behavioral signature using continual learning, capturing individual habits without catastrophic forgetting of general driving patterns.
3) Adaptive Prediction: The Prediction Agent fuses micro-pattern features with driver-specific context through a temporal attention mechanism, dynamically weighting recent versus historical behavior based on driving conditions.

The principal contributions of this paper are as follows:

- We introduce a three-tier agentic architecture for smartphone-based driver intent prediction that decouples signal processing, driver modeling, and prediction into autonomous, communicating agents.
- We propose the Wavelet-Enhanced Temporal Attention Network (WE-TAN), a novel architecture that performs multi-scale wavelet decomposition within the attention mechanism to amplify weak pre-motion micro-patterns.
- We design a continual learning driver profiling strategy that maintains personalized driving signatures while accommodating behavioral drift over time.
- We outline a future-work validation pipeline describing how the proposed framework can be studied through data collection, annotation, model training, baseline comparison, and deployment-oriented validation.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 presents the proposed methodology including the agentic architecture and WE-TAN model. Section 4 describes a proposed validation pipeline for future study. Section 5 discusses implications and limitations, and Section 6 concludes the paper.

2 Related Work

2.1 Driver Intent Prediction

Driver intent prediction has been extensively studied using vehicle-mounted sensors. Jain et al. [3] demonstrated anticipation of maneuvers using recurrent neural networks (RNNs) fusing camera, GPS, and vehicle dynamics data, achieving prediction horizons of 3.5 seconds for lane changes. Zyner et al. [11] employed LSTM networks on naturalistic driving data to predict driver paths at intersections, highlighting the importance of temporal context. Xing et al. [1] provided a comprehensive survey of lane change intention inference, categorizing approaches into vehicle dynamics-based, driver behavior-based, and traffic context-based methods. Li et al. [12] introduced attention-augmented LSTMs for lane change prediction, demonstrating that selective temporal attention significantly improves prediction accuracy over vanilla recurrent architectures.

A common limitation across these works is the reliance on CAN bus signals (steering angle, throttle position, brake pressure) or external perception sensors (cameras, LiDAR) that are inaccessible to aftermarket or smartphone-based systems. Our work addresses this gap by operating exclusively on smartphone IMU data.

2.2 Smartphone-Based Driving Analysis

The use of smartphones for driving behavior analysis has grown substantially. Johnson and Trivedi [4] were among the first to demonstrate driving style recognition using smartphone accelerometer data. Wang et al. [13] applied machine learning classifiers to smartphone sensor data for driving behavior classification, achieving over 90% accuracy for aggressive versus normal driving detection. Chen et al. [5] showed that smartphone IMUs can estimate vehicle trajectories and detect lane changes with reasonable accuracy. Abdelrahman and Hassanein [14] used deep learning on smartphone sensor streams for real-time driver behavior classification.

However, these works predominantly focus on classification of completed maneuvers rather than prediction of future intent. The critical distinction is temporal: detecting a lane change as it occurs is fundamentally different from anticipating it 2–4 seconds before onset. Ren et al. [6] moved closer to prediction but relied on sliding window approaches that do not explicitly model the subtle pre-motion patterns we target.

2.3 Agent-Based Systems in Intelligent Transportation

Multi-agent systems (MAS) have a rich history in intelligent transportation [15]. Traditional MAS approaches model vehicles, infrastructure, and traffic controllers as interacting agents for traffic optimization and cooperative driving. Recent work on LLM-powered autonomous agents [10] has expanded the agent paradigm to encompass perception, reasoning, and action planning within individual systems.

Our work differs from both traditions. Rather than modeling external entities (vehicles, infrastructure) as agents, we decompose the internal processing pipeline of a single driver monitoring system into specialized agents. This intra-system agent architecture enables modular adaptation, each agent can be independently updated, personalized, or replaced, while maintaining coherent end-to-end prediction through structured inter-agent communication. To our knowledge, this is the first application of an agentic architecture to smartphone-based driver intent prediction.

3 Proposed Methodology

3.1 System Overview

Our system takes as input a continuous stream of smartphone IMU data (tri-axial accelerometer readings $a(t) = [a_x, a_y, a_z]^T$ and tri-axial gyroscope readings $\omega(t) = [\omega_x, \omega_y, \omega_z]^T$ sampled at 100 Hz) and outputs intent predictions $\hat{y} \in$ {braking, lane change, turn, normal} with an anticipation horizon $\Delta t \in [2, 4]$ seconds before maneuver onset.

The architecture comprises three autonomous agents that communicate through a shared message bus, as shown in Figure 1.

1) Signal Processing Agent (A_S): Denoises raw sensor data, extracts pre-motion micro-patterns, and produces a structured feature vector.
2) Driver Profiling Agent (A_P): Maintains a personalized driver model that encodes individual behavioral signatures and adapts over time.
3) Prediction Agent (A_R): Fuses micro-pattern features with the driver profile to generate intent predictions via the WE-TAN model.

Figure 1. Overview of the proposed agentic driver intent prediction framework. The Signal Processing Agent extracts micro-patterns from raw IMU data, the Driver Profiling Agent maintains personalized behavior models, and the Prediction Agent generates anticipatory intent classifications.

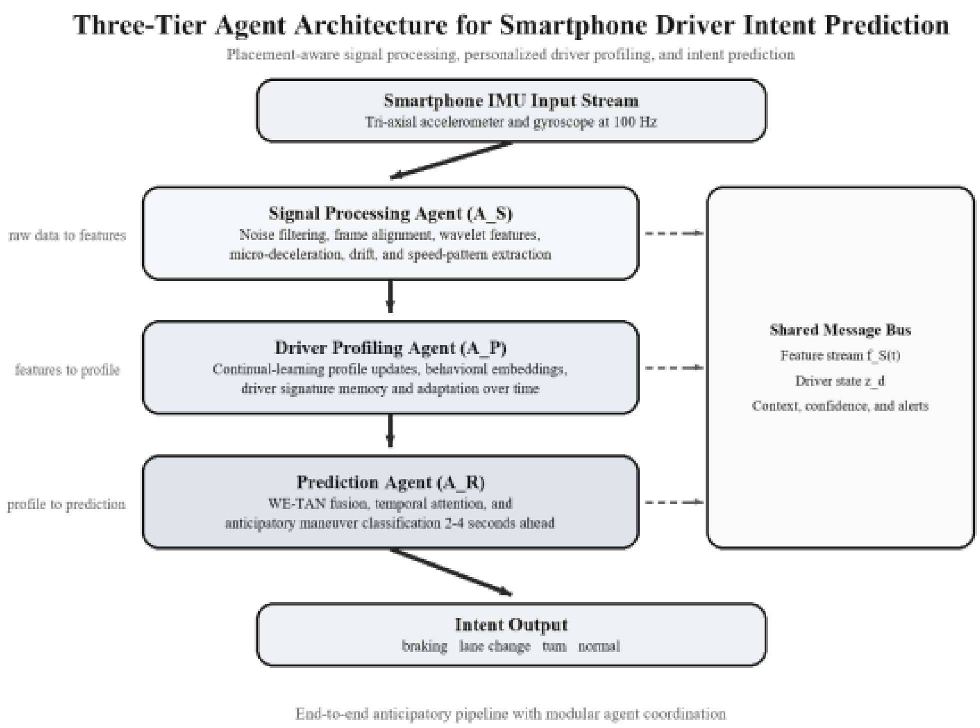

3.2 Signal Processing Agent (A_S)

The Signal Processing Agent addresses the core SNR challenge by extracting three categories of pre-motion micro-patterns from raw IMU data.

1) Preprocessing: Raw accelerometer and gyroscope signals are first reoriented from the device frame to the vehicle frame using a gravity-based rotation estimation [5]. A 4th-order Butterworth low-pass filter

with a 20 Hz cutoff removes high-frequency vibration noise while preserving maneuver-relevant dynamics.

2) Micro-Deceleration Extraction: Micro-decelerations are subtle longitudinal deceleration events ($|a_x| < 0.3$ m/s^2) that precede full braking by 1–3 seconds. We detect these using a continuous wavelet transform (CWT) with a Morlet mother wavelet:

$$W_a(s, \tau) = (1 / \sqrt{s}) \int_{-\infty}^{\infty} a_x(t)\, \psi^*((t - \tau) / s)\, dt \tag{1}$$

where s denotes scale and τ denotes translation. Micro-deceleration events are identified as local maxima in the wavelet coefficient magnitude at scales corresponding to 0.5–3 Hz, capturing the characteristic frequency range of preparatory foot-to-brake transitions.

3) Lateral Drift Detection: Pre-lane-change lateral drift manifests as a slow lateral acceleration (a_y) trend with magnitude 0.05–0.2 m/s^2 sustained over 1–2 seconds. We compute a drift score:

$$D_{lat}(t) = (1 / T_w) \int_{t-T_w}^{t} a_y(\tau)\, d\tau - \mu_{ay} \tag{2}$$

where $T_w = 2$ s is the integration window and μ_{ay} is the running mean lateral acceleration over the past 30 seconds. This baseline subtraction removes road camber and device mounting bias.

4) Speed Variation Patterns: Pre-turn behavior often exhibits a characteristic speed modulation pattern: gradual deceleration followed by a brief stabilization phase. We capture this through the first and second derivatives of the longitudinal acceleration envelope:

$$v_{pattern}(t) = [\dot{a}_x(t), \ddot{a}_x(t), \sigma_{Tw}(t)] \tag{3}$$

where σ_{Tw} is the rolling standard deviation over window T_w. The agent outputs a 24-dimensional feature vector $f_S(t)$ comprising wavelet coefficients (12 features), drift metrics (6 features), and speed variation descriptors (6 features) at each time step.

3.3 Driver Profiling Agent (A_P)

The Driver Profiling Agent constructs and maintains a personalized representation of each driver's behavior, addressing inter-driver variability that degrades population-level models.

1) Driving Signature Representation: Each driver's profile is encoded as a latent vector $z_d \in \mathbb{R}^{32}$ learned through a variational autoencoder (VAE) that compresses historical driving segments into a compact behavioral signature. The VAE encoder processes sequences of micro-pattern features:

$$q(z_d \mid F_d) = N(\mu_\varphi(F_d), \mathrm{diag}(\sigma^2_\varphi(F_d))) \tag{4}$$

where $F_d = \{f^{(1)}, \ldots, f^{(N)}\}$ is the set of historical feature vectors for driver d, and φ denotes encoder parameters.

2) Continual Learning with Elastic Weight Consolidation: Driver behavior evolves over time (e.g., fatigue, route familiarity, weather adaptation). To accommodate this drift without catastrophic forgetting, we employ Elastic Weight Consolidation (EWC):

$$L_{EWC} = L_{task} + \lambda \Sigma_i F_i(\theta_i - \theta^*_i)^2 \quad (5)$$

where F_i is the diagonal Fisher information matrix approximation for parameter θ_i, θ_i are the parameters after learning previous tasks, and λ controls the plasticity-stability trade-off. The agent updates z_d every 15 minutes of driving data, allowing gradual adaptation to changing conditions while preserving the core driving signature.

3.4 Prediction Agent (A_R) and WE-TAN

The Prediction Agent implements the Wavelet-Enhanced Temporal Attention Network (WE-TAN), the core predictive model that fuses micro-pattern features with the driver profile.

1) Wavelet-Enhanced Attention: Standard attention mechanisms [16] operate on raw feature sequences and can be dominated by high-magnitude noise events. WE-TAN performs multi-scale wavelet decomposition within the attention computation:

$$\text{WE-Attn}(Q, K, V) = \text{softmax}(Q_w K_w^T / \sqrt{d_k}) V \quad (6)$$

where $Q_w = DWT(Q)$ and $K_w = DWT(K)$ are the discrete wavelet transforms of the query and key matrices using a Daubechies-4 wavelet. By computing attention in the wavelet domain, the network naturally decomposes temporal patterns across scales, allowing it to attend to slow preparatory trends (low-frequency components) independently of rapid transient events (high-frequency components).

2) Driver-Conditioned Prediction: The driver profile z_d conditions the prediction through a FiLM (Feature-wise Linear Modulation) layer:

$$h_{cond} = \gamma(z_d) \odot h_{attn} + \beta(z_d) \quad (7)$$

where γ and β are learned affine transformations of the driver embedding, and h_{attn} is the output of the WE-Attention layers. This allows the prediction to be modulated by driver-specific parameters without retraining the entire network.

3) Intent Classification: The conditioned representation passes through a two-layer MLP with GELU activation, followed by a softmax classifier:

$$P(\hat{y} \mid f_S, z_d) = \text{softmax}(W_2 \, \text{GELU}(W_1 h_{cond} + b_1) + b_2) \quad (8)$$

The model is trained with focal loss [12] to handle class imbalance, as normal driving segments vastly outnumber maneuver events:

$$L_{focal} = -\alpha_c(1 - p_c)^\gamma \log(p_c) \quad (9)$$

where p_c is the predicted probability for the true class c, α_c is the class weight, and $\gamma = 2$ is the focusing parameter.

4 Proposed Validation Pipeline

Because no new data collection, training, or quantitative testing was performed in the present work, this section reframes the empirical portion of the manuscript as a future validation plan. The purpose is to explain how the proposed agentic pipeline can be implemented and studied rigorously.

4.1 Proposed Data Collection and Labeling Strategy

A future study can collect a naturalistic driving dataset using an Android or iOS application that records smartphone IMU streams during routine driving. A lightweight web-based mobile sensor logging tool can also support early prototyping and instrumentation design before a full native data-collection stack is built [17]. To evaluate the proposed framework credibly, the collection protocol should include:

- Diverse drivers: multiple age groups, driving styles, and route habits to capture inter-driver variability.
- Multiple environments: urban, suburban, and highway segments collected under different traffic densities and times of day.
- Placement diversity: dashboard, windshield, cup holder, seat, pocket, and handheld placement conditions to test robustness.
- Sensor synchronization: accelerometer and gyroscope streams at high sampling rate, with GPS or video used only for alignment and annotation support.
- Ground-truth annotation: maneuver onset labels for braking, lane changes, and turns, ideally annotated from synchronized video by at least two reviewers.

The key design principle is that labels should correspond to maneuver onset rather than maneuver completion, because the objective is anticipation rather than retrospective event recognition.

4.2 Baseline Models and Training Plan

The proposed framework should be compared against a sequence of increasingly strong baselines:

- SVM-SW: Support Vector Machine with sliding-window features [4].
- CNN-1D: One-dimensional convolutional network over raw sensor windows [13].
- LSTM: Two-layer LSTM on raw sensor sequences [18].
- LSTM-Attn: LSTM with temporal attention [12].
- WE-TAN (no profile): proposed wavelet-attention model without personalization.
- WE-TAN + Agent: full agentic framework with driver profiling and adaptive state.

All deep models can be implemented in PyTorch. A reasonable initial training configuration would use hidden dimension 128, four attention heads, a learning rate near 3×10^{-4}, class-balanced loss, and leave-driver-out evaluation. These settings should be treated as starting points for future tuning rather than fixed validated hyperparameters.

4.3 Evaluation Questions and Metrics

The future empirical study should answer four main questions:

1) Does the agentic decomposition improve maneuver anticipation over monolithic baselines?
2) How much does personalization help relative to a population-level model?
3) How robust is the framework to changes in smartphone placement and trip context?
4) What trade-off emerges between anticipation horizon and false positive rate?

The primary evaluation metrics should include macro F1-score, classwise precision and recall, anticipation lead time, false positive rate, and calibration of predicted probabilities.

In addition, ablation studies should isolate the contribution of wavelet attention, driver memory, and individual micro-pattern feature groups.

4.4 Step-by-Step Validation Workflow

Figure 2 summarizes a future validation workflow for this research direction. The workflow can be executed in the following sequence:

1) Instrumentation: build or adapt a mobile logging application [17] and define supported phone placements, consent procedures, and trip metadata.
2) Trip collection: record multi-trip IMU streams across drivers, routes, and operating conditions.
3) Annotation and alignment: synchronize maneuver events using video review or assisted labeling, then mark braking, lane-change, and turn onsets.
4) Preprocessing: transform device-frame signals, denoise the raw measurements, segment anticipation windows, and extract micro-pattern candidates.
5) Model training: train the baseline models first, then train the WE-TAN model with and without the driver-profiling agent.
6) Ablation and robustness checks: evaluate sensitivity to phone placement, missing data, horizon length, and removal of individual agent modules.
7) Deployment-oriented validation: test warning logic, calibration, and edge cases through simulation-in-the-loop workflows before any real-world pilot deployment [19].

Figure 2. Proposed future validation pipeline for smartphone-based agentic driver intent prediction. The workflow spans instrumentation, data collection, labeling, preprocessing, model training, benchmark comparison, robustness analysis, and simulation-in-the-loop validation before deployment.

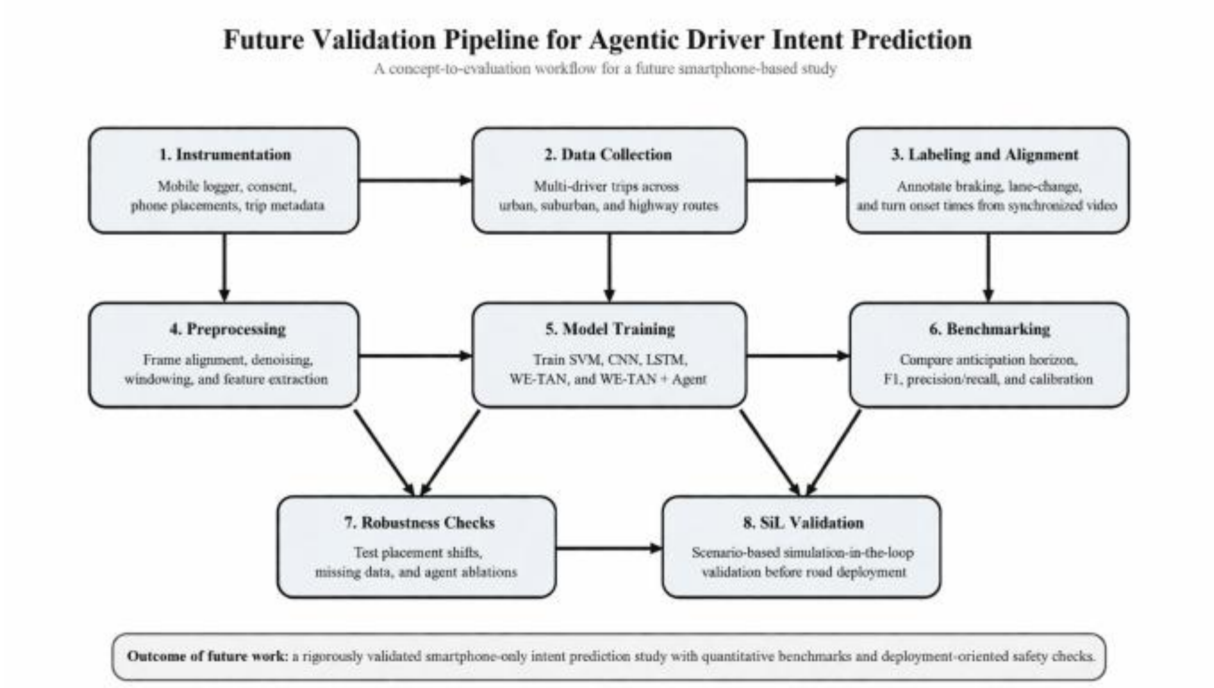

4.5 Expected Outputs of a Future Study

If this pipeline is executed, the resulting paper should report:
• benchmark comparisons between proposed and baseline models,
• classwise anticipation performance at multiple horizons,
• false positive behavior under driver-specific adaptation,
• ablation evidence for the value of each agent, and
• failure modes associated with placement changes, noisy trips, and sparse personalization data.

5 Discussion

5.1 Why Personalization Matters

Personalization is likely to matter because driver intent prediction from weak IMU signals is not only a signal processing problem, but also a modeling problem. Two drivers can produce similar accelerometer magnitudes during normal driving yet exhibit different pre-maneuver signatures. One driver may show a distinct micro-deceleration pattern before braking, while another may prepare more through lateral drift or steering correction. A population-level model must average across these patterns, which can dilute predictive power for any individual.
The agentic architecture addresses this directly. The Driver Profiling Agent is intended to learn which micro-pattern features are most informative for each driver, effectively tuning the signal-to-noise ratio on a per-user basis. This is especially relevant for maneuvers such as lane changes, whose preparatory lateral patterns are often highly idiosyncratic.

5.2 Signal-to-Noise Considerations

The wavelet-enhanced attention mechanism provides a principled approach to the SNR challenge. By decomposing queries and keys into wavelet coefficients before computing attention scores, WE-TAN effectively separates the temporal scales at which pre-motion patterns and noise operate. Pre-motion micro-patterns typically occupy the 0.3–3 Hz band, while device vibration and road roughness noise concentrate above 5 Hz. Standard attention mechanisms that operate in the time domain cannot exploit this spectral separation.
However, the SNR limitation is expected to become more severe at longer prediction horizons ($\Delta t > 4$ s), where pre-motion signals may become vanishingly weak. This suggests a possible physical limit: smartphone IMU sensors may reliably capture intent-related signals only within a temporal window determined by the biomechanics of driving preparation (foot movement to brake pedal, hand positioning on steering wheel, postural adjustments).

5.3 Limitations

Several limitations should be acknowledged. First, this paper does not report new data collection or empirical testing, so the proposed framework remains unvalidated at present. Second, smartphone mounting position and orientation introduce variability that our gravity-based reorientation may only partially address. Third, the current agent communication protocol is described conceptually rather than benchmarked in a real-time implementation. Finally, the system does not yet incorporate external context (traffic density, weather, road type) that could improve predictions in a future deployed system.

6 Conclusion

This paper presented a conceptual agentic framework for predicting driver intent, braking, lane changes, and turns, before execution using only smartphone inertial sensors. By decomposing the prediction pipeline into three specialized agents (Signal Processing, Driver Profiling, and Prediction), the proposed approach targets the twin challenges of weak pre-motion signals and inter-driver variability that have limited prior smartphone-based approaches.

The proposed Wavelet-Enhanced Temporal Attention Network (WE-TAN) is intended to exploit multi-scale wavelet decomposition within the attention mechanism to selectively amplify discriminative micro-patterns buried in noisy IMU data. Combined with agent-based personalization through continual learning driver profiles, it offers a plausible architecture for future smartphone-only intent prediction research.

Rather than claiming validated performance, this paper contributes a structured research direction and an implementation pipeline that explains how such a system can be studied rigorously. The framework provides a principled and modular basis for combining signal processing, personalization, and prediction, and it can be extended to incorporate additional context sources in future work.

Future directions include collecting a naturalistic dataset, integrating GPS-derived road context, exploring federated learning for privacy-preserving cross-driver knowledge transfer, deploying the system in real-time on-device inference scenarios with resource-constrained hardware, and evaluating scenario-based validation workflows in simulation-in-the-loop settings before road deployment [19].

Acknowledgment

The author acknowledges the contributions of industry practitioners, regulatory bodies, and open access researchers whose datasets and evaluation frameworks have informed this work. The author acknowledges the use of OpenAI (GPT-5 model, 2025) and Grammarly to assist in improving the clarity, organization, and structure of this manuscript.

Conflict of Interest

The author declares that there are no conflicts of interest related to the publication of this paper. The author conducted this research entirely independently of their professional duties and responsibilities at their respective employing organization. The research and opinions presented in this paper are solely those of the author and do not represent the views, positions, or policies of their respective employers or affiliations.

References

1. Y. Xing, C. Lv, H. Wang, D. Cao, and E. Velenis, "Driver lane change intention inference for intelligent vehicles: Framework, survey, and challenges," IEEE Transactions on Vehicular Technology, vol. 69, no. 3, pp. 2941–2950, 2020.
2. G. Pokharkar, "Design and evaluation of AI safety mechanisms in ADAS and autonomous vehicle architectures," International Journal of Emerging Trends in Computer Science and Information Technology, pp. 57–67, Sep. 2025. [Online]. Available: https://ijetcsit.org/index.php/ijetcsit/article/view/388
3. A. Jain, H. S. Koppula, B. Raghavan, S. Soh, and A. Saxena, "Car that knows before you do: Anticipating maneuvers via learning temporal driving models," in 2015 IEEE International Conference on Computer Vision (ICCV), 2015, pp. 3182–3190.

4. D. A. Johnson and M. M. Trivedi, "Driving style recognition using a smartphone as a sensor platform," in 2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC), 2011, pp. 1609–1615.
5. Z. Chen, J. Zheng, and J. Yu, "Smartphone-based vehicle trajectory estimation and lane-change detection using inertial sensors," IEEE Sensors Journal, vol. 19, no. 22, pp. 10688–10696, 2019.
6. G. Ren, Y. Liu, and T. Zhang, "Smartphone-based driver behavior recognition using deep learning," IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 3, pp. 2820–2831, 2023.
7. F. L. Mannering, V. Shankar, and C. R. Bhat, "Unobserved heterogeneity and the statistical analysis of highway accident data," Analytic Methods in Accident Research, vol. 11, pp. 1–16, 2016.
8. P. S. Bokare and A. K. Maurya, "Acceleration-deceleration behaviour of various vehicle types," Transportation Research Procedia, vol. 25, pp. 4733–4749, 2017.
9. M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," The Knowledge Engineering Review, vol. 10, no. 2, pp. 115–152, 1995.
10. L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin et al., "A survey on large language model based autonomous agents," Frontiers of Computer Science, vol. 18, no. 6, p. 186345, 2024.
11. A. Zyner, S. Worrall, and E. Nebot, "Naturalistic driver intention and path prediction using recurrent neural networks," IEEE Transactions on Intelligent Transportation Systems, vol. 21, no. 4, pp. 1584–1594, 2020.
12. X. Li, W. Ying, and F. Zhu, "Lane change intent prediction using LSTM with attention mechanism," IEEE Transactions on Intelligent Vehicles, vol. 6, no. 2, pp. 340–349, 2021.
13. W. Wang, J. Xi, and D. Zhao, "Driving behavior analysis and classification based on smartphone sensors using machine learning," IEEE Access, vol. 7, pp. 178366–178379, 2019.
14. A. Abdelrahman and H. S. Hassanein, "A machine learning smartphone-based sensing for driver behavior classification," in 2020 IEEE International Conference on Communications (ICC), 2020, pp. 1–6.
15. A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey," IEEE Access, vol. 6, pp. 28573–28593, 2018.
16. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in Advances in Neural Information Processing Systems 30, 2017, pp. 5998–6008.
17. G. Pokharkar, "Mobile sensor logger (web-based)," https://github.com/gpokhark/sensor-logger, 2026, client-side mobile sensor logging tool for quick research and prototyping.
18. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.
19. G. Pokharkar, "Scenario-based validation for SAE level 2+ features using simulation-in-the-loop (SIL) systems," International Journal of Innovative Research and Creative Technology, vol. 11, no. 4, Jul. 2025. [Online]. Available: https://doi.org/10.5281/zenodo.16883284