

Physics-guided GPS motion features for UAV mmWave beam prediction under high mobility

Do Thuy Linh (Ha Noi, Viet Nam)

Nguyen Binh Thanh (Ha Noi, Viet Nam)

Le Huu Son (Ha Noi, Viet Nam)

Hoang Trong Minh* (Ha Noi, Viet Nam)

Dedicated to Professor Ha Huy Khoai on the occasion of his 80 birthday.

(Received Mar. 26, 2026; accepted July 8, 2026)

Abstract. UAV mmWave systems can exploit multiple information sources to predict beams. Multimodal solutions often achieve higher accuracy, but they also increase sensing requirements, synchronization burden, and deployment complexity. In contrast, GPS-based approaches offer a simpler, lower-cost alternative, but their performance degrades noticeably under high UAV mobility when only position information is used. This paper addresses that limitation from the perspective of input representation design rather than predictor scaling. Specifically, we construct GPS-derived motion features in the Earth-Centered, Earth-Fixed coordinate system and reorganize them according to the UE-BS geometry to form radial, direction-aware tangential, and velocity-angular feature families. Relative to prior GPS-aided beam prediction that uses position coordinates or changes the predictor, this work keeps the predictor and protocol fixed and measures the contribution of each GPS-derived motion component. Experiments on the public DeepSense dataset show that not all physically meaningful motion variables contribute equally to beam prediction. Among the investigated variants, the radial feature provides the most reliable overall improvement, increasing Top-1 accuracy from 73.74% to 74.60% while reducing the mean power loss to 0.1674 dB. The remaining variants exhibit different trade-offs across operating conditions, and the overall results consistently indicate that input feature design has a stronger impact on performance than predictor scaling in GPS-based UAV mmWave beam prediction.

* Corresponding author

Key words and phrases: UAV communications, millimeter wave, beam prediction, GPS-aided learning, motion-aware features, radial velocity, mobility-aware prediction.

2020 Mathematics Subject Classification: 94A12, 94A15, 94A40, 68T07.

1. Introduction

In recent years, millimeter-wave (mmWave) communication combined with unmanned aerial vehicles (UAVs) has emerged as a promising approach for building high-throughput aerial wireless networks that support sensor data collection and maintain reliable connectivity for critical applications. To establish broadband links between UAVs, base stations, and user equipment (UE), mmWave systems generally rely on narrow beams to enhance received power and compensate for severe path loss. However, these narrow beams are highly sensitive to rapid motion and heading changes of the UAV [1–3]. As a result, a beam selected at the current time can quickly become suboptimal, increasing beam training overhead and degrading link quality.

The recent development of learning-based methods has made beam prediction a practical alternative to exhaustive or hierarchical beam search [4, 5]. In this context, UAV systems can exploit multiple information sources to support beam prediction. Multimodal solutions often provide richer environmental and motion information and may achieve higher accuracy, but they also increase sensing requirements, synchronization burden, and deployment complexity [6]. In contrast, GPS-based approaches are particularly attractive because GPS is already available on many UAV platforms and introduces little additional sensing overhead. However, recent studies have shown that raw GPS position sequences alone are still insufficient to maintain reliable beam prediction performance under complex or highly dynamic mobility conditions [7, 8]. This limitation becomes more severe as the UAV moves faster or the prediction horizon is longer. The key difficulty is that a position sequence contains motion information only implicitly. Under higher mobility or longer prediction horizons, the model must infer beam evolution solely from temporal differences in coordinates. This inference is challenging because the same absolute displacement may correspond to very different beam dynamics depending on the relative geometry between the UE (UAV) and the base station (UE-BS). Therefore, instead of adding more sensors or scaling the predictor architecture, this paper revisits the representation of GPS information itself and asks which GPS-derived motion variable is the most useful for beam prediction. Specifically, we construct motion features from GPS trajectories in the Earth-Centered, Earth-Fixed (ECEF) coordinate system and reorganize them according to the UE-BS geometry to form physically interpretable feature families, including radial motion, direction-aware tangential motion, and a simple representation of angular-rate-related behavior. Using the public DeepSense dataset for UAV scenarios, we systematically evaluate these motion variables to answer a central question: among the motion quantities derived from GPS, which representation is most useful for beam prediction when expressed in a geometrically meaningful form?

The resulting analysis not only clarifies the role of each feature family but also provides practical evidence for selecting more effective GPS-based beam prediction strategies in UAV mmWave systems.

The main contributions of this work are as follows. First, we formulate a physics-guided GPS motion representation for UAV mmWave beam prediction by organizing GPS-derived variables according to the UE-BS geometry. This differs from previous GPS-based beam prediction by resolving GPS displacement into UE-BS radial, tangential, and angular-rate components and testing each feature set under the same predictor, split, and metric set. Second, we show that radial motion provides the most reliable overall improvement over the position-only baseline, while the remaining motion variants exhibit more regime-dependent behavior. Third, we demonstrate that input feature design has a stronger impact on performance than predictor scaling in the considered approach. This study focuses on GPS-derived motion-feature representation for UAV mmWave beam prediction, where position, range variation, heading, and speed-related terms are structured to improve beam-index prediction under mobility.

2. Related Work

Previous studies on beam management for mmWave systems have mainly relied on beam sweeping, hierarchical codebooks, or compressed channel estimation techniques [3–5]. These approaches are suitable when link geometry varies slowly. However, as the mobility of connected nodes increases, the cost of beam training and beam updates also rises substantially. This challenge becomes more severe in UAV-enabled mmWave communications, where both the propagation geometry and the optimal beam may change rapidly along the flight trajectory [1, 2]. For this reason, beam prediction has become an important research direction for reducing dependence on repeated beam searches under high-mobility conditions [9].

Several recent studies have focused on modern sensing modalities, such as cameras and LiDAR, to enhance environmental awareness and support beam inference [10–15]. These methods often provide richer spatial context, thereby improving prediction performance. However, these benefits come at the cost of additional sensing hardware, constrained time synchronization across data sources, and higher deployment complexity in practical systems. Therefore, although multimodal sensing is a strong direction for performance, its deployment on UAV platforms remains challenging because UAV systems typically require lightweight configurations and low operational costs. Another approach

Table 1. Positioning of the proposed GPS motion-feature study relative to beam-management and beam-prediction directions.

Technical direction	What work addresses	prior ad-	Remaining limitation	Difference of this paper
Beam sweeping and code-book search	Select beams through direct beam training or hierarchical search.		Training overhead increases when the UAV-BS geometry changes rapidly.	Predict future beam labels from GPS-derived features to reduce repeated beam training.
Multimodal sensing-aided prediction	Use camera, LiDAR, or other sensors to enrich environmental context.		Extra sensors increase payload, synchronization, and deployment complexity on UAV platforms.	Use GPS only, avoiding extra sensing payload and cross-modal synchronization.
GPS-aided beam prediction	Use position information to predict current or future beams.		Position sequences encode motion only implicitly and do not identify which motion variable is responsible for high-mobility gains.	Resolve GPS displacement into UE-BS radial, tangential, and angular-rate features and test them with the same predictor.
Predictor scaling	Improve the learning backbone or sequence model.		Accuracy changes can be caused by both input representation and model capacity.	Keep the comparison controlled so the measured gain is attributed to the feature set, not predictor size.

in this field of study focuses on GPS-based beam prediction, since GPS is already integrated into many UAV platforms and does not require additional sensing devices. Recent studies on the DeepSense 6G dataset have shown the clear potential of position information for beam prediction [7, 8, 16]. However, these studies have shown consistent limitations, such as raw position sequences only indirectly reflecting motion dynamics, and then prediction performance often degrades when mobility becomes faster or requires more time to predict. In other words, GPS is attractive from a deployment perspective, but the way GPS information is represented remains insufficient to maintain accuracy under high-mobility conditions. Recent real-world studies confirm the utility of position information for beam prediction, but they also show that GPS-based performance remains sensitive to mobility, data quality, and feature representation, which leaves open the question of how to exploit GPS information more effectively for prediction [7, 8, 17]. More specifically, we focus on identifying which GPS-derived motion quantity is most valuable for beam prediction. This question is particularly important for UAV systems that operate with GPS only and do not rely on additional sensing modalities [8, 18]. Compared with previous GPS-based studies [7, 8], our goal is not to introduce a new predictive method but to isolate and systematically evaluate the role of physically meaningful

motion variables derived from GPS for enhancing prediction performance. We organize the input variables according to the UE-BS geometry and distinguish three levels of effect: overall improvement, gains under high-speed regimes, and the additional effect of architecture refinement, after identifying the strongest feature family. In this direction, we clarify, through simulation, which motion variable is practically useful for GPS-based UAV mmWave beam prediction.

3. System Model and Problem Formulation

We consider a UAV user equipment (UE) communicating with a ground BS at 60 GHz. The BS uses a beamforming codebook $\mathcal{F} = \{f_m\}_{m=1}^M$, where $f_m \in \mathbb{C}^{Q \times 1}$ is the m th beamforming vector and M is the number of candidate beams.

3.1. Received signal model

At time t , the received signal is written as

$$(3.1) \quad y[t] = h^H[t]f[t]x[t] + n[t],$$

where $h[t] \in \mathbb{C}^{Q \times 1}$ is the channel vector, $f[t] \in \mathcal{F}$ is the selected beam, $x[t]$ is the transmitted symbol, and $n[t]$ is additive noise. The main purpose of this equation is to define the link-level quantity affected by beam mismatch, namely $|h^H[t]f[t]|^2$.

3.2. Optimal beam label

The optimal beam index is

$$(3.2) \quad b^*[t] = \arg \max_{b \in \{1, \dots, M\}} |h^H[t]f_b|^2.$$

This beam index is the supervision target for the learning model. In practice, the model does not observe the channel directly, but instead predicts future beam labels from side information derived from GPS trajectories.

3.3. Prediction task

Let $O[t]$ denote the feature vector observed at time t and let

$$(3.3) \quad O_{t,W} = \{O[t-W+1], \dots, O[t]\}$$

be the observation window of length W . The objective is to predict the beam indices over a future horizon of $V+1$ steps,

$$(3.4) \quad \{b^*[t+\ell]\}_{\ell=0}^V,$$

where ℓ denotes the prediction step. The predictor outputs a beam-probability vector $\hat{p}[t+\ell] \in \mathbb{R}^M$ at each future step, and the predicted beam index is

$$(3.5) \quad \hat{b}[t+\ell] = \arg \max \hat{p}[t+\ell].$$

The main problem addressed in this study is how to design $O[t]$ so that GPS-only inputs better reflect future beam evolution under fast mobilities.




Figure 1. Workflow of the GPS motion-feature pipeline from dataset input to beam-output evaluation.

Figure 1 shows the processing chain used in this paper. GPS samples are converted to ECEF coordinates, projected onto the UE-BS geometry, and converted into candidate motion-feature sequences. The predictor then outputs current and future beam labels, while the evaluation uses Top- K accuracy, mean power loss, and beam-search overhead proxies. This flow makes clear that the proposed contribution is the GPS feature set, while the predictor and evaluation protocol are controlled.

4. Physics-Guided Motion Features

4.1. ECEF representation and relative direction

We use Earth-Centered Earth-Fixed coordinates because they support stable vector operations for relative geometry. Let $r_{\text{UE}}[t] \in \mathbb{R}^3$ denote the UE position at time t and let $r_{\text{BS}} \in \mathbb{R}^3$ denote the BS position. The normalized UE-BS direction is

$$(4.1) \quad u_{\text{UE-BS}}[t] = \frac{r_{\text{UE}}[t] - r_{\text{BS}}}{\|r_{\text{UE}}[t] - r_{\text{BS}}\|}.$$

The purpose of this equation is to capture the line-of-sight direction that is directly connected to beam alignment.

4.2. GPS-derived velocity and geometric decomposition

Assuming sampling interval Δt , the GPS-derived velocity in ECEF coordinates is computed as

$$(4.2) \quad v_{\text{ECEF}}[t] = \frac{r_{\text{UE}}[t] - r_{\text{UE}}[t-1]}{\Delta t}.$$

We then decompose this velocity with respect to the UE-BS direction. The radial component is

$$(4.3) \quad v_{\text{rad}}[t] = v_{\text{ECEF}}[t] \cdot u_{\text{UE-BS}}[t],$$

and the tangential magnitude is

$$(4.4) \quad v_{\text{tan}}[t] = \sqrt{\|v_{\text{ECEF}}[t]\|^2 - v_{\text{rad}}^2[t]}.$$

The purpose of this decomposition is to separate distance-changing motion from direction-changing motion. Radial motion affects how quickly the UAV approaches or recedes from the BS, while tangential motion is more directly associated with the rotation of the UE-BS direction. A first-order angular-rate proxy is

$$(4.5) \quad \omega_{\text{proxy}}[t] \approx \frac{v_{\text{tan}}[t]}{d_{\text{UE-BS}}[t]},$$

where

$$(4.6) \quad d_{\text{UE-BS}}[t] = \|r_{\text{UE}}[t] - r_{\text{BS}}\|.$$

The role of $\omega_{\text{proxy}}[t]$ is not to replace the learning model, but to provide a physically interpretable cue for how quickly the pointing direction may evolve.

To keep the results presentation concise, we use abbreviations for contextual descriptions in Table 2.

Table 2. Model aliases used in the manuscript

Alias	Meaning
Pos	5D position-only model
RADIAL	Pos with radial velocity
DIRTAN	Pos with directed tangential components
VELANG	Pos with velocity and angular cues
RAD-CNN-GRU	Radial-feature CNN-GRU
RAD-2GRU	Radial-feature 2-layer GRU
RAD-TRANS	Radial-feature Transformer

Algorithm 1 Physics-Guided Motion Feature Extraction from GPS

Require: GPS trajectory $\{r_{\text{UE}}[t]\}_{t=0}^T$, BS position r_{BS} , sampling interval Δt

Ensure: Motion-aware feature sequence $\{O[t]\}$

- 1: **for** $t = 1$ to T **do**
 - 2: Compute relative unit direction $u_{\text{UE-BS}}[t]$
 - 3: Compute GPS-derived ECEF velocity $v_{\text{ECEF}}[t]$
 - 4: Compute radial velocity $v_{\text{rad}}[t] \leftarrow v_{\text{ECEF}}[t] \cdot u_{\text{UE-BS}}[t]$
 - 5: Compute tangential magnitude $v_{\text{tan}}[t] \leftarrow \sqrt{\|v_{\text{ECEF}}[t]\|^2 - v_{\text{rad}}^2[t]}$
 - 6: Compute distance $d_{\text{UE-BS}}[t] \leftarrow \|r_{\text{UE}}[t] - r_{\text{BS}}\|$
 - 7: Compute angular-rate proxy $\omega_{\text{proxy}}[t] \leftarrow v_{\text{tan}}[t]/d_{\text{UE-BS}}[t]$
 - 8: Form the desired candidate feature vector by concatenating the 5D position terms with selected motion terms
 - 9: **end for**
-

The methodological emphasis of this paper is on feature design rather than architectural novelty. Accordingly, all core motion-feature comparisons use the same lightweight baseline predictor with approximately 260K parameters. This design choice ensures that improvements observed in the motion-feature study can be attributed to the input representation rather than to model scaling. To assess whether better inputs still benefit from architecture changes, we additionally evaluate a controlled radial-family architecture study on the strongest feature family, namely the 5D radial input. Three predictor variants are considered: the baseline CNN-GRU, a 2-layer GRU encoder, and a Transformer encoder.

5. Experimental Setup

5.1. Dataset and scenario

We evaluate on DeepSense 6G Scenario 23, a real-world dataset for sensing and communication at 60 GHz [16]. DeepSense 6G is a public measurement dataset for sensing-aided wireless communication. Scenario 23 records a drone mmWave communication setup with synchronized UAV-side GPS, visual data, and mmWave beam-training measurements with the ground BS. This paper uses only GPS samples and optimal beam labels derived from beam training; visual data are excluded to keep the evaluation GPS-only. The study follows the same beam-prediction setting used in the same baseline scenario, with codebook size $M = 32$, observation window $W = 8$, and prediction horizon $V = 3$. Each sample uses a GPS-derived observation window to predict the current and three future beam labels. This protocol supports horizon-wise and speed-wise evaluation, which is required for high-mobility beam prediction because beam mismatch increases when the UAV changes range or bearing within a short time interval. The train/validation/test split and preprocessing follow the same project protocol for all compared models so that the reported differences come from the feature family or the predictor variant rather than from data handling changes.

5.2. Training protocol

Unless otherwise stated, all motion-feature experiments use the same baseline predictor and identical training settings. The motion-feature study is reported over 10 independent runs, while the radial-family architecture study is reported over 5 independent runs. All compared models follow the same optimization recipe and early-stopping rule as the project baseline, so the comparisons isolate the effect of feature design and controlled architecture changes rather than changes in the training procedure.

5.3. Beam-prediction metrics

Top- K beam prediction accuracy at prediction step ℓ is defined as

$$(5.1) \quad \text{Top-}K[\ell] = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \mathbf{1}\left\{b_i^*[t + \ell] \in \text{Top}K(\hat{p}_i[t + \ell])\right\}.$$

This metric measures whether the correct beam is contained in the top- K model outputs.

Mean power loss is reported as

$$(5.2) \quad \text{PL}[\ell] = \frac{10}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \log_{10} \left(\frac{P_{f^*}^{(i)} - 0.5P_n}{P_{\hat{f}}^{(i)} - 0.5P_n} \right).$$

The purpose of this metric is to translate classification errors into a link-level degradation measure. Unlike the earlier draft, the final reported PL values are now available from exported beam-rank evaluation and are included directly in the system-level tables.

Beam-training overhead saving at reliability level R is defined as

$$(5.3) \quad \eta_{\text{OH}}^{(R)} = 1 - \frac{b^{(R)}}{M},$$

where $b^{(R)}$ is the minimum number of top-ranked beams required so that the optimal beam is included with probability at least R . This metric connects beam prediction accuracy to operational beam-search reduction.

6. Results and Discussion

6.1. Motion-feature results

Table 3 reports the finalist motion-feature results. The main outcome is that the strongest overall result comes from RADIAL, not from the tangential family. Adding radial velocity to the 5D position baseline improves Top-1 accuracy from 73.74% to 74.60% and gives the tightest confidence band among all finalists. By contrast, DIRTAN reaches a similar mean of 74.54%, but its higher variance makes the gain statistically inconclusive, while VELANG reaches 73.94% and is therefore not a strong overall model. The important interpretation is that physically meaningful motion variables do not contribute equally once GPS noise, finite sampling, and sequence learning are taken into account. In the present setup, the simpler radial component is the most learnable and the most reliable cue.

This significance analysis is central to the paper. RADIAL is the only statistically significant improvement over POS, with a paired p -value of 0.032 and a 95% bootstrap confidence interval of $[+0.20, +1.45]$. By contrast, DIRTAN has




Figure 2. Overall Top-1 accuracy of the finalist feature models.

Table 3. Overall performance of the motion-feature models

Alias	Dim	Top-1 (%)	Top-3 (%)	<i>p</i> vs Pos
POS	5	73.74 ± 0.95	97.53	—
RADIAL	6	74.60 ± 0.41	97.57	0.032
DIRTAN	8	74.54 ± 1.09	96.95	0.132
VELANG	9	73.94 ± 1.15	97.03	0.686

$p = 0.132$ and VELANG has $p = 0.686$. Therefore, only the radial augmentation is strong enough to support the main claim.

Figure 3 shows that RADIAL improves all four prediction steps relative to POS, from $v0$ to $v3$. The gain is therefore not concentrated in a single easy step; instead, the radial feature provides a consistent horizon-wise improvement. The speed-regime analysis reveals a different but complementary pattern. RADIAL is the strongest overall and slow-to-medium model, whereas VELANG gives the best fast-speed accuracy at 62.40%. The correct reading is therefore two-level: radial velocity is the best overall GPS-derived cue, whereas richer velocity-plus-angular cues are more useful when the evaluation emphasizes harder mobility bins.

Table 4. Speed-regime performance of representative models

Alias	Slow	Medium	Fast	$v3$ Top-1
POS	78.36	65.12	62.15	71.86
RADIAL	79.39	66.45	61.71	72.61
VELANG	77.94	67.93	62.40	71.81
RAD-TRANS	79.63	67.85	62.37	73.50




Figure 3. Top-1 accuracy across prediction horizons.




Figure 4. Speed-horizon Top-1 accuracy of representative models.

Figure 4 makes this distinction explicit. VELANG is strongest in the medium-speed regime and remains competitive in the fast regime, whereas RADIAL is stronger in the slow regime and yields the cleanest overall mean. The radial-family Transformer model RAD-TRANS does not alter the main feature-level conclusion, but it does improve the longest-horizon behavior and produces the strongest $v3$ result at 73.50%. This separation is precisely why the architecture study is included: it isolates the dominant effect of feature choice from the smaller effect of backbone refinement.

6.2. Radial-family architecture results

Table 5 summarizes the radial-family architecture study together with the best 8D references. Three findings matter. First, the radial feature family outperforms the corresponding 8D reference family under all three tested architectures. The gain is +0.85 points for the baseline CNN-GRU, +0.67 points




Figure 5. Accuracy–complexity trade-off across feature families and backbones.

Table 5. Accuracy–complexity trade-off in the radial-family architecture study

Alias	Family	Param. (K)	Top-1 (%)	Top-3 (%)	PL (dB)	v3 (%)	Time (ms)
8D-CNN-GRU	8D ref.	261.2	73.85 ± 1.12	97.06	0.1782	71.84	1.14
8D-2GRU	8D ref.	360.3	74.24 ± 1.25	97.30	0.1850	71.92	1.15
8D-TRANS	8D ref.	558.7	74.02 ± 0.58	96.84	0.1730	72.41	1.96
RAD-CNN-GRU	6D radial	260.0	74.70 ± 0.54	97.55	0.1665	72.57	1.07
RAD-2GRU	6D radial	360.0	74.91 ± 0.64	97.17	0.1703	72.52	1.11
RAD-TRANS	6D radial	558.0	75.09 ± 0.98	97.76	0.1623	73.50	1.83

for the 2-layer GRU, and +1.07 points for the Transformer. Second, within the radial feature family, architecture changes provide only small additional gains. RAD-TRANS reaches 75.09% Top-1, but its gain over RAD-CNN-GRU is only +0.39 points and is not statistically significant ($p = 0.491$). Third, the ANOVA result across RAD-CNN-GRU, RAD-2GRU, and RAD-TRANS gives $p = 0.720$, which indicates that backbone differences within the same radial feature family are not statistically significant. These results support the central methodological conclusion of the paper: feature choice dominates architecture. Once the right motion variable is selected, the simple baseline predictor already captures most of the available gain, and larger backbones mainly provide limited refinements on the farthest horizon and on mean power loss.

Table 6 reports the final system-level summary. Mean power loss is more discriminative than overhead saving. All finalist models already achieve Top-3 above 97%, so they require only 2 beams for $R = 0.90$ and 3 beams for $R = 0.95$, which yields identical overhead savings of 93.8% and 90.6%. In contrast, PL still separates the stronger models: RADIAL reduces PL from 0.1768 dB for POS to 0.1674 dB, while RAD-TRANS further reduces PL to 0.1623 dB. The resulting model choice depends on the operating target. If the goal is the most

Table 6. System-level summary of representative models

Alias	Top-1 (%)	Top-3 (%)	PL (dB)	$b^{(0.90)}$	$\text{OH}^{(0.90)}$	$b^{(0.95)}$	$\text{OH}^{(0.95)}$	Role
POS	73.74 ± 0.95	97.53	0.1768	2	93.8%	3	90.6%	Position baseline
RADIAL	74.60 ± 0.41	97.57	0.1674	2	93.8%	3	90.6%	Primary model
VELANG	73.94 ± 1.15	97.03	0.1778	2	93.8%	3	90.6%	Mobility-stressed model
RAD-CNN-GRU	74.70 ± 0.54	97.55	0.1665	2	93.8%	3	90.6%	Radial CNN-GRU ref.
RAD-TRANS	75.09 ± 0.98	97.76	0.1623	2	93.8%	3	90.6%	Radial Transformer ref.

defensible and stable improvement, RADIAL is the primary model because it is the only significant gain over the baseline and has the lowest variance among finalists. If the goal is the best achievable PL and the strongest $v3$ behavior, RAD-TRANS serves as an upper-bound radial-family reference.

Based on the experimental results, we can identify three main claims. First, radial motion is the strongest overall signal in the GPS-only setting. Second, combining velocity and angle parameters can provide a good assessment in speed-stressed situations. Third, changing the feature family has a greater impact than improving the predictor architecture.

The results indicate a direct mechanism for high-mobility UAV mmWave links. Fast UAV motion changes the UE-BS range and bearing between beam updates. A position-only sequence hides these changes in coordinate differences, whereas RADIAL and VELANG expose radial speed and bearing-rate information before prediction. This explains why RADIAL improves Top-1 accuracy and mean power loss overall, while VELANG is useful in the fastest speed group. At the link level, higher Top- K accuracy allows the receiver to test fewer candidate beams, and lower PL indicates less received-power loss when the selected beam is not exactly optimal.

The interpretation of these results is limited by the experimental scope. The evaluation uses only DeepSense 6G Scenario 23, so additional routes, BS sites, carrier frequencies, blockage states, and weather conditions are needed before claiming broad generalization. Because the features are computed from GPS, their reliability also depends on sampling rate, timestamp alignment, and positioning error. The study excludes camera and LiDAR inputs by design; therefore, it should not be read as a comparison against fully multimodal beam prediction. Finally, the system metrics are Top- K accuracy, mean power loss, and beam-search overhead proxies; end-to-end throughput, control signaling, MAC-layer latency, and beam-recovery delay remain outside the present validation.

7. Conclusion

This paper revisited GPS-aided UAV mmWave beam prediction from the perspective of motion-feature engineering. We derived physically interpretable GPS motion descriptors in ECEF coordinates and evaluated them through a motion-feature study and a radial-family architecture study. The results show that a simple radial-velocity augmentation, RADIAL, is the strongest overall GPS-derived feature, improving Top-1 accuracy from 73.74% to 74.60%, reducing mean power loss from 0.1768 dB to 0.1674 dB, and remaining the only statistically significant gain over the position-only baseline. They also show that a richer velocity-plus-angular feature, VELANG, is more useful in the fast-speed situations, where it reaches 62.40% Top-1, although it is not the best overall model. Finally, the radial-family architecture study indicates that architecture changes within the 6D radial family are not statistically significant, even though the radial-family Transformer model provides the best longest-horizon accuracy and the lowest PL. In summary, we see that selecting the appropriate physical motion variable is more efficient than scaling the predictor in GPS-only UAV beam prediction.

References

- [1] **Y. Zeng, Q. Wu, and R. Zhang**, *Accessing from the sky: A tutorial on UAV communications for 5G and beyond*, Proc. IEEE, vol. 107, no. 12, pp. 2327–2375, Dec. 2019.
- [2] **Z. Xiao et al.**, *A survey on millimeter-wave beamforming enabled UAV communications and networking*, IEEE Commun. Surveys Tuts., vol. 24, no. 1, pp. 557–610, 1st Quart., 2022.
- [3] **M. Giordani et al.**, *A tutorial on beam management for 3GPP NR at mmWave frequencies*, IEEE Commun. Surveys Tuts., vol. 21, no. 1, pp. 173–196, 1st Quart., 2019.
- [4] **K. Chen and C. Qi**, *Beam training based on dynamic hierarchical codebook for millimeter wave massive MIMO*, IEEE Commun. Lett., vol. 23, no. 1, pp. 132–135, Jan. 2019.

- [5] **A. Alkhateeb et al.**, *Channel estimation and hybrid precoding for millimeter wave cellular systems*, IEEE J. Sel. Topics Signal Process., vol. 8, no. 5, pp. 831–846, Oct. 2014.
- [6] **R. Mahmoud, D. Castanheira, A. Silva, and A. Gameiro**, *Sensing-assisted communication for mmWave networks: A review of techniques, applications, and future directions*, Electronics, vol. 14, no. 19, Art. no. 3787, Sep. 2025.
- [7] **J. Morais, A. Bchboodi, H. Pezeshki, and A. Alkhateeb**, *Position-aided beam prediction in the real world: How useful GPS locations actually are?*, Proc. IEEE ICC, Jun. 2023, pp. 1824–1829.
- [8] **V. A. Nugroho and B. M. Lee**, *GPS-aided deep learning for beam prediction and tracking in UAV mmWave communication*, IEEE Access, vol. 13, pp. 117065–117077, 2025.
- [9] **J. Tan et al.**, *Beam alignment in mmWave V2X communications: A survey*, IEEE Commun. Surveys Tuts., vol. 26, no. 3, pp. 1676–1709, 3rd Quart., 2024.
- [10] **G. Charan et al.**, *Vision-position multi-modal beam prediction using real millimeter wave datasets*, Proc. IEEE WCNC, Apr. 2022, pp. 2727–2731.
- [11] **S. Jiang, G. Charan, and A. Alkhateeb**, *LiDAR aided future beam prediction in real-world millimeter wave V2I communications*, IEEE Wireless Commun. Lett., vol. 12, no. 2, pp. 212–216, Feb. 2023.
- [12] **G. Charan et al.**, *Towards real-world 6G drone communication: Position and camera aided beam prediction*, Proc. IEEE GLOBECOM, Dec. 2022, pp. 2951–2956.
- [13] **I. Ahmad et al.**, *Vision-assisted beam prediction for real world 6G drone communication*, Proc. IEEE PIMRC, Sep. 2023, pp. 1–7.
- [14] **G. Charan, M. Alrabeiah, T. Osman, and A. Alkhateeb**, *Camera based mmWave beam prediction: Towards multi-candidate real-world scenarios*, IEEE Trans. Veh. Technol., vol. 74, no. 4, pp. 5897–5913, Apr. 2025.
- [15] **G. Charan and A. Alkhateeb**, *Sensing-aided 6G drone communications: Real-world datasets and demonstration*, IEEE Open J. Commun. Soc., vol. 6, pp. 8745–8774, 2025.
- [16] **A. Alkhateeb et al.**, *DeepSense 6G: A large-scale real-world multi-modal sensing and communication dataset*, IEEE Commun. Mag., vol. 61, no. 9, pp. 122–128, Sep. 2023.

- [17] **R. Vladimirov and P. Laskov**, *Challenges of predictive beamforming using geographical positioning: Insights from the DeepSense dataset*, Proc. 16th IFIP Wireless and Mobile Networking Conf. (WMNC), Nov. 2025, pp. 112–119.
- [18] **C. Jiang, X. Zhou, H. Chen, and T. Liu**, *UAV positioning using GNSS: A review of the current status*, Drones, vol. 10, no. 2, Art. no. 91, Feb. 2026.

Do Thuy Linh

Thang Long University
Ha Noi
Viet Nam
linhdt@thanglong.edu.vn

Nguyen Binh Thanh

FPT University
Ha Noi
Viet Nam
thanh14022005@gmail.com

Le Huu Son

Posts and Telecommunications Institute of Technology
Ha Noi
Viet Nam
sonlh.b22dcvt439@stu.ptit.edu.vn

Hoang Trong Minh

Posts and Telecommunications Institute of Technology
Ha Noi
Viet Nam
hoangtrongminh@ptit.edu.vn