Learning-based local routing decisions in sparse aeronautical communication networks

Musab Ahmed, Kevin Nabiev, Konrad Fuger, Koojana Kuladinithi, Andreas Timm-Giel

Angaben zur Veröffentlichung / Publication details:

Ahmed, Musab, Kevin Nabiev, Konrad Fuger, Koojana Kuladinithi, and Andreas Timm-Giel. 2026. "Learning-based local routing decisions in sparse aeronautical communication networks." In *7th KuVS Fachgespräch on Machine Learning in Networking (MaLeNe), University of Augsburg, Augsburg, Germany, March 19-20, 2026: proceedings*, edited by Michael Seufert, Andreas Blenk, and Björn Richerzhagen. Augsburg: Universität Augsburg.

7TH KUVS FACHGESPRÄCH ON
MACHINE LEARNING IN NETWORKING (MaLeNe)
PROCEEDINGS

MARCH 19-20
2026

UNIVERSITY OF AUGSBURG, AUGSBURG, GERMANY

Learning-Based Local Routing Decisions in Sparse Aeronautical Communication Networks

Musab Ahmed, Kevin Nabiev, Konrad Fuger, Koojana Kuladinithi, Andreas Timm-Giel
Hamburg University of Technology, Institute of Communication Networks, Hamburg, Germany
{musab.ahmed, kevin.nabiev, k.fuger, koojana.kuladinithi, timm-giel}@tuhh.de

***Abstract*—Geographic greedy routing provides the scalability required for future L-band Digital Aeronautical Communications System (LDACS) Air-to-Air (A2A) networks. Yet, its topology blindness becomes a critical limitation in sparse airspaces. This blindness induces a path-dependent process, recently characterized as the "memory effect," where suboptimal local choices steer packets toward dead ends multiple hops away. These failures are not arbitrary: they follow learnable patterns embedded in the local network structure. We propose a decentralized, density-adaptive routing scheme that replaces standard geographic progress toward the destination with learned neighbor ranking. The model is trained using the *Topological Advance (TA)* metric as a ground-truth label to indicate whether a chosen next-hop reduces the shortest-path hop distance to the destination. By predicting TA from three-hop local features, nodes improve routing reliability. Evaluations in realistic French airspace scenarios show a success ratio exceeding 0.93 over topologically connected nodes in sparse networks, a 33% improvement over traditional greedy routing, while maintaining a near-optimal hop stretch of 1.03.**

***Index Terms*—LDACS, Air-to-Air communication, geographic greedy routing, betweenness centrality, machine learning.**

I. Introduction

Modernizing air traffic management to handle projected density growth requires L-band Digital Aeronautical Communications System (LDACS), which is currently standardized to replace legacy analog systems with Air-to-Ground (A2G) and Air-to-Air (A2A) links [1–3]. For A2A multi-hop communication, geographic greedy routing (Greedy-1), which forwards to the direct neighbor closest to the destination, is attractive for its scalability [4, 5]. Yet, it fails in sparse airspaces (fewer than ten neighbors) even when topological paths exist [5]. While traditionally attributed to a local minimum, where no reachable neighbor is closer to the destination, this failure is now understood as a path dependency termed the "memory effect" [6]. As shown in Figure 1, a sequence of locally greedy choices can constrain future options, making the eventual local minimum a late symptom of a suboptimal trajectory. Increasing the neighborhood radius (Greedy-k) does not resolve this, as the protocol remains blind to underlying topological constraints [6]. However, the structural nature of this path dependency implies that local routing decisions follow predictable patterns that can be learned.

Key challenges remain when applying machine learning to the routing problem in sparse scenarios. Prior methods often target dense scenarios to optimize delay [7–9] or assume impractical global topology state [10, 11]. Others attempt trap escape reactively after a packet is already stuck [12], a stage

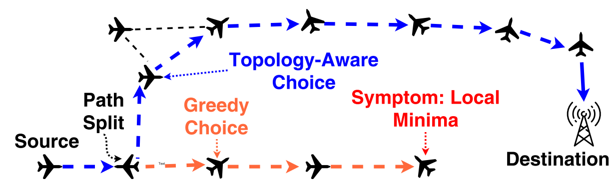

Fig. 1: Memory effect.

where recovery is difficult in sparse networks. A proactive, topology-aware policy for sparse scenarios is currently missing. We bridge this gap with a decentralized, density-adaptive neighbor ranker. Our approach utilizes three-hop local features and is supervised by Topological Advance (TA) metric, a measure of hop-distance progress toward the destination. This enables forwarding decisions that proactively avoid path-dependent traps. Our main contributions are:

- Application of the TA metric as a ground-truth label for neighbor ranking based on shortest-path progress.
- A density-adaptive neighbor-ranking policy for proactive next-hop selection without global topology state.
- A performance evaluation in realistic French airspace scenarios representing different LDACS deployment stages.

II. Ranking-based Framework

The proposed scheme transforms local three-hop topology into a ranking task to select neighbors offering the maximum predicted TA toward the destination. If a topological path exists, at least one immediate neighbor lies on a shortest path toward the destination and thus yields TA. To avoid path-dependent traps, each node constructs a local representation using its three-hop neighborhood. Input features include geometric metrics (e.g., distance and angle to the destination) calculated by each node for its neighbors and structural metrics (e.g., degree and Betweenness Centrality (BC)) calculated by nodes and shared with their neighbors. These structural indicators identify "bridge nodes" that, while geographically distant, are topologically critical for routing. We assume full three-hop neighborhood visibility for BC calculations, which could be obtained and maintained via periodic beacons. Prior research shows this is feasible within LDACS beacon size constraints [13].

Given that routing dynamics differ significantly across network densities, the routing scheme employs a two-phase architecture. First, a lightweight classifier evaluates local structural features to determine the current density condition, using a

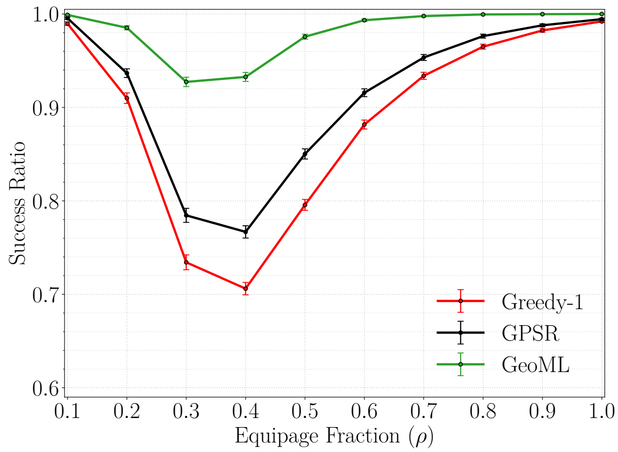

Fig. 2: Average success ratio with varying equipage fraction.

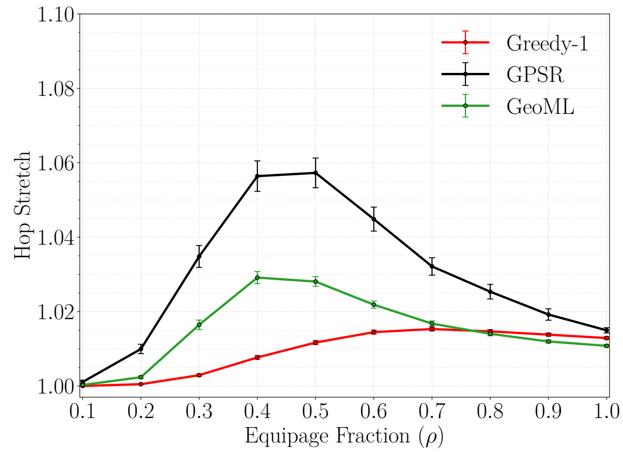

Fig. 3: Average hop stretch with varying equipage fraction.

threshold of ten neighbors to distinguish between sparse and dense environments. Based on this classification, the system selects a specialized ranking model optimized for that density condition. Next-hop selection is formulated as a ranking problem rather than simple classification, treating the decision as a preference task. This allows the model to generalize across varying topologies by ordering neighbors based on their predicted topological advance score. During training, the two rankers are supervised using binary topological advance labels derived from global shortest-path information, where a label of 1 is assigned to neighbors that decrease the shortest-path hop distance to the destination and 0 otherwise. At runtime, the node feeds its current three-hop features into the active ranker to score all candidates. The neighbor with the highest score is selected as the next-hop, promoting hop-by-hop topological progress and bridging the gap between decentralized local heuristics and global shortest-path routing.

III. Results and Discussion

The proposed routing scheme, denoted as GeoML, employs Light Gradient Boosting Machine (LightGBM) [14] for density classification and neighbor ranking. Evaluation is performed on 2000 French airspace test snapshots not used in training (trained on 150 snapshots). Monte Carlo simulations feature 500 aircraft communicating with a ground station. GeoML is benchmarked against Greedy-1 and Greedy Perimeter Stateless Routing (GPSR) (Greedy-1 with perimeter recovery). The equipage fraction (ρ) represents the percentage of aircraft with LDACS capability, where $\rho = 0.4$ corresponds to initial sparse deployment and $\rho = 0.8$ to network maturity. The density classifier labels scenarios with $\rho \geq 0.5$ (average degree ≥ 10) as dense. Performance metrics, reported at 95% confidence intervals, are: the *success ratio* (for a given protocol, the fraction of topologically connected sources that reach the destination) and the *hop stretch factor* (the ratio of the protocol's path length to the global shortest-path length).

At $\rho = 0.4$, Greedy-1 and GPSR success ratios drop to approximately 70% and 77%, respectively (Figure 2). This degradation highlights the fundamental limitation of distance-based heuristics in sparse A2A topologies, consistent with the path dependency discussed in Section I. The $\rho = 0.4$ scenarios are particularly challenging: the network remains connected, yet topological traps are frequent. In contrast, GeoML maintains a success ratio above 93%. Unlike reactive methods like GPSR, GeoML utilizes local three-hop neighborhood structure to identify and avoid routing traps before reaching local minima. The hop stretch factor results show that while GPSR provides a 0.07 absolute increase in success ratio over Greedy-1, it inflates hop stretch factor to around 1.06 due to perimeter detours (Figure 3). GeoML maintains high reliability with a stretch of only 1.03. The lower stretch observed for Greedy-1 in sparse scenarios reflects survivorship bias, as it often fails on complex paths. As the network matures ($\rho > 0.8$), success ratios converge. Furthermore, GeoML slightly outperforms legacy baselines in the hop stretch factor, demonstrating effective adaptation across density scenarios.

IV. Conclusion

In sparse A2A topologies, geographic greedy routing fails due to path-dependent memory effects. The frequently encountered local minima (voids) are often late symptoms of earlier suboptimal decisions rather than isolated local accidents. We introduced a density-adaptive neighbor ranking policy that replaces greedy geographic progress heuristics with topology-aware selection using only local three-hop information. By training with the TA metric as ground truth, nodes learn to proactively avoid routing traps from local observations without global topology knowledge. Evaluated on realistic French airspace snapshots, our GeoML achieves a success ratio exceeding 93% over connected sources with a hop stretch factor of only 1.03. These results outperform both Greedy-1 and GPSR without incurring reactive detours. Grounded in the path-dependency analysis in [6], these findings suggest a practical route toward reliable LDACS A2A networking, particularly in critical early deployment phases where sparse connectivity challenges reliability. Future work will extend the evaluation to additional airspaces to further assess model generalization under different mobility patterns.

REFERENCES

[1] Miguel A. Bellido-Manganell and Michael Schnell. "Towards Modern Air-to-Air Communications: The LDACS A2A Mode". In: *2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC)*. Sept. 2019, pp. 1–10. DOI: 10.1109/DASC43569.2019.9081678.

[2] Eurocontrol. *Aviation Outlook 2050*. Main report. Apr. 13, 2022.

[3] Nils Mäurer, Thomas Gräupl, and Corinna Schmitt. *L-Band Digital Aeronautical Communications System (LDACS)*. Request for Comments RFC 9372. Internet Engineering Task Force, Mar. 2023. 35 pp. DOI: 10.17487/RFC9372.

[4] Daniel Medina et al. "A Geographic Routing Strategy for North Atlantic In-Flight Internet Access Via Airborne Mesh Networking". In: *IEEE/ACM Transactions on Networking* 20.4 (Aug. 2012), pp. 1231–1244. DOI: 10.1109/TNET.2011.2175487.

[5] Musab Ahmed et al. "Enhancing Geographic Greedy Routing in Sparse LDACS Air-to-Air Networks through k-Hop Neighborhood Exploitation". In: *2024 IEEE 49th Conference on Local Computer Networks (LCN)*. Oct. 2024, pp. 1–7. DOI: 10.1109/LCN60385.2024.10639650.

[6] Musab Ahmed et al. "Modeling of Geographic Greedy Routing in Sparse LDACS Air-to-Air Networks Using Absorbing Markov Chains". In: *2025 IEEE 50th Conference on Local Computer Networks (LCN)*. Oct. 2025, pp. 1–9. DOI: 10.1109/LCN65610.2025.11146308.

[7] Zhigang Jin, Qinyi Zhao, and Yishan Su. "RCAR: A Reinforcement-Learning-Based Routing Protocol for Congestion-Avoided Underwater Acoustic Sensor Networks". In: *IEEE Sensors Journal* 19.22 (Nov. 2019), pp. 10881–10891. DOI: 10.1109/JSEN.2019.2932126.

[8] Dong Liu et al. "Deep-Learning-Aided Packet Routing in Aeronautical Ad Hoc Networks Relying on Real Flight Data: From Single-Objective to Near-Pareto Multiobjective Optimization". In: *IEEE Internet of Things Journal* 9.6 (Mar. 2022), pp. 4598–4614. DOI: 10.1109/JIOT.2021.3105357.

[9] Jianmin Liu, Dan Li, and Yongjun Xu. "Deep Distributional Reinforcement Learning-Based Adaptive Routing With Guaranteed Delay Bounds". In: *IEEE/ACM Transactions on Networking* 32.6 (Dec. 2024), pp. 4692–4706. DOI: 10.1109/TNET.2024.3425652.

[10] Shanshan Jiang, Zhitong Huang, and Yuefeng Ji. "Adaptive UAV-Assisted Geographic Routing With Q-Learning in VANET". In: *IEEE Communications Letters* 25.4 (Apr. 2021), pp. 1358–1362. DOI: 10.1109/LCOMM.2020.3048250.

[11] Yaqi Ke et al. "Distributed Routing Optimization Algorithm for FANET Based on Multiagent Reinforcement Learning". In: *IEEE Sensors Journal* 24.15 (Aug. 2024), pp. 24851–24864. DOI: 10.1109/JSEN.2024.3415127.

[12] Chi Wei et al. "QFAGR: A Q-learning-based Fast Adaptive Geographic Routing Protocol for Flying Ad Hoc Networks". In: *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*. Dec. 2023, pp. 4613–4618. DOI: 10.1109/GLOBECOM54140.2023.10437104.

[13] Musab Ahmed et al. "Poster: Can Betweenness Centrality Mitigate Greedy Routing's Sparse Network Problem?" In: *2025 IEEE Vehicular Networking Conference (VNC)*. June 2025, pp. 1–2. DOI: 10.1109/VNC64509.2025.11054179.

[14] Guolin Ke et al. "LightGBM: A Highly Efficient Gradient Boosting Decision Tree". In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017.