GNN-Guided Surrogate-Accelerated Route Optimization for QoS- and QoE-Aware Urban Air Mobility

Yurika Maeda
Ochanomizu University
Graduate School of Humanities and Sciences
Tokyo, Japan
g2440416@edu.cc.ocha.ac.jp

***Abstract*—UAM route planning for aerial tourism must balance passenger experience, communication reliability, and community acceptability. These objectives often conflict and are computationally expensive to evaluate, as they depend on high-resolution urban geometry and geospatial data. We formulate route design as a four-objective optimization problem under flight-time constraints. The objectives are: (1) maximizing along-route landmark-exposure reward as a QoE proxy, (2) maximizing non-residential overflight ratio as a noise-acceptability proxy, (3) maximizing the 10th-percentile RSRP as a QoS robustness proxy, and (4) minimizing route length. To reduce computational cost, we propose a GNN-guided NSGA-III framework in which a graph neural network predicts expensive objectives from a route graph during evolutionary search, and the final non-dominated set is re-evaluated using the original simulator. In a Yokohama waterfront case study with 24 POIs and open 3D city data, the surrogate achieves strong predictive accuracy, with coefficients of determination of 0.94, 0.92, and 0.88 for Scene, Quiet, and Comm, respectively, and reduces expensive simulation calls by approximately 89%. After true re-evaluation, the proposed approach produces a substantially larger true non-dominated set (92 solutions vs. 20 for the baseline) while maintaining comparable best-achieved objective values, enabling richer trade-off exploration for decision support in UAM route design.**

Keywords*— *Urban air mobility, UAV routing, QoS and QoE, multi-objective optimization, NSGA-III, surrogate model, graph neural network, noise-aware planning.

I. Introduction

Urban air mobility (UAM) services are expanding beyond logistics and inspection to include passenger-facing applications such as aerial tourism in dense urban areas. In such scenarios, the route itself becomes part of the service quality. Passengers prefer routes that pass near attractive landmarks. Operators must ensure reliable connectivity for safety-critical communication [1], [2]. In addition, flight paths must avoid noise-sensitive areas to maintain social acceptance [3]. Designing routes that simultaneously satisfy these heterogeneous requirements is challenging because improving one metric (e.g., collecting more reward) can degrade others (e.g., increasing distance, time, or exposure to residential areas).

From a computational perspective, realistic evaluation of user reward and connectivity is expensive because it requires dense spatial sampling along candidate routes. User-centric reward may depend on continuous spatial exposure to landmarks along a trajectory, and connectivity depends on urban blockage and base-station deployment. These components act as black-box functions embedded in a combinatorial optimization problem. Consequently, evolutionary search with full simulation for every candidate becomes prohibitively slow as the solution space grows.

This work addresses these challenges by combining (i) a multi-objective formulation that couples user-oriented reward, network QoS (Quality of Service) robustness, environmental acceptability, and operational cost under resource constraints, and (ii) a learning-based surrogate that accelerates many-objective evolutionary optimization. Although the motivating application is an urban aerial tour, the formulation is generic and applies to reward-collecting UAM services that must remain connected and socially acceptable.

Our main contributions are as follows: (1) We formulate UAM route planning as a four-objective optimization problem that jointly considers QoE (Quality of Experience) reward, noise acceptability, connectivity robustness, and operational cost under flight-time constraints. (2) We propose a GNN (Graph Neural Network)-guided Non-dominated Sorting Genetic Algorithm III (NSGA-III) framework where a graph neural network predicts expensive objectives with an R^2 of over 0.88, reducing simulation calls by 89%. (3) We demonstrate the framework's effectiveness on a realistic Yokohama scenario with 24 POIs using open 3D city data, showing that surrogate guidance discovers high-quality routes with significantly reduced computational cost.

II. Related Work

The UAM routing framework proposed in this study addresses the intersection of multiple technical challenges: ensuring reliable connectivity, mitigating environmental impact, and solving complex multi-objective optimization problems. Accordingly, this section reviews the related literature and background from four key perspectives relevant to our formulation: (A) communication-aware routing strategies; (B) noise and acceptability considerations for low-altitude operations; (C) mathematical models for resource-constrained reward maximization; and (D) surrogate-assisted optimization techniques for computationally expensive objectives.

A. Communication-aware UAV/UAM routing

Cellular-connected UAVs have been extensively studied for 4G/5G networks, highlighting challenges such as altitude-dependent interference and urban blockage [1], [2]. Several works optimize trajectories to improve coverage or reliability under cellular constraints, often using simplified channel models or discrete waypoints [4].

B. Noise- and acceptability-aware planning

Noise footprints and community acceptance are increasingly recognized as critical for low-altitude aerial operations. Prior studies have considered noise-aware path planning using noise assessment models; however, jointly integrating noise, connectivity, and user reward remains underexplored [3].

C. Resource-constrained reward maximization and orienteering

The orienteering problem and its variants model route planning under time/budget constraints to collect spatially distributed rewards [5], [6]. These formulations are relevant to UAV data collection and user-centric routing but typically assume rewards that are inexpensive to evaluate.

D. Surrogate-assisted and learning-based optimization

Surrogate-assisted evolutionary algorithms (SAEAs) are effective for expensive multi-objective optimization by learning approximations of costly objectives [7], [8]. Recently, graph neural networks have been used to represent combinatorial structures, enabling learning on route graphs. Our approach differs from end-to-end neural solvers: we use a GNN as a surrogate within NSGA-III to accelerate exploration while retaining constraint handling and final re-evaluation with the original simulator [9]. Unlike end-to-end neural combinatorial solvers trained with reinforcement learning [9], which often struggle to strictly satisfy hard constraints such as flight time windows and no-fly zones, our proposed hybrid framework leverages the constraint-handling capability of NSGA-III while using GNNs solely for accelerating objective evaluation, ensuring the feasibility of the final solutions.

Table I compares our work with prior studies across five key aspects. To the best of our knowledge, this is the first work to jointly address all five dimensions in a unified optimization framework. Prior UAV communication studies [1], [2], [4] focus on connectivity but neglect user experience and noise. Noise-aware path planning approaches [3] address community acceptance but typically neglect communication constraints. Orienteering formulations [5], [6] model reward collection but assume inexpensive evaluation. Surrogate-assisted methods [7], [8] target generic expensive optimization without domain-specific route representation. Neural combinatorial solvers [9] learn end-to-end policies but lack explicit constraint handling. Our work uniquely combines all five aspects: user-centric QoE reward, communication QoS, noise acceptability, and GNN-based surrogate guidance within a constrained many-objective framework.

Table I Comparison of related work across five key aspects. ✓ indicates the aspect is explicitly addressed; × indicates it is not considered. Our work is the first to integrate all five dimensions within a unified optimization framework.

Study	QoE Reward	QoS (Comm)	Noise Aware	Surrogate Model	Multi-obj Optim.
Zeng et al. [1]	×	✓	×	×	×
3GPP TR 36.777 [2]	×	✓	×	×	×
Tan et al. [3]	×	×	✓	×	×
Behjati et al. [4]	×	✓	×	×	×
Vansteenwegen et al. [5]	✓	×	×	×	×
Gunawan et al. [6]	✓	×	×	×	×
He et al. [7]	×	×	×	×	✓
Díaz-Manríquez et al. [8]	×	×	×	✓	✓
Kool et al. [9]	✓	×	×	✓	×
Ours	✓	✓	✓	✓	✓

III. Problem Formulation

We consider a set of points of interest (POIs) V = {v0, v1, ..., vN}, where v0 is a fixed depot (takeoff/landing site) and the remaining POIs are candidate landmarks. A route is an ordered sequence r = (v0, vi1, ..., viK, v0), where K ∈ [2, 6] is the number of visited POIs (excluding the depot) and no POI is visited more than once.

To capture the sequential nature of user experience, we discretize the polyline of the route r into sample points P(r) with a fixed interval d_{step} (as illustrated in Fig. 1(a)). At each sample point x ∈ P(r), we evaluate (i) a user-side reward (Quality of Experience, QoE proxy), (ii) a non-residential indicator (noise-acceptability proxy), and (iii) communication quality (Quality of Service, QoS proxy).

A. POI-level reward as a QoE proxy

While visual attractiveness is subjective, we construct an interpretable, rule-based decomposition to obtain a reproducible POI reward. Each POI v_i has attributes: iconicity I_i, visual prominence V_i, night-time attractiveness N_i, and waterfront/skyline relevance W_i. The POI weight is defined as:

$$w_i = 0.90 + 0.70 \times (0.40 I_i + 0.25 V_i + 0.15 N_i + 0.10 W_i) \quad (1)$$

which yields $w_i \in [0.90, 1.53]$ under the above normalizations.

B. Route-level QoE reward

For each sample point *x* (see Fig. 1(a)), the instantaneous reward aggregates contributions from nearby POIs with a distance attenuation kernel a(d). As shown in Fig. 1(b), this kernel is defined as:

$$S_{scene}(x) = \sum_i w_i \, a(d(x, v_i)), \quad (2)$$

$$a(d) = \max\{0, 1 - (d/d_{max})^2\} \quad (3)$$

The route-level QoE reward is the average over samples:

$$Scene(r) = \frac{1}{|P(r)|} \sum_{x \in P(r)} S_{scene}(x) \quad (4)$$

which is a discrete approximation of a line integral along the route.

C. Quietness (noise acceptability proxy)

Let Ω_{nonres} denote the union of non-residential polygons (e.g., ports, parks, and waterfront commercial areas). For a sample point x, define q(x) = 1 if x ∈ Ω_{nonres} and 0 otherwise. We define

$$Quiet(r) = \frac{1}{|P(r)|} \sum_{x \in P(r)} q(x), \quad (5)$$

i.e., the non-residential overflight ratio.

D. Connectivity robustness (QoS proxy)

For each sample point x, we compute received signal reference power (RSRP) as

$$\mathrm{RSRP}(x) = max_{\{b\in B\}}[P_{tx} - FSPL(d_{\{b,x\}}) - n(b,x)] (6)$$

where B is the set of base stations, P_{tx} is the transmit power, $d_{\{b,x\}}$ is the 3D distance from base station b to point x, FSPL(·) is the free-space path loss, and η(b,x) is an additional NLOS attenuation applied when the link is obstructed by buildings (determined via ray-intersection with the 3D city model). To emphasize robustness to deep fades and coverage holes, connectivity is evaluated by the 10th percentile along the route:

$$\mathrm{Comm}(r) = P_{10}(\{RSRP(x) | x \in P(r)\}). (7)$$

E. Operational Cost and Constraints

We define the operational cost as the total route distance Dist(r) and minimize it:

$$\mathrm{Dist}(r) = \textstyle\sum_i \mathrm{dist}(v_{i_j}, v_{i_{j+1}}). (8)$$

Flight time is approximated by

$$\mathrm{Time}(r) = \frac{\mathrm{Dist(r)}}{v} + K \times t_{orbit} + t_{\{TO\}} (9)$$

where v is cruise speed set to 60 km/h (1.0 km/min), t_{orbit} is the loiter time per visited POI (2.5 min), and $t_{TO/L}$ is the takeoff/landing overhead (≈1 min). We impose a time-window constraint $T_{min} \leq$Time(r) $\leq T_{max}$ (8–18 min in our experiments).

Finally, we solve a constrained four-objective optimization problem: maximize Scene(r), Quiet(r), and Comm(r) while minimizing Dist(r), subject to the above constraints.

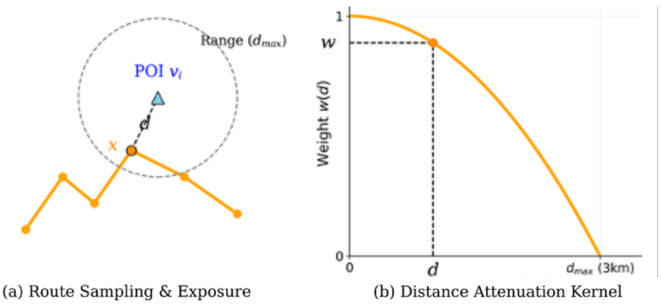

Fig. 1. Illustration of the QoE reward evaluation model. (a) Discrete approximation of the route: Sample points are generated along the route to measure the distance d to nearby POIs. (b) Distance attenuation kernel: The reward weight $w(d)$ decays quadratically with distance, reaching zero at d_{max} = 3 km.

IV. Proposed Method

To efficiently solve the resource-constrained many-objective optimization problem defined in Section III, we propose a framework that combines evolutionary computation with learning-based acceleration. Since direct evaluation of QoE and connectivity is computationally expensive, reliance on raw simulation during the search process is impractical. Therefore, this section details our approach in three parts. First, we describe the baseline NSGA-III algorithm for many-objective optimization. Second, we introduce a GNN surrogate model that predicts costly objectives from route structures. Third, we present the surrogate-guided framework that integrates both components to accelerate the search while preserving solution validity. The overall architecture of this surrogate-guided optimization framework is illustrated in Fig. 2.

A. Baseline: NSGA-III

We adopt NSGA-III as a baseline many-objective evolutionary algorithm due to its reference-direction-based selection, which is effective when optimizing three or more objectives [10]. Individuals encode a variable-length route as a random-key vector; a decoding step converts it to a POI sequence without repetition while enforcing start/end at the depot.

B. GNN Surrogate Model

Evaluating Scene(r), Quiet(r), and Comm(r) requires spatial simulation over many sample points, which becomes the main computational bottleneck. We learn a surrogate model $f_\theta(r)$ that predicts these expensive objectives from the route structure. While graph attention networks [11] offer adaptive weighting, we adopt GCN for computational efficiency.

A route is represented as a graph G=(V,E) where nodes correspond to visited POIs. Each node vi has a feature vector x_i=[lat_i,lon_i,I_i, V_i,N_i, W_i] (normalized coordinates and POI attributes). Directed edges follow the visiting order, forming a closed tour with self-loops. The adjacency matrix is symmetrically normalized: $\hat{A}$=$D^{-1/2}AD^{-1/2}$.

We employ a two-layer Graph Convolutional Network (GCN) [13] with hidden dimension 64: $H^{(l+1)}$ = ReLU ($\hat{A}$ $H^{(l)}$ $W^{(l)}$). A binary mask handles variable-length routes with K ∈ [2,6] visited POIs (excluding the depot), i.e., k = K+1 ∈ [3,7] nodes including the depot. Mean pooling over valid nodes produces a fixed-size route embedding $\mathbb{R}^{64}$, which a two-layer MLP maps to the three predicted objectives ($\hat{S}$,$\hat{Q}$,$\hat{C}$).

The surrogate was trained offline on 5,000 randomly generated routes evaluated by the true simulator. We used Adam optimizer (learning rate 10^{-3}) for 100 epochs, minimizing MSE with feature-wise normalization. To validate surrogate accuracy, we held out 500 routes for testing. Evaluation results showed that the GNN achieves R^2 = 0.94 for Scene, R^2 = 0.92 for Quiet, and R^2 = 0.88 for Comm, confirming sufficient accuracy for guiding evolutionary search.

C. Surrogate-guided NSGA-III

During evolutionary search, we compute Dist(r) and Time(r) exactly (computationally inexpensive), while replacing expensive objectives with surrogate predictions: (Ŝcene, Q̂uiet, Ĉomm)=f_θ(r). Infeasible routes violating the time-window constraint are penalized using constraint violation values following standard practice in constrained multi-objective optimization [12]. NSGA-III selection is performed based on the predicted objective vector and the exact distance objective. After optimization, we re-evaluate the final non-dominated set with the original simulator to report true objective values. This hybrid approach preserves constraint handling and solution validity while significantly reducing expensive simulations inside the evolutionary loop.

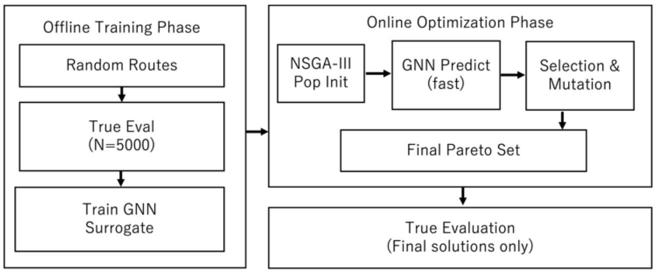

Fig. 2. Overview of the proposed GNN-guided NSGA-III framework. Left: offline training phase where the GNN surrogate is trained on 5,000 randomly sampled routes. Right: online optimization phase where the surrogate predicts expensive objectives (Scene, Quiet, Comm) during evolutionary search, and only the final Pareto set undergoes true evaluation.

V. EXPERIMENTAL SETUP

A. Urban scenario and data

We evaluate the proposed method in a realistic urban environment: the Yokohama waterfront district. The area contains dense high-rise buildings and mixed land-use (commercial/waterfront vs residential blocks), making it suitable to study the trade-off between user reward, connectivity robustness, and noise acceptability. We prepared a set of 24 POIs, including the depot at Landmark Tower and major landmarks (e.g., Red Brick Warehouse, Osanbashi Pier, Yamashita Park, marine structures, and high-rise buildings). POI attributes (Ii,Vi,Ni,Wi) were computed based on explicitly defined rules described in Section III-A. Building geometry was obtained from open 3D city model data (e.g., CityGML/PLATEAU), and land-use polygons were used to define non-residential regions [14]. Base-station locations were synthetically generated following typical urban macro-cell deployment patterns (inter-site distance 500 m, antenna height 25 m).

B. Parameters

Unless stated otherwise, we used d_{step}= 100 m for route sampling, d_{max}= 3 km for the QoE kernel, cruise speed v= 60 km/h, t_{orbit} = 2.5 min per POI and $t_{TO/L}\approx$ 1.0 min. The number of visited POIs was constrained to K $\in$ [2, 6]. The time constraint was set to 8 ≤ Time(r) ≤ 18 minutes. The flight altitude was fixed at 150 m above ground level.

For the communication model, we assumed a carrier frequency of 2.1 GHz and a base station transmit power of 43 dBm. The path loss was calculated using the Free Space Path Loss (FSPL) model with additional attenuation for non-line-of-sight (NLOS) conditions caused by building blockages.

NSGA-III was configured with a population size of 250 and run for 200 generations. We utilized Simulated Binary Crossover (SBX) with a probability of p_c= 1.0 and a distribution index of η_c= 20. Polynomial Mutation was applied with a probability of $p_m=1/L$ (where L is the number of variables) and a distribution index of η_m= 20. Reference directions were generated using the Das–Dennis structured method with division parameter p = 12; because the number of directions can exceed the population size, the niching procedure reuses reference directions as needed. The surrogate was queried for Scene, Quiet, and Comm predictions, while Dist and Time were computed exactly.

C. Rationale for baseline selection

We compare the proposed GNN-guided approach against standard NSGA-III with true evaluation rather than other surrogate methods (e.g., GP-based SAEA) for the following reasons: (1) Isolation of contribution: Using the same NSGA-III framework for both methods isolates the effect of the GNN surrogate, enabling fair assessment of surrogate guidance without confounding algorithmic differences. (2) Practical relevance: NSGA-III with true evaluation represents the realistic baseline that practitioners would use without surrogate assistance, making our comparison directly relevant to operational deployment. (3) Computational focus: Our primary claim is efficiency improvement through surrogate prediction; comparing against the same algorithm with/without surrogate directly validates this claim.

D. Evaluation protocol

After the optimization, we extracted the non-dominated solutions and evaluated them with the true simulator. We report the number of non-dominated solutions, best-achieved objective values, and representative routes.

VI. RESULTS AND DISCUSSION

A. Quantitative comparison of Pareto solutions

Table Ⅱ summarizes the non-dominated solutions obtained by each method after true re-evaluation. The proposed GNN-guided approach yields a substantially

Table Ⅱ Performance comparison of Pareto-optimal solutions after true evaluation. Arrows indicate optimization direction (↑: higher is better, ↓: lower is better). Bold values indicate superior performance.

Method	#ND	Best QoE (Scene)↑	Best Quiet↑	Best QoS (Comm)↑	Min Dist [km]↓	Avg Time [min]
Baseline	20	23.7	1	-45.05	0.97	17.33
Proposed (GNN-guided)	92	23.72	1	-45.27	0.97	17.5

Table Ⅲ Summary of objective statistics computed from the true non-dominated sets (provided by the experiment outputs).

method	n_nd	scene_max	quiet_max	comm_max	dist_min	scene_mean	quiet_mean	comm_mean	dist_mean	time_min_min	time_min_max	time_min_mean
baseline	20.00	23.70	1.00	-45.05	0.97	23.46	1.00	-48.30	1.33	16.97	17.71	17.33
guided	92.00	23.72	1.00	-45.27	0.97	23.44	1.00	-51.44	1.51	16.97	18.00	17.51

larger true non-dominated set (92 solutions) compared to the baseline (20 solutions). This indicates that surrogate guidance can expand the set of candidate trade-off routes available for decision makers, even when the final reporting is done on true objective values.

In terms of best-achieved objective values, the two methods are comparable: both reach the maximum quietness value (1.0) and the same minimum distance (0.97 km). The proposed method slightly improves the best Scene reward (23.72 vs. 23.70), while the baseline achieves a marginally better best communication value (−45.05 dBm vs. −45.27 dBm; higher is better). Average flight time remains similar (17.50 vs. 17.33 min). Overall, these results suggest that the surrogate-guided search can maintain solution quality while providing a richer Pareto set.

Table III summarizes the statistics of the obtained non-dominated solutions. It is noteworthy that while the Baseline solutions are clustered in a narrow region (resulting in lower variance and seemingly better averages for distance), the Proposed method exhibits a much broader distribution of objective values. This statistical difference indicates that the GNN-guided surrogate successfully prevents premature convergence and enables the exploration of diverse trade-off regions—such as routes that maximize quietness or rewards at the cost of longer distances—that the Baseline failed to discover. Consequently, the proposed framework provides decision makers with a significantly more comprehensive set of options (92 vs. 20 solutions).

Computational cost is where the proposed method provides a clear advantage. We define a "true evaluation" as a call to the expensive simulator that computes Scene, Quiet, and Comm for a route. The baseline required approximately $N_{pop} \times N_{gen} = 250 \times 200 = 50{,}000$ true evaluations. In contrast, the proposed method required 5,000 true evaluations for surrogate training, plus N_{final} true evaluations to re-evaluate predicted non-dominated candidates after removing duplicates ($N_{final}=400$ in our setup), totaling 5,400 true evaluations (≈89% reduction).

Fig. 3 visualizes this reduction, highlighting the efficiency benefit of surrogate-guided search.

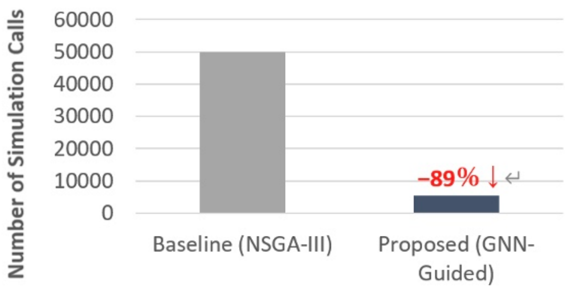

Fig. 3. Comparison of computational cost. The proposed GNN-guided method reduces the number of expensive simulation calls by approximately 89% compared to the baseline.

B. Representative routes

Fig. 4 visualizes representative routes for each method in (a) and (b), and summarizes the route overview in (c). The surrogate-guided solutions tend to concentrate on a compact high-reward corridor in the waterfront area, consistent with the spatial distribution of high-weight landmarks. In contrast, the baseline often explores more spatially dispersed routes. This qualitative difference is useful for decision support: corridor-focused routes can be attractive when planners prefer concentrated flight paths (e.g., for monitoring, public communication, or operational regularity), whereas more dispersed routes may be preferred when emphasizing other objectives such as connectivity robustness.

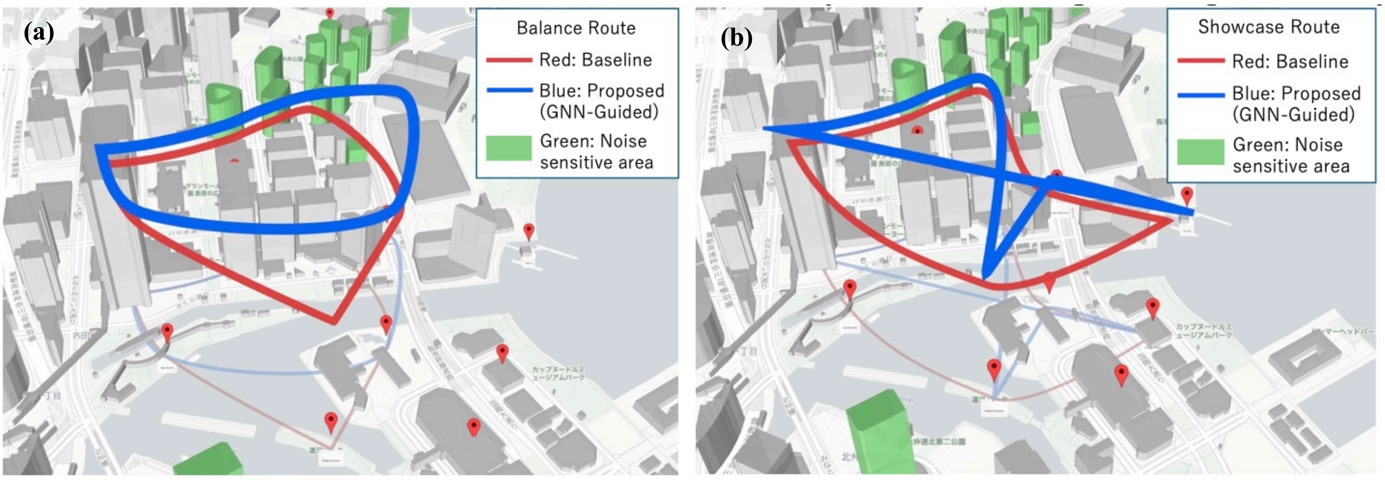

(c)

method	rep	n_poi	route	scene	quiet	dist_km	comm	time_min
baseline	showcase	6.00	LM -> QA -> QB -> QC -> CC -> CN -> KP -> LM	23.70	1.00	1.69	-45.57	17.69
baseline	balanced	6.00	LM -> QA -> QB -> QC -> CC -> KP -> NM -> LM	23.63	1.00	1.40	-45.49	17.40
guided	showcase	6.00	LM -> QA -> QB -> QC -> KP -> CC -> CN -> LM	23.72	1.00	1.96	-45.55	17.96
guided	balanced	6.00	LM -> QA -> QB -> QC -> IC -> CC -> NM -> LM	23.57	1.00	1.35	-49.04	17.35

Fig. 4. Representative routes in the Yokohama district (baseline vs proposed). (a) Balance-route example (red: baseline, blue: proposed, green: noise-sensitive area). (b) Showcase-route example (legend as in (a)). (c) Route sequences and the corresponding evaluation metrics for the representative routes.

C. Trade-off analysis

In the studied district, user reward and quietness are often aligned because attractive landmarks are located along the waterfront where non-residential overflight is possible.

However, connectivity robustness can conflict with distance and reward when the best-connected region (close to base stations) does not coincide with the highest-reward corridor. This highlights the need for true multi-objective optimization rather than single-metric routing, and motivates the use of Pareto sets for operational decision making.

D. Limitations and future work

Our study has several limitations that suggest directions for future research. First, the quietness metric is a land-use proxy and does not model acoustic propagation; integrating physics-based noise footprint models (e.g., helicopter noise contours) would strengthen environmental validity. Second, the QoE reward uses a distance-based kernel; incorporating 3D viewshed analysis with building occlusion would better capture visual exposure. Third, while the GNN surrogate accelerates search, it may reduce Pareto diversity; uncertainty-aware active learning could mitigate this. Finally, our evaluation uses a single urban scenario with 24 POIs; larger-scale studies across multiple cities would improve generalizability.

VII. Conclusion

This paper introduced a QoS- and QoE-aware route-planning framework for reward-collecting UAM services in dense urban environments, where passenger experience, connectivity robustness, and community acceptability must be balanced under strict flight-time constraints. We formulated route design as a constrained four-objective optimization problem—maximizing scenic exposure reward (QoE), non-residential overflight ratio (noise-acceptability proxy), and 10th-percentile RSRP (QoS robustness) while minimizing route length—and proposed a GNN-guided NSGA-III framework to reduce the cost of expensive route-level simulations.

In a realistic Yokohama waterfront case study with 24 POIs and open 3D city data, the learned GNN surrogate achieved strong predictive accuracy ($R^2 \geq 0.88$ for the expensive objectives) and reduced expensive simulator calls by approximately 89%. After true re-evaluation, the proposed method produced a substantially larger true non-dominated set (92 vs. 20 routes) while maintaining comparable best-achieved objective values. These results indicate that surrogate-guided many-objective optimization can serve as a practical decision support approach for designing UAM routes that are attractive to users and operationally feasible without incurring prohibitive simulation cost.

Future work will incorporate higher-fidelity 3D viewshed and physics-based noise propagation models, and extend surrogate guidance with uncertainty-aware active learning to improve robustness and scalability to larger urban instances and operational settings.

References

[1] Y. Zeng, J. Lyu, and R. Zhang, "Accessing from the sky: A tutorial on UAV communications for 5G and beyond," *Proceedings of the IEEE*, vol. 107, no. 12, pp. 2327–2375, Dec. 2019.

[2] 3GPP, "Enhanced LTE support for aerial vehicles," 3rd Generation Partnership Project (3GPP), TR 36.777, v15.0.0, Dec. 2017.

[3] Q. Tan, Y. Li, H. Wu, P. Zhou, H. K. Lo, S. Zhong, and X. Zhang, "Enhancing sustainable urban air transportation: Low-noise UAS flight planning using noise assessment simulator," Aerospace Science and Technology, vol. 147, Art. 109071, 2024, doi: 10.1016/j.ast.2024.109071.

[4] M. Behjati, R. Nordin, M. A. Zulkifley, and N. F. Abdullah, "3D Global Path Planning Optimization for Cellular-Connected UAVs under Link Reliability Constraint," Sensors, vol. 22, no. 22, Art. no. 8957, 2022, doi: 10.3390/s22228957.

[5] P. Vansteenwegen, W. Souffriau, and D. Van Oudheusden, "The orienteering problem: A survey," *European Journal of Operational Research*, vol. 209, no. 1, pp. 1–10, Feb. 2011.

[6] A. Gunawan, H. C. Lau, and P. Vansteenwegen, "Orienteering problem: A survey of recent variants, solution approaches and applications," *European Journal of Operational Research*, vol. 255, no. 2, pp. 315–332, Dec. 2016.

[7] C. He, Y. Zhang, D. Gong, and X. Ji, "A review of surrogate-assisted evolutionary algorithms for expensive optimization problems," Expert Systems with Applications, vol. 217, Art. 119495, 2023, doi: 10.1016/j.eswa.2022.119495.

[8] J. D. Díaz-Manríquez, G. Toscano, J. H. Martínez, and C. A. Coello Coello, "A review of surrogate-assisted evolutionary optimization," in *Proc. IEEE Congress on Evolutionary Computation (CEC)*, Vancouver, BC, Canada, Jul. 2016, pp. 2432–2439.

[9] W. Kool, H. van Hoof, and M. Welling, "Attention, learn to solve routing problems!," in *Proc. International Conference on Learning Representations (ICLR)*, New Orleans, LA, USA, May 2019.

[10] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, Aug. 2014.

[11] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *Proc. International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, Apr. 2018.

[12] H. Jain and K. Deb, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part II: Handling constraints and extending to an adaptive approach," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 602–622, Aug. 2014.

[13] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. International Conference on Learning Representations (ICLR)*, Toulon, France, Apr. 2017.

[14] The Ministry of Land, Infrastructure, Transport and Tourism (MLIT), Japan. "Project PLATEAU: Open 3D city models." Accessed: Dec. 12, 2025. [Online]. Available: https://www.mlit.go.jp/plateau