Supplementary Information
ARIES-Mission: A Vision-Language System for Efficient UAV Mission Generation

Junhao Wei[1], Haochen Li[1], Yifu Zhao[1], Dexing Yao[1], Yanxiao Li[1], Baili Lu[1], Sio-Kei Im[2], Xu Yang[1], Yapeng Wang[1,*]

[1]Faculty of Applied Sciences, Macao Polytechnic University, Macao 999078, China

[2]Macao Polytechnic University, Macao 999078, China

[*]Correspondence: yapengwang@mpu.edu.mo

Scope and isolation of the evaluated method

This study evaluates one fixed method only. ARIES-Mission always constructs exactly four terminal candidates: Raw recognition order, ACO, DE and PSO. It validates and scores all four under the same closed-loop Haversine objective and returns the shortest. It does not contain adaptive operator selection, learned route construction, post-hoc local-search polishing, energy modelling, multi-UAV allocation or dynamic replanning. Code and results from other ARIES-Mission variants were not imported into the experiment.

Per-task route lengths

Supplementary Table S1 reports seed-averaged route lengths for stochastic methods. Raw and HumanPlan are deterministic. HumanPlan is not a controlled routing baseline because its waypoint set can differ from the VLM output.

Supplementary Table S1: Per-task closed-loop trajectory length in kilometres.

Task	N	Raw	ACO	DE	PSO	Portfolio	Human
1	11	2.079	1.774	1.789	1.845	1.774	1.879
2	11	3.438	2.280	2.286	2.442	2.280	2.341
3	16	2.704	2.558	2.673	3.011	2.556	2.534
4	12	2.395	2.134	2.144	2.296	2.134	2.409
5	12	2.560	1.865	1.906	2.020	1.865	2.459
6	11	2.362	1.756	1.756	1.849	1.756	1.903
7	8	1.952	1.842	1.842	1.855	1.842	1.944
8	19	4.334	2.432	2.586	3.380	2.432	2.433
9	13	3.310	2.412	2.441	2.558	2.412	3.445
10	11	3.548	2.296	2.297	2.429	2.286	2.629
11	11	2.900	2.760	2.703	2.823	2.697	3.554
12	15	2.214	1.436	1.486	1.581	1.433	1.982

Task	N	Raw	ACO	DE	PSO	Portfolio	Human
13	11	2.624	2.038	2.035	2.151	2.029	2.604
14	5	1.293	1.067	1.067	1.067	1.067	3.711
15	9	2.435	2.113	2.113	2.138	2.113	2.772
16	17	2.437	2.161	2.168	2.467	2.160	2.537
17	7	2.626	2.010	2.010	2.037	2.010	2.601
18	17	3.956	2.705	2.935	3.209	2.705	3.408
19	6	1.961	1.908	1.908	1.908	1.908	2.168
20	10	2.327	2.170	2.175	2.205	2.170	2.402
21	7	2.499	2.207	2.209	2.232	2.207	3.381
22	42	0.666	0.650	1.380	1.657	0.650	1.133
23	4	2.011	2.011	2.011	2.011	2.011	2.045
24	4	1.554	1.351	1.351	1.351	1.351	1.250
25	6	2.177	1.941	1.941	1.941	1.941	1.976
26	10	3.508	2.218	2.235	2.398	2.218	2.965
27	5	2.511	1.900	1.900	1.901	1.900	2.702
28	10	3.118	2.027	2.036	2.125	2.026	2.982
29	20	5.907	2.504	2.734	3.384	2.503	3.755
30	12	2.250	1.935	1.941	2.140	1.933	2.446

Sequence and path-similarity metrics

We computed transparent path-similarity diagnostics from the archived coordinates. Each closed route, including the return to Home, was resampled to 256 equally spaced arc-length positions in a local metric coordinate system. Sequence RMSE is the root mean squared Euclidean separation between corresponding resampled positions. DTW RMSE is the root mean squared separation along the minimum-cost dynamic-time-warping alignment of those resampled routes. The reference is the HumanPlan route for the same scene. The one-way grounding distance is the mean nearest-neighbour distance from each VLM target to the HumanPlan waypoint set; it is independent of route order.

These diagnostics use seed 0 and should not be interpreted as accuracy against ground-truth building centres: VLM and HumanPlan waypoint sets differ in size and semantics. Portfolio reduced mean sequence RMSE from 271.69 to 248.44 m and mean DTW RMSE from 191.11 to 173.53 m relative to Raw. DE obtained the lowest mean DTW RMSE (149.75 m), despite not being the shortest-distance candidate in most runs. Thus minimising route length and matching a human route are distinct objectives.

Supplementary Table S2: Seed-0 path similarity to HumanPlan. Lower is better. The grounding column is identical across routing methods because the targets are frozen before routing.

Method	Grounding KNN (m)	Sequence RMSE (m)	DTW RMSE (m)
Raw	33.26	271.69	191.11
ACO	33.26	228.69	158.13
DE	33.26	230.38	149.75
PSO	33.26	258.49	166.19
Portfolio	33.26	248.44	173.53

Solver convergence

Figure 1 shows best-so-far route length as a function of objective evaluations for the three stochastic solvers, expressed as the optimality gap to the Held–Karp exact optimum and averaged over the 27 tasks with $N \leq 18$ and their 20 seeds. All three solvers stabilise well before the 7,500-evaluation budget is exhausted, and ACO both converges fastest and settles at the lowest residual gap, which is consistent with its dominant selection frequency in the main experiment.

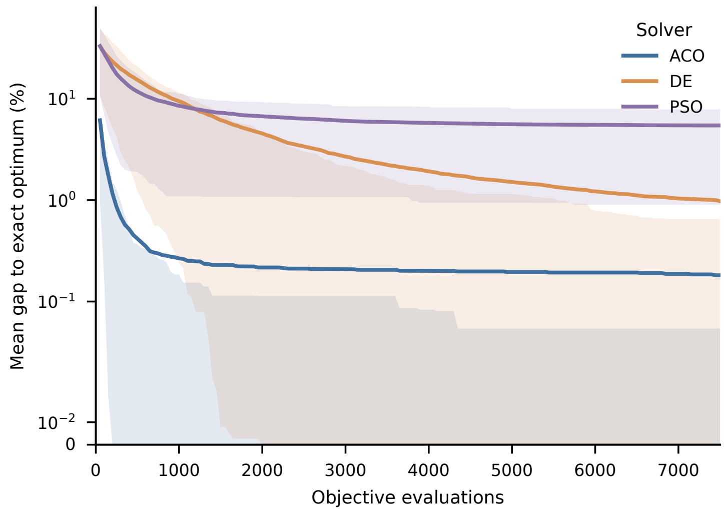

Figure 1: Convergence of ACO, DE and PSO. Lines are the mean gap to the Held–Karp optimum across the 27 exactly solvable tasks; shaded bands span the interquartile range across tasks. The vertical axis uses a symmetric-logarithmic scale.

Exact references

Held–Karp dynamic programming was used when the number of targets was at most 18, yielding exact optima for 27 tasks. Tasks 8, 22 and 29 exceeded this threshold and were excluded from exact-gap summaries rather than being compared with an unverified surrogate optimum. Every reported optimality gap uses the same floating-point Haversine evaluator as the candidate methods.

Reproducibility controls

All stochastic methods used population size 50, 7,500 complete-route evaluations and seeds 0–19. ACO established its incumbent only from sampled ant tours rather than from Raw. Each stored route was checked to be a permutation of every detected target exactly once. The Portfolio selector was also checked programmatically to be no longer than any of its four candidates on every task–seed pair. Unit tests cover coordinate parsing, closed-loop cost, route validity, ACO budget accounting, exact TSP behaviour and terminal selection. Serial timing was collected separately from the parallel experiment to avoid worker-contention artefacts.