

Proceeding Paper

Energy-Aware Route Planning for Differential Drive Mobile Robots: Feasibility First GA and PSO Benchmarking Against A* in Dense Urban Environments [†]

Vanessa Botero-Gómez ¹, Cristian M. Hernández ¹, Juan C. Tejada ², Luis Fernando Grisales-Noreña ³
and Daniel Sanin-Villa ^{4,*}

¹ Facultad de Ingenierías, Institución Universitaria ITM, Medellín 050036, Colombia; vanessabotero@itm.edu.co (V.B.-G.); cristianhernandez@itm.edu.co (C.M.H.)

² Department of Engineering Studies for Innovation, Universidad Iberoamericana, Mexico City 01219, Mexico; juan.tejada@ibero.mx

³ Grupo de Investigación en Alta Tensión, GRALTA, Escuela de Ingeniería Eléctrica y Electrónica, Facultad de Ingeniería, Universidad del Valle, Cali 760015, Colombia; grisales.luis@correounivalle.edu.co

⁴ Área de Industria, Materiales y Energía, Universidad EAFIT, Medellín 050022, Colombia

* Correspondence: dsaninv2@eafit.edu.co

[†] Presented at the 3rd International Online Conference on Energies (IOCEN 2026), 12–15 May 2026; Available online: <https://sciforum.net/event/IOCEN2026>.

Abstract

Urban service robots require route planners that are not only collision-free but also consistent with the energetic behavior of differential drive locomotion. Conventional grid planners such as A* are efficient and reliable for geometric navigation, but their usual cost structure prioritizes path length and does not explicitly account for heading changes, concentrated turns, or localization risk near obstacles. This study presents an energy-aware route-planning formulation for differential-drive mobile robots operating in dense polygonal urban environments. The path is encoded through twelve continuous internal waypoints and is evaluated using an interpretable energy proxy that combines translational distance, cumulative absolute rotation, squared rotation, and a clearance-dependent localization risk term. Collision avoidance, boundary compliance, and maximum turn feasibility are handled through a feasibility-first dominance rule, and the resulting constrained problem is solved using a Genetic Algorithm and Particle Swarm Optimization. A* with clearance inflated occupancy grids is included as a deterministic baseline. The final experiments used a dense urban scenario with sixteen polygonal obstacles, an A* grid resolution of 0.10 m, a robot radius of 0.20 m, a safety clearance of 0.10 m, 80 GA individuals, 80 PSO particles, 100 iterations, and 10 independent runs. All methods achieved a 100% feasibility rate. PSO obtained the lowest average energy proxy, 18.505, compared with 18.645 for A* and 18.645 for GA, and reduced average rotation from 4.136 rad to 4.079 rad. However, A* remained much faster, 5.860 s on average, compared with 174.574 s for GA and 175.869 s for PSO. The ranking analysis shows that GA produced the best aggregate score when route quality, time, clearance, and feasibility were weighted equally, while PSO produced the best route quality. External perturbation tests indicate that open-loop execution in narrow corridors is sensitive to bias and waypoint noise, which motivates closed-loop tracking, online replanning, and physical validation in future work.

Keywords: differential drive robot; energy-aware planning; A* baseline; polygonal obstacle avoidance; constrained optimization; particle swarm optimization; genetic algorithm; smart city robotics



Academic Editor: Marco Pasetti

Published: 13 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Autonomous mobile robots are increasingly used in warehouses, hospitals, campuses, inspection routes, and smart city services, where safe navigation must be achieved under limited onboard energy resources [1,2]. Classical path planning methods usually formulate the problem as the search for a collision-free route between a start point and a target point, using geometric length, travel time, or grid traversal cost as the dominant criterion [2,3]. This view is appropriate for many navigation tasks, but it is incomplete for differential-drive robots because two trajectories with similar Euclidean length may require different levels of in-place rotation, heading correction, wheel-slip compensation, and motor effort.

For differential drive locomotion, the energetic burden of a route is shaped by both translation and rotation. A path that appears short in Cartesian space can contain abrupt heading changes that require stop-turn-start maneuvers or high angular-velocity commands. These maneuvers increase electrical demand, reduce tracking accuracy, and may amplify localization uncertainty near obstacles [4,5]. Therefore, energy-aware planning for this robot class should not only avoid collisions but also penalize accumulated rotational effort and concentrated turns.

Conventional graph search algorithms, especially A*, are widely used because they provide fast and reproducible collision-free routes when the environment can be discretized into an occupancy grid [6]. Nevertheless, their standard implementation optimizes grid traversal cost and path length, whereas the energy implications of heading changes are usually considered only indirectly through post-processing, smoothing, or controller-level compensation. In dense urban settings with narrow passages and dead ends, this limitation becomes relevant because the planner may produce a safe and short path that still imposes unnecessary rotational effort on the robot. This observation motivates the explicit comparison with A* in the present study.

Metaheuristic algorithms have been used in diverse complex problems [7–9], and can provide an alternative for continuous waypoint optimization in nonlinear and constrained spaces [10–12]. Genetic Algorithms can preserve exploration through selection, crossover, and mutation, while Particle Swarm Optimization can refine waypoint coordinates through collective attraction toward personal and global best solutions [13–15]. These mechanisms are attractive for energy-aware route planning because they can optimize nonconvex route quality metrics that combine distance, rotation, and clearance-related risk.

The contribution of this paper is a feasibility first energy-aware planning framework for differential drive robots in polygonal urban environments. The proposed method integrates: (i) a continuous waypoint representation with exact segment to polygon safety checks, (ii) an energy proxy composed of translational distance, cumulative rotation, squared rotation, and localization risk, (iii) a feasibility first dominance rule that avoids manually tuned collision penalty weights, (iv) a benchmark against A* in a dense scenario with narrow passages and a U shaped dead end, and (v) repeatability, ranking, weight sensitivity, and external perturbation analyses. The rest of the paper is organized as follows. Section 2 presents the mathematical model and algorithms. Section 3 reports and discusses the results. Section 4 summarizes the main findings, limitations, and future work.

2. Methods

This section describes the environment model, the energy proxy, the constraint handling strategy, the A* baseline, and the continuous metaheuristic optimizers. The final experimental configuration was selected to address the requirements for benchmarking, scenario complexity, parameter reporting, sensitivity, and robustness.

2.1. Robot, Workspace, and Path Representation

The workspace is a two dimensional bounded map $\mathcal{W} = [0, 10] \times [0, 10]$ m² containing a set of polygonal obstacles $\mathcal{O} = \{O_1, O_2, \dots, O_m\}$. The start and goal positions are $S = (0.7, 0.7)$ m and $G = (9.3, 9.6)$ m, respectively. The mobile robot is represented by a circular footprint of radius $r = 0.20$ m, and a safety clearance $c = 0.10$ m is imposed around each obstacle. Therefore, a path segment is admissible only if its distance to every polygon is at least $r + c = 0.30$ m.

A candidate route is represented as a polyline $P = \{p_0, p_1, \dots, p_{M+1}\}$, where $p_0 = S$, $p_{M+1} = G$, and the $M = 12$ internal waypoints are optimization variables. The decision vector is $\mathbf{x} = [x_1, y_1, \dots, x_M, y_M]^T \in \mathbb{R}^{2M}$, with each coordinate bounded by the workspace limits. This continuous encoding allows GA and PSO to refine route geometry beyond the fixed grid used by A*.

2.2. Energy Proxy and Physical Interpretation

For each segment i , the translational length is $\ell_i = \|p_{i+1} - p_i\|_2$. The total path length is

$$L(P) = \sum_{i=0}^M \ell_i. \quad (1)$$

The heading of segment i is $\theta_i = \text{atan2}(y_{i+1} - y_i, x_{i+1} - x_i)$, and the wrapped heading change is

$$\Delta\theta_i = \text{wrap}_\pi(\theta_i - \theta_{i-1}), \quad i = 1, \dots, M. \quad (2)$$

Two rotational indicators are considered. The term $R_1(P) = \sum_{i=1}^M |\Delta\theta_i|$ approximates the cumulative angular effort required by the differential drive system, while $R_2(P) = \sum_{i=1}^M (\Delta\theta_i)^2$ penalizes concentrated turns more strongly than distributed smooth heading corrections. A clearance-dependent localization risk term is also included,

$$R_{\text{loc}}(P) = \sum_{i=0}^M \sum_{j=1}^m \frac{\sigma_{\text{loc}}^2}{\max(d(p_i, p_{i+1}, O_j) - (r + c), 10^{-4})}, \quad (3)$$

where $d(p_i, p_{i+1}, O_j)$ is the minimum distance between segment (p_i, p_{i+1}) and obstacle O_j , and $\sigma_{\text{loc}} = 0.05$ m is the nominal localization uncertainty. The final energy proxy is

$$J(P) = \alpha L(P) + \beta R_1(P) + \gamma R_2(P) + \eta R_{\text{loc}}(P). \quad (4)$$

The adopted coefficients are $\alpha = 1.00$, $\beta = 0.60$, $\gamma = 0.20$, and $\eta = 0.02$. This cost is an interpretable energy proxy rather than a calibrated electrical energy model in joules. Its purpose is to rank feasible trajectories according to translation, rotation, concentrated turning, and proximity risk. The value $\beta = 0.60$ was selected after a sensitivity analysis over $\beta \in \{0.0, 0.1, 0.2, 0.4, 0.6, 0.8\}$, reported in Section 3.6.

2.3. Constraint Handling Through Feasibility First Dominance

The total constraint violation combines collision clearance, boundary compliance, and maximum turn feasibility. The collision component is

$$V_{\text{obs}}(P) = \sum_{i=0}^M \sum_{j=1}^m \max(0, r + c - d(p_i, p_{i+1}, O_j))^2. \quad (5)$$

The boundary violation $V_{\text{box}}(P)$ penalizes waypoints outside $\mathcal{W}$, and the turn violation $V_{\text{turn}}(P)$ penalizes turns above 100° . The total violation is

$$V(P) = V_{\text{obs}}(P) + V_{\text{box}}(P) + V_{\text{turn}}(P). \quad (6)$$

A path is feasible when $V(P) \leq 10^{-10}$. Candidate solutions are compared using a feasibility-first rule: a feasible solution dominates any infeasible solution; between two feasible solutions, the lower $J(P)$ is preferred; and between two infeasible solutions, the lower $V(P)$ is preferred. This rule reduces dependence on manually adjusted obstacle penalty coefficients.

2.4. A* Baseline

A* was implemented as a deterministic grid-based baseline. The workspace was discretized with a resolution of 0.10 m, obstacles were inflated through the robot radius and safety clearance, and an eight-connected neighborhood was used. The heuristic was the Euclidean distance to the goal. After the grid path was obtained, a shortcut smoothing step removed intermediate vertices when a direct line of sight preserved the required clearance. This baseline represents a classical, fast, and reproducible planner against which the energy-aware continuous optimizers can be compared.

2.5. GA and PSO Implementation

The Genetic Algorithm encodes each path as the decision vector $\mathbf{x}$. Tournament selection with tournament size 3 was used, crossover was performed as a convex combination of two parents, and mutation was applied using Gaussian perturbations with an initial scale equal to 8% of the coordinate range. The final GA configuration used a population of 80 individuals, 3 elite individuals, crossover probability $p_c = 0.85$, mutation probability $p_m = 0.20$, and 100 generations.

In PSO, each particle also represents the full waypoint vector. The velocity update used inertia $\omega = 0.72$, acceleration constants $c_1 = c_2 = 1.60$, and a maximum velocity equal to 20% of the coordinate range. The final swarm size was 80 particles and the number of iterations was 100. Both GA and PSO were seeded around the A* reference path to ensure feasible initial search regions while still allowing stochastic refinement of waypoint coordinates. Full parametrization of the algorithms and test scenario are presented in Table 1.

Table 1. Algorithmic and scenario parameters.

Parameter	Value
Workspace	$[0, 10] \times [0, 10] \text{ m}^2$
Start point	$S = (0.7, 0.7) \text{ m}$
Goal point	$G = (9.3, 9.6) \text{ m}$
Robot radius, r	0.20 m
Safety clearance, c	0.10 m
Required obstacle margin, $r + c$	0.30 m
Internal waypoints, M	12
A* grid resolution	0.10 m
GA population size	80
PSO swarm size	80
Maximum iterations	100
Independent runs	10
Energy weights	$\alpha = 1.00, \beta = 0.60, \gamma = 0.20, \eta = 0.02$

2.6. Evaluation Metrics and Ranking

The evaluation considered the energy proxy $J(P)$, Euclidean distance $L(P)$, cumulative rotation $R_1(P)$, squared rotation $R_2(P)$, number of turns above 5° , minimum clearance, feasibility rate, and runtime. For ranking, lower values were preferred for energy, distance, rotation, number of turns, and runtime, whereas higher values were preferred for clearance and success rate. The global ranking score was computed as the arithmetic mean of the metric ranks. This ranking does not replace the primary energy objective; instead, it provides a multi-criteria view of practical tradeoffs.

3. Results and Discussion

The experiments used the dense urban scenario shown in Figure 1, composed of sixteen polygonal obstacles, including narrow passages, blocks, rotated polygons, and a U-shaped obstacle near the upper right part of the map. This scenario is more demanding than the original sparse environment because it forces the route planner to traverse constrained corridors and to balance length, rotation, and clearance.




Figure 1. Dense urban scenario and representative routes obtained with A*, GA, and PSO. All routes are feasible, but PSO slightly modifies the A* seeded route to reduce energy and rotation.

3.1. Single Representative Comparison

Table 2 reports a representative run with seed 42. A* obtained an energy proxy of 18.645 and a runtime of 5.767 s. GA converged to a route with almost the same geometry as A*, with an energy proxy of 18.648 and a runtime of 183.938 s. PSO reduced the energy proxy to 18.458, corresponding to a 1.00% reduction relative to A*, and also reduced distance from 15.490 m to 15.435 m and rotation from 4.136 rad to 4.045 rad. The result confirms that the continuous waypoint optimizer can refine a feasible A* route, but it also shows a clear computational overhead.

Table 2. Representative comparison for the dense urban scenario, seed 42.

Algorithm	Energy	Distance	Rotation	R_2	Turns	Clearance	Time	Success
A*	18.645	15.490	4.136	3.211	7	0.110	5.767	100.0
GA	18.648	15.490	4.136	3.211	7	0.110	183.938	100.0
PSO	18.458	15.435	4.045	2.699	7	0.103	175.113	100.0

Figure 2 shows that GA rapidly stabilized near the seeded A* route and did not produce relevant improvement after the early iterations. PSO also remained nearly unchanged during the first half of the run, but after approximately iteration 55 it progressively reduced the best score and ended near 18.43. This pattern suggests that the PSO velocity update can perform delayed exploitation once a direction of improvement is found in the continuous waypoint space.




Figure 2. Convergence behavior of GA and PSO in the representative dense urban case. GA remained close to the A* seeded solution, while PSO produced late but substantial objective improvement after approximately 55 iterations.

3.2. Repeatability over Ten Independent Runs

Table 3 summarizes the results over ten independent runs of the real-time algorithm. A* is deterministic and is repeated only for table symmetry. All methods achieved a 100% feasibility rate. PSO achieved the lowest mean energy proxy, 18.505, compared with 18.645 for A* and 18.645 for GA. The average reduction relative to A* was 0.75%. PSO also obtained the shortest distance, 15.445 m, and the lowest cumulative rotation, 4.079 rad, representing a 1.37% rotation reduction relative to A*. However, the average PSO time was 175.869 s, while A* required only 5.860 s. Therefore, PSO improves route quality but is not a direct replacement for A* when strict real time response is required.

Table 3. Repeatability metrics over 10 independent runs in the dense urban scenario.

Algorithm	Energy Avg.	Energy Std.	Dist. Avg.	Dist. Std.	Rot. Avg.	Turns Avg.	Time Avg.	Time Std.	Clear. Avg.	Success (%)
A*	18.645	0.000	15.490	0.000	4.136	7.0	5.860	0.000	0.110	100.0
GA	18.645	0.006	15.488	0.005	4.136	7.0	174.574	3.269	0.110	100.0
PSO	18.505	0.042	15.445	0.013	4.079	7.1	175.869	2.984	0.105	100.0

The boxplots in Figure 3 confirm the deterministic character of A* and the near deterministic response of GA under the adopted seeding strategy. PSO exhibits a wider distribution because it explores more strongly around the A* reference path. This dispersion is not necessarily negative, since the best and average PSO solutions have lower energy and lower rotation than the other two alternatives.




Figure 3. Distribution of energy and cumulative rotation over the repeated experiments. A* is deterministic; GA remains very close to the A*-seeded route; and PSO shows greater dispersion while achieving lower route-quality metrics.

3.3. Ranking Analysis

Table 4 presents a criterion-based ranking of the evaluated methods using the repeatability averages. For each performance criterion, the methods were assigned integer ranks, where rank 1 denotes the best method for that criterion, rank 2 the second-best method, and rank 3 the third-best method. Energy, distance, rotation, number of turns, and computation time were ranked in ascending order because lower values are preferred. Minimum clearance and success rate were ranked in descending order because larger clearance margins and higher feasibility rates are preferred. When two methods obtained the same value for a criterion, the same best available rank was assigned to both methods. This tie rule explains why A* and GA both receive rank 1 for the number of turns, and why all methods receive rank 1 for success rate, since all of them achieved 100% feasible solutions.

Table 4. Explicit criterion-based ranking of the evaluated methods.

Algorithm	Energy	Distance	Rotation	Turns	Time	Clearance	Success	Score	Global
A*	3	3	3	1	1	1	1	1.857	2
GA	2	2	2	1	2	2	1	1.714	1
PSO	1	1	1	3	3	3	1	1.857	2

Note: Ranks 1, 2, and 3 indicate best, second best, and third best performance for each criterion, respectively. Energy, distance, rotation, turns, and time are ranked in ascending order, whereas clearance and success are ranked in descending order. The score is the arithmetic mean of the seven criterion ranks, and the global rank is obtained from this score.

The aggregate score was computed as the arithmetic mean of the seven criterion ranks,

$$R_a = \frac{1}{N_c} \sum_{j=1}^{N_c} r_{a,j}, \quad (7)$$

where R_a is the aggregate rank score of algorithm a , $N_c = 7$ is the number of ranking criteria, and $r_{a,j}$ is the rank assigned to algorithm a under criterion j . Therefore, a lower value of R_a indicates better-balanced performance across all criteria. The final global rank was then obtained by ranking the aggregate scores in ascending order.

The ranking shows that PSO provides the best route-quality metrics, since it ranks first in energy, distance, and rotation. This confirms that PSO is the strongest alternative when the primary objective is to minimize the motion-related cost of the planned route. However, PSO receives lower ranks in number of turns, computation time, and clearance, which penalizes its aggregate score under the equal-weight ranking rule. A* obtains the best ranks in computation time and clearance, confirming its value as a fast and conservative classical planner. GA obtains the best aggregate score, $R_a = 1.714$, because it improves energy, distance, and rotation with respect to A*, keeps the same number of turns as A*, and avoids the larger time and clearance penalties observed for PSO. Thus, the ranking should not be interpreted as evidence of a single dominant method under all operating priorities. Instead,

it quantifies the engineering tradeoff between route quality, computational burden, and safety margin.

3.4. Practical Significance and Computational Deployment

The computational results show that the proposed energy-aware refinement should not be interpreted as a direct replacement for A* in rapid replanning tasks. In the dense urban scenario, A* required an average computation time of 5.860 s, whereas GA and PSO required 174.574 s and 175.869 s, respectively. Therefore, the current GA and PSO implementations are approximately 30 times slower than A*. At the same time, PSO reduced the average energy proxy from 18.6446 to 18.5046, approximately 0.75%, while reducing the average path length from 15.4901 m to 15.4455 m and the average cumulative rotation from 4.1360 rad to 4.0794 rad. These gains are modest for a single route execution, but they can become significant when the same route is repeatedly executed by battery-constrained robots over long operating periods.

The practical value of the proposed method is therefore associated with supervisory or offline route refinement in structured environments where the map is known in advance, and the optimized route can be reused many times. Representative examples include autonomous mobile robots in warehouses, hospital logistics robots, airport service robots, industrial inspection platforms, agricultural monitoring robots, and security patrol robots that operate along repeated routes. In these scenarios, a route may be planned before a shift, mission, or service cycle, and then executed tens or hundreds of times. Under these conditions, even a small reduction in motion-related energy can translate into longer operating time, fewer charging interruptions, reduced battery cycling, and lower mechanical stress from unnecessary rotations. This interpretation is particularly relevant to differential-drive robots, where turn-intensive paths can increase wheel slip, tire wear, actuator effort, and localization uncertainty.

Based on these results, the recommended deployment strategy is hierarchical. A* should be retained as the primary planner for fast route generation, emergency replanning, and unexpected obstacle avoidance. The GA or PSO-based energy-aware module should operate as a supervisory refinement layer when enough computation time is available.

3.5. Computational Burden Reduction Strategies

Several implementation and structural strategies can reduce the computational burden observed in the current version. First, the search can be restricted to a corridor around the A* route rather than allowing each waypoint to move across the full map. This would reduce the feasible search space and decrease the probability of generating infeasible candidates. Second, the number of optimized waypoints can be made adaptive. Long nearly straight segments can be represented with few variables, whereas only geometrically complex regions near obstacles require additional degrees of freedom. Third, early stopping can be used when the best feasible cost does not improve by more than a prescribed tolerance during a fixed number of generations. Fourth, the population size can be decreased after convergence begins, which would preserve exploration during the first iterations while reducing repeated evaluations in the final refinement stage.

The highest computational cost in the current implementation is associated with repeated path evaluation, especially collision checking, clearance estimation, and constraint verification for all agents and all iterations. Therefore, additional acceleration can be obtained through vectorized distance calculations, spatial indexing of obstacles, caching of repeated segment evaluations, and parallel evaluation of individuals. These modifications are compatible with GA and PSO because each candidate path can be evaluated independently within each generation. A further reduction can be obtained through hybrid local

refinement, where PSO or GA performs a limited global search initialized from A^* , and a deterministic smoother then improves the best feasible route. This hybrid structure would retain the route-quality advantage of the energy-aware formulation while reducing the runtime penalty observed in the present experiments.

3.6. Sensitivity to the Rotational Weight

Table 5 presents the sensitivity analysis for the PSO planner. Because changing β changes the definition of the objective itself, the energy proxy values across rows should not be interpreted as direct energy savings. The relevant observation is the response of route geometry to the rotational weight. Increasing β from 0.0 to 0.6 reduced average cumulative rotation from 4.245 rad to 4.061 rad while maintaining a similar distance, near 15.44 m. The number of turns also decreased from 7.9 to 7.3. At $\beta = 0.8$, the average number of turns decreased to 7.0, but cumulative rotation slightly increased to 4.076 rad. Therefore, $\beta = 0.60$ was retained as a balanced setting that minimizes rotation without producing the highest objective weight or excessive sensitivity.

Table 5. Sensitivity analysis for the PSO planner under different rotational weights.

β	Energy Avg.	Energy Std.	Distance Avg.	Rotation Avg.	Turns Avg.	Success (%)
0.000	16.016	0.044	15.457	4.245	7.9	100.0
0.1	16.434	0.037	15.447	4.187	7.8	100.0
0.2	16.859	0.025	15.445	4.136	7.3	100.0
0.4	17.672	0.045	15.450	4.096	7.2	100.0
0.6	18.451	0.045	15.442	4.061	7.3	100.0
0.8	19.305	0.057	15.438	4.076	7.0	100.0

3.7. External Perturbation Robustness

The external perturbation test evaluates the sensitivity of the representative planned routes to open-loop execution errors. Random waypoint perturbations with standard deviation $\sigma \in \{0.00, 0.02, 0.05, 0.08\}$ m and deterministic bias terms $(w_x, w_y) \in \{(0, 0), (0.03, 0), (0, 0.03), (0.03, 0.03)\}$ m were applied to the internal waypoints of each route. The start and goal locations were kept fixed. For each perturbation condition, 100 perturbed trajectories were generated and evaluated using the same collision, clearance, and energy metrics used in the nominal planning problem. Therefore, the success rate represents the percentage of perturbed trajectories that remained feasible after the imposed execution error. These tests should be interpreted as open-loop robustness tests, not as closed-loop tracking experiments.

Table 6 reports the zero-bias cases, namely $w_x = w_y = 0$, for three representative random perturbation levels. The rows with $\sigma = 0.00$ correspond to the unperturbed reference trajectories, which is why the perturbation terms and the energy degradation are equal to zero. In this reference condition, all methods remain feasible, and their nominal clearance margins are preserved. When random waypoint noise of $\sigma = 0.02$ m is introduced, the success rate decreases to 37% for A^* , 32% for GA, and 17% for PSO. This reduction is expected because the dense scenario contains narrow corridors, and the nominal paths operate with clearance margins close to the imposed safety buffer. Under stronger perturbations, $\sigma = 0.05$ m, feasible execution becomes rare for all methods, with success rates of 4%, 8%, and 4% for A^* , GA, and PSO, respectively.

The complete perturbation dataset, including deterministic-bias cases, is reported in Appendix A. Those additional cases are intentionally demanding because even a systematic displacement of 0.03 m can move a route segment across the safety margin in the narrowest passages. Consequently, the robustness analysis does not invalidate the global planner. Instead, it shows that the planned route should be executed with local feedback control,

online clearance monitoring, and replanning triggers when the estimated tracking error approaches the available obstacle margin.

Table 6. Selected zero-bias external perturbation results for the dense urban scenario. The reported cases use $w_x = w_y = 0$ and isolate the effect of random waypoint noise.

Algorithm	σ [m]	w_x [m]	w_y [m]	Mean Energy	Degradation [%]	Success [%]	Mean Clearance [m]
A*	0.00	0.00	0.00	18.645	0.00	100.0	0.110
GA	0.00	0.00	0.00	18.648	0.00	100.0	0.110
PSO	0.00	0.00	0.00	18.458	0.00	100.0	0.103
A*	0.02	0.00	0.00	19.234	3.16	37.0	0.095
GA	0.02	0.00	0.00	19.395	4.01	32.0	0.094
PSO	0.02	0.00	0.00	19.630	6.35	17.0	0.087
A*	0.05	0.00	0.00	20.239	8.55	4.0	0.062
GA	0.05	0.00	0.00	20.464	9.74	8.0	0.060
PSO	0.05	0.00	0.00	20.938	13.44	4.0	0.052

3.8. Practical Integration into a Robot Control Architecture

The proposed planner is intended as a global route-generation layer rather than a low-level motion controller. In a practical robot architecture, the perception module updates polygonal obstacle representations from lidar, depth cameras, or map data; the global planner computes a feasible waypoint route using A*, GA, or PSO depending on time availability; a local trajectory tracker converts waypoints into linear and angular velocity commands; and a safety supervisor monitors obstacle distance, localization confidence, and actuator limits. If the predicted clearance falls below a predefined threshold or if localization uncertainty increases, the supervisor should trigger replanning or switch to a conservative A* route. In this architecture, A* is suitable for fast replanning, while PSO is suitable when route quality is prioritized and sufficient computational time is available.

4. Conclusions

This paper presented an energy-aware route planning framework for differential drive mobile robots in dense polygonal urban environments. The formulation combines translational length, cumulative rotation, squared rotation, and localization risk in a single interpretable energy proxy, while safety is enforced using a feasibility-first dominance rule. The benchmark includes a deterministic A* baseline, a denser urban scenario with sixteen obstacles, ten independent runs, explicit ranking, rotational-weight sensitivity, and external-perturbation tests.

The results show that PSO provides the best route quality among the evaluated methods. Across ten runs, PSO obtained the lowest average energy proxy, 18.505, the shortest average distance, 15.445 m, and the lowest average cumulative rotation, 4.079 rad. These values improve the A* baseline by 0.75% in energy proxy, 0.29% in distance, and 1.37% in rotation. However, the computational cost of PSO is much higher than that of A*, with average times of 175.869 s and 5.860 s, respectively. Therefore, the results support the use of PSO for offline or supervisory route refinement, while A* remains preferable for rapid replanning.

The equal weight ranking produced a different interpretation: GA obtained the best aggregate score because it stayed close to the A* solution while slightly improving some quality metrics and avoiding the lowest clearance and longest runtime ranks assigned to PSO. This confirms that the best planner depends on the decision criterion. If route quality is dominant, PSO is preferable; if fast response and clearance are dominant, A* is preferable; and if balanced rank across the selected metrics is required, GA is competitive.

The external perturbation tests show that open-loop execution in narrow corridors is highly sensitive to deterministic bias and waypoint noise. For this reason, the proposed

global planner should be integrated with closed-loop trajectory tracking, online clearance monitoring, and replanning triggers when the estimated tracking error approaches the available obstacle margin. The main limitation of the present study is the absence of physical prototype validation. Future work will therefore address two complementary directions. First, the computational implementation will be improved through corridor-constrained optimization around the A* reference route, adaptive waypoint selection, parallel objective evaluation, and online replanning integration, with the aim of reducing computational time while preserving the route-quality improvements observed in the dense urban benchmark. Second, the energy proxy will be calibrated against measured motor current and battery discharge, and the planner will be tested on a real differential-drive platform under localization uncertainty, actuator constraints, tracking errors, and dynamic obstacles.

Author Contributions: Conceptualization, V.B.-G., C.M.H. and D.S.-V.; methodology, V.B.-G., C.M.H. and D.S.-V.; software, C.M.H.; validation, V.B.-G. and D.S.-V.; formal analysis, J.C.T. and L.F.G.-N.; investigation, V.B.-G., C.M.H., J.C.T., L.F.G.-N. and D.S.-V.; data curation, C.M.H.; writing—original draft preparation, V.B.-G., C.M.H. and D.S.-V.; writing—review and editing, V.B.-G., J.C.T., L.F.G.-N. and D.S.-V.; visualization, C.M.H.; supervision, V.B.-G. and D.S.-V.; project administration, D.S.-V. All authors have read and agreed to the published version of the manuscript.

Funding: The authors acknowledge the institutional support provided by Universidad EAFIT through research project 819728: “Diseño de una metodología para la implementación de tecnologías 4.0 en la industria automotriz”.

Data Availability Statement: The numerical results reported in the tables are included directly in the manuscript. The scripts and raw outputs are available from the corresponding author upon reasonable request.

Acknowledgments: The authors acknowledge the support of their respective institutions during the development of this research.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Complete External Perturbation Results

Table A1 reports the complete external perturbation dataset used in the robustness discussion.

Table A1. Complete external perturbation robustness dataset.

Algorithm	σ	w_x	w_y	Energy	Degradation (%)	Success (%)	Clearance
A*	0.000	0.000	0.000	18.645	0.000	100.0	0.110
A*	0.000	0.03	0.000	19.609	5.17	0.000	0.084
A*	0.000	0.000	0.03	19.667	5.49	0.000	0.099
A*	0.000	0.03	0.03	19.144	2.68	0.000	0.086
A*	0.02	0.000	0.000	19.234	3.16	37.0	0.095
A*	0.02	0.03	0.000	19.895	6.71	8.0	0.078
A*	0.02	0.000	0.03	19.377	3.93	22.0	0.088
A*	0.02	0.03	0.03	19.409	4.10	10.0	0.082
A*	0.05	0.000	0.000	20.239	8.55	4.0	0.062
A*	0.05	0.03	0.000	20.539	10.16	4.0	0.051
A*	0.05	0.000	0.03	19.918	6.83	10.0	0.063
A*	0.05	0.03	0.03	20.169	8.17	7.0	0.056
A*	0.08	0.000	0.000	20.994	12.60	0.000	0.026
A*	0.08	0.03	0.000	21.197	13.69	2.0	0.016
A*	0.08	0.000	0.03	20.720	11.13	1.0	0.029
A*	0.08	0.03	0.03	20.915	12.18	1.0	0.021
GA	0.000	0.000	0.000	18.648	0.000	100.0	0.110
GA	0.000	0.03	0.000	19.634	5.29	0.000	0.084

Table A1. Cont.

Algorithm	σ	w_x	w_y	Energy	Degradation (%)	Success (%)	Clearance
GA	0.000	0.000	0.03	19.681	5.54	0.000	0.099
GA	0.000	0.03	0.03	19.173	2.82	0.000	0.086
GA	0.02	0.000	0.000	19.395	4.01	32.0	0.094
GA	0.02	0.03	0.000	20.107	7.83	2.0	0.078
GA	0.02	0.000	0.03	19.450	4.30	28.0	0.091
GA	0.02	0.03	0.03	19.512	4.63	8.0	0.083
GA	0.05	0.000	0.000	20.464	9.74	8.0	0.060
GA	0.05	0.03	0.000	20.829	11.70	2.0	0.049
GA	0.05	0.000	0.03	20.231	8.49	12.0	0.065
GA	0.05	0.03	0.03	20.394	9.37	4.0	0.057
GA	0.08	0.000	0.000	21.495	15.27	3.0	0.024
GA	0.08	0.03	0.000	21.773	16.76	1.0	0.013
GA	0.08	0.000	0.03	21.138	13.36	6.0	0.029
GA	0.08	0.03	0.03	21.418	14.86	2.0	0.022
PSO	0.000	0.000	0.000	18.458	0.000	100.0	0.103
PSO	0.000	0.03	0.000	19.935	8.00	0.000	0.073
PSO	0.000	0.000	0.03	18.980	2.83	0.000	0.082
PSO	0.000	0.03	0.03	19.458	5.42	0.000	0.077
PSO	0.02	0.000	0.000	19.630	6.35	17.0	0.087
PSO	0.02	0.03	0.000	20.491	11.01	0.000	0.068
PSO	0.02	0.000	0.03	19.432	5.28	8.0	0.079
PSO	0.02	0.03	0.03	19.778	7.15	0.000	0.072
PSO	0.05	0.000	0.000	20.938	13.44	4.0	0.052
PSO	0.05	0.03	0.000	21.226	15.00	1.0	0.040
PSO	0.05	0.000	0.03	20.651	11.88	7.0	0.055
PSO	0.05	0.03	0.03	20.841	12.91	1.0	0.048
PSO	0.08	0.000	0.000	22.137	19.93	1.0	0.014
PSO	0.08	0.03	0.000	22.381	21.26	0.000	0.004
PSO	0.08	0.000	0.03	21.846	18.36	2.0	0.019
PSO	0.08	0.03	0.03	22.104	19.76	1.0	0.012

References

- Köhler, J.; Zhang, D.; Soloperto, R.; Carron, A.; Zeilinger, M. An MPC framework for efficient navigation of mobile robots in cluttered environments. *arXiv* **2025**, arXiv:2509.15917.
- Bao, H.; Fang, J.; Fang, M.; Zhang, J.; Zhang, Z.; Cai, H. An Enhanced A*-DWA Fusion Algorithm for Robot Navigation in Complex Environments. *Biomimetics* **2026**, *11*, 138. [CrossRef] [PubMed]
- Jiang, C.Y. The Impact of Autonomous Delivery Robots on Urban Logistics. 2025. Available online: <https://www.theseus.fi/handle/10024/882698> (accessed on 6 July 2026).
- Beaver, L.E.; Tron, R.; Cassandras, C.G. A graph-based approach to generate energy-optimal robot trajectories in polygonal environments. *IFAC-PapersOnLine* **2023**, *56*, 7000–7005. [CrossRef]
- Tang, Y.; Zakaria, M.A.; Younas, M. Path planning trends for autonomous mobile robot navigation: A review. *Sensors* **2025**, *25*, 1206. [CrossRef] [PubMed]
- Liu, L.; Johnson, P.; Reina, D.G. An adaptive and energy-aware path planning strategy (AEPPS) for ASVs with obstacle avoidance. In Proceedings of the 2025 International Conference on Machine Intelligence and Smart Innovation (ICMISI), Alexandria, Egypt, 10–12 May 2025; pp. 487–494.
- Sanin-Villa, D.; Montoya, O.D.; Gil-González, W.; Grisales-Noreña, L.F.; Perea-Moreno, A.J. Parameter estimation of a thermoelectric generator by using salps search algorithm. *Energies* **2023**, *16*, 4304. [CrossRef]
- Caballero-Mora, J.A.; Sanin-Villa, D.; Girón-Nieto, H.; Botero-Gómez, V.; Portillo-Vélez, R.d.J.; López-Romero, J.C.; Tejada, J.C. UR3 Collaborative Robot Inverse Kinematics Using Metaheuristic Optimization: A Unified Comparative and Experimental Evaluation. *Appl. Syst. Innov.* **2026**, *9*, 140. [CrossRef]
- Grisales-Noreña, L.F.; Botero-Gómez, V.; Bolaños, R.I.; Moreno-Gamboa, F.; Sanin-Villa, D. An effective parameter estimation on thermoelectric devices for power generation based on multiverse optimization algorithm. *Results Eng.* **2025**, *25*, 104408. [CrossRef]
- Katona, K.; Neamah, H.A.; Korondi, P. Obstacle avoidance and path planning methods for autonomous navigation of mobile robot. *Sensors* **2024**, *24*, 3573. [CrossRef] [PubMed]

11. Neri, F.; Palmieri, G.; Callegari, M. Non-Holonomic Mobile Manipulator Obstacle Avoidance with Adaptive Prioritization. *Robotics* **2025**, *14*, 52. [[CrossRef](#)]
12. Debnath, D.; Vanegas, F.; Sandino, J.; Hawary, A.F.; Gonzalez, F. A review of UAV path-planning algorithms and obstacle avoidance methods for remote sensing applications. *Remote Sens.* **2024**, *16*, 4019. [[CrossRef](#)]
13. Abu, N.; Bukhari, W.; Adli, M.; Ma'arif, A. Optimization of an autonomous mobile robot path planning based on improved genetic algorithms. *J. Robot. Control (JRC)* **2023**, *4*, 557–571. [[CrossRef](#)]
14. Amr, M.; Bahgat, A.; Rashad, H.; Ibrahim, A.; Youssef, A. Optimizing Navigation in Mobile Robots: Modified Particle Swarm Optimization and Genetic Algorithms for Effective Path Planning. *Algorithms* **2025**, *18*, 719. [[CrossRef](#)]
15. Altun, G.; Aydın, İ. Optimizing Unmanned Vehicle Navigation: A Hybrid PSO-GWO Algorithm for Efficient Route Planning. *Firat Univ. J. Exp. Comput. Eng.* **2025**, *4*, 100–114. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.