

Hybrid Path Planning for Industrial Autonomous Vehicles: Integrating Neural Heuristics with Energy-Constrained Search

Lijun Xiao¹, Xinwu Jian¹, Zijian Wang¹, Sisi Zhou², and Kuanching Li¹

¹ School of Computer Science and Engineering, Hunan University of Science and Technology
Xiangtan 411201, China

ljxiao@hnust.edu.cn, 1435115512@qq.com, 19555725865@163.com,
aliric@hnust.edu.cn (corresponding author)

² School of Computer Science and Engineering, Hunan Women's University
Changsha 410004, China
sisizhou@mail.hnust.edu.cn (corresponding author)

Abstract. Energy-aware path planning for industrial autonomous vehicles requires more than a geometric shortest path: terrain, payload, drivetrain efficiency, turning, and a hard task-level energy budget jointly affect feasibility. We present Neuro-PathNet-EC, a physics-aware constrained planner with learned multimodal guidance. The planner operates on heading-aware grid states and evaluates directed edge energy from friction and elevation. Exact Pareto labels enforce the hard energy budget, while a goal-conditioned ResNet-18 network fuses four raster channels with vehicle semantics and predicts an eight-channel dense cost-to-go map. The learned heuristic is used only for guidance; it is not assumed admissible or consistent. Exact Energy-A* obtains its guarantee from an analytic admissible anchor, and a FOCAL extension provides the bound $C \leq wC^*$. On a layout-disjoint synthetic industrial dataset, five training seeds achieve a test relative heuristic MAE of 0.0289 ± 0.0014 and Spearman correlation of 0.9963 ± 0.0004 . Across 512 paired binding-budget queries per seed (2,560 seed–query runs), full multimodal neural-constrained search returns the exact constrained objective in every run while reducing expanded labels by 10.73% relative to Energy-A*. The semantic branch also reduces labels relative to an otherwise matched image-only model. FOCAL satisfies its declared bound in all evaluated cases, although its current correctness-first OPEN management does not improve runtime. The evaluation is based on synthetic layouts and does not constitute public-benchmark or real-vehicle validation.

Keywords: Industrial Autonomous Vehicles, Energy-Constrained Path Planning, Neural Heuristic, Resource-Constrained Search, Multimodal Fusion, Bounded-Suboptimal Search.

1. Introduction

Industrial autonomous vehicles, including automated guided vehicles and unmanned transport platforms, repeatedly navigate warehouses, workshops, and structured outdoor facilities. Recent reviews show that graph search remains central to mobile-robot navigation

even as sampling, optimization, and learning-based planners diversify [1,2,3,4,5,6]. Classical graph search provides transparent and reliable path construction when the state and edge cost are well specified [7,8]. In industrial operation, however, distance alone is insufficient. Payload, rolling resistance, elevation change, drivetrain efficiency, and steering can materially change the energy required by a route. Moreover, a route with a favorable weighted objective may still be unacceptable if it exceeds the energy assigned to the current task.

Learning-based planners can extract guidance from raster maps, graphs, and operating context [9,10,11,12,13]. Yet a learned heuristic creates a theoretical difficulty: low regression error does not imply admissibility or consistency. Treating such a heuristic as if it preserved A* optimality confounds empirical guidance quality with a search guarantee. Hard energy budgets create a second difficulty. The problem is resource-constrained, so keeping only one spatial label can discard a higher-objective but lower-energy prefix that is necessary for a feasible optimum.

We therefore position Neuro-PathNet-EC as *physics-aware energy-constrained search with learned multimodal guidance*. The physical model and hard budget define feasibility and path cost. Search retains nondominated objective–energy labels at each heading-aware state. A dense goal-conditioned network predicts unconstrained cost-to-go for all cells and incoming headings in one pass, but its output is not granted theoretical admissibility. Guarantees instead come from an analytic anchor used by exact Energy-A* or by a bounded FOCAL rule.

The contributions are as follows.

- We formulate directed, heading-aware energy-constrained planning with elevation-derived slope, nonnegative traction energy, turning energy, a dimensionless distance–energy objective, and Pareto labels that enforce a hard task-level budget.
- We design a goal-conditioned multimodal network that maps a four-channel numerical raster and physically relevant vehicle semantics to a dense tensor $\hat{H} \in \mathbb{R}^{8 \times H \times W}$. Exact directed reverse shortest paths provide the supervised target.
- We separate empirical neural guidance from guarantees: Energy-A* is exact under an analytic anchor, Safe-Neural is an exact tie-break variant, and FOCAL-Neural is a w -bounded extension whose bound does not depend on learned admissibility.
- We provide a reproducible five-seed evaluation on layout-disjoint synthetic industrial maps. The evaluation includes exact-solver cross-checks, ablation, paired statistics, physical response, binding budgets, sensitivity, robustness stress tests, and censored scalability measurements.

The remainder of the paper reviews related work, formulates the problem, presents the method and theoretical properties, describes the experimental protocol, discusses results and limitations, and concludes with evidence-supported future work.

2. Related Work

2.1. Classical and resource-constrained search

Dijkstra and A* are foundational shortest-path methods [7,8]. Weighted, anytime, and multi-heuristic extensions trade search effort against bounded solution quality [14,15,16].

Their behavior depends on the state representation, edge objective, and heuristic. Recent mobile-robot surveys and an RRT-focused review provide broader comparisons with sampling, evolutionary, and reactive planners [1,3,4,17].

A hard consumable-resource constraint changes the search problem fundamentally. Modern exact algorithms use Pareto labeling and indexing [18,19]. Other accelerations include bidirectional dynamic programming, backtracking, and decomposition [20,21,22]. Related formulations address stochastic constraints, constrained tours, and time resources [23,24,25]. These results support retaining multiple nondominated labels: under an energy budget, a higher-objective but lower-energy prefix can be essential to the feasible optimum. Our formulation specializes this principle to a heading-aware physical grid and never treats the budget as a post hoc filter.

2.2. Learned planning guidance

Neural A* differentiates through grid search to learn spatial guidance maps [9], while GraphMP and PHIL learn graph-structured search guidance [10,11]. Recent work further studies local learned heuristics and losses that optimize search ranking rather than only cost-to-go regression [12,13]. Multimodal perception-aware path planning also demonstrates the value of combining heterogeneous information rather than relying on a single raster modality [26]. Our encoder uses residual spatial features [27], but differs from these studies by predicting all eight incoming-heading cost-to-go maps and coupling them to resource-constrained labels.

Such guidance may reduce exploration, but prediction accuracy alone does not establish A* admissibility. Learned policies and learned heuristics have therefore been combined with FOCAL or bounded-suboptimal search [28,29]. Contemporary bounded-search studies analyze alternative queue rules, error handling, bidirectional bounds, and tunable suboptimality [30,31,32,33,34]. In this work, the learned map is evaluated empirically through absolute error, rank correlation, admissible fraction, directed consistency, and expansion reduction. Optimal or bounded-suboptimal claims derive only from the anchor search rules.

2.3. Energy-aware industrial navigation

Recent AGV research models transport energy directly and constructs route-efficiency benchmarks [35,36]. Related formulations jointly consider productivity, conflict avoidance, multiple loads, or heterogeneous fleets [37,38,39]. Energy-aware A*-type planning has also been developed for Ackermann-steered vehicles, while industrial studies combine global and local planning or optimize pickup-and-delivery routes [40,41,42,43]. These works establish that industrial route energy depends on vehicle state and route geometry, but they do not provide the same dense heading-conditioned learned guidance over a hard Pareto-label budget.

Terrain-aware planning further motivates our directed elevation representation. Recent studies consider resource-aware agricultural autonomy, complex-terrain optimization, and sandy-terrain traversability [44,45,46]. Others address mountainous planning and control, three-dimensional ground-vehicle energy, and large-scale rough terrain [47,48,49]. In parallel, multi-objective planners balance distance, safety, smoothness, or coverage

using graph search and population-based optimization [50,51,52,53,54]. Our distinction is the combination of (i) directed elevation-based energy, (ii) a hard task budget, (iii) heading-aware Pareto labels, (iv) dense multimodal guidance, and (v) an anchor/FOCAL framework that states precisely which component carries each guarantee. We do not claim quantitative comparison with external learned, evolutionary, or sampling-based planners that were not reproduced in the present evaluation.

3. Problem Formulation

3.1. Heading-aware state graph

Let an $H \times W$ grid contain binary occupancy O , friction μ , and elevation z . Motion uses an eight-connected neighborhood. A search state is

$$s = (y, x, q), \quad q \in \{0, \dots, 7\}, \quad (1)$$

where q is the incoming heading. A sample specifies a start state (y_s, x_s, q_s) and a goal cell $G = (y_g, x_g)$; goal orientation is unconstrained. Diagonal corner cutting through occupied cells is forbidden.

For a forward transition $i \rightarrow j$, horizontal displacement is d_{xy} and

$$\Delta z_{ij} = z_j - z_i, \quad \phi_{ij} = \text{atan2}(\Delta z_{ij}, d_{xy}). \quad (2)$$

Consequently, the grade is directed: $\phi_{ji} = -\phi_{ij}$.

3.2. Physical edge energy

Let m_v and m_p denote vehicle and payload mass, $m = m_v + m_p$, η drivetrain efficiency, v speed, a acceleration, and d edge distance. The force components are

$$F_{\text{roll}} = \mu_{ij} mg \cos \phi_{ij}, \quad (3)$$

$$F_{\text{grade}} = mg \sin \phi_{ij}, \quad (4)$$

$$F_{\text{acc}} = m_{\text{eff}} a, \quad (5)$$

$$F_{\text{aero}} = \frac{1}{2} \rho C_d A v^2. \quad (6)$$

The physical work and nonnegative traction energy are

$$E_{\text{phys}} = \frac{d}{\eta} (F_{\text{roll}} + F_{\text{grade}} + F_{\text{acc}} + F_{\text{aero}}), \quad (7)$$

$$E_{\text{traction}} = \max(0, E_{\text{phys}}). \quad (8)$$

No regenerative braking is modeled; a downhill edge therefore receives no negative energy credit. Aerodynamic drag remains in the formulation but is disabled in the low-speed experiments. If q_i and q_j are the incoming and outgoing headings, the wrapped angle satisfies $0 \leq |\Delta\psi| \leq \pi$ and

$$E_{\text{turn}} = k_{\text{turn}} |\Delta\psi|, \quad \Delta E_{ij} = E_{\text{traction}} + E_{\text{turn}}. \quad (9)$$

Here k_{turn} has units J/rad. Its experimental value is a nominal synthetic coefficient, not a hardware-calibrated parameter.

3.3. Objective and hard budget

Adding metres to joules is dimensionally invalid. We therefore define

$$c_{ij} = \alpha \frac{d_{ij}}{d_{\text{ref}}} + \beta \frac{\Delta E_{ij}}{E_{\text{ref}}}, \quad \alpha, \beta \geq 0, \quad \alpha + \beta = 1. \quad (10)$$

For path P , $C(P) = \sum_{(i,j) \in P} c_{ij}$ and $E(P) = \sum \Delta E_{ij}$. The planning problem is

$$P^* = \arg \min_P C(P) \quad \text{s.t. } P : s_s \rightsquigarrow G, \quad E(P) \leq E_{\text{task,max}}. \quad (11)$$

The task budget is distinguished from total battery availability. If nominal battery energy is given in Wh, physical availability may be written as $3600 E_{\text{nominal}} \text{SOC SOH } r_{\text{reserve}}$ joules, while $E_{\text{task,max}}$ allocates no more than this amount to the current route.

3.4. Pareto labels

A label is

$$\ell = (y, x, q, g, E, \text{parent}), \quad (12)$$

where g is accumulated objective and E is accumulated energy. A successor with $E > E_{\text{task,max}}$ is rejected. For labels at the same (y, x, q) , ℓ_a dominates ℓ_b iff

$$g_a \leq g_b, \quad E_a \leq E_b, \quad \text{and at least one inequality is strict.} \quad (13)$$

Dominance is never applied across headings because future turning cost differs.

4. Neuro-PathNet-EC

4.1. Goal-conditioned multimodal input

The spatial input has four numerical channels: binary occupancy, friction, elevation, and a Gaussian goal map. Occupancy and goal encoding remain in $[0, 1]$; friction and elevation are standardized using training-split statistics only. The semantic vector is

$$\mathbf{s}_{\text{sem}} = [m_v, m_p, \eta, v]. \quad (14)$$

SOC and SOH do not enter this unconstrained cost-to-go network because they do not directly change mechanical edge energy in the adopted model; they determine resource availability. This separation avoids assigning unsupported physical influence to semantic variables.

4.2. Dense eight-heading architecture

A ResNet-18-style encoder receives $\mathbf{X} \in \mathbb{R}^{B \times 4 \times 128 \times 128}$. Its first convolution has four input channels and weights are initialized randomly because the raster is not natural RGB imagery. The bottleneck is $\mathbf{F} \in \mathbb{R}^{B \times 512 \times 4 \times 4}$. A semantic MLP $4 \rightarrow 128 \rightarrow 128 \rightarrow 64$ produces $\mathbf{e}_{\text{sem}}$. Spatial broadcasting and concatenation yield $B \times 576 \times 4 \times 4$, followed by a 1×1 projection to 512 channels.




Fig. 1. The proposed Neuro-PathNet-EC framework. The neural network predicts a dense heading-conditioned cost-to-go map. Exact or bounded properties arise from the analytic anchor and label-search rule, not from assuming that h_{NN} is admissible.

No global average pooling is used. A bilinear upsampling decoder with convolutional widths $512 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow 32 \rightarrow 8$ restores the grid resolution. Softplus produces nonnegative normalized predictions

$$\hat{H}_{\text{norm}} \in \mathbb{R}^{B \times 8 \times 128 \times 128}. \quad (15)$$

The planner uses $\hat{H} = H_{\text{scale}} \hat{H}_{\text{norm}}$ and imposes the known boundary $\hat{H}[q, y_g, x_g] = 0$ for every heading.

4.3. Exact dense supervision

The target

$$H^*[q, y, x] = \min_{P: (y, x, q) \rightsquigarrow G} C(P) \quad (16)$$

is the exact unconstrained objective-to-go; hard resource feasibility remains the planner's responsibility. Directed reverse Dijkstra starts from all eight goal headings. When reverse search is at current state v and enumerates predecessor u , it relaxes with the *forward* edge cost $c(u, v)$. Random forward exact solves verified this construction.

Only finite reachable free states contribute to the masked SmoothL1 loss:

$$\mathcal{L} = \frac{1}{|\mathcal{V}|} \sum_{(q, y, x) \in \mathcal{V}} \text{SmoothL1} \left(\hat{H}_{\text{norm}}, H^* / H_{\text{scale}} \right). \quad (17)$$

H_{scale} is the training-split 95th percentile of finite positive H^* .

4.4. Search variants

The analytic anchor is

$$h_{\text{anc}}(s) = \alpha \frac{\|p(s) - G\|_2}{d_{\text{ref}}} + \beta \frac{\|p(s) - G\|_2 e_{\min}}{E_{\text{ref}}}. \quad (18)$$

In our no-regeneration terrain family, a universally safe $e_{\min}$ may be zero; we prefer a weak valid bound to an unsupported positive one.

All constrained variants retain identical feasibility and Pareto dominance. They differ only in queue ordering:

- **Energy-A***: priority $g + h_{\text{anc}}$; exact constrained baseline.
- **Full Neural-Constrained**: priority $g + h_{NN}$; learned, potentially inadmissible guidance with stale-label handling. Exactness is evaluated empirically, not claimed theoretically.
- **Safe-Neural**: primary key $g + h_{\text{anc}}$, secondary key h_{NN} ; exact, but neural influence occurs only at anchor ties.
- **FOCAL-Neural**: let $f_{\min} = \min_{\ell \in \text{OPEN}} (g + h_{\text{anc}})$. Choose from $\{\ell : f_{\text{anc}}(\ell) \leq wf_{\min}\}$ by (h_{NN}, E, tie) .

5. Theoretical Properties

5.1. What the learned heuristic does not guarantee

Strong MAE or correlation does not imply $h_{NN}(s) \leq H^*(s)$ or $h_{NN}(s) \leq c(s, s') + h_{NN}(s')$. We measure both properties, but do not assume them. The boundary $h(G) = 0$ is known exactly and is enforced without using unknown path cost.

5.2. Anchor validity under a resource constraint

Let $\mathcal{P}_{\text{con}}(s)$ be feasible continuations satisfying the remaining budget and $\mathcal{P}_{\text{unc}}(s)$ all obstacle-feasible continuations. Since $\mathcal{P}_{\text{con}}(s) \subseteq \mathcal{P}_{\text{unc}}(s)$, the optimal unconstrained continuation cost is no greater than the constrained continuation cost. Therefore any lower bound on the unconstrained cost, including Eq. (18), is also a lower bound for the constrained problem. A zero turning lower bound is valid because turn energy is nonnegative.

5.3. FOCAL bounded-suboptimality

Proposition 1. Consider the finite heading-grid feasible label graph generated under a hard energy budget. Assume nonnegative edge objective, safe Pareto dominance within identical heading states, an admissible anchor, correct OPEN maintenance, and termination when a selected FOCAL label is a goal. Then FOCAL-Neural with $w \geq 1$ returns cost $C_{\text{ret}} \leq wC^*$ whenever a feasible solution exists.

Proof. Before an optimal goal label is selected, OPEN contains a surviving prefix label on an optimal feasible path. Admissibility implies its anchor key is at most C^* , hence $f_{\min} \leq C^*$. Every selected FOCAL label satisfies $f_{\text{anc}} \leq wf_{\min}$. At a selected goal, $h_{\text{anc}} = 0$, so $C_{\text{ret}} = g \leq wf_{\min} \leq wC^*$. Safe Pareto dominance cannot remove every optimal continuation because a dominating label has no greater objective and energy at the same heading state. The finite grid and finite set of generated simple nondominated path labels under positive movement cost ensure termination. $\square$

The learned score affects which eligible FOCAL label is chosen, but not the bound. For $w = 1$, FOCAL retains the optimal anchor threshold, with learned ordering only among equal-minimum anchor labels.

Table 1. Frozen implementation and training settings.

Item	Setting
Hardware	NVIDIA L4, 22.03 GiB VRAM
Software	Python 3.12.13; PyTorch 2.11.0+cu128
CUDA / cuDNN	12.8 / 9.1.9
Input / output	$4 \times 128 \times 128 / 8 \times 128 \times 128$
Full parameters	13,070,984
Optimizer	AdamW, lr 3×10^{-4} , wd 10^{-2}
Scheduler	ReduceLRonPlateau, factor 0.5, patience 3
Batch / epochs	16 / maximum 50
Early stopping	patience 8 on validation loss
Seeds	0, 1, 2, 3, 4
Split samples	2,048 / 256 / 512
Split layouts	128 / 16 / 32, disjoint
α, β	0.3, 0.7
E_{ref}	543.4587 J, training edge median
H_{scale}	128.8968, training H^* p95
k_{turn}	25 J/rad, nominal synthetic

6. Experimental Setup

6.1. Implementation details

All neural training and dense inference use the GPU. Search is an ordinary CPU implementation; no CUDA search kernel or GPU-parallel expansion is used. Checkpoints are selected only by minimum validation normalized SmoothL1.

6.2. Synthetic industrial data

The generator produces structured warehouse-like rack layouts, factory-like aisles, and industrial-park-like open/structured layouts. Each numerical sample stores occupancy, spatially structured friction, continuous elevation, start, goal, start heading, vehicle state, layout metadata, and exact dense H^* . Layout identities are disjoint across train, validation, and test. Continuous normalization uses training statistics only.

The binding benchmark contains 512 three-route queries (Fig. 2). Route A is short and energy intensive, Route B is balanced, and Route C is longer and lower energy. Every accepted query verifies $d_A < d_B < d_C$, $E_A > E_B > E_C$, and $E_{\text{obj}} > E_{\text{min}}$. For tightness ρ ,

$$E_{\text{max}}(\rho) = E_{\text{min}} + \rho(E_{\text{obj}} - E_{\text{min}}), \quad (19)$$

with $\rho \in \{0.05, 0.25, 0.50, 0.75, 0.95\}$.

6.3. Baselines, variants, and metrics

We evaluate Dijkstra-style exact label setting, Classical A* for geometric context, Energy-A*, Image-only and Full Neural-Constrained search, Safe-Neural, FOCAL with $w = 1.05$




Fig. 2. Binding-budget benchmark. Left: a representative three-route layout. Right: all 512 generated queries satisfy the intended distance–energy ordering. These are synthetic stress-test maps, not real-site telemetry.

Table 2. Main comparison on 512 paired binding queries. Values are mean $\pm$ standard deviation of five seed-level means. Objective and energy are deterministic for exact methods.

Method	Objective	Energy (J)	Labels	Runtime (ms)	Property
Dijkstra	48.4649	18171.3	412.50	6.76 $\pm$ 0.95	Exact
Classical A*	59.0000	21815.5	60.00	0.52 $\pm$ 0.07	Infeasible
Energy-A*	48.4649	18171.3	406.53	6.94 $\pm$ 0.91	Exact anchor
Image-only Neural	48.4649	18171.3	383.69 $\pm$ 22.16	6.36 $\pm$ 0.94	Empirical
Full Multimodal	48.4649	18171.3	362.92 $\pm$ 21.09	6.14 $\pm$ 0.99	Empirical
Safe-Neural	48.4649	18171.3	406.53	7.06 $\pm$ 1.05	Exact tie-break
FOCAL $w = 1.05$	48.5597 $\pm$ 0.0650	18061.5 $\pm$ 71.2	406.74 $\pm$ 4.29	22.07 $\pm$ 3.00	$C \leq 1.05C^*$
FOCAL $w = 1.10$	48.5149 $\pm$ 0.0504	18127.2 $\pm$ 37.5	403.90 $\pm$ 6.02	23.00 $\pm$ 3.27	$C \leq 1.10C^*$
Oracle- H^*	48.4649	18171.3	254.28	4.45 $\pm$ 0.61	Diagnostic

and 1.10, and diagnostic Oracle- H^* and spatially shuffled- H controls. Classical A* does not enforce the energy budget and is not treated as a feasible constrained competitor.

Heuristic metrics are MAE, RMSE, relative MAE, Pearson, Spearman, R^2 , empirical admissible fraction, and directed consistency violation. Planner metrics are objective, energy, path length, expanded labels, runtime, success, and gap to exact constrained search. Statistical comparisons use the paired query as the unit: each method is first averaged across five seeds per query, followed by two-sided Wilcoxon signed-rank, paired Cohen’s d , 10,000-sample query bootstrap confidence intervals, and Holm correction.


7. Results and Discussion

7.1. Main comparison

Full Multimodal Neural-Constrained search returns the exact constrained objective and energy in all 2,560 seed–query runs. Relative to Energy-A*, it reduces labels by 10.73%

Table 3. Ablation of guidance mechanisms.

Variant	Semantic Neural FOCAL			Labels
Energy-A*	–	–	–	406.53
Image-only	–	✓	–	383.69 ± 22.16
Full	✓	✓	–	362.92 ± 21.09
Safe-Neural	✓	tie	–	406.53
FOCAL 1.05	✓	✓	✓	406.74 ± 4.29
FOCAL 1.10	✓	✓	✓	403.90 ± 6.02

**Fig. 3.** Guidance diagnostic. Spatially meaningful multimodal predictions reduce expansions relative to zero and shuffled controls; Oracle- H^* is a diagnostic upper reference.

and mean runtime by 0.80 ms. This exact agreement is an experimental finding, not a consequence of learned admissibility. Classical A* finds the shortest geometric corridor but violates the binding budget in all 512 queries.

7.2. Ablation and guidance contribution

The semantic branch reduces labels by a further 5.41% relative to Image-only. Safe-Neural matches Energy-A* because learned guidance acts only at exact anchor ties. Oracle guidance shows that additional expansion reduction remains possible, while shuffling destroys most of the learned benefit (Fig. 3). The empirical guidance recovery ratio is 0.313 relative to Zero/Dijkstra and Oracle endpoints.

7.3. Heuristic quality

The dense map has low relative error and strong spatial ordering (Fig. 5). Nevertheless, only about 47% of predictions are empirically admissible and about 14.7% of directed




Fig. 4. Five-seed ablation. Full multimodal guidance reduces labels and modestly reduces runtime, whereas correctness-first FOCAL incurs substantial OPEN-management overhead.

Table 4. Full-model heuristic properties over five seeds.

Metric	Mean $\pm$ SD
MAE	1.9370 ± 0.0917
RMSE	3.1145 ± 0.1716
Relative MAE	0.02888 ± 0.00137
Pearson	0.99592 ± 0.00044
Spearman	0.99628 ± 0.00042
R^2	0.99182 ± 0.00090
Admissible fraction	0.47034 ± 0.02663
Consistency violation rate	0.14681 ± 0.00829

edge tests violate consistency. These measurements directly motivate the analytic-anchor separation.

7.4. Physical semantic response

At fixed map, start, goal, and heading, increasing payload from 0 to 800 kg raises exact start cost-to-go from 82.85 to 115.49. Full prediction follows from 82.48 to 114.13, while Image-only remains 94.33 because it cannot observe payload. Similarly, increasing η from 0.75 to 0.90 lowers exact cost from 106.25 to 94.19 and Full prediction from 105.86 to 93.03; Image-only again remains unchanged. This controlled response provides mechanism-level evidence for the semantic branch.

7.5. Statistical analysis

The label reduction is statistically and practically meaningful. Conversely, no objective or energy difference exists between Full and Energy-A*, and no path-quality improvement is claimed. FOCAL is significantly slower in this implementation.




Fig. 5. Representative exact and predicted cost-to-go maps for the sample’s start heading, with a subsampled cell-wise correlation plot. Aggregate five-seed statistics are reported in Table 4.




Fig. 6. Controlled payload and drivetrain-efficiency response. Multimodal predictions follow the exact physical target, whereas the matched image-only model cannot respond to hidden vehicle semantics.

7.6. Binding budgets and FOCAL bounds

For each ρ , Energy-A* and exact label setting agree in every one of 200 checked queries. Binding changes the path in 100% of cases. Mean energy reduction relative to the unconstrained objective optimum decreases from 4010.7 J at $\rho = 0.05$ to 1694.5 J at $\rho = 0.95$. Across 5,000 FOCAL cases (200 queries, five budgets, five w values), the bound violation count is zero. Maximum gaps for $w = 1.02, 1.05, 1.10, 1.20$ are 0.00246, 0.04855, 0.04453, and 0.06367. Runtime, however, increases with w because the current implementation reconstructs the eligible FOCAL subset by scanning OPEN.

7.7. Parameter sensitivity

At nominal $\alpha = 0.3$, mean normalized distance and energy contributions are 24.804 and 23.778, respectively; $\beta = 0.7$ therefore does not make energy dominate by orders of magnitude. Under the fixed tight budget, decreasing efficiency by 20% or increasing friction by 20% makes all 100 cases infeasible. Increasing payload by 20% changes 26% of paths

Table 5. Selected paired statistical comparisons over 512 query units after averaging training seeds. CI is a query-bootstrap 95% interval; p_{Holm} is Holm-adjusted.

Method A	Method B	Metric	Mean diff.	95% CI	p_{Holm}	Paired d
Full	Energy-A*	Objective	0	[0, 0]	1.0	0
Full	Energy-A*	Labels	-43.618	[-48.184, -39.178]	2.15×10^{-84}	-0.833
Full	Energy-A*	Runtime (ms)	-0.799	[-0.975, -0.619]	4.25×10^{-49}	-0.390
Full	Image-only	Labels	-20.773	[-22.491, -19.104]	2.15×10^{-84}	-1.068
FOCAL 1.05	Energy-A*	Runtime (ms)	15.129	[14.952, 15.316]	2.15×10^{-84}	7.250

Table 6. Objective-weight sensitivity on 100 paired queries.

α	β	Objective	Energy (J)	Length (m)
0.1	0.9	38.4839	17984.8	87.00
0.3	0.7	48.5823	18460.8	82.68
0.5	0.5	58.3245	18460.8	82.68
0.7	0.3	68.0667	18460.8	82.68
0.9	0.1	77.8089	18460.8	82.68

and makes 26% of nominal paths exceed the budget; replanning returns no violating successful path. The nominal turn coefficient has a smaller effect, changing 1–5% of paths under $\pm 20\%$ perturbation.

7.8. Robustness stress tests

These deliberately severe tests expose failure modes rather than establish industrial robustness. A 10% budget overestimate preserves planner success but violates the nominal true budget in 74% of cases, illustrating why uncertainty-aware budgeting is important. Occupancy and localization perturbations are synthetic and are not calibrated sensor-error distributions.

7.9. Scalability

Energy-A* mean runtime is 2.43 s at 128 and 9.97 s at 256. At 512 it completes two of four settings. Higher obstacle density can reduce runtime by shrinking the reachable label set. FOCAL’s correctness-first linear OPEN scan is the dominant limitation; these results do not support a large-scale acceleration claim.

8. Limitations

The present evidence has clear boundaries. First, all maps are generated synthetic industrial layouts; there is no Gazebo experiment, public benchmark, industrial telemetry, or real-vehicle deployment. Layout-disjoint testing measures generalization within this generator only. Second, the learned heuristic is neither theoretically admissible nor consistent. Its empirical exact-objective behavior does not replace the analytic-anchor proof. Third,




Fig. 7. Energy–distance response to the frozen α grid ($\beta = 1 - \alpha$). The discrete Route-C choices produce plateaus rather than a smooth continuum.

Table 7. Synthetic robustness stress tests (50 seed-0 queries).

Condition	Levels	Success rates
Occupancy flips (%)	0/1/3/5/10	1/.48/.12/.02/0
Localization shift (m)	0/.05/.10/.20	1/1/1/.52
Budget error (%)	-10/0/+10	.48/1/1
SOH (%)	100/90/80/70	1/.48/0/0

k_{turn} is nominal and has not been calibrated against hardware. Fourth, FOCAL uses a correctness-first OPEN scan and is not runtime optimized. Fifth, robustness perturbations are synthetic stress tests rather than calibrated sensing or battery uncertainty. Sixth, scalability uses one generated map per size–density cell and includes censored runs. Seventh, external RRT*, CNN-A*, or GNN-based planners were not reproduced; no external state-of-the-art claim is made. Finally, GPU use is limited to neural training and inference, not parallel search.

9. Conclusion

We presented Neuro-PathNet-EC as physics-aware energy-constrained search with multi-modal neural guidance. Directed elevation, friction, payload, efficiency, and turning define a dimensionless objective and hard resource constraint over heading-aware Pareto labels. A goal-conditioned dense network predicts eight heading-specific cost-to-go maps, while exact or bounded properties are supplied separately by analytic-anchor Energy-A*, Safe-Neural, and FOCAL.

Five-seed experiments show strong dense regression quality and exact constrained objective agreement for Full Neural-Constrained search across 2,560 runs. Full guidance

Table 8. Completion within a 30 s per-method threshold. Four density settings are evaluated at each size.

Grid Energy-A* FOCAL 1.05 Peak RSS (GiB)			
128	4/4	1/4	1.04
256	4/4	0/4	1.19
512	2/4	0/4	1.39




Fig. 8. Censored scalability results. Runtime values include the 30 s censor time. The current FOCAL data structure prioritizes correctness and does not scale competitively.

reduces expanded labels by 10.73% relative to Energy-A* and improves on a matched image-only network; controlled payload and efficiency experiments confirm that the semantic branch responds to physical variables. At the same time, learned admissibility and consistency do not hold, and FOCAL's present data structure is slower and poorly scalable. These negative results are part of the method's current scope.

Future work should validate on public benchmarks and real vehicles, calibrate turning and drivetrain parameters from measurements, replace linear FOCAL scans with efficient OPEN/FOCAL structures, and model uncertainty explicitly. Remaining-budget-aware guidance is also a natural extension because two Pareto labels at the same heading state may have different useful continuations.

Acknowledgment. This work is supported by the Natural Science Foundation of Hunan Province under Grant 2025JJ50399.

References

1. José Ricardo Sánchez-Ibáñez, Carlos J. Pérez-del Pulgar, and Alfonso García-Cerezo. Path planning for autonomous mobile robots: A review. *Sensors*, 21(23):7898, Nov 2021.
2. Shiwei Lin, Ang Liu, Jianguo Wang, and Xiaoying Kong. A review of path-planning approaches for multiple mobile robots. *Machines*, 10(9):773, Sept 2022.
3. Jaafar Ahmed Abdulsahab and Dheyaa Jasim Kadhim. Classical and heuristic approaches for mobile robot path planning: A survey. *Robotics*, 12(4):93, June 2023.

4. Liwei Yang, Ping Li, Song Qian, He Quan, Jinchao Miao, Mengqi Liu, Yanpei Hu, and Erexidin Memetimin. Path planning technique for mobile robots: A review. *Machines*, 11(10):980, Oct 2023.
5. Kornél Katona, Husam A. Neamah, and Péter Korondi. Obstacle avoidance and path planning methods for autonomous navigation of mobile robot. *Sensors*, 24(11):3573, June 2024.
6. Yuexia Tang, Muhammad Aizzat Zakaria, and Maryam Younas. Path planning trends for autonomous mobile robot navigation: A review. *Sensors*, 25(4):1206, Feb 2025.
7. Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
8. Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
9. Ryo Yonetani, Tatsunori Tanai, Mohammadamin Barekatin, Mai Nishimura, and Asako Kanezaki. Path planning using neural A* search. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 12029–12039. PMLR, 2021.
10. Xiao Zang, Miao Yin, Jinqi Xiao, Saman Zonouz, and Bo Yuan. GraphMP: Graph neural network-based motion planning with efficient graph search. In *Advances in Neural Information Processing Systems*, volume 36, pages 3131–3142, 2023.
11. Michal Pándy, Weikang Qiu, Gabriele Corso, Petar Veličković, Zhitao Ying, Jure Leskovec, and Pietro Lio. Learning graph search heuristics. In *Proceedings of the First Learning on Graphs Conference*, volume 198 of *Proceedings of Machine Learning Research*, pages 10:1–10:13. PMLR, 2022.
12. Rishi Veerapaneni, Muhammad Suhail Saleem, and Maxim Likhachev. Learning local heuristics for search-based navigation planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, pages 406–414, 2023.
13. Leah Chrestien, Stefan Edelkamp, Antonin Komenda, and Tomáš Pevný. Optimize planning heuristics to rank, not to estimate cost-to-goal. *Advances in Neural Information Processing Systems*, 36, 2023.
14. Judea Pearl and Jin H. Kim. Studies in semi-admissible heuristics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-4(4):392–399, 1982.
15. Maxim Likhachev, Geoffrey J. Gordon, and Sebastian Thrun. ARA*: Anytime A* with provable bounds on sub-optimality. In *Advances in Neural Information Processing Systems*, volume 16, 2004.
16. Sandip Aine, Siddharth Swaminathan, Venkatraman Narayanan, Victor Hwang, and Maxim Likhachev. Multi-heuristic A*. *The International Journal of Robotics Research*, 35(1–3):224–243, 2016.
17. Tong Xu. Recent advances in rapidly-exploring random tree: A review. *Heliyon*, 10(11):e32451, June 2024.
18. Saman Ahmadi, Guido Tack, Daniel D. Harabor, and Philip Kilby. A fast exact algorithm for the resource constrained shortest path problem. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12217–12224, May 2021.
19. Libin Wang and Raymond Chi-Wing Wong. QHl: A fast algorithm for exact constrained shortest path search on road networks. *Proceedings of the ACM on Management of Data*, 1(2):1–25, June 2023.
20. Matteo Salani, Saverio Basso, and Giovanni Righini. Enhanced bi-directional dynamic programming algorithm for the resource constrained shortest path problem. *Transportation Research Procedia*, 78:361–368, 2024.
21. Bryce Ford and Mrinal Kumar. Analysis of backtracking a* for resource constrained shortest path problems. In *2024 American Control Conference (ACC)*, pages 3545–3550. IEEE, July 2024.

22. Xuefeng Jin and Shunzheng Yu. A two-stage method for doubly resource constrained elementary shortest path problems. *Knowledge-Based Systems*, 293:111661, June 2024.
23. Sungkweon Hong and Brian C. Williams. An anytime algorithm for constrained stochastic shortest path problems with deterministic policies. *Artificial Intelligence*, 316:103846, Mar 2023.
24. Sébastien Martin, Youcef Magnouche, Corentin Juvigny, and Jérémie Leguay. Constrained shortest path tour problem: Branch-and-price algorithm. *Computers & Operations Research*, 144:105819, Aug 2022.
25. Pan Liu and Wulan Huang. A graph algorithm for the time constrained shortest path. *Connection Science*, 34(1):1500–1518, May 2022.
26. Zhenyu Li, Aiguo Zhou, Jiakun Pu, and Jiangyang Yu. Multi-modal neural feature fusion for automatic driving through perception-aware path planning. *IEEE Access*, 9:142782–142794, 2021.
27. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
28. Matias Greco. Exploiting learned policies and learned heuristics in bounded-suboptimal search. *Proceedings of the International Symposium on Combinatorial Search*, 12(1):219–221, July 2021.
29. Matias Greco and Jorge A. Baier. Avoiding errors in learned heuristics in bounded-suboptimal search. *Proceedings of the International Symposium on Combinatorial Search*, 15(1):336–338, July 2022.
30. Maximilian Fickert, Tianyi Gu, and Wheeler Ruml. New results in bounded-suboptimal search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9):10166–10173, June 2022.
31. Ryo Kuroiwa and J. Christopher Beck. Biased exploration for satisficing heuristic search. *Proceedings of the International Conference on Automated Planning and Scheduling*, 32:213–221, June 2022.
32. Dor Atzmon, Shahaf S. Shperberg, Netanel Sabah, Ariel Felner, and Nathan R. Sturtevant. W-restrained bidirectional bounded-suboptimal heuristic search. *Proceedings of the International Conference on Automated Planning and Scheduling*, 33(1):26–30, July 2023.
33. Stephen Wissow, Fanhao Yu, and Wheeler Ruml. Tunable suboptimal heuristic search. *Proceedings of the International Symposium on Combinatorial Search*, 17:170–178, June 2024.
34. Jaemin Lim and Panagiotis Tsiotras. Cbs-budget (cbsb): A complete and bounded suboptimal search for multi-agent path finding. *Artificial Intelligence*, 346:104349, Sept 2025.
35. Zhongwei Zhang, Lihui Wu, Wenqiang Zhang, Tao Peng, and Jun Zheng. Energy-efficient path planning for a single-load automated guided vehicle in a manufacturing workshop. *Computers & Industrial Engineering*, 158:107397, Aug 2021.
36. Luoke Hu, Xiaoliang Zhao, Weipeng Liu, Wei Cai, Kangkang Xu, and Zhongwei Zhang. Energy benchmark for energy-efficient path planning of the automated guided vehicle. *Science of The Total Environment*, 857:159613, Jan 2023.
37. Jianbin Xin, Liuqian Wei, Andrea D’Ariano, Fangfang Zhang, and Rudy Negenborn. Flexible time–space network formulation and hybrid metaheuristic for conflict-free and energy-efficient path planning of automated guided vehicles. *Journal of Cleaner Production*, 398:136472, Apr 2023.
38. Zhongwei Zhang, Lihui Wu, Boqiang Zhang, Shun Jia, Weipeng Liu, and Tao Peng. Energy-efficient path planning for a multi-load automated guided vehicle executing multiple transport tasks in a manufacturing workshop environment. *Environmental Science and Pollution Research*, Mar 2024.
39. Jia Gao, Xiaojun Zheng, Feng Gao, Xiaoying Tong, and Qiaomei Han. Heterogeneous multi-type fleet green vehicle path planning of automated guided vehicle with time windows in flexible manufacturing system. *Machines*, 10(3):197, Mar 2022.

40. Haojie Zhang, Yudong Zhang, Chuankai Liu, and Zuoyu Zhang. Energy efficient path planning for autonomous ground vehicles with ackermann steering. *Robotics and Autonomous Systems*, 162:104366, Apr 2023.
41. Te Wang, Aijuan Li, Dongjin Guo, Guangkai Du, and Weikai He. Global dynamic path planning of agv based on fusion of improved a* algorithm and dynamic window method. *Sensors*, 24(6):2011, Mar 2024.
42. Zhengrui Jiang, Wang Chen, Xiaojun Zheng, and Feng Gao. Research on vehicle path planning of automated guided vehicle with simultaneous pickup and delivery with mixed time windows. *IET Collaborative Intelligent Manufacturing*, 6(2), May 2024.
43. Meltem Eyuboglu and Gokhan Atali. A novel collaborative path planning algorithm for 3-wheel omnidirectional autonomous mobile robot. *Robotics and Autonomous Systems*, 169:104527, Nov 2023.
44. Stuart Eiffert, Nathan D. Wallace, He Kong, Navid Pirmarzashti, and Salah Sukkarieh. Resource and response aware path planning for long-term autonomy of ground robots in agriculture. *Field Robotics*, 2:1–33, Mar 2022.
45. Sunil Kumar and Afzal Sikander. An intelligent optimize path planner for efficient mobile robot path planning in a complex terrain. *Microsystem Technologies*, 29(4):469–487, June 2022.
46. Go Sakayori and Genya Ishigami. Modeling of slip rate-dependent traversability for path planning of wheeled mobile robot in sandy terrain. *Frontiers in Robotics and AI*, 11, Jan 2024.
47. Se-Wook Yoo, E-In Son, and Seung-Woo Seo. Traversability-aware adaptive optimization for path planning and control in mountainous terrain. *IEEE Robotics and Automation Letters*, 9(6):5078–5085, June 2024.
48. Omer Melih Gul. Energy-aware 3d path planning by autonomous ground vehicle in wireless sensor networks. *World Electric Vehicle Journal*, 15(9):383, Aug 2024.
49. Yuxiang Li, Kun Chen, Yifei Wang, Weifan Zhang, Jiancheng Wang, Haoyao Chen, and Yunhui Liu. Real-time multilevel terrain-aware path planning for ground mobile robots in large-scale rough terrains. *IEEE Transactions on Robotics*, 41:4159–4179, 2025.
50. Dan Xiang, Hanxi Lin, Jian Ouyang, and Dan Huang. Combined improved a* and greedy algorithm for path planning of multi-objective mobile robot. *Scientific Reports*, 12(1), Aug 2022.
51. K. S. Suresh, R. Venkatesan, and S. Venugopal. Mobile robot path planning using multi-objective genetic algorithm in industrial automation. *Soft Computing*, 26(15):7387–7400, June 2022.
52. V. Sathiya, M. Chinnadurai, and S. Ramabalan. Mobile robot path planning using fuzzy enhanced improved multi-objective particle swarm optimization (fimopso). *Expert Systems with Applications*, 198:116875, July 2022.
53. Senthil Kumar Angappamudaliar Palanisamy, Dinesh Selvaraj, and SivaBalaKrishnan Ramasamy. Hybrid multi-objective optimization approach intended for mobile robot path planning model. *Journal of Intelligent & Fuzzy Systems*, 42(3):2681–2693, Feb 2022.
54. Monex Sharma and Hari Kumar Voruganti. Multi-objective optimization approach for coverage path planning of mobile robot. *Robotica*, 42(7):2125–2149, Mar 2024.

Lijun Xiao received the Ph.D. degree in mechanical engineering from Hunan University, China. He is currently a Professor with the School of Computer Science and Engineering, Hunan University of Science and Technology. His research interests include artificial intelligence, machine learning, and intelligent unmanned systems, with a current focus on motion planning and cooperative optimization for autonomous vehicles and multi-UAV systems.

Xinwu Jian is currently working toward the M.Eng. degree in software engineering with the School of Computer Science and Engineering, Hunan University of Science and Technology, China. His research interests include sampling-based motion planning, UAV path planning, and reinforcement learning.

Zijian Wang is currently working toward the M.Eng. degree in software engineering with the School of Computer Science and Engineering, Hunan University of Science and Technology, China. His research interests include combinatorial optimization, metaheuristic algorithms, and UAV–UGV cooperation.

Sisi Zhou is currently with Hunan Women's University. She received the B.S., M.S., and Ph.D. degrees in 2009, 2012, and 2026, respectively. Her research interests include privacy protection and information security, with a focus on federated learning and blockchain technology.

Kuanching Li (Senior Member, IEEE) is a Professor at Hunan University of Science and Technology. He received his Ph.D. degree from San Pablo Main in 2001, and has published papers in high-ranked journals and conferences and served as associate and guest editor for several scientific journals. His research interests include cloud and edge computing, big data, and blockchain.

Received: April 25, 2026; Accepted: August 8, 2026.

