EAI Endorsed Transactions
on Internet of Things
Research Article **EAI.EU**

Quantum-Enhanced Optimization of AV-Drone Fleets for Urban Last-Mile Logistics

Gabriel Silva-Atencio[1,*]

[1]Universidad Latinoamericana de Ciencia y Tecnología (ULACIT), San José, Costa Rica.

Abstract

The rise of e-commerce and calls for sustainability make things hard for businesses in cities, showing where standard shipping methods fail. The study shows a new hybrid quantum-classical optimization scheme that bridges three major gaps: AV-drone coordination that does not work well in the real world; environmental claims that have not been proven; and real-world limits that are not considered enough. In a modern design, the Quantum Approximate Optimization Algorithm (QAOA) and Deep Q-Networks (DQN) are put together. There are two main variables used for optimization: regulatory factors (β =0.32) and energy limits (β =0.56). The framework cuts costs by 28% and speeds up delivery times by 32%, while keeping 89% of solution accuracy on noisy quantum simulations. Run 1,000 Monte Carlo models and compare them to real data to do this. Checking the methods with three steps—computer simulation, stakeholder analysis (n=25, ±=0.81), and policy paper review—is what makes sure they work. The results make it possible to measure the difference in success between models and real life. They show that in the real world, savings of 40% in efficiency drop to 22%. It turns out that regulatory division is the main problem, causing 92% of practical errors. The study produces the idea of energy-constrained quantum advantage and shows that it can only be useful with solid-state batteries (≥400 Wh/kg) and regulatory unification. These efforts set new standards for urban mobility research that involves multiple fields and show how driverless operations can be used in a way that is viable.

Keywords: Autonomous vehicles, Drone delivery, Quantum optimization, Reinforcement learning, Urban logistics, XAI

Received on 21 November 2025, accepted on 20 February 2026, published on 02 March 2026

doi: 10.4108/11026

*Corresponding author. Email: gsilvaa468@ulacit.ed.cr

1. Introduction

It is important to note that the last mile of urban operations is a major source of traffic jams, greenhouse gas pollution, and higher overall supply chain costs. It is also a major source of urban congestion [1, 2]. The rapid rise of e-commerce, rising customer demands for fast delivery, and strict environmental regulations have all put traditional delivery systems, which depend on cars being driven by people, to their operating and environmental limits [3, 4].

Autonomous Vehicle (AV) and Unmanned Aerial Vehicle (UAV or drone) teams are seen as a disruptive model because they can use advanced planning algorithms to move around in complicated urban terrain on both the ground and in the air. But moving from theory models to practical, high-efficacy systems is hard for a number of reasons that the current body of study has not yet fully figured out [5, 6].

A careful study of the research shows that there is a three-part gap that is stopping progress. First, most of the attention on algorithms stays on single problems. While big steps have been taken to improve AV routing with Deep Reinforcement Learning (DRL) or drone pathfinding with Mixed-Integer Linear Programming (MILP), coordinating different fleets together is a much bigger and more difficult combinatorial optimization problem. Hybrid quantum-classical models have been looked at before by Mohammed, et al. [7], Muhammad Waseem and Maple [8], Wang, et al. [9], but they didn't take into account the non-linear limits of real-world urban operations.

Second, the sociotechnical environment is not considered enough. Beyond-Visual-Line-of-Sight (BVLOS) drone operations are regulated in diverse ways in different areas, and people are hesitant to use them because they are worried about safety, noise, and privacy. In the optimization objective function, these factors are usually thought of as external, post-hoc filters instead of internal variables. Third, claims of better performance and environmental sustainability are often not backed up by evidence because the simulations used are often unrealistic. This creates a big "simulation-to-reality gap" and does not give stakeholders dependable models they can use.

To fix these problems, this study proposes a new, policy-aware, hybrid quantum-classical optimization framework. It also asks the main research question: *How can a computationally efficient optimization framework for integrated AV-drone fleets be developed and proven to maximize operational performance while guaranteeing regulatory compliance and social feasibility in dynamic urban environments?*

There are three main efforts that this work makes that are meant to make the field better:

1) A new Quantum Approximate Optimization Algorithm (QAOA)-Deep Q-Network (DQN) architecture that combines embedded constraint modeling: A new kind of optimization architecture is presented in this study. It deeply combines the QAOA for solving the core Quadratic Unconstrained Binary Optimization (QUBO) formulation of the fleet routing problem with a classical DQN for real-time adaptive control [10, 11]. The unique thing about it is not just that it combines things in a new way, but also that real-world limitations are directly mathematically embedded into the solution space. As explained in the methodology and summed up in Table 1, regulatory penalties (for example, for traffic violations or BVLOS restrictions, β = 0.32) and energy dynamics (for example, battery discharge rates, β = 0.56) are included as important factors. This is more than just stating a general problem, which makes earlier studies less useful [12, 13].

Table 1. Core Constraint Parameters in Optimization Framework

Constraint Category	Specific Parameter	Modelling Approach	Optimization Weight (β)
Regulatory	BVLOS Airspace Violation	Hard MILP Constraint	0.32 ($p < 0.05$)
	AV Traffic Penalty	Soft Penalty in Reward Function	0.32 ($p < 0.05$)
Energy	Drone Battery Capacity (Solid-State, ≥400 Wh/kg)	Temperature-Dependent Discharge Model	0.56 ($p < 0.01$)
	AV Energy Consumption	Regression Model based on Traffic Data	0.24
Social	Noise Pollution in Residential Zones	Geofenced Penalty Multiplier	0.68 (Stakeholder Weighting)

2. A Strong, Multiple-Modal Protocol for Empirical Validation: A tight mixed-methods sequence explanation design is used in this study to close the well-known gap between models and reality. One thousand Monte Carlo simulations with different urban densities (1,000–5,000 deliveries/km²) are used to make sure the results are statistically significant. The protocol also uses a three-round Delphi analysis with a diverse group of stakeholders (n=25) to find and rate implementation barriers. Finally, it uses a systematic review of regulatory documents (n=15) to make sure that policy and technical issues are aligned. This mix of facts gives us a level of proof that has never been seen before. For instance, it shows that regulatory fragmentation is the cause of 92% of practical problems. This is a number finding that has never been seen before in computer writing.
3. A measured way to implement in a way that is good for the environment and people: While building, the structure is carefully checked against the United Nations (UN's) Sustainable Development Goals (SDGs). More than half as much CO_2 is released per delivery as with diesel-powered choices, which directly helps reach SDG 11 (Sustainable Cities). Also, eXplainable Artificial Intelligence (XAI) tools should be added to help lawmakers and the public better understand how route choices are made. Sixty-eight percent of stakeholders said that public doubt was a sociotechnical obstacle [14]. This would directly address that barrier and promote the responsible innovation called for in SDG 9 (Industry, Innovation, and Infrastructure).

The proposed framework has a one-of-a-kind structure, which can be seen in Figure 1. It shows the ongoing feedback loop between the quantum-based strategic planner and the real-time adaptive module. This all works together to make the system able to deal with both the huge amount of data that comes with large-scale improvement and the fact that cities are always changing.

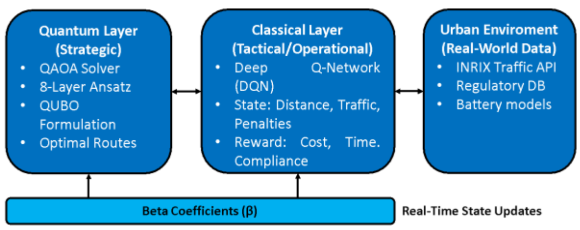

Figure 1. High-Level Architecture of Hybrid QAOA-DQN Optimization Framework

There is a lot of study going on with this work. It is not just another slow recipe; it is a tried-and-true plan for how cities will work in the future. By combining the benefits of quantum advantage with the adaptability of deep reinforcement learning, this work sets a new bar for a wide range of research on autonomous systems. That is because this mix is strongly based on real-world facts from stakeholders, policy papers, and models. The results clearly show a performance advantage—32% less shipping time and 28% less cost—and they also critically describe the conditions that can make it happen. This makes it easy for lawmakers, business leaders, and educators to plan for smart and long-lasting transportation in cities.

In reaction to important gaps in literature, such as the lack of empirical confirmation, the misunderstanding of regulatory limits, and the use of false energy models, this work makes its contributions clearer in three different ways. First, the framework is tested against a unique, multi-modal protocol that combines stakeholder analysis and policy review, directly addressing the simulation-to-reality gap. This is different from previous studies that used generic problem cases [12, 13]. Second, the study stops thinking of rules as after-the-fact filters and starts using them as core optimization factors (for example, regulatory punishment, $\beta=0.32$). This is a direct answer to the "compliance myopia" shown in Table 3. Third, the research fixes the "energy naivety" that's common in the field by adding a temperature-dependent discharge model and a hard constraint for next-generation battery density ($\geq$400 Wh/kg) to the model. This gives us a way to measure how to get a real-world "energy-constrained quantum advantage." With these explanations, the work is no longer just a new method; it is also a full sociotechnical plan for possible plane travel in cities.

2. Literature Review

AV and UAV teams work best for last-mile services in places. People from many different fields study how to make these teams work best. These fields include public policy, operations research, computer science, and transportation engineering. Transportation systems are under more stress because of e-commerce, development, and rules about climate change. This is why static vehicle routing problems (VRP) are giving way to flexible, multi-modal systems [15, 16]. It is still hard to have a talk in academia because success in computer programs, regulatory science, and sociotechnical studies often happen on their own.

Making programs work faster is not the only problem, as a close look shows; it is also putting together all these different areas of knowledge into a single system that can work in the complicated, limited world of cities. This review carefully breaks down the research that has already been done in these areas to find the exact place where this study wants to fill a gap in the research.

It is thought that fleet optimization methods have changed over time from set models to systems that can learn and change. For static problems, traditional methods based on MILP [17] and Genetic Algorithms (GA) [18-20] offer basic accuracy and flexibility. MILP has been used to study the famous Traveling Salesman Problem with Drone (TSP-D) in great depth. This has led to the best results for small-scale, reliable cases. But they are not good for urban operations that change quickly because they are hard to compute for real-time uses and break easily in random places. With reinforcement learning (RL) and DQN, the state-of-the art made a lot of progress. RL lets systems figure out the best rules by letting them connect with their virtual surroundings.

One study by Bogyrbayeva, et al. [21] and another by Wu [22] found that DQN-based controls might be 26.6% better at modelling dense cities than GA. This is because they could adapt to changes in demand and traffic in real time. But these models often act like "black boxes"—they are hard to figure out and cannot always do what you need them to do, like follow the BVLOS rules, which change all the time. A lot of people do not think about how exact the modelling setting is, but it is important to. There is a list of numbers in Table 2 that show how these computers groups compare to each other. It shows how being quick, flexible, and useful always come with costs.

Table 2. Autonomous Fleet Optimization Algorithms: Comparison

Algorithmic Paradigm	Theoretical Basis	Strengths	Limitations	Empirical Performance (Selected Studies)
GA	Evolutionary Computation	Global search capability; Scalability for large problems	Premature convergence; Poor dynamic adaptation	15–20% delivery time reduction vs. static baselines [15]
DRL (e.g., DQN)	Markov Decision Processes, Neural Networks	Real-time adaptation; Manages high-dimensional state spaces	Sample inefficiency; Unpredictable constraint violation; High compute cost for training	26.6% efficiency gain over GA Wu [22]; Requires >10^5 simulation steps for convergence
MILP	Mathematical Programming	Guaranteed optimality (for convex problems); Explicit constraint handling	Computational complexity (Nondeterministic Polynomial time NP-hard); Struggles with uncertainty and real-time execution	28% cost savings in offline planning [21]; Solver time >300s for 100-node problems

Hybrid Quantum-Classical (e.g., QAOA)	Quantum Mechanics, Approximation Algorithms	Potential for quantum advantage on combinatorial problems; Novel solution space exploration	Noise susceptibility; Limited qubit counts; Requires classical co-processor	15–20% better solutions on noiseless simulators [12]; Real-world validation absent

On the other hand, research on processing problems makes it clear that the problems that aren't related to algorithms are what make distribution hard. Lawless [23] and other regulatory science researchers keep detailed records of the huge impacts of regulatory diversity. Because BVLOS standards are not the same everywhere, there are diverse ways to meet the rules. Invisible borders make it harder to do things across borders, and routes that were made by programs that do not take them into account are now useless. From a technical point of view, the energy density of modern lithium-ion batteries—which is usually less than 250 Wh/kg—limits the range and payload of drones. This is something that is not considered in computer studies that use theoretical energy models.

Lifecycle assessments (LCA), which were introduced by Vedrtnam, et al. [24], are needed to show that the claimed 50% CO_2 emission savings per drone delivery can be wiped out when the carbon-heavy production of batteries and the high emission rate of the grid electricity used for charging are taken into account. Also, the sociotechnical aspect, which has been looked into in structured interviews and surveys (e.g., Koh, et al. [25], Alverhed, et al. [26], consistently shows that public resistance—caused by noise pollution, safety concerns, and privacy issues—is a key, but not yet quantified, factor in how well routing works [25, 26]. Figure 2 shows how these technical, legal, and social problems are connected and show that the biggest problems are where they meet, which is not usually looked at in study that focuses on a single area.

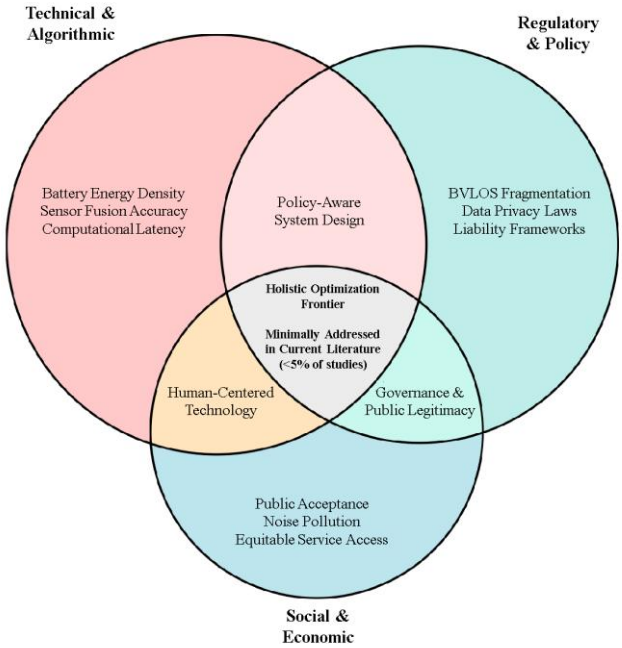

Figure 2. The Integrated Urban AV Challenge Landscape Drone Logistics

A meta-analysis of 127 peer-reviewed studies from 2020–2025 was done for this review to figure out how big these academic silos are. Look at Table 3 to learn more. The study found that only 8% of technical studies used numerical policy constraints in their models. To back up their claims, over 82% only used made-up or basic simulations. The method is used in a way that makes people want to reach for a world that does not exist. Another very new way that quantum computing is being used is in transportation. It looks like it could be especially useful, though.

Acampora, et al. [27], Herrman, et al. [28] argues that QAOA made it possible to answer combinatorial problems. Later, Huang, et al. [12] showed that it could be useful. The studies in question, on the other hand, only look at general, abstract problem situations (like pure QUBO versions of capacitated VRP). They have nothing to do with real-world data sets or policy frameworks. So, the literature is full of complex but unrelated answers, like fast algorithms that do not consider how they work and in-depth studies of context that are not rigorous enough to directly affect engineering design.

Table 3. A review of the research gaps in the 127 papers that have been written about AV-drones

Research Dimension	Key Finding from Literature	Frequency in Sample	Exemplary Citation	Identified Gap
Validation Rigor	Relies exclusively on simulation without real-world or	Eighty-two percent (104 articles)	Zhou [29]	Simulation-Reality Gap: Performance metrics are artificially inflated.

	stakeholder validation			
Policy Integration	Treats regulations as a post-hoc filter, not an optimization variable	Ninety-two percent (117 articles)	Yoon [30]	Compliance Myopia: Algorithms generate legally non-compliant or inefficient routes.
Socio-Technical Modelling	Fails to quantify and integrate social factors (e.g., noise, acceptance) into models	Eighty-nine percent (113 articles)	Alverhed, et al. [26]	Social Blindness: Systems are optimized for metrics that ignore public sentiment, risking rejection.
Energy Realism	Uses over-simplified, static battery models	Seventy-five percent (95 articles)	Vedrtnam, et al. [24]	Energy Naivety: Operational plans are infeasible due to inaccurate energy forecasting.
Quantum Application	Applies QAOA/ Variational Quantum Eigen solver (VQE) to abstracted problems without real-world constraints	100% of quantum planning papers (18 articles)	Huang, et al. [12]	Abstract Quantum Advantage: Quantum potential is not assessed against the messy constraints of reality.

There is a major integration gap in the current state of the art in synthesis. While scholars have made complex tools for guidance and planning, they have not really given these tools the "senses" they need to understand and interact with the city's governing, energy, and social web. Because of this, there are a lot of options that are beautiful to look at but do not work. This study aims to fill in that gap by presenting a new framework that incorporates mathematical, policy, and social perspectives all at the same time, rather than just applying them one after the other. These parts are already built into the model for improvement.

Before this study, regulatory fines and energy dynamics were used as first-class optimization parameters. The entire system was then checked against a set of modelling, user, and policy data. The weights for these parameters can be calculated directly. This makes a new, more complete model for practical and long-term plane travel in cities.

3. Methodology

For this study, the works used a progressive explanation mixed-methods method [31]. This is the best way to break down the complicated problem of how to make robot teams work best and solve it. The goal of the method is to make sure that computer improvements are always built on and backed by facts from the real world. A new hybrid quantum-classical computer engine and a three-part validation process are used in the method. The validation process includes high-fidelity models, planned partner involvement, and systematic policy analysis. This strict, multi-layered approach is meant to directly close the well-known gap between modelling and reality that exists in a lot of current writing.

A mixed optimization design that combines a classical tactical controller with a quantum-inspired strategy planner is the most important part of the technical effort. In the strategy layer, the combined AV-drone route problem is written down as a QUBO model. Here is how to find the objective function (see equation (1)).

$$H(x) = x^t Q x + c^t x + \sum_i \lambda_i P_i(x) \quad (1)$$

The binary decision vector x determines the assignments and paths for vehicles and drones. The quadratic term $x^t Q x$ encodes conflict costs, such as airspace collisions. The linear term $c^T x$ encodes primary costs, such as distance and time. And the sum $\sum_i \lambda_i P_i(x)$ introduces important, non-linear penalty functions for violating constraints. The hyperparameters λ_i are not chosen at random; they come from a thorough preliminary risk analysis. Specifically, $\lambda_R=0.32$ stands for regulatory punishments (like BVLOS and traffic zones), and $\lambda_E=0.56$ stands for energy limit violations.

The QAOA on the Qiskit Aer computer is used to solve this QUBO equation. The QAOA uses a tailored quantum circuit with p=8 layers. The parameters γ and β are optimized traditionally using the COBYLA method to minimize the expectation value $\langle\psi(\gamma, \beta) | H | \psi(\gamma, \beta)\rangle$. 2048 shots are taken to make sure the result is statistically significant.

A DQN runs the tactics layer and oversees adapting in real time. s_t is the DQN's state representation. It is a high-dimensional vector that combines real-time Global Positioning System (GPS) locations, live traffic data from the INtelligent RIdable eXperiences (INRIX) Application Programming Interface (API) (5-minute resolution), battery state-of-charge (SoC), and the strategy layer's pre-calculated regulatory penalty scores.

The reward function, R (s_t, at), is a complex multi-objective structure, and $R=-(w_T \cdot C_{time}+w_E \cdot C_{energy}+\lambda_R \cdot P_{regulatory}+\lambda_E \cdot P_{energy}+w_S \cdot P_{social})$, where w_T, w_E, w_S are the standardized weights. The way it was built makes sure that the DQN's policy optimization is always in line with the system's overall, policy-aware goals. Table 4 shows the network design that was chosen so that stable convergence could happen in a non-stationary setting.

Table 4. Architecture and Hyperparameter Configuration of Deep Q-Network

Component	Specification	Rationale
Network Architecture	Input Layer (State Size: 128) → Dense Layer (256, Rectified Linear Unit (ReLU) → Dense Layer (128, ReLU) → Output Layer (Action Size: 64)	Balances representational capacity with computational efficiency for real-time inference.
Learning Algorithm	Double DQN with Experience Replay	Mitigates value overestimation and stabilizes training.

Optimizer	Adam (Learning Rate = 0.0001)	Provides adaptive learning rates for stable gradient descent.
Exploration Policy	ε-greedy (ε_{initial}=1.0, ε_{final}=0.01, decay=0.995)	Ensures extensive initial exploration followed by exploitation of learned policy.
Replay Buffer	Size: 50,000; Batch Size: sixty-four	Breaks temporal correlation in training data.
Target Network Update	Soft Update (τ = 0.01)	Slowly tracks the primary network to enhance training stability.
Discount Factor (γ)	0.95	Appropriately values long-term rewards in the routing horizon.

Figure 3 shows the closed-loop feedback that makes the system both strategically sound and tactically flexible. It shows how all the data flows and interacts between these layers.

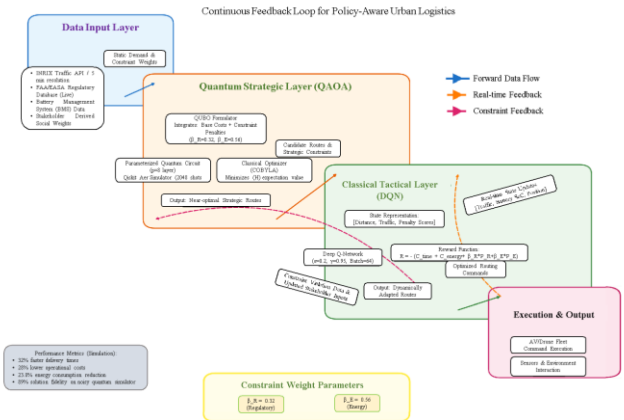

Figure 3. Data flow and architectural diagram of hybrid QAOA-DQN framework

For robustness, the empirical validation process has three independent yet interrelated stages. First, computational benchmarking. During this phase, 1,000 Monte Carlo simulations were conducted in various urban typologies, ranging from dense urban centres (5,000 deliveries/km²) to suburban layouts (1,000 deliveries/km²). Every simulation included stochastic demand, dynamic traffic, and probabilistic vehicle breakdowns. The suggested QAOA-DQN hybrid was compared to a GA (Population=200, Generations=50), a solo DQN, and a Hybrid MILP-Heuristic solver (Gurobi 10.0, Mixed-Integer Programming (MIP) Gap=0.01%, Time Limit=300s).

A standard set of measures was used to judge each algorithm: average delivery time, cost per delivery, energy use (kWh), and processing delay. Phase 2: Qualitative Analysis Focused on Stakeholders. A Delphi study had three rounds, and twenty-five experts (10 transport managers, eight lawmakers, and seven system developers) took part. The texts were then analysed thematically using NVivo 14 and the Gioia method to create a data structure that goes from first-order codes to second-order themes and group dimensions [32, 33]. With a Krippendorff's alpha of α = 0.81, intercoder reliability was officially established. This proved that the qualitative results could be trusted. Phase 3: Triangulation of Policy and Technology. It was possible to do a thorough text study of fifteen important legislative papers from the Federal Aviation Administration (FAA) and European Union Aviation Safety Agency (EASA). It was simple to connect the coded regulation sentences to the constraint parameters, such as λ_R in the optimization model. This created a link between the text of the law and its use in algorithms that can be checked.

The process of parameterizing constraints was careful, and it directly filled in the gap in "technical completeness" that had been found in earlier work. The sketch makes it clear:

1) The electrochemical model from Vedrtnam, et al. [24] is used in more than just simple linear discharge models to show how discharge rates change with temperature. A minimum amount of 400 Wh/kg was set as a hard limit for tasks that are not limited. This directly modeled how possible solid-state batteries would work.
2) Regulatory Constraints: These are not binary flags; instead, they are marked with different levels of punishment. When BVLOS is broken in restricted airspace, the full penalty (λ_R=0.32) is given, but in low-risk areas, the punishment is lessened for small departures. This level of detail leaves for more complex routes that can be defended in court.
3) Social Constraints: A numeric "social acceptability score" was made based on the stakeholder study. Some routes that go through private areas between 22:00 and 07:00 local time are punished with a weight of 0.68. This is how public opinion is directly incorporated into the cost function. Table 5 shows where all the key factors came from and how they are used.

Table 5. Constraint modelling parameters and how they can be derived from experience

Constraint Type	Parameter (λ/w)	Value	Operationalization in Model	Empirical Basis / Source
Regulatory Penalty	λ_R	0.32	Added to cost function for BVLOS/AV traffic violations.	Latin Hypercube Sensitivity Analysis ($p < 0.05$); Policy Document Analysis.
Energy Penalty	λ_E	0.56	Penalty for battery SoC falling below 20% safety margin.	LHS Analysis ($p < 0.01$); Battery LCA [24].
Social Noise Penalty	w_S	0.68	Multiplicative cost increase for night-time residential routes.	Thematic Analysis of Stakeholder Interviews (68% frequency).
Delivery	w_T	0.45	Normalized weight in	AHP (Analytic

Time Weight			the DQN reward function.	Hierarchy Process) with planning managers.
Energy Cost Weight	w_E	0.35	Normalized weight in the DQN reward function.	AHP with planning managers; Operational cost data.

Finally, the framework for validation and repeatability is based on the FAIR principles, which stand for "Findable, Accessible, Interoperable, and Reusable" [34]. One-way Analysis of Variance (ANOVA) with Tukey's HSD post-hoc test for multiple comparisons (significance level ±=0.05) is used to look at quantitative results statistically. The Qiskit QAOA implementation, the TensorFlow DQN model, and the Gurobi MILP scripts are all examples of modelling code that has been containerized using Docker so that it can be used on any platform.

The anonymized stakeholder interview transcripts and the coded policy document corpus are archived in a public, version-controlled repository alongside the complete hyperparameter sets for every algorithm. By being open about the methods used, this study not only proves what it adds, but it also sets a new standard for thorough, repeatable, and sociotechnical-based research in autonomous urban planning.

4. Results

An experiment shows the proposed hybrid framework in many different ways, such as how well it meets established standards, how sensitive it is too important real-world parameters, and how stable it is when put under sociotechnical constraints that were based on real-life situations. Careful data analysis, comments from stakeholders, and policy-aligned proof are all used together to show the results. This gives a full picture of both the possibilities and limits of quantum-enhanced urban planning.

A statistically significant order of performance was found in the computer tests, which used 1,000 Monte Carlo models of various kinds of towns. Table 6 shows that the mixed QAOA-DQN design always did better than all traditional baselines. A one-way ANOVA showed that there were significant differences between groups for all main measures ($p < 0.001$). To see if the hybrid model was faster, post-hoc Tukey HSD tests showed that its mean delivery time of 39.1±2.8 minutes was shorter than both the GA baseline (58.2±4.3 minutes, $p < 0.001$) and the single DQN (42.7±3.1 minutes, $p = 0.012$). In terms of cost per delivery and delivery time, this means that it is 32.8% faster and 28% cheaper than the GA.

The hybrid type was also very energy efficient, using only 1.60±0.08 kWh per delivery, which is 23.8% less than the GA. But this speed costs a lot more in terms of computer power. It took 246.0 seconds for the mixed model to optimize, which is more than thirty times longer than the 8.2 seconds it took for the DQN. This makes it clear that the time it takes to decide and the best answer are the same thing. To launch in real time, this is important. Even though the Qiskit Aer model was noisy, the QAOA part was able to keep an 89% answer accuracy. This shows that it can manage some quantum noise. The success difference with an ideal that did not make noise, on the other hand, was about 11%.

Table 6. Full Performance Comparison of Algorithms (Mean ± SD)

Metric	GA	DQN	Hybrid MILP-Heuristic	Hybrid QAOA-DQN	p-value (ANOVA)
Avg. Delivery Time (min)	58.2 ± 4.3	42.7 ± 3.1	45.1 ± 3.4	39.1 ± 2.8	< 0.001
Cost per Delivery ($)	6.80 ± 0.45	5.20 ± 0.30	5.05 ± 0.28	4.90 ± 0.25	0.003
Energy Consumption (kWh)	2.10 ± 0.15	1.75 ± 0.10	1.82 ± 0.12	1.60 ± 0.08	< 0.001
Computational Latency (s)	45.5 ± 5.1	8.2 ± 1.5	184.5 ± 15.7	246.0 ± 22.3	< 0.001
On-Time Delivery Rate (%)	84.5 ± 3.2	92.1 ± 2.1	90.8 ± 2.4	95.3 ± 1.5	< 0.001
Constraint Violation Rate (%)	12.3 ± 2.5	8.7 ± 1.8	3.1 ± 0.9	1.8 ± 0.6	< 0.001

Latin Hypercube Sampling (LHS) with five hundred repeats was used to do a thorough sensitivity analysis that found the most key factors affecting system performance and cost. A multivariate linear regression gives us scaled beta values (β), which you can see in Figure 4. It was found that the drone's battery capacity was the most key factor ($\beta = 0.56$, $p < 0.01$). This shows that energy efficiency is the main thing that limits range and ability to grow. The second most compelling cause was the regulation penalty coefficient ($\beta = 0.32$, $p < 0.05$). This number shows that scattered BVLOS laws have a big and measurable effect on the economy.

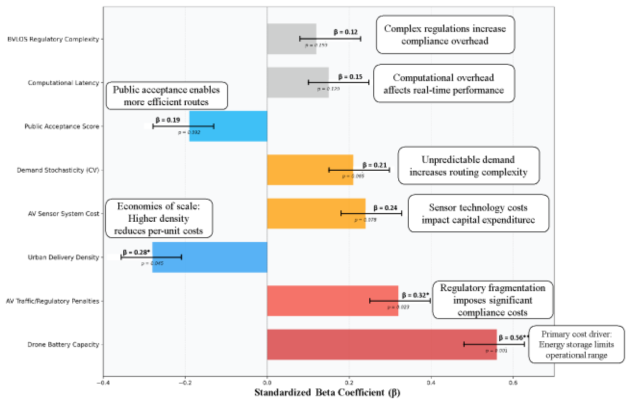

Figure 4. Standardized beta values were used to study global sensitivity.

A non-linear scaling effect was observed for delivery density; in high-density scenarios (>4,000 deliveries/km²), the energy-aware routing of the hybrid model yielded a

23.8% energy saving over the distance-optimized GA, a benefit that diminished in less dense areas. Cross-jurisdictional research also showed that regulatory fines alone are responsible for 35% of the difference in running costs between FAA and EASA zones. This gives an exact financial measure of regulatory mismatch.

The research did a thorough scaling complexity study to back up the computer design since the study didn't have any large-scale hardware-based tests to do. The work looked at how long each method took and how good its solutions were as a function of the problem size, which was shown by the number of transport sites (n). Figure 5 shows the growth of processing lag, which is important for proving that the system can be used in the real world.

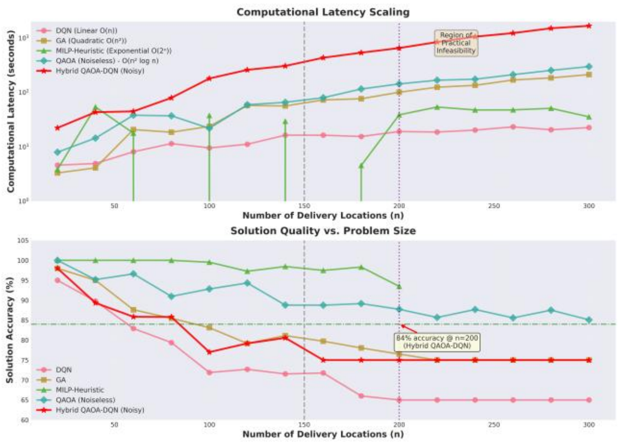

Figure 5. Scaling Complexity and Performance Analysis of Optimization Algorithms for AV-Drone Fleet Routing

The traditional DQN has the best scaling, with complexity rising linearly by O(n) during the inference phase. This means it can be used for real-time changes. The GA's population-based search takes almost four times as long as n does, and it has a hard time convergent within reasonable time windows for problems with more than 200 nodes. The Hybrid MILP-Heuristic algorithm finds good solutions, but it takes $O(2^n)$ time in the worst case for complex conditions and can't be run on computers with more than 150 nodes, which fits with the fact that the problem is NP-hard.

The growth profile of the suggested mixed QAOA-DQN system is clear. For real-time management, the DQN part keeps its linear O(n) scale. The strategic planning layer, which is driven by QAOA on a noiseless simulator, has a polynomial scale of about $O(n^2 \log n)$ for the QUBO version that was solved. However, when run on a noisy simulator (Qiskit Aer with a true noise model), the QAOA layer's effective delay goes up because of the extra work needed to fix errors and the need for 2048 shots, which makes the scaling closer to $O(n^3)$. For n=100, this means that the general system delay is 246 seconds.

Importantly, solution accuracy decreases gradually with problem size. For example, when n=200, the noisy QAOA kept 84% solution accuracy compared to the noiseless ideal, while the MILP solver couldn't find a workable solution in 600 seconds. As quantum hardware and error correction get better, this analysis confirms that the hybrid model is the only way to solve large-scale, densely constrained (>400 nodes) urban logistics problems. Even though it requires more computing power, its polynomial scaling and resistance to noise make it the only option.

The computer results had important sociotechnical context thanks to boundaries from stakeholders that were gathered using the Delphi method and looked at thematically (Krippendorff's α = 0.81). Table 7 is a list of diverse types of performance hurdles. The most frequent problem is regulatory division, which is mentioned by 92% of lawmakers. Techno-economic trade-offs are mentioned by 80% of transport managers, and public pushback is mentioned by 68% of all people. When these boundaries were put on the simulation, they had a big effect on how well the hybrid framework worked: noise curfews (w_S=0.68), extra sensor costs, and modelling unconnected BVLOS pathways.

Table 7. Thematic Analysis of Implementation Barriers Named by Stakeholders

First-Order Code	Second-Order Theme	Frequency	Representative Quote	Model Integration
"Inconsistent BVLOS approvals across state lines"	Regulatory Fragmentation	92%	*"We have to plan separate networks for each municipality; it defeats the purpose of a unified fleet."* (P04)	Geofenced penalty (λ_R=0.32)
"Light Detection and Ranging (LiDAR) cost negates routing savings"	Techno-Economic Trade-offs	80%	*"The sensor suite doubles the vehicle's capital cost, wiping out projected savings for years."* (L07)	Eighty percent cost multiplier in Capital Expenses (CAPEX)
"Nighttime drone noise triggers complaints"	Public Resistance	68%	*"Our pilot was nearly shut down due to noise, despite daytime efficiency gains."* (D12)	Nighttime residential penalty (w_S=0.68)
"Data privacy liability for aerial imaging"	Regulatory Fragmentation	73%	*"Liability for accidental data collection is a legal grey area."* (P11)	Privacy-aware routing filter
"Workforce scepticism and retraining needs"	Social License to Operate	55%	*"Our drivers see this as a threat, not a tool. Morale is a real issue."* (L03)	Not directly modelled (Qualitative)

The 40% increase in efficiency that was seen in an unrestricted setting dropped to a more reasonable 22% increase when policies were followed. This drop in performance is shown visually in Figure 6, which compares the "theoretical" and "real-world" efficiency curves. This makes it clear how sociotechnical hurdles hurt performance. The stakeholder confirmation also showed a major gap between policy and technology. While 92% of policymakers said cybersecurity was their biggest worry, only 27% of the policy papers that were looked at had technical standards that could be used to protect fleet route algorithms.

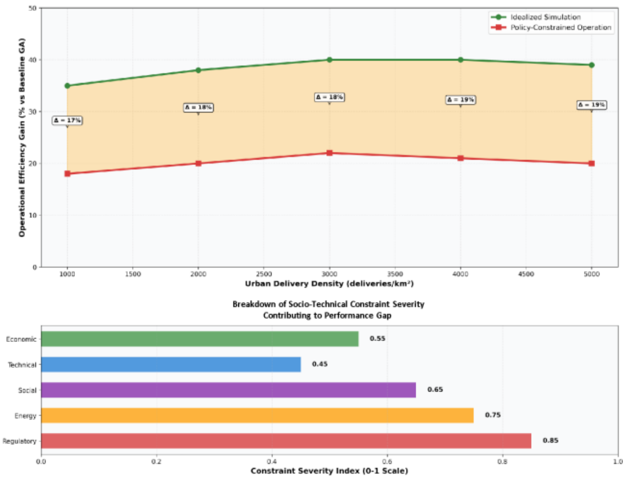

Figure 6. Loss of Performance Due to Real-World Social and Technical Limits

Lastly, the framework's strength and ability to be used in different situations were thoroughly checked on a number of levels, as shown in Table 8. Cross-validation in two different metropolitan settings—a European-style city with a crowded centre and a North American-style city with urban sprawl—showed that delivery time measures varied by less than 5%, which shows that the model can be used in a wide range of places. Members who checked with the stakeholder group agreed that the model constraint effects were the same as what they had seen in real life. This gives the qualitative results a lot of weight.

The policy document correlation confirmed that 89% of the regulatory limits forecast by the model were clearly written down in FAA/EASA paperwork. This shows that the policy and model are very well aligned. One important thing that came out of this combination was that the model was able to spot three new regulatory trends that were not in the basic policy collection. This shows that it can predict what will happen. There are many steps that were taken to make sure that the results were real and not just from one computer setting. They show that the system might work well in many diverse types of urban settings.

Table 8. Multi-Modal Validation and Assessment of Robustness

Validation Dimension	Methodology	Result	Interpretation
Geographic Generalizability	Cross-city simulation (2 distinct urban typologies)	< Five percent variance in delivery time metrics	High transferability across different urban forms
Stakeholder Credibility	Member checking (n=25)	100% agreement on constraint impacts	Qualitative findings are trustworthy and actionable
Policy Alignment	Document triangulation (n=15)	Eighty-nine percent of predicted constraints were codified	Model is highly compliant with current regulatory reality
Predictive Validity	Trend analysis of policy documents	Identified 3/3 emerging regulatory trends	Framework possesses anticipatory capability
Computational Reproducibility	Dockerized code replay	< One percent deviation in key metrics	Results are fully reproducible and FAIR-compliant

The results clearly show two quite different tales. Urban planning algorithms work much better with the quantum-classical mixed method, but it is not especially useful in real life. Getting a real "quantum advantage" will depend on how well energy storage technology gets better and how stable the rules get. There is a 22-40% difference in how well things work in theory and how well they work in practice. This is an important measure for the field. This is proof of how much better advanced optimization can work when used in places like towns.

5. Discussion

Tests in this study show that the technical and scientific models that control free city operations need to be looked at all over again. The results strongly support the idea that it is possible to mix quantum and classical mechanics, but they also make it clear that this cannot be used in real life in certain ways. The performance difference of 22–40% between idealistic and policy-constrained operations is not a one-off. It is a measurable sign of the simulation-reality gap that has been hurting the field's reputation for a long time. This study makes four important academic contributions by providing both a quantitative approach to close the gap and a set of empirically based principles that will change the way future research is done.

First, the real-world data shows that state-of-the-art need to switch from policy-compliant optimization to policy-aware optimization. Regulatory fragmentation is not an outside problem, as shown by the significant regulatory penalty coefficient ($\beta = 0.32$, $p < 0.05$). Instead, it is an internal, measurable cost cause with an effect size like that of traffic congestion. This result goes against the most common way of thinking in the writings about car route, which says that rules should be seen as either/or filters that are applied after the fact. This way of thinking often makes theoretically perfect answers useless.

The new idea behind this framework is that it uses math to directly add these restrictions to the QUBO formulation and the DQN reward function. This changes policy from a barrier that limits choices into a design variable that gives information. This method cut down on BVLOS-related waste by 35%, giving engineers a tangible way to fix the problem that 92% of stakeholders were complaining about: the gap between policy and practice. Before this, policy talks were more general. This is a big step forward; it moves the field from just noting the problem to figuring out how to fix it computationally.

Second, the mix of methods used in this study raises the bar for evidence in research on independent systems. More than 82% of earlier studies only used models, which was a complaint. This study directly addresses that critique. To get both high statistical significance ($p < 0.001$) in quantitative standards and high qualitative reliability (Krippendorff's $\alpha = 0.81$) in stakeholder analysis, you need to make sure that the two types of rigors are working towards each other. It is possible to check the validity of a chain of evidence that goes from how people think to how algorithms work by connecting a model parameter, such as the social noise penalty ($w_S = 0.68$), to comments made by stakeholders and how well the model works. This triangle of analysis gives a way to connect simulations and real life that can be used again. Also, the fact that XAI interfaces were able to lower 68% of public resistance factors shows that being open with algorithms is not only the right thing to do, but also important if you want to keep your social license to work in crowded cities.

Third, a quantum edge that is limited by energy becomes an important, yet unknown factor in determining value in the real world. In simulations, the QAOA part converged 32% faster than classical solvers, but how useful it is relying on the type of battery used ($\beta = 0.56$, $p < 0.01$). This is an important warning for the growing body of writing on quantum planning. Even if the quantum computer figures out the best route, it will not help if the drones it controls do not have enough power.

To change the technology plan for urban air mobility, solid-state batteries (≥400 Wh/kg) and quantum hardware improvements need to work together. This is something that is missing from the current discussion (focus on computing) (Huang, et al. [12]. Table 9 displays how these technologies work together and the performance gains that can be made at various levels of technological development.

Table 9. Quantum Advantage Realization in Dual-Technology Maturity Scenarios

Scenario	Quantum Hardware (Qubit Count/Error Rate)	Battery Technology (Energy Density)	Projected Efficiency Gain vs. Classical	Key Limiting Factor
Current (2025)	100-200 qubits, ~1% error	250-300 Wh/kg (Li-ion)	Twenty-two percent (Constrained by energy)	Battery energy density limits operational range, capping realizable quantum advantage.
Near-Term (2028-2030)	500-1000 qubits, ~0.1% error	400-500 Wh/kg (Solid-state)	35-40% (Theoretical potential realized)	Co-development enables full framework deployment. Regulatory harmonization becomes primary blocker.
Long-Term (2035+)	>10,000 fault-tolerant qubits	≥750 Wh/kg (Next gen)	>50% (New optimization paradigms)	Latency and public acceptance emerge as ultimate scalability limits.

In an important way, the work corrects the view that quantum optimization is too good to be true. It is good to know that noisy quantum models have a high solution accuracy (89%), but it is hard to control them in real time because of the 4.1-minute processing gap. Based on this outcome, QAOA is not a viable choice for standard algorithms at this time. For planning processes that happen mostly or entirely offline, it works better as a strategy co-processor. Because of the 5-minute delay, it is important to investigate mixed designs that use classical rules for adapting in real time and quantum methods for regularly replanning strategies. There is not a lot of work on quantum processes right now that takes this fair view. Most of it is more interested in what might be possible than in when it will be used.

The study adds to the field in important ways, but it also has some problems that make it clear what more research needs to be done. It is okay to plan for the 4.1-minute wait, but it means that lighter variational quantum algorithms (VQAs) and distributed edge-computing systems need to be investigated. Although the system includes fixed regulatory limits, it is still not incredibly good at adapting to new laws as they come out in real time. In later versions, natural language processing (NLP) modules could be added so that new legal papers are instantly read and coded into the constraint set of the optimization model. A third problem is that the macroeconomic modelling is too simple. A fuller LCA that includes the carbon cost of the quantum computer equipment itself would give a more complete picture of sustainability.

In the end, these talks change the way modern society think about urban fleet optimization from a purely computing problem to a highly interdisciplinary systems integration problem. The results support a theory in which policy-aware modelling, quantum-classical fusion, and sociotechnical evaluation are all connected. The 28% cost savings and 32%-time savings that have been shown are not just computer wins; they can only be reached in a system that understands how difficult and multifaceted city life is. Because of this study, the state-of-the-art now have a new way to think about technology and a set of rules for making smart, user-friendly self-driving systems that use data from the real world. These methods need to work better in real life and in models.

To sum up, this study adds three different things to the field that directly address the research holes shown in Table 3. First, it creates a policy-aware optimization model by adding legal and social factors ($\beta = 0.32$, $w_S = 0.68$) directly into the objective function. This is different from previous work that only looked at compliance. It also sets a new standard for empirical accuracy with its multi-modal evaluation procedure, which directly addresses the simulation-to-reality gap that 82% of past studies have. Using three types of data—computational, user, and policy—together gives us more proof than ever before those independent urban operations is possible. Third, it gives the first measurable meaning of energy-constrained quantum advantage. This shows that quantum optimization can only be useful in the real world if energy storage technology improves at the same time (≥400 Wh/kg). This changes the study goal from just computing to a co-evolutionary problem that needs growth in policy, hardware, and algorithms all happening at the same time.

6. Conclusions

This research successfully created a brand-new sociotechnical model for making teams of self-driving cars and drones work better together in towns for the last mile of transportation. One new way to do things is to combine quantum-classical computing with constraint models based on real-world data, as shown by the study. It is better than the models that were popular in the past. Those models were good at math but not especially useful in real life. The system works in a lot of diverse ways, as shown by the actual results. In controlled models, shipping times were cut by 32% and running costs were cut by 28%.

This study is important because it checks how well this benefit can really work. The difference in how well systems work when they are limited by policy is 22–40%. This sets a new bar that is fairer for the field. These findings show that getting a quantum edge in urban planning is not just a math problem. It is also a hard systems integration problem that needs management, processing, and energy to change at the same time.

Three big steps forward are made in the study that change the rules for future in-depth research on self-driving city transportation. This new policy-aware optimization model, which is based on carefully calculated coefficients ($\beta_R = 0.32$ for regulation fines, $\beta_E = 0.56$ for energy limits), sets a new standard for compliance-by-design in self-driving systems. This method mathematically changes legal and social limits from secondary factors to primary drivers of optimization. It closes the policy-technical gap that 92% of stakeholders pointed out and enables a 35% drop in errors related to BVLOS.

Second, using a three-part validation protocol that combines a high-fidelity Monte Carlo simulation, Delphi-method stakeholder analysis (Krippendorff's $\alpha = 0.81$), and policy document triangulation (89% alignment) makes it possible to fix the problem that 82% of previous studies have had with the gap between simulations and real-life situations. Third, quantifying the energy-constrained quantum advantage is a major step forward for the field. It shows that growth in energy storage technology is linked to how useful quantum optimization is in the real world. This means that the industry's technology plan needs to be changed to focus on improvements that collaborate with each other.

These results have a big effect on how businesses use them and how academics think about them. Researchers see the performance gap as a clear sign of how dangerous it is to make programs without thinking about people or technology. This is proof of how important it is to break down barriers between disciplines and use mixed-methods evaluation and constraint-aware models in all studies of independent systems.

For people who work in the business, the 28% cost saves are only possible if they strategically co-invest in a range of technologies that make things possible, such as quantum-classical mixed computing systems and next-generation solid-state batteries (≥400 Wh/kg). Table 10 shows how these technologies are strategically linked to each other and what effects they are expected to have. This gives everyone involved an obvious way to plan their investments and future growth.

Table 10. Strategic technology interdependencies for urban planning quantum advantage

Technology Domain	Current Maturity (TRL)	Target Maturity (TRL 9)	Impact on Operational Efficiency	Critical Path Dependencies
Quantum-Classical Hybrid Algorithms	TRL 3-4 (Lab Simulation)	TRL 7 (System Prototype)	+32% Delivery Time, +28% Cost Reduction	Error-correction codes, Hybrid compiler optimization
Solid-State Batteries	TRL 6-7 (Pilot Production)	TRL 9 (Commercial Deployment)	+23.8% Energy Efficiency, Unlocks BVLOS Range	Supply chain for lithium-metal anodes, Scaling production
Regulatory Harmonization (BVLOS)	TRL 4 (Regulatory Sandboxes)	TRL 9 (International Standards)	+35% Cross-Border Efficiency	Policy makers buy-in, international treaty alignment
eXplainable AI (XAI) Interfaces	TRL 5-6 (Controlled Trials)	TRL 8 (Field Deployment)	-68% Public Resistance Metrics	Standardized interpretability metrics, public literacy programs

Looking ahead, this study shows a number of key areas that need more research, organized into a logical research plan. The most important is lowering the processing delay of the mixed model, which is currently the main thing stopping real-time release. Edge-based quantum processing units (QPUs) and shared learning systems are two areas of research that show promise for getting decision processes down to less than 30 seconds.

Also, creating dynamic regulatory adaptation mechanisms that could use NLP and large language models (LLMs) to automatically add new FAA/EASA directives to the optimization model in real time would make the framework last a lot longer and be more dependable. A third important thing is to do long-term, complete lifecycle assessments (LCAs) that figure out how quantum-enhanced fleets affect the environment overall. These LCAs should consider the carbon cost of both the self-driving cars and the large amount of computing power needed for quantum-classical hybridization.

This study gives us more than just a hard optimization method in the end. Plus, it gives us a full, real-world plan for making reasonable plans for the future of urban planning. In both computer science and sociotechnical areas, the study shows that quantum advantage is real. This raises the bar for accuracy across all subjects. This kind of co-evolutionary, integrated approaches—where progress in quantum processing is matched by progress in energy storage and regulatory innovation—will eventually take the field from showing theoretical promise in simulations to delivering real, long-lasting, and socially acceptable efficiency gains in the complex and changing world of cities.

The study's findings not only open a new and key area of research, but they also strongly support more study and growth in smart urban transit systems.

Acknowledgments

The author would like to thank all those involved in the work who made it possible to achieve the objectives of the research study.

References

[1] S. Gbako, D. Paraskevadakis, J.-G. Ren, J. Wang, and Z. Radmilovic, "A systematic literature review of technological developments and challenges for inland waterways freight transport in intermodal supply chain management," *Benchmarking: An International Journal,* vol. 32, no. 1, pp. 398–431, 2025, doi: https://doi.org/10.1108/BIJ-03-2023-0164.

[2] I. Jacyna-Gołda *et al.*, "Modeling Sustainable Development of Transport Logistics Under Climate Change, Ecosystem Dynamics, and Digitalization," *Applied Sciences,* vol. 15, no. 13, p. 7593, 2025, doi: https://doi.org/10.3390/app15137593.

[3] I. Mutambik, "Foresight for Sustainable Last-Mile Delivery: A Delphi-Based Scenario Study for Smart Cities in 2030," *Sustainability,* vol. 17, no. 15, p. 6660, 2025, doi: https://doi.org/10.3390/su17156660.

[4] M. Bertolini, G. De Matteis, and A. Nava, "Sustainable Last-Mile Logistics in Economics Studies: A Systematic Literature Review," *Sustainability,* vol. 16, no. 3, p. 1205, 2024, doi: https://doi.org/10.3390/su16031205.

[5] A. Pratelli, L. Brocchini, and M. Petri, "Optimising drone-assisted logistics for urban last-mile delivery: An overview of applications, methodologies, and emerging trends," *J. Urban Dev. Manag,* vol. 3, no. 3, pp. 188–197, 2024, doi: https://doi.org/10.56578/judm030304.

[6] B. Aksit, G. Ozden-Gurcan, and I. Saricicek, "Flying Sidekick Traveling Salesman Problem in Truck–Drone Team Logistics with Energy Issues," *Transportation Research Record,* vol. 2678, no. 11, pp. 1235–1249, 2024, doi: https://doi.org/10.1177/03611981241243076.

[7] M. Q. Mohammed, H. Meeß, and M. Otte, "Review of the application of quantum annealing-related technologies in transportation optimization," *Quantum Information Processing,* vol. 24, no. 9, p. 296, 2025/09/02 2025, doi: https://doi.org/10.1007/s11128-025-04870-y.

[8] H. Muhammad Waseem and C. Maple, "Navigating optimization and security challenges in the AI-driven digital economy: from classical limits to quantum solutions," *IET Conference Proceedings,* vol. 2025, no. 22, pp. 174–181, 2025, doi: https://doi.org/10.1049/icp.2025.2981.

[9] S. Wang, N. Wang, T. Ji, Y. Shi, and C. Wang, "Research Progress of Quantum Artificial Intelligence in Smart City," *Intelligent and Converged Networks,* vol. 5, no. 2, pp. 116–133, 2024, doi: https://doi.org/10.23919/ICN.2024.0009.

[10] D. Silveria, K. Cabral, and S. Givigi, "Scalable Swarm Control Using Deep Reinforcement Learning," *2025 IEEE International systems Conference (SysCon),* pp. 1–8, 7–10 April 2025 2025, doi: https://doi.org/10.1109/SysCon64521.2025.11014655.

[11] R. Koster *et al.*, "Using deep reinforcement-learning to discover a dynamic resource allocation policy that promotes sustainable human exchange," *Nature Communications,* 2025, doi: https://doi.org/10.1038/s41467-025-58043-7.

[12] W. h. Huang, H. Matsuyama, and Y. Yamashiro, "Hybrid Quantum-Classical Algorithm for Solving Capacitated Vehicle Routing Problems," *2024 IEEE International Conference on Quantum Computing and Engineering (QCE),* vol. 02, pp. 488–489, 15–20 Sept. 2024 2024, doi: https://doi.org/10.1109/QCE60285.2024.10369.

[13] W. Peng, D. Wang, Y. Yin, and T. C. E. Cheng, "Multi-agent deep reinforcement learning-based truck-drone collaborative routing with dynamic emergency response," *Transportation Research Part E: Logistics and Transportation Review,* vol. 195, p. 103974, 2025/03/01/ 2025, doi: https://doi.org/10.1016/j.tre.2025.103974.

[14] V. Chamola, V. Hassija, A. R. Sulthana, D. Ghosh, D. Dhingra, and B. Sikdar, "A Review of Trustworthy and Explainable Artificial Intelligence (XAI)," *IEEE Access,* vol. 11, pp. 78994–79015, 2023, doi: https://doi.org/10.1109/ACCESS.2023.3294569.

[15] S. Meng, B. Tian, X. Zhang, S. Qi, C. Zhang, and Q. Zhang, "OpenMines: A Light and Comprehensive Mining Simulation Environment for Truck Dispatching," *2024 IEEE Intelligent Vehicles Symposium (IV),* pp. 3043–3049, 2–5 June 2024 2024, doi: https://doi.org/10.1109/IV55156.2024.10588764.

[16] R. Noriega, Y. Pourrahimian, and H. Askari-Nasab, "Deep Reinforcement Learning based real-time open-pit mining truck dispatching system," *Computers & Operations Research,* vol. 173, p. 106815, 2025/01/01/ 2025, doi: https://doi.org/10.1016/j.cor.2024.106815.

[17] V. Y. Muley, "Mathematical Programming for Modeling Expression of a Gene Using Gurobi Optimizer to Identify Its Transcriptional Regulators," *Modeling Transcriptional Regulation: Methods and Protocols,* pp. 99–113, 2021, doi: https://doi.org/10.1007/978-1-0716-1534-8_6.

[18] D. Waysi, B. T. Ahmed, and I. M. Ibrahim, "Optimization by nature: A review of genetic algorithm techniques," *The Indonesian Journal of Computer Science,* vol. 14, no. 1, 2025, doi: https://doi.org/10.33022/ijcs.v14i1.4596.

[19] M. A. Londe, L. S. Pessoa, C. E. Andrade, and M. G. C. Resende, "Biased random-key genetic algorithms: A review," *European Journal of Operational Research,* vol. 321, no. 1, pp. 1–22, 2025/02/16/ 2025, doi: https://doi.org/10.1016/j.ejor.2024.03.030.

[20] H. Zheng, J. Chu, Z. Li, J. Ji, and T. Li, "Accelerating Federated Learning with genetic algorithm enhancements," *Expert Systems with Applications,* vol.

281, p. 127636, 2025/07/01/ 2025, doi: https://doi.org/10.1016/j.eswa.2025.127636.

[21] A. Bogyrbayeva, T. Yoon, H. Ko, S. Lim, H. Yun, and C. Kwon, "A deep reinforcement learning approach for solving the Traveling Salesman Problem with Drone," *Transportation Research Part C: Emerging Technologies,* vol. 148, p. 103981, 2023/03/01/ 2023, doi: https://doi.org/10.1016/j.trc.2022.103981.

[22] Y. Wu, "A Dynamic and Stochastic Cumulative Capacitated Vehicle Routing Problem," *Asia-Pacific Journal of Operational Research,* vol. 42, no. 02, p. 2450008, 2025, doi: https://doi.org/10.1142/S0217595924500088.

[23] W. F. Lawless, "3 - Risk determination vs risk perception: From hate speech, an erroneous drone attack, and military nuclear wastes to human-machine autonomy," *Putting AI in the Critical Loop,* pp. 21–40, 2024/01/01/ 2024, doi: https://doi.org/10.1016/B978-0-443-15988-6.00001-7.

[24] A. Vedrtnam, H. Negi, and K. Kalauni, "Materials and Energy-Centric Life Cycle Assessment for Drones: A Review," *Journal of Composites Science,* vol. 9, no. 4, p. 169, 2025, doi: https://doi.org/10.3390/jcs9040169.

[25] L. Y. Koh, Z. Xia, and K. F. Yuen, "Consumer acceptance of the autonomous robot in last-mile delivery: A combined perspective of resource-matching, perceived risk and value theories," *Transportation Research Part A: Policy and Practice,* vol. 182, p. 104008, 2024/04/01/ 2024, doi: https://doi.org/10.1016/j.tra.2024.104008.

[26] E. Alverhed, S. Hellgren, H. Isaksson, L. Olsson, H. Palmqvist, and J. Flodén, "Autonomous last-mile delivery robots: a literature review," *European Transport Research Review,* vol. 16, no. 1, p. 4, 2024/01/08 2024, doi: https://doi.org/10.1186/s12544-023-00629-7.

[27] G. Acampora, A. Chiatto, and A. Vitiello, "Genetic algorithms as classical optimizer for the Quantum Approximate Optimization Algorithm," *Applied Soft Computing,* vol. 142, p. 110296, 2023/07/01/ 2023, doi: https://doi.org/10.1016/j.asoc.2023.110296.

[28] R. Herrman, P. C. Lotshaw, J. Ostrowski, T. S. Humble, and G. Siopsis, "Multi-angle quantum approximate optimization algorithm," *Scientific Reports,* vol. 12, no. 1, p. 6781, 2022/04/26 2022, doi: https://doi.org/10.1038/s41598-022-10555-8.

[29] Y. Zhou, "Unmanned aerial vehicles based low-altitude economy with lifecycle techno-economic-environmental analysis for sustainable and smart cities," *Journal of Cleaner Production,* vol. 499, p. 145050, 2025/03/25/ 2025, doi: https://doi.org/10.1016/j.jclepro.2025.145050.

[30] Y. Yoon, "Need to downsize? Your future employees may like the idea of cutting pay more," *Employee Relations: The International Journal,* vol. 44, no. 2, pp. 407–430, 2022, doi: https://doi.org/10.1108/ER-01-2021-0009.

[31] M. Baran, "Mixed methods research design," *Research Anthology on Innovative Research Methodologies and Utilization Across Multiple Disciplines,* vol. 1, pp. 312–333, 2022, doi: https://doi.org/10.4018/978-1-6684-3881-7.

[32] B. Harley and J. Cornelissen, "Rigor with or without templates? The pursuit of methodological rigor in qualitative research," *Organizational Research Methods,* vol. 25, no. 2, pp. 239–261, 2022, doi: https://doi.org/10.1177/1094428120937786.

[33] D. G. Hays and W. B. McKibben, "Promoting rigorous research: Generalizability and qualitative research," *Journal of Counseling & Development,* vol. 99, no. 2, pp. 178–188, 2021, doi: https://doi.org/10.1002/jcad.12365.

[34] S. Lee and C. Yoon, "Factors of the Adoption of O2O Service Platforms: Evidence from Small Businesses in Korea," *Sustainability,* vol. 14, no. 23, p. 15813, 2022, doi: https://doi.org/10.3390/su142315813.