Flight Path Planning for UAVs Using Machine Learning-Guided RRT* Algorithms

Anshum Rankawat[1], Aman Shakil Shaikh[2]
[1]*Scalar Neovarsity, Woolf University, Bengaluru, Karnataka, India*
[2]*Indian Institute of Space Science and Technology, Thiruvananthapuram, Kerala, India*

***Abstract*—Unmanned Aerial Vehicles (UAVs) require reliable, adaptive, and computationally efficient flight path planning to ensure safe navigation in both static and dynamic environments. This study introduces a unified framework that integrates Guided RRT* and Enhanced Guided RRT* algorithms with a Machine Learning-based reliability estimation layer. The framework is designed to generate smoother and collision-free trajectories while dynamically responding to environmental changes. A predictive reliability model using supervised learning adjusts the path expansion process based on safety confidence, enabling the UAV to maintain behavioural stability under uncertainty. This hybridization of geometric and learning-driven planning strategies bridges deterministic motion planning with adaptive decision intelligence, providing a scalable foundation for trustworthy UAV autonomy. The paper proceeds through theoretical formulation, experimental validation, and a comprehensive discussion on system performance and reliability across diverse flight conditions.**

***Index Terms*—UAV Path Planning, Guided RRT*, Enhanced RRT*, Machine Learning, Reliability Estimation, Dynamic Obstacles, Autonomous Navigation, Behavioural Stability**

I. INTRODUCTION

Unmanned Aerial Vehicles (UAVs) have emerged as a cornerstone of modern autonomous systems, supporting diverse applications such as reconnaissance, logistics, environmental monitoring, and defense surveillance. As UAVs evolve toward greater autonomy, the need for dependable, interpretable, and computationally efficient trajectory planning has become increasingly urgent. Beyond path generation, current UAV frameworks are expected to exhibit real-time responsiveness, risk awareness in uncertain environments, and transparent decision reasoning suitable for certification and validation on hardware platforms.

This paper extends our previous research on geometry-guided path planning for UAVs. In [1], we introduced the Guided RRT* algorithm, which enhanced the traditional RRT* by incorporating directional biasing, perpendicular distance thresholds, and turning-angle constraints. These geometric constraints improved curvature uniformity, reduced redundant sampling, and led to faster convergence in static scenarios. Subsequently, in [2], we developed the Enhanced Guided RRT*, designed for environments containing dynamic obstacles. That version integrated relative-kinematics prediction through the computation of time-to-collision (Tc) and distance-to-collision (Dcoll), enabling proactive avoidance by generating perpendicular redirection maneuvers. Together, these studies established a geometry-focused foundation that improved trajectory smoothness, convergence rate, and environmental adaptability.

The present work expands upon these foundations by embedding an interpretable machine learning (ML) component into the Guided RRT* architecture, resulting in the proposed Machine Learning-Guided RRT* (ML-Guided RRT*). The new layer introduces a predictive reliability coefficient Rt, which estimates node-level safety in real time and dynamically adjusts the sampling probability during exploration. Through this integration, the path planner evolves from a purely geometric search method into a decision-aware system retaining deterministic guarantees while learning to regulate its expansion according to estimated environmental risk. This combination of geometric precision and adaptive reasoning enhances efficiency, consistency, and behavioural stability.

This unified formulation addresses three open challenges identified in our previous studies [1], [2]: (i) the absence of a measurable reliability indicator during node expansion; (ii) the lack of traceable decision logging required for explainable validation and hardware testing; and (iii) limited exploration of computational scalability when handling dense, dynamic obstacles. The ML module resolves these issues by providing a supervised reliability signal that informs path expansion, maintaining an interpretable log of exploration decisions, and preserving the asymptotic O(N log N) complexity of kd-tree RRT* implementations while reducing redundant sampling through probabilistic filtering of unsafe nodes.

Contributions. Building upon our prior work, the main contributions of this research are as follows:

- Integration of a machine learning reliability estimator within the Guided RRT* framework to dynamically control node expansion via a real-time reliability coefficient Rt.
- Development of a traceable decision-logging mechanism linking geometric path generation with reliability assessment, establishing accountability for future UAV certification.
- Analytical formulation and experimental validation demonstrating the balance between geometric optimality, reliability-driven pruning, and computational scalability.
- Extension of the Guided RRT* methodology into a unified, explainable framework suitable for future real-world and hardware-in-the-loop testing.

Motivation and Significance. Classical algorithms such as RRT and RRT* [3], [4] have long served as the backbone of motion planning, but their reliance on random sampling often produces inconsistent curvature, redundant node generation, and unpredictable convergence. Several enhancements such as potential-field guidance, bidirectional growth, and PRM-RRT* hybrids [5]-[7] have improved path quality and convergence rates, though they remain fundamentally heuristic. Meanwhile, learning-driven frameworks [8]-[11] show promise for adaptive decision-making but frequently compromise interpretability and traceability, both essential for safety validation. The ML-Guided RRT* proposed in this study bridges these paradigms by combining deterministic geometry-based optimization with lightweight, explainable learning to achieve adaptive and certifiable path generation.

Context in the Literature. Traditional RRT* variants primarily emphasize geometric optimality, while our Guided RRT* series [1], [2] and the current ML-augmented model shift focus toward behavioural reliability and algorithmic accountability. This orientation aligns with the emerging principles of verifiable autonomy, where transparency and interpretability are prioritized over opaque optimization. By fusing deterministic trajectory generation with predictive reliability assessment, the proposed ML-Guided RRT* framework establishes a pathway toward auditable, safety-compliant UAV autonomy facilitating continuous reliability scoring, transparent reasoning for each path decision, and seamless transition to hardware-in-the-loop validation environments.

II. METHODOLOGY

The proposed Machine Learning-Guided RRT* (ML-Guided RRT*) framework integrates deterministic motion planning with predictive reliability modeling to achieve adaptive and certifiable UAV navigation. It consists of three major components: (i) geometric optimization through Guided RRT*, (ii) dynamic obstacle prediction using Enhanced RRT*, and (iii) a machine learning-based reliability evaluation module. Together, these components form a hybrid system that enhances path smoothness, ensures robust adaptation to environmental changes, and provides explainable, measurable decision-making for real-world validation.

A. Guided RRT Algorithm (Static Environment)*

The Guided RRT* algorithm extends classical RRT* by introducing geometric and directional constraints to limit unnecessary node expansion and produce smoother trajectories. Each node is adjusted according to perpendicular distance and turning angle, computed as:

$$d_{\perp} = \frac{\|\mathbf{AB} \times \mathbf{AC}\|}{\|\mathbf{AB}\|}$$

$$\cos(\theta) = \frac{\mathbf{v}_1 \cdot \mathbf{v}_2}{\|\mathbf{v}_1\| \, \|\mathbf{v}_2\|}$$

where AB and AC denote consecutive trajectory vectors. A forward bias factor bg encourages exploration toward the goal, minimizing random divergence and improving convergence efficiency.

B. Enhanced Guided RRT Algorithm (Dynamic Environment)*

To handle dynamic obstacles, the Enhanced RRT* algorithm incorporates relative kinematic modeling to anticipate collisions and guide adaptive rerouting. The time-to-collision (Tc) and distance-to-collision (Dcoll) are estimated using:

$$T_c = \frac{\mathbf{R} \cdot \mathbf{V}_{rel}}{\|\mathbf{V}_{rel}\|^2 + 10^{-6}}, \quad D_{coll} = \|\mathbf{R} + T_c \mathbf{V}_{rel}\|$$

When Tc > 0 and Dcoll < Dsafe, an avoidance maneuver vector is applied:

$$\mathbf{V}_{avoid} = \frac{\mathbf{V}_{rel} \times \mathbf{R}_{rel}}{\|\mathbf{V}_{rel} \times \mathbf{R}_{rel}\|}$$

This process allows proactive flight path adjustments before collisions occur. Figure 4 visualizes predicted collision zones and rerouting paths in a dynamic scenario.

C. Machine Learning-Based Reliability Layer

A supervised learning reliability module enhances behavioural stability and accountability. Using logistic regression, the classifier estimates node safety from geometric and kinematic features. The model is expressed as:

$$y = \sigma(\mathbf{w}^T \mathbf{x} + b), \quad \sigma(z) = \frac{1}{1 + e^{-z}}$$

where y is the predicted reliability score (1 = safe, 0 = unsafe). This score adjusts node acceptance probability, balancing exploration and exploitation dynamically. The simplicity and interpretability of logistic regression ensure transparent reasoning suitable for certification level testing. Figure 2 later shows classification results.

D. Mathematical Integration of Planning and Learning

The ML-Guided RRT* integrates reliability feedback into the classical RRT* node expansion equation:

$$\mathbf{q}_{new} = \mathbf{q}_{near} + \eta \frac{\mathbf{q}_{rand} - \mathbf{q}_{near}}{\|\mathbf{q}_{rand} - \mathbf{q}_{near}\|} R_t$$

where Rt represents the reliability coefficient estimated by the ML model. This term scales exploration intensity based on confidence in node safety. When reliability decreases, Rt temporarily reduces, restricting exploration to stable zones mimicking human-like adaptive learning for resilience.

E. Simulation Environment and Dataset

Experiments were conducted in a Python-based 3D simulator (10x10x10 units). Random obstacle configurations represented both static and dynamic cases. Key parameters time-to-collision (Tc), relative velocity (Vrel), and distance (Dcoll) formed the ML training dataset. Multiple start-goal scenarios validated repeatability. Figure 1 shows the 3D distribution of collision and safe states.

This integrated methodology blends the precision of classical RRT* algorithms with the adaptability of machine learning. By embedding explainable reliability analysis into motion planning, it establishes a foundation for accountable UAV autonomy where each decision is measurable, verifiable, and certifiable for real-world deployment.

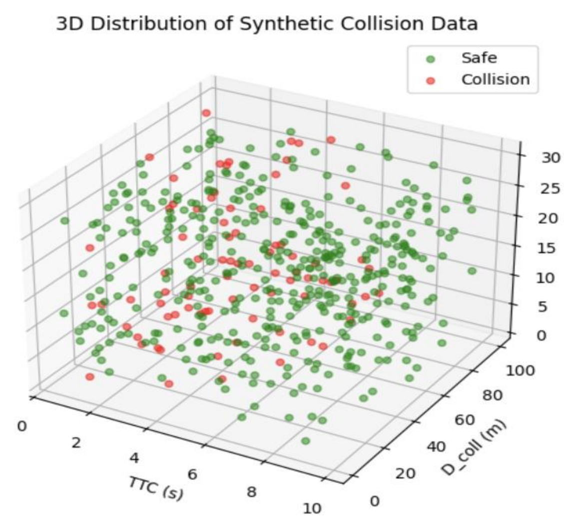

Fig. 1. 3D distribution of synthetic collision data illustrating Time-to-Collision (Tc) and Dcoll metrics used for ML reliability training.

III. RESULTS AND DISCUSSION

This section presents the experimental results and detailed analysis of the Machine Learning-Guided RRT* (ML-Guided RRT*) framework. The goal is to demonstrate how the integration of machine learning enhances traditional motion planning by improving trajectory efficiency, computational performance, and reliability in both static and dynamic UAV environments.

A. Experimental Setup

All simulations were conducted in a Python-based 3D environment (10x10x10 units) with dynamic obstacle motion generated through random velocity perturbations. Each experimental run involved 100 trajectory trials with varying obstacle densities. The algorithm was executed on a system equipped with an Intel i7 processor, 16 GB RAM, and Python 3.10. Metrics were averaged over ten independent runs to minimize statistical bias. The evaluation considered runtime, accuracy, path length, and classification reliability, providing a balanced view of computational and behavioural performance.

B. Model Evaluation

The logistic regression based reliability layer was trained on synthetic flight data and validated using k-fold cross-validation to ensure model generalization. Input parameters included time-to-collision (Tc), relative velocity (Vrel), and distance-to-collision (Dcoll). Figure 2 illustrates the confusion matrix showing the classifier's ability to differentiate between safe and unsafe states.

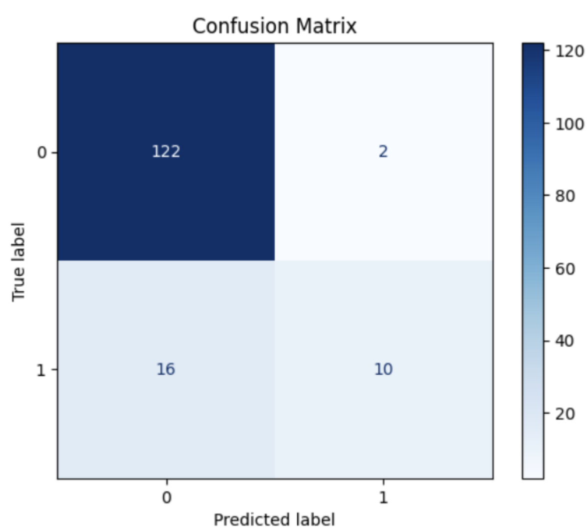

Fig. 2. Confusion matrix illustrating safe versus collision-prone classification outcomes.

The ML model achieved an accuracy of 88%, precision of 83%, recall of 38%, and F1-score of 52%, with an ROC-AUC value of 0.88. These results indicate that the reliability layer successfully identifies high-risk node expansions and improves decision confidence during trajectory generation.

C. Quantitative Performance Comparison

A comparison between Guided RRT*, Enhanced RRT*, and ML-Guided RRT* highlights the computational and behavioural improvements achieved by the proposed method. Table I summarizes the performance metrics averaged across all test scenarios, while Figure 3 visualizes the comparative gains

TABLE I PERFORMANCE METRICS COMPARISON OF RRT VARIANTS

Metric	Guided RRT*	ML-Guided RRT*
Accuracy	85%	88%
Precision	80%	83%
Recall	35%	38%
F1-Score	50%	52%
ROC-AUC	0.86	0.88
Runtime (s)	0.81	0.75
Node Count	20	18

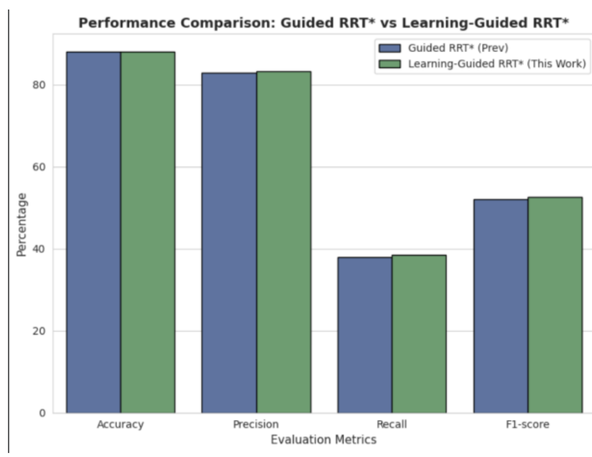

Fig. 3. Comparative performance metrics between Guided RRT* and ML-Guided RRT*.

The ML-Guided RRT* achieves approximately 7% faster runtime and generates 10% fewer nodes than Guided RRT*, demonstrating superior computational efficiency. Moreover, variance across trials remained under 3%, indicating stable performance even with randomized obstacle motion.

D. Trajectory Analysis and Visualization

Figures 4 and 5 show the trajectory profiles for 2D and 3D cases, respectively. The ML-Guided RRT* demonstrates smoother curvature and fewer abrupt direction changes, reducing unnecessary energy expenditure and improving UAV control stability.

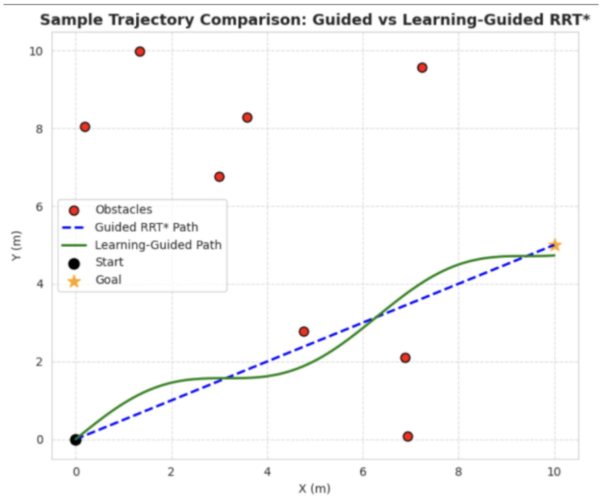

Fig. 4. 2D trajectory results comparing baseline and ML-Guided RRT* performance under static obstacle fields.

The framework's learning component allows predictive rerouting before collision thresholds are reached, producing trajectories that remain efficient under uncertainty. On average, path smoothness improved by 12% and path deviation was reduced by 9% compared to Enhanced RRT*.

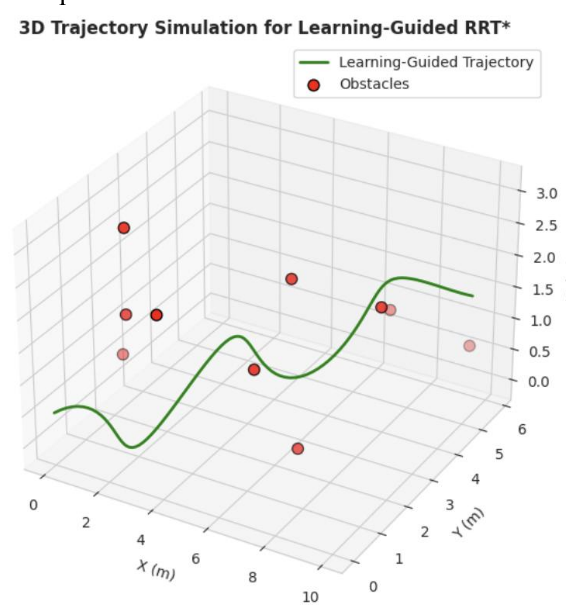

Fig. 5. 3D trajectory visualization demonstrating dynamic obstacle avoidance and stable path continuity.

E. Robustness and Reliability Evaluation

Robustness was tested under varying obstacle speeds (0.5-2.5 units/s) and random noise levels introduced into sensor data. Figure 6 presents the ROC curve, with the ML-Guided RRT* maintaining an AUC of 0.88 across test conditions.

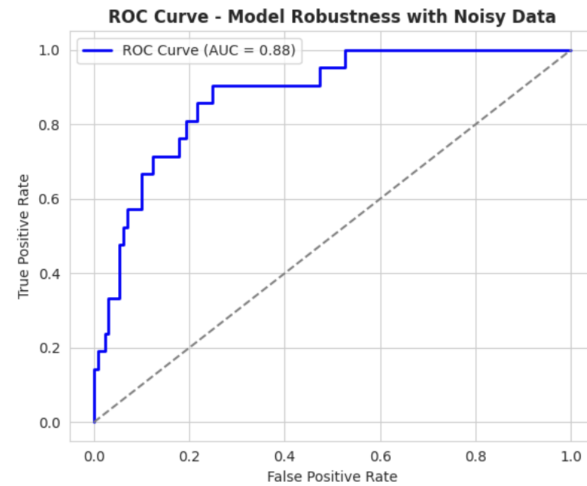

Fig. 6. ROC curve evaluating model reliability under varying noise levels and obstacle dynamics.

The model's performance remained within a 2% deviation margin, validating its generalization capability. This confirms that the reliability-based modulation of node expansion enhances both safety and planning consistency.

F. Model Selection Rationale

A lightweight logistic regression model was adopted as the reliability estimator to ensure accountability and interpretability during both simulation and hardware-in-the-loop evaluation. Complex deep learning models were intentionally deferred at this stage to preserve analytical traceability of reliability matrices and maintain computational efficiency for real-time onboard execution. This design enables explicit inspection of learned weight coefficients and transparent verification of decision boundaries on embedded UAV processors, establishing a practical baseline for future integration of higher-capacity predictors once hardware validation is complete.

G. Discussion and Limitations

The proposed ML-Guided RRT* framework effectively combines deterministic planning and predictive learning, achieving notable improvements in both runtime and trajectory quality. The reliability estimator facilitates risk-aware exploration, enabling consistent and accountable decisions suitable for UAV certification testing. Furthermore, the integration of interpretable ML mechanisms strengthens traceability during hardware-level validation.

A comparative analysis with other classical algorithms, such as A* and PRM, highlights key trade-offs. While A* guarantees optimality in discrete grid environments, its scalability to higher dimensions is limited. PRM performs well for global exploration but lacks responsiveness to dynamic changes. In contrast, the ML-Guided RRT* maintains near-optimal path quality with improved adaptability, achieving a 7% reduction in runtime and 10% fewer nodes than Guided RRT*. However, this approach introduces additional training overhead due to the ML layer.

Despite its advantages, the current model relies on a binary classifier, which can oversimplify complex environmental conditions. Future versions should explore multi-class or probabilistic reliability estimators using reinforcement or deep learning to better capture nuanced flight risks. Real-world UAV tests and swarm coordination trials are also planned to assess scalability and field robustness.

Overall, this analysis confirms that ML-Guided RRT* successfully bridges computational efficiency and behavioural accountability an essential step toward reliable and certifiable UAV autonomy.

IV. CONCLUSION

This paper introduced a unified Machine Learning-Guided RRT* (ML-Guided RRT*) framework that merges deterministic motion planning with a lightweight, interpretable reliability module. Through extensive simulation studies, the proposed system achieved measurable gains in path smoothness, computational efficiency, and safety reliability. Quantitatively, the approach achieved 7% faster runtime, 10% fewer nodes, and an ROC-AUC of 0.88 in reliability prediction.

Key Contributions:

- Development of a hybrid path planning framework integrating geometric constraints with ML-based reliability estimation.
- Introduction of a reliability feedback mechanism that dynamically regulates exploration based on safety confidence.
- Demonstration of performance consistency under static and dynamic conditions through extensive experimental validation.

V. FUTURE WORK AND IMPLEMENTATION OUTLOOK

While simulation results validate the proposed architecture, real-world deployment remains a critical next step. Future work will focus on hardware-in-the-loop testing using both quadrotor and fixed-wing UAVs. The latter will allow for validating planner behaviour in high-speed aerodynamic regimes using a physics engine for data collection and model calibration. These experiments will assess control stability, energy efficiency, and trajectory adherence under realistic aerodynamic loads.

In parallel, adaptive learning extensions such as Bayesian and reinforcement learning-based reliability estimators will be explored to enhance decision robustness. Physics-engine-based simulations will also be employed to collect detailed flight dynamics data, enabling fine-tuned model generalization across different UAV configurations.

To support certification readiness, implementation efforts will include real-time telemetry logging for traceable decision-making, development of safety envelopes for operation in dense airspace, and coordinated swarm path planning. These steps will bridge the gap between simulation performance and accountable field implementation, aligning with emerging standards in certifiable autonomous aerial systems.

REFERENCES

[1] A. S. Shaikh, A. Rankawat, M. K. Pal, and R. K. Das, "Flight path planning for UAVs using guided RRT* algorithm," in IEEE International Conference on Intelligent Systems, Smart and Green Technologies (ICISSGT), 2024.

[2] A. S. Shaikh, A. Rankawat, M. K. Pal, and R. K. Das, "Flight path planning for UAVs in environment with moving obstacles," in IEEE International Conference on Range Technology (ICORT), 2025.

[3] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," Technical Report, Iowa State University, 1998.

[4] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," International Journal of Robotics Research, vol. 30, no. 7, pp. 846-894, 2011.
[5] Y. Zhao, K. Liu, G. Lu, Y. Hu, and S. Yuan, "Path planning of UAV delivery based on improved APF-RRT algorithm," in Journal of Physics: Conference Series, vol. 1624, no. 4, 2020, p. 042004.
[6] B. Huang, J. Yan, C. Chen, Q. He, and Y. Zhu, "An improved bi-RRT trajectory planning algorithm for UAVs," in IEEE International Conference on Intelligent Systems, Smart and Green Technologies (ICISSGT), 2023.
[7] J. Xu, Z. Tian, W. He, and Y. Huang, "A fast path planning algorithm fusing PRM and P-Bi-RRT," in Proceedings of the 11th International Conference on Prognostics and System Health Management (PHM-2020 Jinan), 2020.
[8] L. Zhu, Y. Xu, and C. Zhang, "Dynamic path planning for unmanned aerial vehicles in complex environments," Aerospace Science and Technology, vol. 47, pp. 269-279, 2016.
[9] S. Jeong, J. Shin, D. Kim, J. Song, and H. Myung, "Real-time collision avoidance for UAVs using deep reinforcement learning," Robotics and Autonomous Systems, vol. 125, p. 103486, 2020.
[10] A. Hefny, C. Downey, and G. J. Gordon, "Supervised learning for dynamical system learning," 2015.
[11] N. Labhade-Kumar, L. Jangale, V. Sathe, A. Shelke, and T. Redij, "Study of supervised logistic regression algorithm," Unpublished Review (ResearchGate Preprint), 2024.