ADVANCES
IN GEOMATICS

VOLUME 04 • ISSUE 01 • APRIL 2026
e-ISSN: 3023-4980

ORIGINAL ARTICLE

Submission: *02/02/2026*
Accepted: *27/04/2026*

MACHINE LEARNING-BASED TRAFFIC PREDICTION AND TRAFFIC-AWARE VEHICLE ROUTING OPTIMIZATION

Gökçe TEMEL[1]*

[1] Department of Geomatics, Engineering Faculty, Kocaeli University, Kocaeli, Türkiye.
* Corresponding Author: G. Temel, ✉ gokce.temel.142@gmail.com iD 0009-0002-5996-8043

ABSTRACT

Fixed-speed planning approaches widely used in urban logistics operations lead to unreliable delivery commitments by disregarding the dynamic traffic conditions of large cities. This study aims to develop an integrated pipeline that connects traffic speed prediction directly to road network cost calculation and vehicle routing optimization. A CatBoost gradient boosting algorithm was trained using approximately 19 million rows of 2024 traffic data obtained from the Istanbul Metropolitan Municipality Open Data Portal. The developed M2 model achieved R^2=0.856 in temporal testing, while in the spatial holdout test — where the baseline model collapsed with R^2=–0.340 — it maintained functionality with R^2=0.418, demonstrating that geographic features contributed +0.758 to spatial generalization capacity. Consistent performance in the range of R^2=0.863–0.884 was achieved in independent external validation conducted with data from January 2025. Model predictions were transferred to the Istanbul road network of 506,667 edges via Dijkstra's algorithm and provided as input to the Vehicle Routing Problem with Time Windows solved using OR-Tools. The dynamic validation analysis revealed a critical finding: while the static model predicted reaching 28 out of 30 points on time, simulation of this plan with traffic predictions showed that only 19 points could be reached. This 32% error rate concretely demonstrates that traffic-aware planning is indispensable for operational reliability.

Keywords: Vehicle Routing Problem, Vehicle Routing Problem with Time Windows, Traffic Speed Prediction, CatBoost, Machine Learning, Spatial Generalization.

Cited As:

Temel, G. (2026). Machine Learning-Based Traffic Prediction and Traffic-Aware Vehicle Routing Optimization, Advances in Geomatics, 4(1), 106-122. https://doi.org/10.5281/zenodo.20561364

1. INTRODUCTION

The growing global e-commerce volume and increasing urban population density are transforming city logistics operations into an increasingly complex planning problem. The Vehicle Routing Problem (VRP) constitutes the mathematical framework of this complexity and has remained one of the central research topics in combinatorial optimization literature for decades (Dantzig and Ramser, 1959; Cordeau et al., 2002). However, the majority of existing routing approaches rely on fixed travel time assumptions. This assumption deviates significantly from reality in megacities such as Istanbul, where traffic density varies dramatically on an hourly and regional basis.

Istanbul has a unique transportation dynamic with a population exceeding 15 million and a geographical structure divided in two by the Bosphorus. Systematic congestion at bridge crossings can double travel times during morning peak hours compared to nighttime conditions. Routing systems that exclude this level of variability from the planning process not only produce inefficient routes but also systematically violate delivery commitments made to customers.

In recent years, significant advances have been made in traffic speed prediction with the development of machine learning methods. Boukerche and Wang (2020) systematically reviewed machine learning-based traffic prediction models, while Bratsas et al. (2020) comparatively evaluated the contribution of different algorithms to urban traffic prediction. In the Turkish context, Subaşı (2025) demonstrated that the GWO-XGBoost combination achieves high prediction accuracy on Istanbul traffic data. In the vehicle routing literature, Ichoua et al. (2003) and Figliozzi (2012) examined the effect of time-dependent travel times on routing decisions; however, these studies relied on synthetic speed profiles rather than real data. Among studies directly combining machine learning with VRP, Shahbazian et al. (2024) demonstrated that hybrid systems produce more realistic results compared to static approaches.

Three fundamental gaps are evident when the literature is evaluated. First, traffic prediction studies focus on model accuracy without addressing the conversion of predictions into route cost functions. Second, routing studies largely rely on fixed travel time assumptions or synthetic speed profiles not derived from real data. Third, an integrated pipeline approach combining traffic prediction and vehicle routing within a single decision process remains quite limited in the literature.

This study addresses these gaps with the following contributions. A CatBoost model trained on approximately 19 million rows of real Istanbul traffic data was used as input to route cost calculation. Predictions were transferred to the Istanbul road network of 506,667 edges. A spatial holdout test conducted with geographic features measured the model's generalization capacity at locations it had never seen and demonstrated the operational significance of this capacity. Through dynamic validation, the performance of static plans under traffic conditions was simulated and the operational cost of planning that disregards traffic information was concretely measured.

2. METHODOLOGY

2.1 Data Sources

Three different data sources were utilized in this study. Hourly traffic speed data for 2024 was obtained from the Istanbul Metropolitan Municipality Open Data Portal (IMM, 2024). This dataset consists of approximately 19 million rows containing geographic coordinates, timestamps, and average speed information. January 2025 traffic data was reserved as an external validation set to test the model's generalizability to an independent time period. The Istanbul road network was obtained from OpenStreetMap data through the OSMnx library developed by Boeing (2017), and road type and legal speed limit information were compiled for each road segment in this network consisting of 192,716 nodes and 506,667 edges (OpenStreetMap Contributors, 2024). Hourly average temperature and precipitation data were obtained by averaging the three nearest meteorological stations to Istanbul through the Meteostat library (Meteostat, 2024).

2.2 Traffic Speed Prediction Model

The variable predicted by the model was defined as speed ratio rather than raw speed: SPEED_RATIO = Average Speed / Speed Limit.The ratio was bounded between 0.05 and 2.0. The lower bound of 0.05 prevents division-by-zero errors and eliminates physically implausible near-zero speed recordings likely caused by sensor noise. The upper bound of 2.0 reflects the empirical observation that vehicles rarely exceed twice the posted speed limit under any real-world conditions; values above this threshold were treated as sensor anomalies. In the 2024 dataset, only 0.3% of raw observations exceeded this upper bound, confirming that the clipping operation has negligible impact on the data distribution while eliminating outliers that would otherwise distort model training.

Four categories of features were used. Within temporal features, hour and month variables were encoded using sine-cosine transformation to preserve cyclical continuity (Bratsas et al., 2020). Within geographic features, Haversine distances from each geohash cell to Taksim, Kadıköy, and the nearest Bosphorus bridge were calculated. Within road features, road type obtained from OSM and speed limit information derived through a hybrid strategy were used. Within weather features, hourly temperature and precipitation amounts were incorporated into the model.

Traffic data was encoded using the geohash method developed by Niemeyer (2008) at precision=6 level, corresponding to cells of approximately 1.2 km x 0.6 km. A total of 2,463 unique geohash cells were observed across Istanbul in the 2024 dataset.

Four different models were developed using the CatBoost gradient boosting algorithm developed by Prokhorenkova et al. (2018). CatBoost prevents target leakage by processing categorical variables through the ordered target statistics method and eliminates prediction shift through the ordered boosting mechanism. The M1 baseline model contains only temporal, road, and GEOHASH identity features. The M2 model was constructed by adding geographic distance features to M1. The M3 model was developed by adding weekly lagged observations inspired by Ichoua et al. (2003) and Figliozzi (2012), while the M4 model incorporated hourly lagged observations. The M2 model was selected as the operational model; M3 and M4 require access to historical observations and therefore cannot be used in real operational scenarios such as new locations or system interruptions.

All models were trained with identical hyperparameters: iterations=2000, depth=8, learning_rate=0.05, early_stopping_rounds=200, random_seed=42. This standardization guarantees that performance differences between models arise solely from the feature set.

Model performance was evaluated using two different strategies: temporal and spatial. In the temporal split, the training set covers January–September 2024, the validation set covers October–November 2024, and the test set covers December 2024. In the spatial holdout test, drawing on the spatial generalization requirement emphasized by Boukerche and Wang (2020), 20% of the 2,463 geohash cells were completely excluded from the training set, forcing the model to produce predictions for these locations without having seen them.

2.3 Graph-Based Route Cost

A travel time weight based on the speed prediction produced by the M2 model was assigned to each road network edge:

$$t_{edge} = \frac{L}{\left(\frac{v_{predicted}}{3.6}\right)} \quad (1)$$

Where L represents edge length in meters and vpredicted represents the speed in km/h obtained by multiplying the model's predicted speed ratio by the speed limit. Rather than making separate predictions for 506,667 edges, a geohash-based architecture was adopted; predictions were made at the level of 2,463 unique geohash cells and results were transferred to the corresponding edges. The shortest-time routes between point pairs were determined using Dijkstra's (1959) algorithm and the resulting travel times formed the time matrix provided as input to the VRPTW solver. The static model used for comparison assigns a fixed speed of 45 km/h to all edges.

2.4 Vehicle Routing Problem with Time Windows

Drawing on the VRPTW framework standardized by Solomon (1987), the problem was modeled as an optimization problem in which three vehicles departing from the Yenibosna depot visit 30 customer points. A service time of 15 minutes and an operation window of 09:00–19:00 were defined at each point. Customer points were selected to cover both sides of Istanbul; points close to the depot received early and narrow windows, while distant points received late windows. Three different scales containing 5, 15, and 30 customer points were compared to examine how the advantage of traffic-aware planning varies with problem size. The 5 and 15-point scenarios are subsets of the 30-point scenario, isolating the scale effect independently from point selection. Additionally, the most complex scenario — the 30-point problem — was run for two different dates, weekday and weekend, to concretize the traffic difference between working days and holidays. Optimization was performed using the Google OR-Tools library developed by Perron and Furnon (2023); the initial solution produced by the PATH_CHEAPEST_ARC strategy was improved using the Guided Local Search metaheuristic proposed by Voudouris and Tsang (1999).

2.5 Dynamic Validation

Dynamic validation is the process of re-evaluating the route plan determined by the static model using M2 model traffic predictions instead of the fixed speed assumption. Inspired by the hybrid simulation approach proposed by Shahbazian et al. (2024), the route sequence determined by the static model was kept fixed while the travel time of each edge was recalculated using the speed prediction produced by the M2 model for that hour. The arrival time at each customer point was calculated cumulatively and compared with the time window. This analysis does not use actual sensor data, but it references M2 estimates validated with January 2025 data.

The use of M2 predictions rather than observed sensor data for dynamic validation is a deliberate methodological choice driven by data availability constraints. The 2025 sensor dataset does not provide uniform spatial coverage across all 30 customer points in the scenario; Istanbul's sensor network is concentrated on major arterials and central districts, leaving peripheral locations without direct measurements. M2 predictions, by contrast, can generate travel time estimates for any location within the road network by leveraging geographic proximity features, thereby enabling a consistent and spatially complete simulation environment. While this introduces model-based uncertainty into the validation, the M2 model's demonstrated accuracy of $R^2=0.863–0.884$ on January 2025 sensor data provides a reasonable empirical basis for its use as a reference.

3. RESULTS

3.1 Model Performance

Four models were evaluated on the December 2024 test set. Results are presented in Table 1. Although only geographic distance features were added in the transition from M1 to M2, R^2 increased from 0.820 to 0.856 and RMSE dropped from 7.69 to 6.82 km/h. The additional gain achieved by M3 and M4 with historical observations is an expected result; traffic conditions from the same hour of the previous week serve as a strong predictor. However, these models were not used in this study due to operational constraints. The standard deviation ratio calculated to measure the model's tendency toward mean regression was found to be 0.910 for M2, indicating that the model can represent the extreme conditions of traffic density with sufficient variation.

Table 1. Performance of M1–M4 models on the December 2024 test set.

Model	R^2	RMSE (km/h)	MAE (km/h)
M1 Baseline	0.820	7.69	5.26
M2 Geographic	0.856	6.82	4.58
M3 Historical	0.895	5.85	3.75
M4 Real-Time	0.901	5.69	3.68

The M1 model failed completely in the spatial holdout test with $R^2=-0.340$. This result clearly demonstrates that M1 learns location information only through the geohash identity variable and cannot produce meaningful predictions for a geohash it has not seen during training. The M2 model achieved $R^2=0.418$ in the same test. The spatial generalization capacity difference between the M1 and M2 models was observed as 0.758 (Figure 1). This finding indicates that the integration of geographic features substantially enhances the model's capability to generalize across different spatial zones within the studied network. In real-world logistics scenarios, such as a vehicle fleet expanding into new depots or areas with limited sensor coverage, the M2 model demonstrates the potential to mitigate spatial prediction degradation by leveraging location-based relationships, thereby providing more robust estimations than the baseline model.

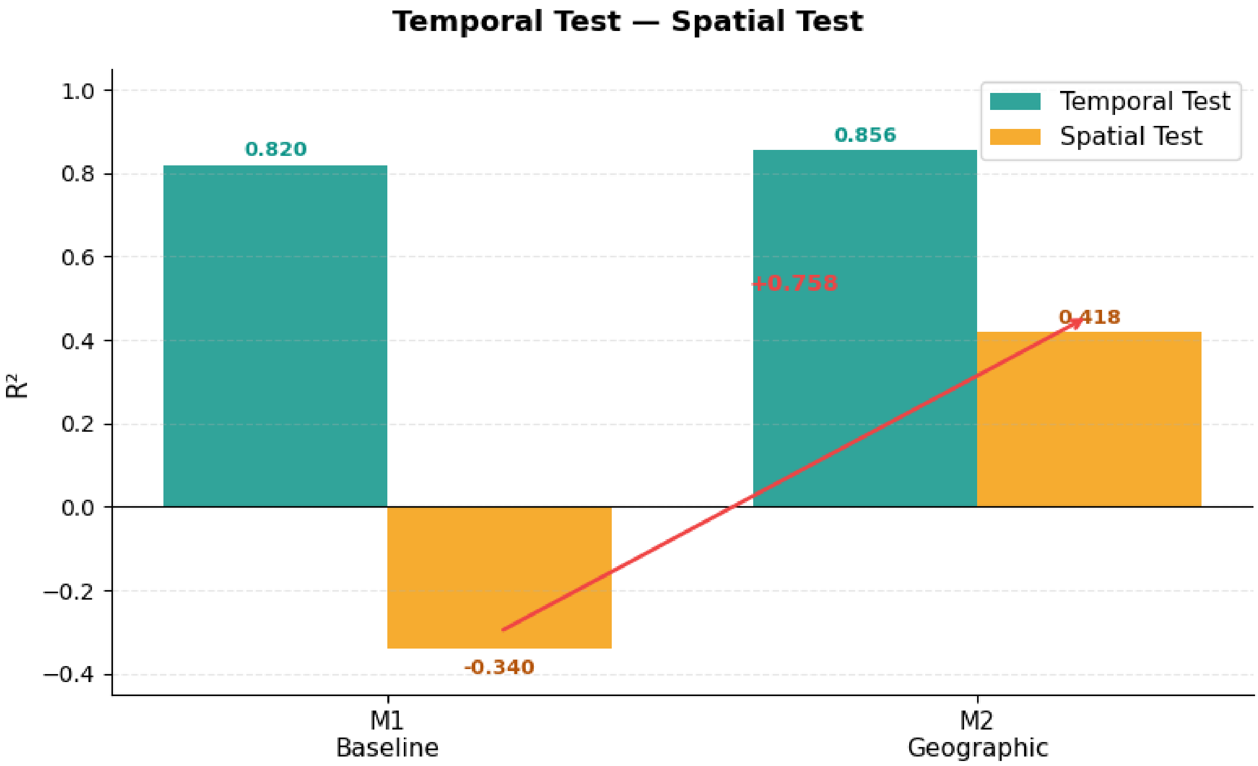

Figure 1. Temporal and spatial validation comparison of M1 and M2 models.

The external validation of the model with January 2025 data is presented in Figure 2. R^2=0.863 was achieved during morning peak hours (08:00), R^2=0.884 at midday (12:00), and R^2=0.847 at nighttime (23:00). The preservation of this performance on data from January 2025 demonstrates that the model has strong generalization capacity against seasonal traffic variability. The highest accuracy at midday reflects the relatively stable nature of traffic patterns at this hour, while the relatively lower performance at nighttime reflects the more variable character of nighttime traffic.

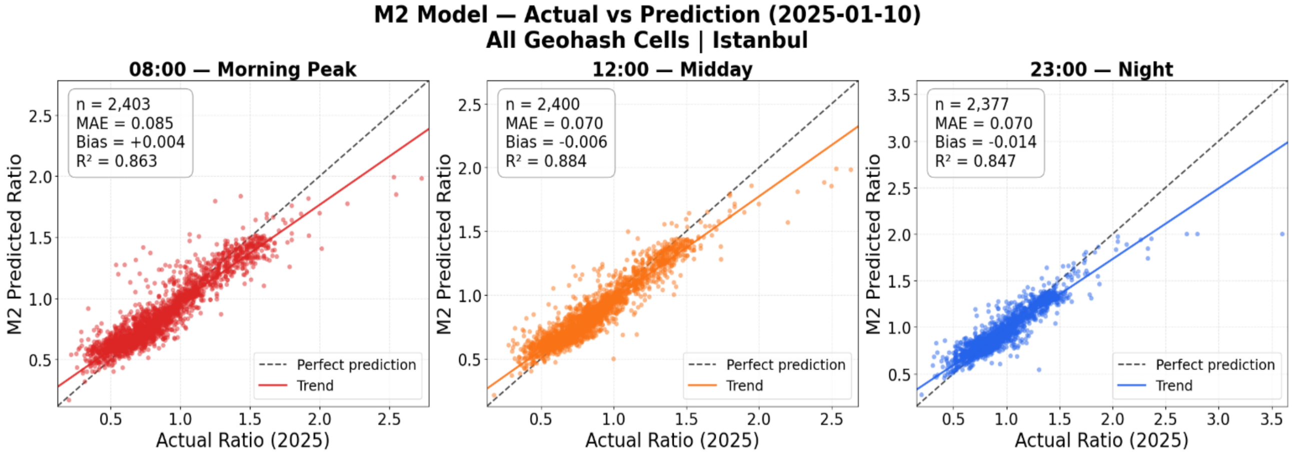

Figure 2. Comparison of the M2 model with real data in geohash cells on 2025-01-10.

Additionally, point-based validation was conducted at the Taksim (sxk97m) and Kadıköy (sxk9hq) geohash cells to evaluate whether the model captures the daily traffic pattern. Results are presented in Figure 3.

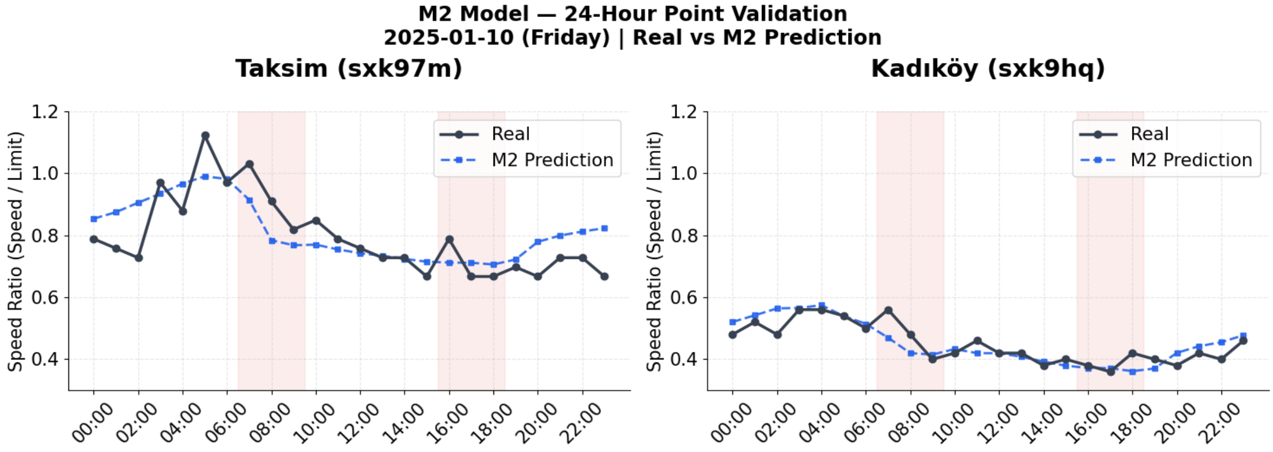

Figure 3. 24 Hour point validation.

When the figure is examined, it can be seen that the model accurately captures the daily traffic pattern at both points. The decrease in speed during morning and evening rush hours and the increase during nighttime hours were predicted by the model.

3.2 Segment Analysis

The two models diverged significantly in the comparative analysis conducted on the Esenyurt–Üsküdar route. As shown in Table 2, dynamic validation simulating the route that the static model planned at 53.7 minutes in the M2 environment revealed that it actually takes 64.8 minutes. The M2 model produced a more realistic plan of 59.1 minutes under the same conditions. The comparative map is presented in Figure 4.

Table 2. Esenyurt–Üsküdar route comparison at 08:00 and 23:00.

Scenario	Model	Planned Duration	Distance	Actual Duration	Error
08:00	M2	59.1 min	40.6 km	—	—
08:00	Static	53.7 min	40.3 km	64.8 min	+11.1 min
23:00	M2	45.5 min	41.5 km	—	—
23:00	Static	53.7 min	40.3 km	53.9 min	+0.2 min

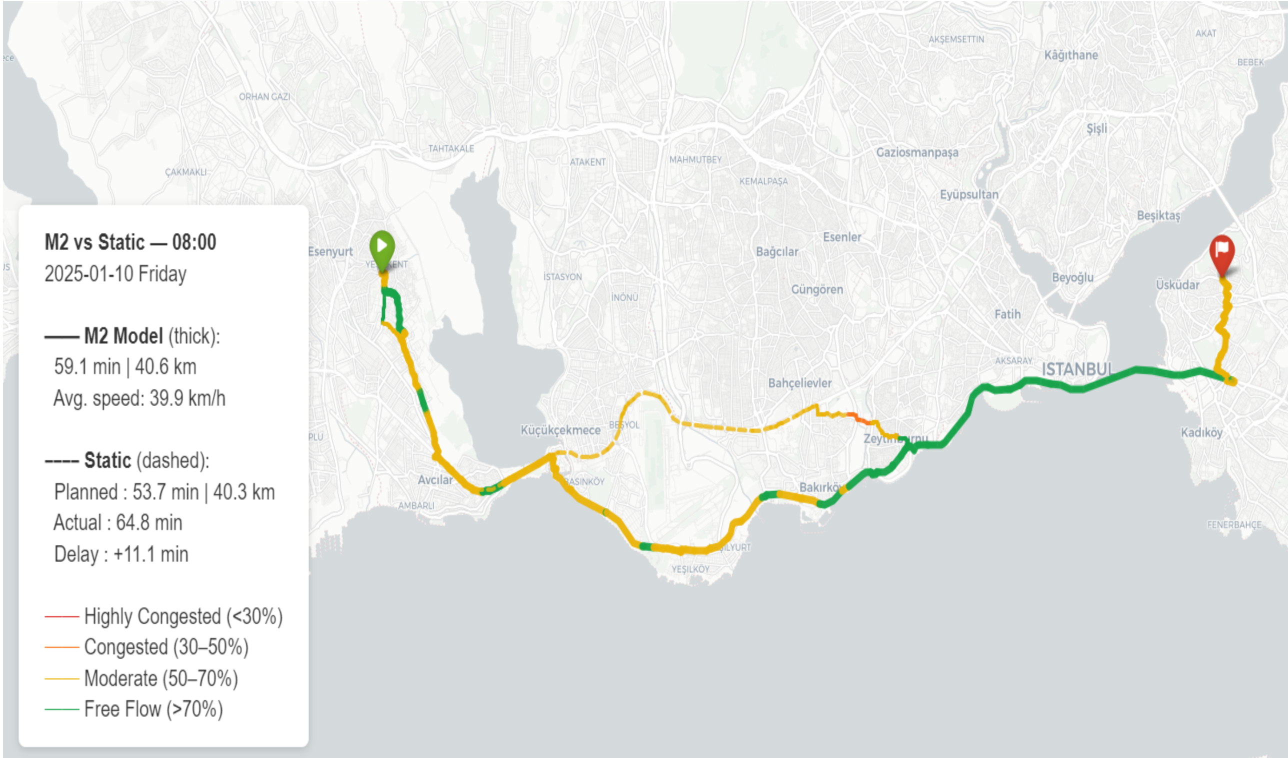

Figure 4. Segment analysis map at 08:00.

Under nighttime 23:00 conditions, the route changes substantially (Figure 5). While the M2 model produces a plan of 45.5 minutes by leveraging traffic fluidity, the static model again yields the same 53.7 minutes. Dynamic validation measured the static plan's duration in the M2 environment at nighttime as 53.9 minutes; this 0.2-minute deviation demonstrates that the fixed speed assumption converges to reality during nighttime hours. It was observed that the two models propose not only different durations but also different routes, demonstrating that traffic information contributes to planning not only in terms of duration but also in terms of geographic decision-making.

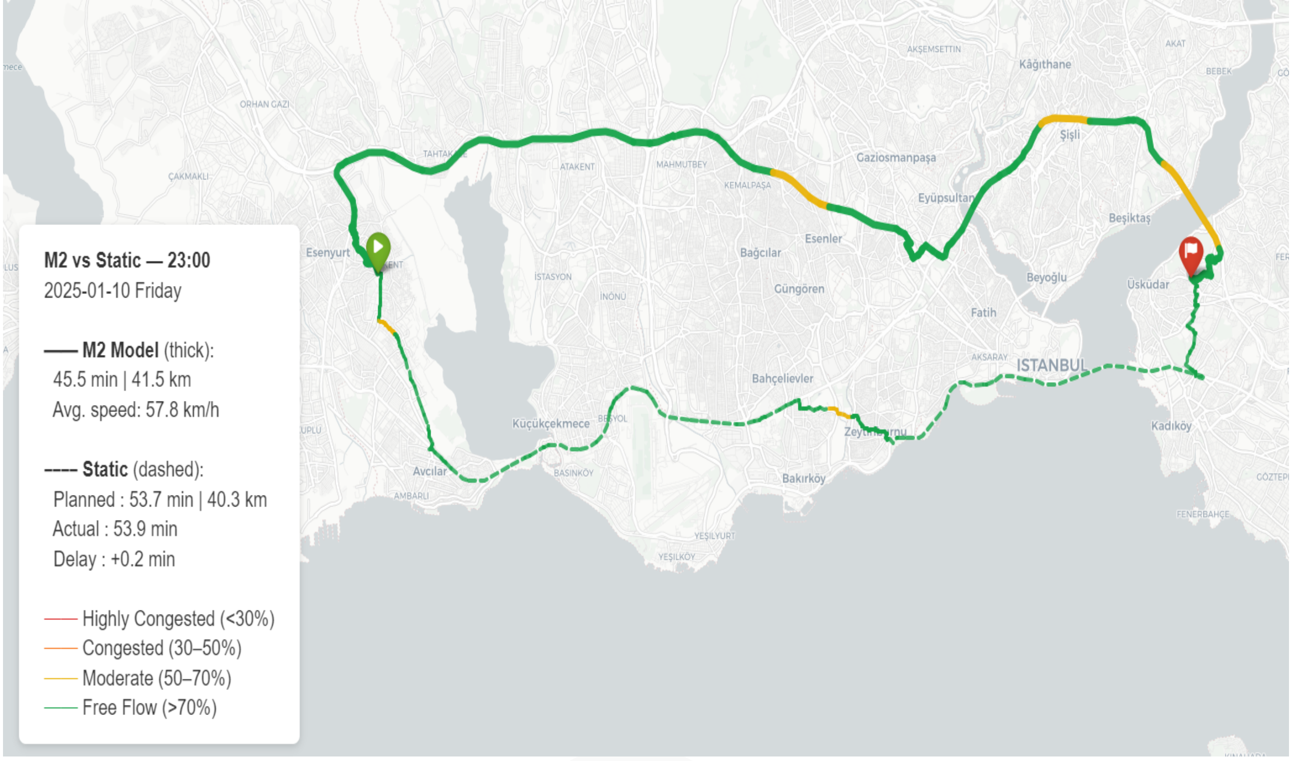

Figure 5. Segment analysis map at 23:00.

3.3 Vehicle Routing Results

Scale analysis results conducted with soft time windows under morning 08:00 peak conditions are presented in Table 3.

Table 3. Comparison of M2 and static models at different problem scales (Soft window, 08:00 am).

Scale	M2 (min)	Static (min)	Difference (min)	Difference (%)
5 points	280.5	259.5	+21.0	+8.1
15 points	638.9	543.8	+95.1	+17.5
30 points	1035.8	917.4	+118.4	+12.9

The absolute difference between M2 and the static model increases as problem scale grows. The 21-minute difference in the 5-point scenario reaches 118 minutes in the 30-point scenario. Each additional customer point adds a new travel segment and each segment contributes to cumulative delay.

Hour-based analysis results for the 30-customer-point scenario are presented in Table 4.

Table 4. Hour-based soft window comparison for 30 point scenario.

Hour	M2 (min)	Static (min)	Difference (min)	Difference (%)
08:00	1035.8	917.4	+118.4	+12.9
12:00	1011.9	917.4	+94.5	+10.3
23:00	916.0	917.4	−1.4	−0.2

The static model producing the same total duration at all three hours is the inevitable consequence of the fixed speed assumption. The −1.4- minute difference at 23:00 demonstrates that the M2 model is capable of not only cautious planning during peak hours but also efficient planning during fluid traffic conditions.

Under hard time window conditions, the static model plans to reach 28/30 points on time while the M2 model plans 26/30. The distribution of the 4 points excluded from the route in the M2 model is shown in Figure 6. Dynamic validation results obtained by simulating the static plan in the M2 environment are presented in Table 5.

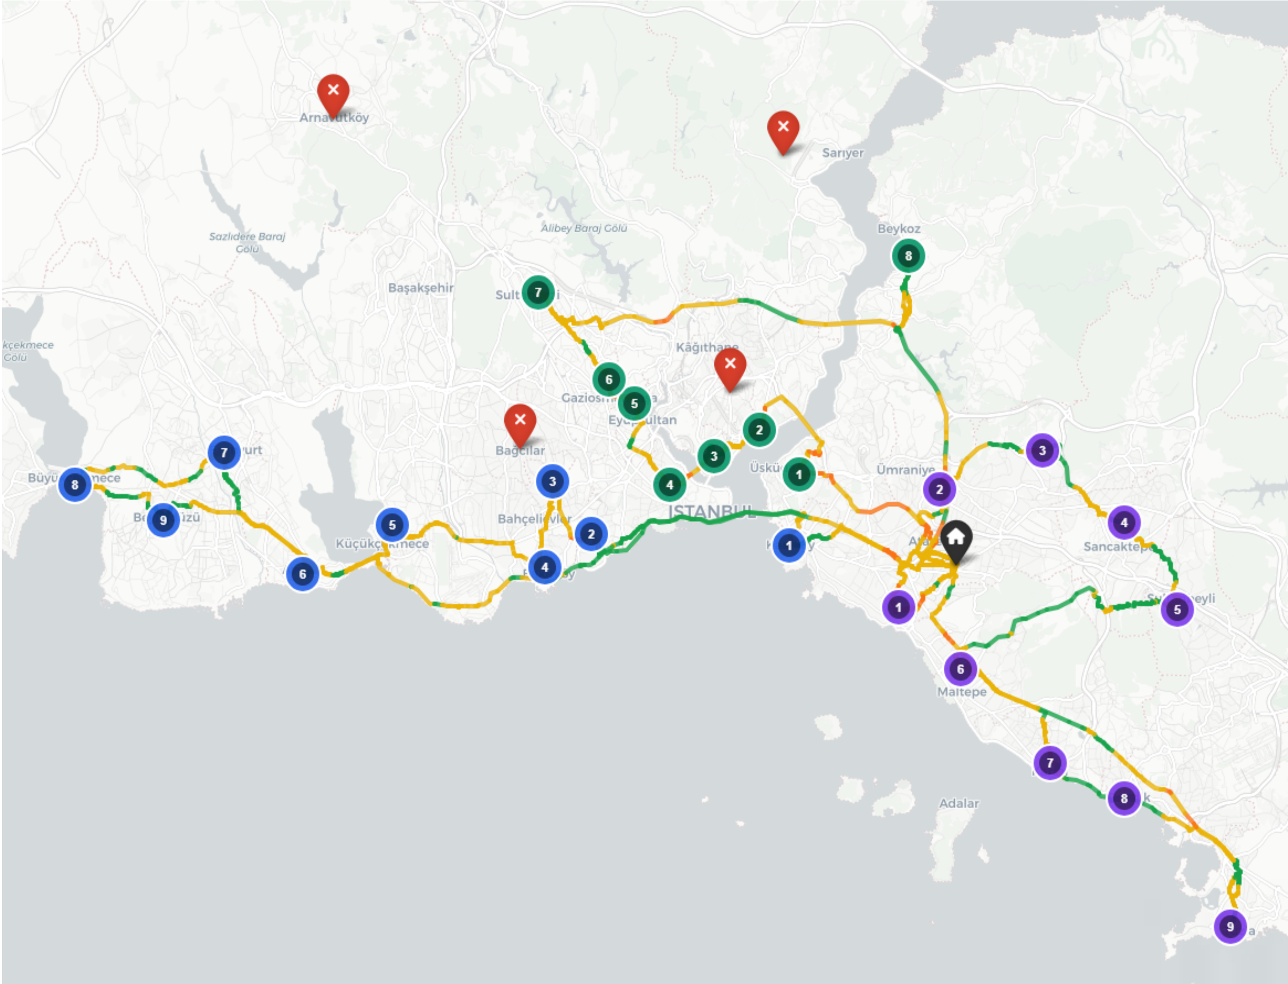

Figure 6. Solution of the M2 model.

Table 5. Dynamic validation of static plan in M2 environment.

Hour	On Time	Missed	Total Delay (min)
08:00	19/30	11/30	1236
12:00	19/30	11/30	1214
23:00	19/30	11/30	975

The difference between the 28 planned by the static model and the 19 actually achieved corresponds to a 32% error rate. The highest delays were observed at bridge-dependent points — Üsküdar, Ümraniye, and Beşiktaş. The lower total delay in the nighttime scenario confirms that the static model's error diminishes as traffic load decreases. The route on which the static plan was applied within the scope of dynamic validation is presented in Figure 7.

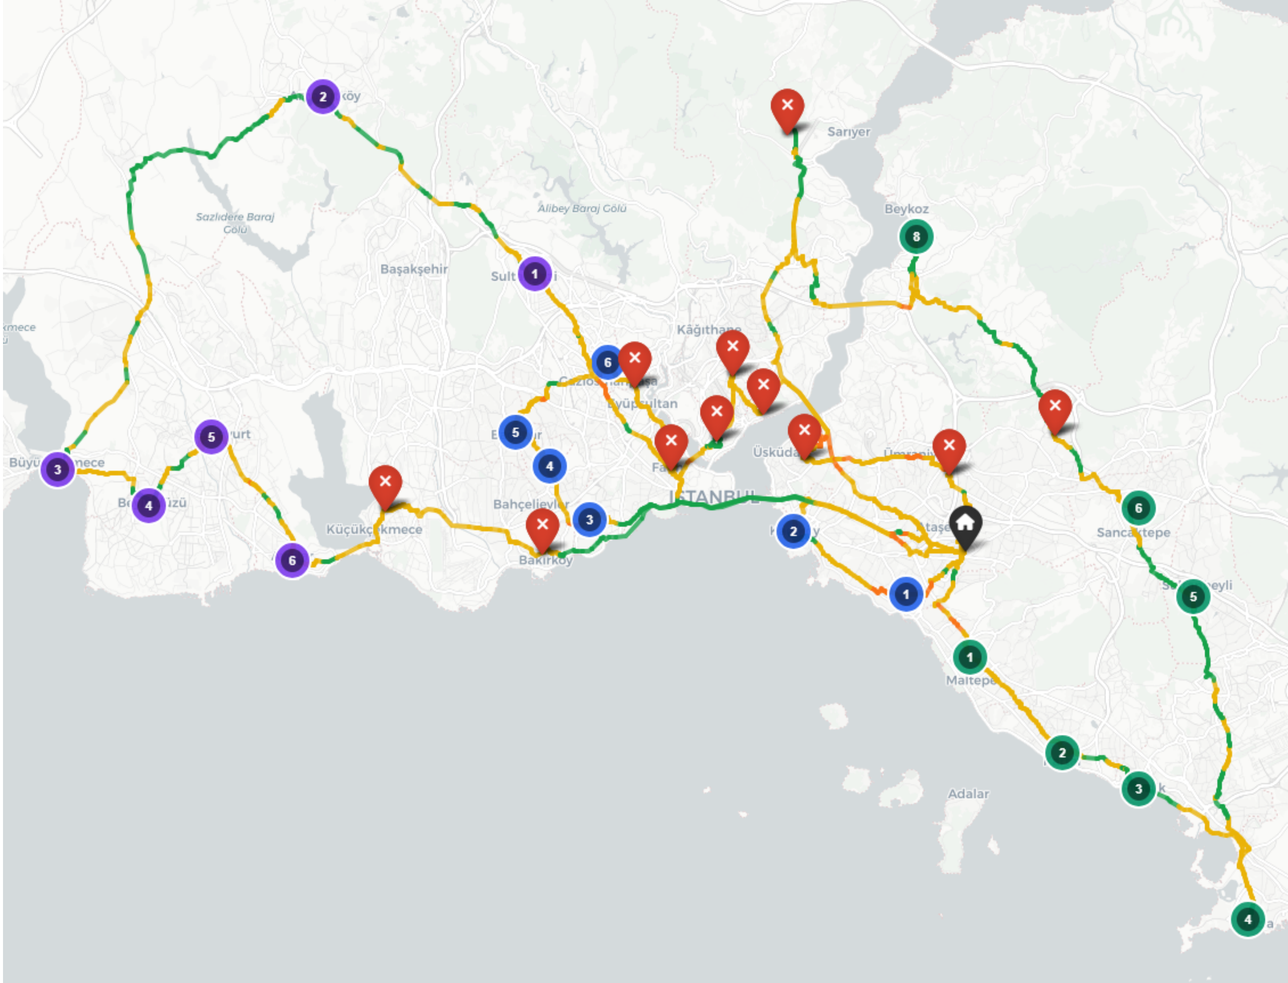

Figure 7. Implementation of the static plan in M2 traffic.

The comparison conducted for Friday (2025-01-10) and Saturday (2025-01-11) revealed that the advantage provided by the M2 model is directly related to traffic density. The M2–Static difference of 12.9% on Friday morning dropped to 3.5% on Saturday morning. When dynamic validation results are examined, it is observed that the number of missed points remained at similar levels for both days, while total delay dropped significantly on Saturday: the total delay of 1,236 minutes on Friday morning fell to 797 minutes on Saturday morning. Four comparative charts are presented together in Figure 8.

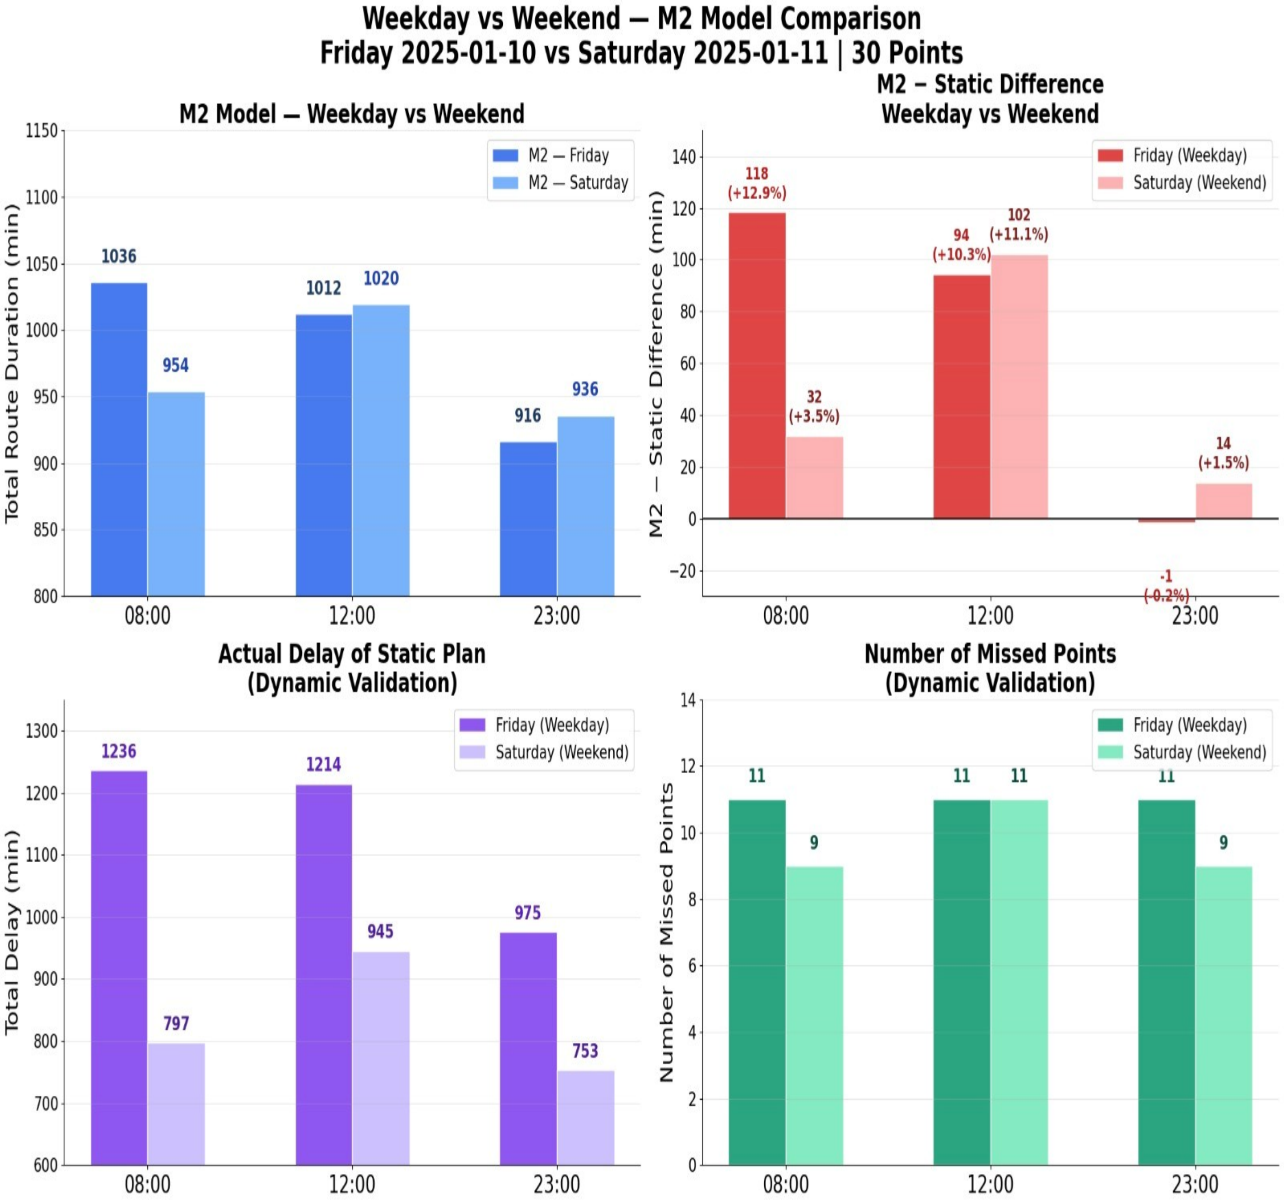

Figure 8. Weekday vs. weekend model comparison.

4. DISCUSSION

The findings obtained in this study concretize the impact of a pipeline integrating traffic prediction with vehicle routing optimization at three different levels.

4.1 Evaluation of Model Performance

When model performance is evaluated, achieving $R^2=0.856$ accuracy with the M2 model is a strong result when considered in isolation. However, what is truly decisive in this study is not the absolute accuracy value but the spatial generalization capacity that geographic features contribute to the model. While the temporal test difference between M1 and M2 is only 0.036, the difference in the spatial holdout test reaches +0.758. This finding raises a distinction frequently overlooked in traffic prediction literature: a good temporal test score does not guarantee that the model will be functional at new locations. Boukerche and Wang (2020) emphasized that spatial generalization capacity is a critical criterion in evaluating machine learning-based traffic prediction models; the spatial holdout test findings of this study concretely confirm this emphasis.

The higher accuracy achieved by M3 and M4 models is an expected and explainable result. Traffic conditions from the same hour of the previous week serve as a strong predictor in a city like Istanbul with repeating weekly patterns. However, this performance advantage comes with an operational cost. The selection of M2 is therefore not only a technical but a strategic decision. A less accurate model that works under all conditions is operationally superior to a more accurate but fragile model that works only under certain conditions. Shahbazian et al. (2024) used a similar framework when evaluating the operational applicability of hybrid systems and demonstrated that generalization capacity is as decisive as accuracy in model selection.

The preservation of performance in the range of $R^2=0.847$–0.884 with January 2025 data is particularly noteworthy. Subaşı (2025) achieved $R^2=0.8486$ with the GWO-XGBoost combination on Istanbul traffic data. The M2 model in this study remaining within a similar accuracy range with data from January 2025 confirms that gradient boosting-based approaches can model Istanbul traffic dynamics in a stable manner.

4.2 Evaluation of Routing Findings

When routing findings are evaluated, the fact that the M2 model produces longer plans than the static model under peak hour conditions appears to be a disadvantage at first glance. However, this impression reflects a common misconception. The dynamic validation analysis concretely revealed this error. The static model claims to reach 28/30 points under hard time windows, while in reality only 19 can be reached. Figliozzi (2012) quantitatively examined the effect of traffic density on time windows and demonstrated that the fixed speed assumption leads to serious planning errors particularly during morning peak hours. The dynamic validation findings of this study measure the magnitude of this error in the Istanbul context at 32%.

The increase in M2's advantage as problem scale grows is also an important finding. While traffic effects accumulate in a limited number of segments in small-scale operations, in the 30-point scenario each vehicle passes through dozens of segments consecutively and the delay in each segment cumulatively transfers to the final points. This finding is consistent with the cumulative delay effect demonstrated by Li et al. (2010) in their stochastic travel time VRPTW study.

The weekday and weekend comparison is perhaps the analysis that most clearly demonstrates M2's value. The drop from 12.9% on Friday morning to 3.5% on Saturday morning proves that the model's advantage is directly linked to traffic density. Kim et al. (2015), while examining the dynamic effects of traffic conditions in urban routing problems, reached a similar conclusion and emphasized that traffic-aware planning plays a decisive role particularly in weekday morning operations.

4.3 Limitations

The methodological and data-related limitations of this study should be explicitly addressed. The traffic data covers only locations with sensors, and underrepresentation is an issue in Istanbul's peripheral districts. The model was trained only on 2024 data; regular updates will be required to capture traffic dynamics that change over the years. Furthermore, the single depot and homogeneous fleet assumptions do not fully reflect the complexity of real logistics operations.

Dynamic validation in this study is based on model-predicted travel times rather than observed travel times from real sensors. While the M2 model was validated against January 2025 sensor data at $R^2=0.863–0.884$, predicted travel times are inherently smooth estimates that may underrepresent the stochastic variability of real traffic conditions — sudden incidents, road closures, or weather events that are not captured in the training data. Consequently, the reported error rates and delay figures should be interpreted as conservative estimates; the actual performance gap between static and traffic-aware planning under real operational conditions may be larger than reported here.

5. CONCLUSION

This study combined a machine learning prediction system modeling Istanbul-specific traffic dynamics with the Vehicle Routing Problem with Time Windows within a single integrated pipeline. Travel times derived from traffic predictions were transferred to road network edges via Dijkstra's algorithm and converted directly into routing decisions. The effectiveness of the developed framework was evaluated comparatively with a static planning approach based on fixed speed assumptions at the levels of model accuracy, segment analysis, and fleet planning.

At the model level, the most decisive finding relates to spatial generalization capacity. While the baseline model collapsed with R²=–0.340 at locations it had never seen, the M2 model achieved R²=0.418, and the contribution of geographic features to spatial generalization capacity was measured at +0.758. Performance in the range of R²=0.863–0.884 was preserved in independent external validation conducted with January 2025 data. These results demonstrate that the model has a generalizing rather than memorizing structure and can operate reliably across different time periods and new locations.

At the segment level, the contribution of traffic information to route decisions was concretized. The Esenyurt–Üsküdar route, which the static model planned at 53.7 minutes during morning peak hours, took 64.8 minutes in the M2 environment. Under nighttime conditions, the M2 model leveraged traffic fluidity to fall 8.2 minutes below the static model. It was observed that the two models propose not only different durations but also different routes, demonstrating that traffic information shapes planning not only in terms of duration but also in terms of geographic decision-making.

At the fleet level, the most critical finding of the study was obtained from the dynamic validation analysis. While the static model planned to reach 28 out of 30 points on time under hard time windows, simulating this plan in the M2 environment showed that only 19 points could be reached. This 32% error rate proves that plans produced without traffic information render customer commitments unfulfillable from the outset. This gap widens as problem scale grows and traffic density increases.

When all these findings are evaluated together, the fundamental contribution of this study can be summarized as follows: by developing a workable pipeline bridging traffic prediction and vehicle routing, operational outputs that neither field could produce independently were obtained. In cities like Istanbul that are dependent on bridge crossings and where traffic density changes dramatically on an hourly basis, the operational counterpart of this transformation is concrete and measurable. The integration of traffic information into logistics planning processes should be regarded in this context not merely as a performance improvement but as a prerequisite for being able to make reliable service commitments.

REFERENCES

Boeing, G. (2017). OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. Computers, Environment and Urban Systems, 65, 126-139.

Boukerche, A., & Wang, J. (2020). Machine learning-based traffic prediction models for intelligent transportation systems. Computer Networks, 181, 107530.

Bratsas, C., Koupidis, K., Salanova, J.-M., Giannakopoulos, K., Kaloudis, A., & Aifadopoulou, G. (2020). A comparison of machine learning methods for the prediction of traffic speed in urban places. Sustainability, 12(1), 142.

Cordeau, J. F., Desaulniers, G., Desrosiers, J., Solomon, M. M., & Soumis, F. (2002). VRP with time windows. In P. Toth & D. Vigo (Eds.), The Vehicle Routing Problem (pp. 157-193). Philadelphia: SIAM.

Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. Management Science, 6(1), 80-91.

Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. Numerische Mathematik, 1(1), 269-271.

Figliozzi, M. A. (2012). The time dependent vehicle routing problem with time windows: Benchmark problems, an efficient solution algorithm, and solution characteristics. Transportation Research Part E, 48(3), 616-636.

Ichoua, S., Gendreau, M., & Potvin, J. Y. (2003). Vehicle dispatching with time-dependent travel times. European Journal of Operational Research, 144(2), 379-396.

Istanbul Metropolitan Municipality [IMM]. (2024). Hourly traffic density data set [Dataset]. Retrieved October 2024, from https://data.ibb.gov.tr/dataset/hourly-traffic-density-data-set

Kim, G., Ong, Y. S., Heng, C. K., Tan, P. S., & Zhang, N. A. (2015). City vehicle routing problem (City VRP): A review. IEEE Transactions on Intelligent Transportation Systems, 16(4), 1654-1666.

Li, X., Tian, P., & Leung, S. C. H. (2010). Vehicle routing problems with time windows and stochastic travel and service times: Models and algorithm. International Journal of Production Economics, 125(1), 137-145.

Meteostat. (2024). Meteostat: Historical weather and climate data [Software]. Retrieved October 2024, from https://meteostat.net

Niemeyer, G. (2008). Geohash [Software]. Retrieved October 2024, from https://en.wikipedia.org/wiki/Geohash

OpenStreetMap Contributors. (2024). OpenStreetMap [Geographic database]. Retrieved October 2024, from https://www.openstreetmap.org

Perron, L., & Furnon, V. (2023). OR-Tools [Software]. Google. Retrieved October 2024, from https://developers.google.com/optimization

Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., & Gulin, A. (2018). CatBoost: Unbiased boosting with categorical features. Advances in Neural Information Processing Systems, 31, 6638-6648.

Shahbazian, R., Pugliese, L. D. P., Guerriero, F., & Macrina, G. (2024). Integrating machine learning into vehicle routing problem: Methods and applications. IEEE Access, 12, 93087-93115.

Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. Operations Research, 35(2), 254-265.

Subaşı, N. (2025). Optimization of Istanbul urban traffic data with grey wolf optimization-assisted machine learning techniques: Comparative analysis of long short-term memory and eXtreme gradient boosting. ELECTRICA, 25(1), 1-15.

Voudouris, C., & Tsang, E. P. K. (1999). Guided local search and its application to the traveling salesman problem. European Journal of Operational Research, 113(2), 469-499.