A Data-Driven Framework for Fleet Electrification Planning: Integrating Total Cost of Ownership Analysis for University Facilities Management

Nick DiCintio[1], Tracy Hua[1], Sean Kerrigan[1], Liam Tuohy[1], Emily Spradley[1], Michael Duffy[2], and B. Brian Park[1]
**pvt2fb@virginia.edu, gpz9ad@virginia.edu, njn4gb@virginia.edu, dyc3jf@virginia.edu*
xar6dp@virginia.edu, med7p@virginia.edu, bp6v@virginia.edu

[1] *Department of Systems and Information Engineering*
University of Virginia
Charlottesville, VA, USA

[2] *UVA Facilities Management*
University of Virginia
Charlottesville, VA, USA

***Abstract*—Fleet electrification represents an important strategy for reducing greenhouse gas emissions while achieving long-term cost savings. Large fleets often include vehicles with diverse usage patterns, which makes uniform electrification decisions inefficient and financially risky. Decision makers therefore require analytical tools that account for these variations to evaluate economic and environmental outcomes when planning fleet transitions. This work developed a data-driven decision-support framework to evaluate electrification strategies for UVA Facilities Management (FM), which oversees over 400 vehicles. The framework identified which vehicles were suitable candidates for replacement with plug-in hybrid vehicles or battery electric vehicles and quantified trade-offs between cost and emissions. The methodology combined vehicle-level telematics data with a supporting machine learning based modeling approach that learned patterns in vehicle usage and energy consumption from historical telematics data for future fuel and energy estimates. Sensitivity analysis was used to evaluate uncertainty in key parameters such as fuel prices, electricity rates, charging availability, and vehicle utilization. Results indicated that a subset of high-emitting vehicles, often characterized by repeated excessive idling, contributed disproportionately to overall fleet emissions but were not always suitable candidates for electrification based on their operational profiles. Case studies of real FM replacement decisions demonstrated that while electric and hybrid alternatives offer meaningful emissions reductions of 69–82%, their higher upfront acquisition costs mean the total cost of ownership premium of electrification varies substantially by vehicle utilization, and targeted replacement strategies are more cost-effective than uniform fleet-wide electrification. These findings demonstrate that the framework provides FM with a structured, data-grounded tool to evaluate when electrification is economically justified and when operational emissions reduction should take priority.**

***Keywords*—Fleet electrification, machine learning, telematics data, decision support systems, total cost of ownership, data-driven decision making**

I. Introduction

Transportation emissions represent a significant and growing share of greenhouse gas (GHG) output in the United States, and institutional fleet operators face increasing pressure to reduce their environmental footprint while maintaining operational efficiency and fiscal responsibility. Fleet electrification—replacing internal combustion engine (ICE) vehicles with battery electric vehicles (BEVs) or plug-in hybrid electric vehicles (PHEVs)—presents a promising path toward meeting both sustainability and cost goals. However, the viability of electrification varies substantially across individual vehicles within a fleet, depending on factors such as daily mileage, duty cycle, idle behavior, and available charging infrastructure [1].

The University of Virginia (UVA) Facilities Management (FM) department operates a diverse fleet of over 400 vehicles, comprising 13 diesel, 13 electric, 280 gasoline, and 17 hybrid assets, with gasoline-powered vehicles accounting for approximately 84% of all fleet vehicles. UVA has expressed institutional sustainability commitments that include reducing fleet-related emissions, yet the heterogeneous nature of the fleet—spanning vehicle classes from compact passenger cars to pickup trucks, with usage patterns ranging from campus shuttle routes to field maintenance operations—makes blanket electrification decisions both economically and operationally risky.

Prior work by the National Renewable Energy Laboratory (NREL) has demonstrated the value of telematics-based analysis for university fleet electrification planning [1], and telematics data more broadly have been shown to offer actionable insight into fleet energy efficiency when combined with data science methods [2]. However, existing tools such as NREL's ezEV Suitability Assessment rely primarily on range-compatibility thresholds and are applied prospectively without evaluating whether past decisions delivered expected savings. FM currently lacks a structured method to audit prior vehicle replacements against real operational data or to generate cost projections grounded in observed, vehicle-specific driving behavior rather than manufacturer estimates.

This paper presents a data-driven decision-support framework with two complementary components. The first is a retrospective audit of vehicle replacement decisions already made by UVA FM, evaluating each historical substitution from both an economic or an emissions perspective using actual fuel consumption data drawn from Geotab, a fleet telematics platform. This audit quantifies the cost and carbon savings realized by past transitions to hybrid and electric alternatives and identifies where FM's purchasing decisions have been most and least cost-effective. The second component is a forward-looking prototype that trains a machine learning model on historical telematics records to predict fuel and energy consumption for vehicles not yet in the fleet, enabling prospective TCO comparisons for vehicles under consideration for future replacement. Specifically, the vehicle replacements focused on are Class 1 and Class 2 vehicles, as those classes are typically the ones feasible for electrification. Both Class 1 and Class 2 are grouped as light-duty vehicles, with Class 1 including vehicles up to 6,000 lbs and Class 2 including vehicles from 6,001 to 10,000 lbs [14]. Together, these components form a total cost of ownership (TCO) framework that incorporates acquisition cost, fuel and energy costs, maintenance, insurance, salvage value, and the social cost of carbon (SCC) [5], producing a vehicle-by-vehicle comparison tool that FM leadership can use to evaluate electrification scenarios under varying assumptions.

The remainder of this paper is organized as follows. Section II reviews relevant literature. Section III describes the data sources and fleet composition. Section IV details the TCO framework and the ML modeling approach. Section V presents the retrospective audit findings and ML model results. Section VI discusses implications and limitations, and Section VII concludes the paper.

II. Background

Prior work by NREL established data-driven methods for assessing EV suitability in institutional fleets using GPS and mileage data [2]. A multi-university NREL study applied this approach to six campuses including UVA, analyzing 1.13 million trips and 7.29 million vehicle miles traveled, with UVA contributing 247,000 trips and 930,000 miles [1]. That work and subsequent campus fleet studies [6] demonstrated that university environments—with centralized parking, predictable daily routes, and overnight charging opportunities—are well-suited for electrification, but that economic feasibility depends heavily on vehicle utilization and the degree to which fuel savings offset higher acquisition costs [7].

Prediction of vehicle energy consumption has been approached through analytical, statistical, and machine learning methods. Ullah et al. compared XGBoost, LightGBM, MLR, and ANN for EV energy prediction using GPS data from 68 vehicles across 38,362 trips, finding that gradient boosting methods consistently outperformed statistical and neural network baselines [3], [4]. Feature importance analysis identified speed, trip distance, road gradient, and ambient temperature as the dominant predictors—findings that directly inform the feature design of the ML component in this study. TCO analysis for fleet replacement typically incorporates acquisition cost, fuel cost, maintenance, insurance, and salvage value, discounted over a holding period. This paper extends the standard TCO framework by including the social cost of carbon (SCC), monetized using the interagency methodology described in Greenstone et al. [5].

III. Data

This study grounded all cost and emissions analysis in vehicle-level operational records rather than generic manufacturer estimates, drawing on three categories of data: financial records, fleet telematics, and supplementary external datasets.

A. Financial and Fleet Records

Of the 323 vehicles in the UVA FM fleet, those analyzed fall within Class 1 (GVWR $\leq$6,000 lbs, including passenger cars, light-duty pickups, and minivans) and Class 2 (GVWR 6,001–10,000 lbs, including full-size pickups, SUVs, and vans), which together constitute the light-duty segment of the fleet and the primary candidates for electrification [14]. Of these, 95 had complete financial records—confirmed purchase price, multi-year maintenance cost history, and FY2025 insurance premium—and were used for full TCO modeling (78 ICE, 8 FHEV, 6 BEV, 3 PHEV). Purchase prices not available in FM records were supplemented by Kelly Blue Book, and salvage values were estimated at 20% of original MSRP [15].

B. Fleet Telematics Data

Vehicle operations data were sourced from Geotab, which logged GPS coordinates, instantaneous speed, ambient temperature, and odometer readings at approximately 5-second intervals across the entire fleet. Fuel economy records captured total fuel (liters) and energy (kWh) consumed at roughly five timestamps per day; a separate fill-up log recorded refueling events, and charge event logs captured BEV recharging sessions. A validation pipeline was applied to the raw fuel data, checking for duplicate timestamps, negative values, non-monotonic cumulative totals, and mismatches between interval and cumulative consumption (tolerance: 0.25 liters or kWh).

C. Supplementary Data Sources

Road grade was computed by matching each GPS coordinate to the USGS 1 arc-second digital elevation model ($\approx$30 m resolution) for the Charlottesville region. Ambient temperature was sourced from the NOAA daily min/max dataset recorded at the Leander McCormick Observatory on the UVA grounds (January 2020–present); daytime average temperature was derived using the sinusoidal model of Wann et al. [8]: $T_{\text{avg}} = T_{\min} + \frac{2}{\pi}(T_{\max} - T_{\min})$. Vehicle specifications (curb weight, engine displacement, drag coefficient, factory MPG/MPGe) were collected via an automated API pipeline.

IV. Methodology

A. Data Preprocessing Pipeline

Raw telematics data were ingested and transformed through a multi-step preprocessing pipeline implemented in Python. For each telematics record, the pipeline computed link distance from consecutive GPS coordinate pairs using the Haversine formula, retrieved road elevation from USGS tiles to calculate link grade, assigned daytime average temperature from the

NOAA dataset (using the Wann et al. sinusoidal model [8]), and aggregated all records into observation windows aligned with each vehicle's fuel reporting intervals. Each row in the resulting modeling dataset represents one fuel-reporting interval for one vehicle.

B. Total Cost of Ownership Framework

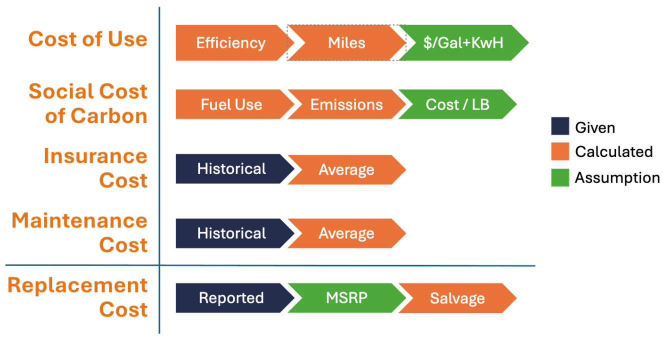

Fig. 1. TCO Methodology Framework

The methodology is centered on the TCO framework, as seen in Fig. 1, which evaluates past and prospective vehicle replacement decisions using observed and predicted fuel consumption data. Machine learning is incorporated as a supporting component to estimate future fuel and energy usage for vehicles not yet in the fleet. The TCO model estimated the net present value (NPV) of the total cost of ownership of the vehicle over a 10-year holding period, discounted at 7%, for both the incumbent vehicle and its proposed replacement. The TCO for a given vehicle is:

$$\text{TCO} = C_{\text{acq}} - C_{\text{salvage}} + \sum_{t=1}^{T} \frac{C_{\text{fuel},t} + C_{\text{maint},t} + C_{\text{ins},t} + C_{\text{carbon},t}}{(1+r)^t} \tag{1}$$

where C_{acq} is acquisition cost, C_{salvage} is the discounted end-of-period salvage value, $C_{\text{fuel},t}$ is annual fuel or electricity cost, $C_{\text{maint},t}$ is annual maintenance, $C_{\text{ins},t}$ is annual insurance, $C_{\text{carbon},t}$ is the annual social cost of carbon emissions, and r is the discount rate.

Fuel cost. For retrospective audits of vehicle replacement decisions already made by FM, annual fuel use was drawn directly from each vehicle's Geotab telematics records, providing observed consumption figures rather than estimates. For prospective evaluation of vehicles not yet in the fleet, annual fuel use is estimated by the trained ML model described in Section IV-C. In both cases, fuel cost is calculated by multiplying annual consumption by the unit cost: $3.20 per gallon of gasoline [9] and $0.1169 per kWh of electricity [10], scaled to 150 driving days per year based on FM's three-day operational pattern.

Carbon cost. The annual carbon cost is computed as:

$$C_{\text{carbon},t} = E_t \cdot f_{\text{CO}_2} \cdot \text{SCC} \tag{2}$$

where E_t is annual fuel or energy consumption, f_{CO_2} is the emissions intensity (0.009796 metric tons CO_2/gal for ICE [11]; 0.000328 metric tons CO_2/kWh for the Virginia grid [12]), and SCC is $190 per metric ton (US EPA interim estimate [5]). BEVs are assigned zero tailpipe carbon cost; upstream grid emissions are incorporated directly using the Virginia eGRID emissions factor.

Acquisition and salvage. Vehicle purchase prices are sourced from FM financial records where available, supplemented by Kelly Blue Book. Future-year replacement vehicle prices are projected using a 3.7% compound annual growth rate [13]. Salvage value is estimated at 20% of original MSRP for all powertrain types [15].

Maintenance and insurance. Annual maintenance costs are drawn from FM's historical cost records (service events only; accident repairs excluded). Insurance premiums are sourced from FY2025 invoices, with BEV maintenance costs discounted by 30% relative to equivalent ICE vehicles, reflecting the reduced service requirements of electric drivetrains [16].

Sensitivity analysis. Four parameters were varied to bound uncertainty: gasoline price, SCC, salvage rate, and discount rate.

TABLE I
TCO MODEL ASSUMPTIONS

Parameter	Assumed Value	Source
Gasoline price	$3.20 / gal	[9]
Electricity rate	$0.1169 / kWh	[10]
Social cost of carbon	$190 / metric ton	[5]
CO_2 factor (ICE)	0.009796 t / gal	[11]
CO_2 factor (BEV)	0.000328 t / kWh	[12]
Annual driving days	150 days / year	FM operational data
Vehicle holding period	10 years	FM policy
Discount rate	7%	UVA endowment rate
Vehicle price inflation	3.7% CAGR	[13]
Salvage value	20% of MSRP	[15]
BEV maintenance savings	30% vs. ICE	[16]

Recommendation scoring. A vehicle replacement recommendation score was generated using a rule-based grading system that assessed each vehicle's TCO differential, operational compatibility with an EV replacement, charging availability, and vehicle condition. Vehicles categorized as structurally irreplaceable (e.g., specialized equipment with no EV equivalent) were excluded from the replacement pool.

C. Machine Learning Model for Forward-Looking Fuel Prediction

For vehicles not yet owned by UVA FM, observed Geotab fuel data are unavailable. The ML model fills this gap by predicting fuel or energy consumption from driving behavior features extracted from telematics data of comparable incumbent vehicles. The model was trained and evaluated following the comparative framework established by Ullah et al. [3], which identified gradient boosting methods as consistently superior to linear and neural network approaches for vehicle energy prediction tasks.

Features. Input features capture the primary physical drivers of energy consumption: total distance, idle ratio, speed distribution across three regimes (0–40, 40–80, and 80+ km/h), average squared speed (a proxy for aerodynamic drag), road grade distribution across five slope buckets, ambient temperature, and vehicle-level static attributes (curb weight, engine displacement, drag coefficient, and manufacturer fuel economy rating). Target variables were gallons consumed per fuel-reporting interval (ICE and FHEV) and kWh consumed per interval (BEV).

Models evaluated. Four models were trained and compared: multiple linear regression (MLR) [17] as a statistical baseline, an artificial neural network (ANN) [18] as a flexible nonparametric comparator, XGBoost [19] as a regularized gradient boosting ensemble, and LightGBM [20] as a histogram-based gradient boosting method optimized for large datasets. Models were evaluated on held-out test data using R^2, RMSE, and MAE, with stratified k-fold cross-validation used to assess generalizability.

V. Results

A. Retrospective Audit of Vehicle Replacement Decisions

The primary results of this work are a set of retrospective TCO analyses covering vehicle replacement decisions already made by UVA FM. For each evaluated pair, the incumbent ICE vehicle's actual Geotab fuel consumption was used to calculate its realized annual fuel and carbon cost, which was then compared against the actual observed consumption of the replacement vehicle over the same holding period. This provides a data-grounded audit of whether past electrification decisions delivered the cost and emissions reductions that a simple manufacturer-MPG comparison would have predicted.

Preliminary analysis identified a subset of high-idling vehicles that contributed disproportionately to fleet fuel use and emissions. Table II lists vehicles flagged for excessive idling, defined as an idle fraction exceeding 20% of total engine hours. Several vehicles exceed 40–48%, consistent with extended warm-up periods or stationary auxiliary use.

TABLE II
VEHICLES FLAGGED FOR EXCESSIVE IDLING

Vehicle Type	Idle Fraction	Excess vs. Fleet Avg.
ICE Vehicle A	40.5%	+24.9%
ICE Vehicle B	48.8%	+33.1%
ICE Vehicle C	24.9%	+9.3%
ICE Vehicle D	29.6%	+14.0%
ICE Vehicle E	46.3%	+30.7%

Despite their high emissions, these vehicles are not necessarily strong electrification candidates: their high idle-to-drive ratios are not well-matched to BEV or PHEV operational profiles, as sustained idling in a PHEV draws from the combustion engine rather than the battery, limiting the emissions benefit.

B. Machine Learning Model Results

Four models were evaluated for the ICE/FHEV and BEV prediction tasks. Across both tasks, gradient boosting methods outperformed MLR and ANN, consistent with the findings of Ullah et al. [3]. XGBoost was chosen for the ICE prediction task and LightGBM for the BEV task based on performance metrics, as seen in Table III. The fuel model demonstrates more consistent performance between training and testing compared to the energy model, which exhibits a larger drop in performance, suggesting comparatively weaker generalization for the energy predictions. This disparity reflects the imbalance in training data: the fleet contains 280 ICE and FHEV vehicles compared to only 6 BEVs. Feature importance analysis from the ICE model identified distance, speed profiles, trip duration, and vehicle characteristics as the strongest predictors, consistent with Ullah et al. [3]. The trained models are deployed as the fuel estimation component of the forward-looking TCO tool.

TABLE III
CHOSEN MODEL PERFORMANCE COMPARISON FOR FUEL AND ENERGY PREDICTION

Dataset	Model	Split	MAE	RMSE	R^2
Fuel	XGBoost	Train	0.1608	0.3628	0.8309
Fuel	XGBoost	Test	0.1793	0.4297	0.8146
Energy	LightGBM	Train	0.0848	0.1488	0.9456
Energy	LightGBM	Test	0.1379	0.2173	0.8841

C. Case Studies of Historical Vehicle Replacements

Two replacement pairs were selected to validate the TCO model against real FM decisions, covering the two primary electrification scenarios: an ICE-to-FHEV substitution for a light-duty passenger vehicle and an ICE-to-BEV substitution for a high-utilization cargo van.

TABLE IV
CASE 1: 2016 JEEP COMPASS (ICE) → 2022 FORD ESCAPE HYBRID (FHEV)

Component	ICE	FHEV
Vehicle year	2016	2022
Age (2025)	9 yrs	3 yrs
Purchase (2025$)	$27,521	$35,481
Insurance/yr	$441	$603
Maintenance/yr	$845	$401
Fuel cost/yr	$320	$57
Carbon cost/yr	$186	$33
Total opex/yr	$1,793	$1,094
10yr Opex NPV	$12,591	$7,683
Salvage value (NPV)	$2,798	$3,607
TOTAL 10yr TCO	$37,314	$39,557
CO_2 emissions/yr	0.981 t	0.174 t
CO_2 reduction	—	82% less
Fuel used (FY25)	100 gal	18 gal
→ ICE cheaper by $2,243 over 10 years		

The first case, shown in Table IV, examines the replacement of a 2016 Jeep Compass (ICE) with a 2022 Ford Escape Hybrid (FHEV). The hybrid's lower fuel and maintenance costs reduce annual operating expenditure by $699, but its higher purchase price means the 10-year TCO of $39,557 modestly exceeds the incumbent's $37,314—a gap of $2,243. The FHEV delivers an 82% reduction in CO_2 emissions, from

0.981 to 0.174 metric tons per year, demonstrating that near-term emissions benefits can be achieved at a relatively small cost premium for comparable light-duty vehicles.

TABLE V
CASE 2: 2017 CHEVROLET CITY EXPRESS (ICE) → 2023 FORD E-TRANSIT 350 (BEV)

Component	ICE	BEV
Vehicle year	2017	2023
Age (2025)	8 yrs	2 yrs
Purchase (2025$)	$28,692	$65,070
Insurance/yr	$575	$996
Maintenance/yr	$435	$51
Fuel cost/yr	$535	$183
Carbon cost/yr	$311	$98
Total opex/yr	$1,857	$1,327
10yr Opex NPV	$13,039	$9,320
Salvage value (NPV)	$2,917	$6,616
TOTAL 10yr TCO	$38,814	$67,774
CO_2 emissions/yr	1.638 t	0.514 t
CO_2 reduction	—	69% less
Fuel source (FY25)	167 gal	1,564 kWh

→ ICE cheaper by $28,960 over 10 years

The second case, shown in Table V, examines the replacement of a 2017 Chevrolet City Express cargo van (ICE) with a 2023 Ford E-Transit 350 (BEV). The E-Transit's substantially lower maintenance and fuel costs reduce annual operating expenditure by $530, but the $36,378 purchase price premium results in a 10-year TCO of $67,774—$28,960 more than the incumbent's $38,814. The BEV achieves a 69% reduction in CO_2 emissions. At the vehicle's current utilization level, the operating savings do not offset the upfront cost premium within the 10-year holding period, indicating that this replacement is not justified on TCO alone and that FM accepted an economic cost in exchange for emissions benefits.

VI. DISCUSSION

The case studies reveal a nuanced picture that reinforces the need for vehicle-level analysis. Case 1 shows that FHEV substitution can deliver large emissions reductions (82%) at a modest 10-year cost premium of $2,243—a trade-off many institutions would consider favorable. Case 2 shows that full BEV electrification at current acquisition prices carries a significant TCO penalty ($28,960) for lower-utilization vehicles, even with a 69% emissions reduction. Vehicles with higher annual mileage would see a more favorable break-even point, underscoring that utilization is the critical variable driving whether electrification is cost-justified.

The dual-component design of the framework—retrospective audit plus forward-looking ML prediction—reflects the practical needs of an institutional fleet manager. The audit component provides accountability by evaluating whether past purchasing decisions delivered on their projected savings, grounded in actual Geotab consumption data rather than manufacturer estimates. As shown in the case studies, grounding cost projections in real operational data surfaces meaningful differences between projected and actual savings that a manufacturer-MPG-based analysis would miss. The ML component extends this capability to vehicles not yet owned, enabling FM to generate data-driven TCO comparisons at the point of purchase decisions for new vehicles.

Several limitations should be acknowledged. First, the financial data coverage is incomplete: only 95 of 424 vehicles had confirmed purchase prices and full cost histories, which restricted the breadth of the retrospective TCO analysis. Second, the framework does not currently account for infrastructure costs (charging station installation) or additional training requirements for maintenance staff. Third, the BEV ML model's performance is constrained by the small number of BEVs in the training set (6 vehicles), a limitation that will diminish as the fleet continues to electrify and more BEV telematics data accumulate.

Regarding the social cost of carbon: the $190/ton SCC adopted in this study reflects the US EPA interim estimate [5]. The sensitivity analysis bounds this across a range of ±$100 per metric ton, reflecting the policy and methodological uncertainty documented in the SCC literature. Future work should incorporate lifecycle analysis extending to vehicle manufacturing and end-of-life disposal to provide a more complete emissions accounting.

The framework is designed to be updatable: as fleet usage patterns evolve, new telematics data can be re-ingested and the TCO analysis rerun without structural changes to the pipeline.

A. Scalability and Applicability to Other University Fleets

Although developed for UVA FM, the framework was designed for portability. The external data dependencies—USGS elevation tiles, NOAA weather records, and public vehicle specification APIs—are available nationwide, requiring only a geographic region and local weather station identifier to reconfigure for a different campus. The NREL multi-university study confirmed that university fleets share structural characteristics—centralized parking, predictable duty cycles, sustainability commitments—that make the electrification problem broadly similar across institutions [1]. Geotab is widely used in university and municipal fleets, and the preprocessing pipeline can be adapted to other telematics providers by remapping column names and unit conversions. The modular design allows the TCO model to be deployed with manufacturer MPG estimates as a fallback while telematics data accumulate, lowering the barrier to entry for fleets at earlier stages of data maturity.

VII. CONCLUSION

This paper presented a data-driven decision-support framework for fleet electrification planning at UVA Facilities Management, combining a retrospective audit of past vehicle replacement decisions with a forward-looking machine learning-based fuel prediction tool. The framework integrated vehicle-level telematics data, financial cost records, and the social cost of carbon into a total cost of ownership model that provides vehicle-by-vehicle electrification guidance grounded in observed operational data rather than manufacturer estimates.

Retrospective case studies demonstrated that hybrid and battery electric replacements deliver meaningful emissions reductions of 69–82%, but that the economic justification for electrification depends heavily on vehicle utilization. For the FHEV case, the replacement was nearly cost-neutral over 10 years ($2,243 premium) with substantial emissions benefit. For the BEV cargo van, the current acquisition cost premium of $28,960 over 10 years was not offset by operating savings at the vehicle's observed utilization level. These findings confirm that targeted, utilization-aware replacement strategies are substantially more cost-effective than uniform fleet-wide electrification.

The forward-looking ML component extends this capability to future purchasing decisions, enabling FM to estimate energy costs for vehicles not yet owned using driving behavior patterns extracted from comparable incumbent vehicles. Future work should expand the case study set to include PHEV replacements, incorporate sensitivity analysis outputs to identify the mileage thresholds at which BEV total cost becomes competitive, and account for charging infrastructure costs as the fleet continues to electrify. The modular, configurable design of the framework positions it as a replicable template for other university fleets seeking rigorous, data-driven electrification planning tools.

Acknowledgment

The authors thank UVA Facilities Management for providing the vehicle data that made this work possible.

References

[1] S. Booth, J. Bennett, M. Helm, D. Arnold, B. Baker, R. Clay, M. Till, and T. Sears, "Identifying electric vehicles to best serve university fleet needs and support sustainability goals," National Renewable Energy Laboratory, Golden, CO, Tech. Rep. NREL/TP-5400-81596, Feb. 2022. [Online]. Available: https://www.nrel.gov/publications.

[2] R. Daley and M. Helm, "Telematics and data science: Informing energy-efficient mobility," National Renewable Energy Laboratory, Golden, CO, Subcontract Rep. NREL/SR-5400-73661, Aug. 2019. [Online]. Available: https://www.nrel.gov/publications.

[3] I. Ullah, K. Liu, T. Yamamoto, R. E. Al Mamlook, and A. Jamal, "A comparative performance of machine learning algorithm to predict electric vehicles energy consumption: A path towards sustainability," *Sustainable Computing: Informatics and Systems*, 2022.

[4] I. Ullah, K. Liu, T. Yamamoto, R. E. Al Mamlook, and A. Jamal, "A comparative performance of machine learning algorithm to predict electric vehicles energy consumption: A path towards sustainability," *Energy Reports*, vol. 8, pp. 278–285, 2022.

[5] M. Greenstone, E. Kopits, and A. Wolverton, "Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation," *Review of Environmental Economics and Policy*, vol. 7, no. 1, pp. 23–46, 2013.

[6] T. Yoon and C. R. Cherry, "Migrating towards using electric vehicles in campus: Proposed methods for fleet optimization," *Sustainability*, vol. 10, no. 2, p. 285, Jan. 2018.

[7] J. Moon, H. Anwar, M. Villani, M. Q. Fahim, P. Jain, K. Ramakrishnan, and Q. Ahmed, "Electrified fleet and infrastructure aware energy efficient routing," *eTransportation*, 2023.

[8] M. Wann, D. Yen, and H. Gold, "Evaluation and calibration of three models for daily cycle of air temperature," *Agricultural and Forest Meteorology*, vol. 34, pp. 121–128, 1985.

[9] U.S. Energy Information Administration, "Gasoline and diesel fuel update," [Online]. Available: https://www.eia.gov/petroleum/gasdiesel/. Accessed: 2025.

[10] U.S. Energy Information Administration, "Electric power monthly," [Online]. Available: https://www.eia.gov/electricity/monthly/. Accessed: 2025.

[11] U.S. Environmental Protection Agency, "Greenhouse gas emissions from a typical passenger vehicle," EPA-420-F-23-014, 2023. [Online]. Available: https://www.epa.gov/greenvehicles.

[12] U.S. Environmental Protection Agency, "Emissions & generation resource integrated database (eGRID)," [Online]. Available: https://www.epa.gov/egrid. Accessed: 2025.

[13] U.S. Bureau of Labor Statistics, "Consumer price index: New vehicles," [Online]. Available: https://www.bls.gov/cpi/. Accessed: 2025.

[14] U.S. Department of Energy, Alternative Fuels Data Center, "Vehicle weight classifications for light-duty trucks," [Online]. Available: https://afdc.energy.gov/data/10380. Accessed: 2025.

[15] Kelley Blue Book, "Car depreciation: How much is your new car worth after 5 years?" [Online]. Available: https://www.kbb.com. Accessed: 2025.

[16] A. Burnham, M. Wang, and Y. Wu, "Development and applications of GREET 2.7—the transportation vehicle-cycle model," Argonne National Laboratory, ANL/ESD/06-5, 2006.

[17] D. C. Montgomery, E. A. Peck, and G. G. Vining, *Introduction to Linear Regression Analysis*, 5th ed. Hoboken, NJ: Wiley, 2012.

[18] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016. [Online]. Available: https://www.deeplearningbook.org.

[19] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, San Francisco, CA, 2016, pp. 785–794.

[20] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "LightGBM: A highly efficient gradient boosting decision tree," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.