

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre

Bonacina, Monica; Demir, Mert; Sileo, Antonio; Zanoni, Angela

Working Paper

What hinders electric vehicle diffusion? Insights from a neural network approach

Working Paper, No. 16.2025

Provided in Cooperation with:

Fondazione Eni Enrico Mattei (FEEM)

Suggested Citation: Bonacina, Monica; Demir, Mert; Sileo, Antonio; Zanoni, Angela (2025): What hinders electric vehicle diffusion? Insights from a neural network approach, Working Paper, No. 16.2025, Fondazione Eni Enrico Mattei (FEEM), Milano

This Version is available at: https://hdl.handle.net/10419/331435

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



August 2025



Working Paper

16.2025

What Hinders
Electric Vehicle
Diffusion? Insights
from a Neural
Network Approach

Monica Bonacina, Mert Demir, Antonio Sileo, Angela Zanoni

What Hinders Electric Vehicle Diffusion? Insights from a Neural Network Approach

Monica Bonacina (Fondazione Eni Enrico Mattei, Università degli Studi di Milano), Mert Demir (Fondazione Eni Enrico Mattei), Antonio Sileo (Fondazione Eni Enrico Mattei, GREEN - Università Bocconi), Angela Zanoni (Fondazione Eni Enrico Mattei, Università di Roma La Sapienza, Research Institute for Sustainable Economic Growth - National Research Council)

Summary

The transition to a zero-emission vehicle fleet represents a pivotal element of Europe's decarbonization strategy, with Italy's participation being particularly significant given the size of its automotive market. This study investigates the potential for battery electric cars (BEVs) to drive decarbonization of Italy's passenger vehicle fleet, focusing on the feasibility of targets set in the National Integrated Plan for Energy and Climate (PNIEC). Leveraging artificial neural networks, we integrate macroeconomic indicators, market-specific variables, and policy instruments to predict fleet dynamics and identify key factors influencing BEV adoption. We forecast that while BEV registrations will continue growing through 2030, the growth rate is projected to decelerate, presenting challenges for meeting ambitious policy targets. Our feature importance analysis demonstrates that BEV adoption is driven by an interconnected set of economic, infrastructural, and behavioral factors. Specifically, our model highlights that hybrid vehicle registrations and the vehicle purchase index exert the strongest influence on BEV registrations, suggesting that policy interventions should prioritize these areas to maximize impact. By offering data-driven insights and methodological innovations, our findings contribute to more effective policy design for accelerating sustainable mobility adoption while accounting for market realities and consumer behavior.

Keywords: sustainable mobility, electric vehicle, neural networks, shap interpretation

JEL Classification: N74, Q55, Q58, R40, C45

Corresponding Author:

Angela Zanoni Fondazione Eni Enrico Mattei, Via del Commercio, 9/11 - 00154 Roma angela.zanoni@feem.it

What Hinders Electric Vehicle Diffusion? Insights from a Neural Network Approach

Monica Bonacina^{a,b}, Mert Demir^a, Antonio Sileo^{a,c}, Angela Zanoni*a,d, e

- ^a Fondazione Eni Enrico Mattei, Milan
- ^b Università degli Studi di Milano, Milan
- ^c GREEN Università Bocconi, Milan
- ^d Università di Roma la Sapienza, Rome

Abstract

The transition to a zero-emission vehicle fleet represents a pivotal element of Europe's decarbonization strategy, with Italy's participation being particularly significant given the size of its automotive market. This study investigates the potential for battery electric cars (BEVs) to drive decarbonization of Italy's passenger vehicle fleet, focusing on the feasibility of targets set in the National Integrated Plan for Energy and Climate (PNIEC). Leveraging artificial neural networks, we integrate macroeconomic indicators, market-specific variables, and policy instruments to predict fleet dynamics and identify key factors influencing BEV adoption. We forecast that while BEV registrations will continue growing through 2030, the growth rate is projected to decelerate, presenting challenges for meeting ambitious policy targets. Our feature importance analysis demonstrates that BEV adoption is driven by an interconnected set of economic, infrastructural, and behavioral factors. Specifically, our model highlights that hybrid vehicle registrations and the vehicle purchase index exert the strongest influence on BEV registrations, suggesting that policy interventions should prioritize these areas to maximize impact. By offering data-driven insights and methodological innovations, our findings contribute to more effective policy design for accelerating sustainable mobility adoption while accounting for market realities and consumer behavior.

^e Research Institute for Sustainable Economic Growth - National Research Council, Rome, Rome

^{*}Corresponding author. angela.zanoni@feem.it

1 Introduction

Transport sector decarbonization has emerged as a central pillar in both EU-wide and national climate change policies, with the diffusion of battery electric cars (BEVs) frequently presented as a key solution. Following this idea, the EU has established a ban on internal combustion engine vehicles by 2035, while national governments have made commitments to increase electric vehicle penetration in their fleets. However, the rate of fleet substitution has not matched expectations, either in Europe or worldwide (International Energy Agency (2025); Rietmann et al. (2020)).

The complexities inherent in electric mobility deployment have generated substantial research interest. A well-established body of literature now addresses the expected uptake of battery-electric cars, examining both the demand side—focusing on consumer choices (Ziegler (2012); van Rijnsoever et al. (2013); Plötz et al. (2014b))—and the supply side - concerned with the quality and quantity of electric models produced by manufacturers (Al-Alawi and Bradley, 2013a; Domarchi and Cherchi, 2023a). Parallel to this, a growing debate investigates the actual environmental benefits of electric vehicles (Rietmann et al., 2020; Ziegler, 2012). Despite uncertainties surrounding these topics, the policy direction, particularly at the European level, remains steadfast.

Nevertheless, data suggest that consumers exhibit exceptionally low propensity for BEV purchase (Bonacina and Sileo (2024)). This reality poses significant challenges for policymakers, adding uncertainty to the already debated issue of environmental benefits. What factors are hindering BEV uptake? Is it a matter of technical characteristics? Is it about price? Can increased governmental incentives stimulate purchases? As will be presented in Section 2, researchers have found evidence supporting all these explanations. However, few studies (Jensen et al. (2016b)) attempt to identify the optimal areas for policy intervention.

In this landscape, Italy represents a particularly meaningful case study as the country with the second-largest vehicle fleet in Europe. While the adoption of electric vehicles in Italy has gained momentum in recent years, it still lags behind other European front-runners. Studies suggest that limited infrastructure availability, high capital expenses, and cultural affinity for traditional vehicles have slowed the diffusion process (Bonacina and Sileo (2024)). Nevertheless, Italy's National Integrated Plan on Energy and Climate has established ambitious targets for sustainable mobility, aiming to have over 4 million battery electric cars (BEVs) and more than 2 million plug-in hybrid electric vehicles (PHECs) on Italian roads by 2030. This strategy relies on fiscal incentives (e.g.,

Ecobonus), investments in public charging infrastructure, and various local-level measures (Bonacina and Sileo (2024); van Rijnsoever et al. (2013)).

The following research questions guide our analysis. First, can a data-driven, machine learning-based approach provide reliable predictions of Italy's BEV fleet evolution compared to previous works, and how do these align to the existing policy goals? This focuses on assessing the predictive power and reliability of the machine learning framework, especially in contrast to existing policy-target-based models Siskos et al. (2019), which may not fully capture the complexities of the Italian automotive market. Second, what factors are most critical in shaping the future adoption of BEVs in Italy, and how can this knowledge be leveraged to design more effective policy interventions? This question aims to identify the key drivers of BEV adoption—such as market dynamics, consumer behavior, and scrappage rates—and explore how this understanding can inform more targeted and realistic policy measures.

Our work contributes to the literature by proposing a framework to identify the most influential factors determining BEV uptake. To this end, we train a machine learning model to forecast new BEV registrations in Italy, leveraging comprehensive data on market conditions, consumer behavior, and policy interventions. We demonstrate that our model not only performs exceptionally well in out-of-sample prediction but also yields forecasts aligned with real-world developments. Our methodological approach draws inspiration from recent developments in machine learning applications for market forecasting. In particular, we adopt a similar framework to that employed in Afandizadeh (2023), who applied machine learning to forecast electric vehicles' penetration in the United States. Their work highlights the potential of data-driven models to uncover hidden patterns in complex systems, motivating us to explore whether a different set of algorithms and input features can still lead to convincing predictive performances for the Italian case and yield policy relevance in this domain. While this is not the first work employing ML to study BEV uptake, we broaden the scope of modeling by incorporating a wider pool of machine learning algorithms, including both foundational architectures and cutting-edge models. This enables us to benchmark performance across diverse techniques and identify architectures best suited for the prediction. In addition, we define our prediction task using a more comprehensive set of variables compared to previous works. Our inputs range from macroeconomic indicators (GDP, inflation, purchasing power) to specific market factors (fuel prices, energy indices, consumer sentiment), allowing for a more nuanced modeling of adoption drivers. To ensure the policy relevance of our findings, we integrate feature importance analysis within our framework. In this way, we

make a concerted effort to ensure the interpretability of our results, providing researchers and policymakers with a convincing, data-driven guide to understand the feasibility of policy goals and to frame grounded interventions.

The remainder of this study is organized as follows. Section 2 reviews the relevant scientific literature on vehicle stock evolution and electric vehicle penetration. Section 3 outlines our theoretical framework and empirical methodology. Section 4 describes our data sources, their characteristics, strengths, and limitations. Section 5 presents the results of our analysis and concludes by offering policy recommendations.

2 Literature Review

To collocate our work in the existing research, we need to refer to two separate strands of research.

2.1 Studies on factors driving BEV adoption

The decision-making process for electric vehicle (EV) adoption represents a complex interplay of multiple factors that have been extensively studied across various disciplines.

Economic factors constitute perhaps the most extensively researched domain in EV adoption literature. While EVs typically command higher initial purchase prices than internal combustion engine (ICE) vehicles, their total cost of ownership (TCO)—incorporating fuel, maintenance, insurance, and residual value—can be lower over the vehicle's lifetime. Government incentives, including direct subsidies, tax exemptions, and reduced registration fees, play a crucial role in offsetting initial cost barriers (Van et al., 2023).

The deterrent effect of high upfront costs has been consistently documented across multiple studies. Lévay et al. (2017), Mohsin et al. (2023), and Kamal and Mansur (2024) all identify initial purchase price as a significant barrier to adoption. Regional variations in economic sensitivity are substantial, with Rotaris et al. (2021) finding that Italian consumers' EV preferences are strongly influenced by economic factors and charging infrastructure availability, contributing to relatively slow uptake despite growing interest. Similarly, Buhmann et al. (2023) identified distinct segments of German EV buyers with varying price sensitivities, suggesting the need for targeted economic strategies.

Policy effectiveness emerges as a critical theme in this strand of literature. Plötz et al. (2014a) demonstrated that purchase subsidies significantly stimulated EV adoption in Germany, while comprehensive reviews by Mohsin et al. (2023) and Kamal and Mansur

(2024) consistently confirm the vital role of supportive financial policies. However, the universal efficacy of incentives remains contested. Wang et al. (2017) found no significant effect of financial incentives on Chinese consumers' EV adoption intentions, a finding corroborated by Sierzchula et al. (2014), suggesting that incentive effectiveness may be context-dependent.

The importance of policy design emerges from studies by Lutsey and Nicholas (2019) and Held and Gerrits (2019), who emphasize the need for comprehensive, consistent, and long-term policy frameworks. Fragmented or abrupt policy changes prove counterproductive to market development. Notably, Springel (2021) found that subsidies for charging infrastructure could be more cost-effective than direct vehicle purchase subsidies in driving initial EV uptake, attributing this to network externalities that benefit all EV users.

Technological limitations and performance concerns represent persistent barriers to EV adoption. The literature consistently identifies three primary technological barriers: limited driving range, extended charging times, and inadequate charging infrastructure. "Range anxiety"—the fear of running out of battery during travel—and "charging tolerance"—the inconvenience of longer charging times compared to rapid refueling—remain paramount concerns for potential adopters.

Empirical studies highlight the salience of these concerns. Rotaris et al. (2021) identified driving range as a key factor influencing Italian consumers' EV preferences, while Hidrue et al. (2011) quantified consumers' willingness to pay for improved range and reduced charging times. The psychological dimension of range anxiety is illuminated by Franke and Krems (2013), who demonstrated that users' preferred range often exceeds their actual driving needs, highlighting the formation of "psychological range levels" that influence behavior.

The evolution of range anxiety discourse reflects technological progress. While absolute battery range has increased significantly, Papa et al. (2020) note that focus has shifted toward range variability and predictability under real-world conditions. This evolution has transformed the challenge from simply providing more range to ensuring more predictable and reliable range performance (Mohsin et al., 2023).

Consumer experience can mitigate technological concerns. Jensen et al. (2013) demonstrated that direct EV interaction significantly improves consumer attitudes and preferences. However, perceived limitations often lead to preferences for hybrid technologies as bridging solutions. Axsen and Kurani (2013) found that consumers frequently favor plug-in hybrid electric vehicles (PHEVs) over battery electric vehicles (BEVs) due to

range and cost concerns, suggesting a phased approach to electrification.

Significantly, economic factors have increasingly overshadowed technological concerns. S&P Global Mobility (2023) revealed that EV affordability has become a greater barrier than charging or range anxiety, contributing to demand growth deceleration. This finding suggests that as technological barriers diminish, economic considerations become more prominent in consumer decision-making.

Environmental consciousness significantly impacts EV adoption decisions, though it is often underestimated by policymakers (Tran et al., 2012). Consumers with heightened environmental awareness and climate change concerns demonstrate greater inclination toward BEV adoption, perceiving this choice as contributing to greenhouse gas reduction.

The environmental benefits of EVs are well-documented. Farajnezhad et al. (2024) provide a comprehensive review of environmental factors influencing EV adoption, while Mohsin et al. (2023) and Kamal and Mansur (2024) confirm through systematic reviews that environmental benefits, particularly greenhouse gas reduction when charged with renewable energy, serve as crucial adoption drivers. Xia et al. (2020) demonstrate that EV adoption in China directly reduces CO₂ emissions and generates positive spatial spillover effects.

Social factors extend beyond environmental concerns to encompass status and peer influence. EVs are increasingly perceived as symbols of social status and technological progressiveness, conferring positive social identity on owners. Buhmann and Criado (2023) reveal that status and reputation influence EV adoption, with some consumers purchasing more expensive EVs as status symbols. This finding aligns with Noppers et al. (2014), who found that symbolic and environmental attributes were stronger predictors of sustainable innovation adoption than purely instrumental factors.

The theoretical foundation for environmental and social influences is well-established. Ajzen (1991) Theory of Planned Behavior posits that behavioral intentions, shaped by attitudes, subjective norms, and perceived behavioral control, strongly predict actual behavior. Rezvani et al. (2017) complement this by demonstrating that anticipated emotions and personal moral norms significantly influence pro-environmental consumer behavior.

Psychological and behavioral factors often outweigh or modify the impact of economic and technical considerations in EV adoption decisions. Consumer attitudes toward new technology significantly influence EV acceptance, with openness to technological innovation serving as an important precursor to adoption. The innovation diffusion literature suggests that early adopters are typically more receptive to new technologies, while later adopters exhibit greater skepticism.

Key psychological constructs impacting adoption include perceived behavioral control—the perceived ease or difficulty of EV adoption—and subjective norms—the perceived social pressure to adopt. Moral norms, stemming from personal environmental responsibility, can directly influence purchase intentions. Prior EV experience, through test drives or ride-sharing, can significantly mitigate psychological barriers by building familiarity and confidence.

Established behavioral patterns present significant obstacles to EV adoption. Habitual refueling behaviors and car usage patterns prove difficult to alter, with the need to plan charging and adapt to new routines representing substantial barriers for many potential adopters. The literature on these topics continues to expand, with recent research such as Myers et al. (2025) emphasizing the importance of integrating detailed driver behavior understanding into charging infrastructure design and deployment.

Demographic characteristics consistently correlate with EV adoption decisions. Regional socio-economic development and digitalization levels positively correlate with EV adoption propensity (Popescu et al., 2024). Individual factors including income, education, age, gender, household composition, existing vehicle ownership, geographical location, and residence type all influence EV adoption choices.

Income emerges as the most consistently correlated demographic factor. Higher income levels provide the financial capacity to manage elevated EV purchase prices and are often associated with home ownership, facilitating dedicated charging infrastructure installation. Educational attainment also correlates positively with EV adoption, potentially reflecting greater environmental awareness, technological openness, and improved ability to process complex information regarding long-term savings and incentive structures.

Age patterns in EV adoption show interesting variations. Many studies indicate that BEV buyers tend to be middle-aged, potentially reflecting a balance between technological openness and environmental sensitivity common in younger generations, combined with the financial stability typically achieved in middle age. Geographically, BEV penetration is higher in urban and suburban areas, characterized by shorter travel distances and more developed charging infrastructure.

Theoretical frameworks such as Ajzen (1991) Theory of Planned Behavior help explain these demographic patterns. Pamidimukkala et al. (2025) demonstrate how various motivational factors influence EV acceptance in India, with gender playing a moderating role in these dynamics, illustrating the complex interplay between demographic characteristics and adoption decisions.

The competitive automotive landscape also significantly influences BEV adoption patterns. The continuous improvement and diversification of ICE models, coupled with inherently lower upfront costs, maintains substantial market appeal. In 2024, approximately 25% of European ICE vehicles were priced below €30,000, compared to only 5% of BEVs (International Energy Agency, 2025).

Hybrid technologies serve as compelling alternatives, bridging conventional and fully electric powertrains. These bridging technologies experience robust growth and capture market share potentially at BEV expense. By April 2025, hybrid electric vehicles (HEVs) commanded 35.3% of the EU market share, compared to 15.3% for BEVs and 7.9% for PHEVs (European Automobile Manufacturers' Association, 2025).

The entry of Chinese EV manufacturers profoundly disrupts this competitive land-scape through aggressive pricing and rapid technological advancement. China's domestic market, where nearly two-thirds of BEVs were cheaper than ICE equivalents in 2024 (International Energy Agency, 2025), fuels this global trend. This influx creates strong downward pressure on new EV prices and affects used vehicle residual values (Pao, 2025; Roberts, 2025). While this intensifies competition and enhances EV affordability, long-term impacts on global automotive profitability and brand positioning remain uncertain.

2.2 Studies employing ML to predict electric vehicle diffusion

The application of machine learning techniques to electric vehicle market prediction represents a paradigmatic shift from traditional econometric modeling toward data-driven approaches capable of capturing complex, non-linear relationships in adoption patterns. This evolution reflects both the increasing availability of high-dimensional datasets and the recognition that EV adoption involves intricate interactions between economic, technological, behavioral, and policy factors that may be difficult to model using conventional methods.

2.2.1 Evolution and methodological foundations

The transition from classical econometric models to machine learning approaches in EV forecasting can be traced to the early 2010s, when researchers began recognizing the limitations of linear models in capturing the complex dynamics of technology adoption. Traditional approaches, while theoretically grounded, often struggled with the non-linearities inherent in early adoption phases, the heterogeneity of consumer preferences, and the dynamic nature of policy interventions (Al-Alawi and Bradley, 2013b).

Hülsmann et al. (2012) provided one of the earliest systematic comparisons of traditional econometric methods with machine learning approaches for automobile sales forecasting, demonstrating that decision trees could deliver both accuracy and explainability. Their work on German and US markets established decision trees as particularly suitable for automotive forecasting due to their ability to capture threshold effects while maintaining interpretability—a crucial requirement for policy applications.

The methodological landscape has since evolved to encompass a broad spectrum of techniques, from classical time series models enhanced with machine learning components to sophisticated deep learning architectures. This evolution reflects not only technological advances but also the growing recognition that EV adoption involves multiple interacting systems that require flexible modeling approaches.

2.2.2 Time series and classical ML approaches

Traditional time series methods have formed the foundation for many EV forecasting efforts, often serving as benchmarks against which more sophisticated approaches are evaluated. Zhang et al. (2017) conducted a comprehensive comparison between univariate time series models (singular spectrum analysis) and multivariate approaches (vector autoregressive models) for Chinese EV sales forecasting. Their findings demonstrated that incorporating exogenous variables—particularly economic indicators and policy measures—significantly improved prediction accuracy, with the VAR model achieving superior performance in both short-term (up to December 2017) and long-term (up to 2020) forecasting horizons.

The integration of economic indicators has proven particularly valuable. Wang et al. (2011) developed a sophisticated forecasting methodology for Taiwanese automobile sales that incorporated exchange rates, GDP, consumer confidence indices, and consumer price indices. Their adaptive neuro-fuzzy inference system (ANFIS) outperformed both traditional ARIMA models and standard artificial neural networks, achieving superior accuracy by effectively combining the pattern recognition capabilities of neural networks with the interpretability of fuzzy logic systems.

Ensemble methods have gained prominence for their ability to combine multiple models' strengths while mitigating individual weaknesses. Tiwari et al. (2024) conducted an extensive comparison of eleven machine learning models for volatility forecasting, including Linear Regression, Ridge Regression, Lasso Regression, Random Forest, Support Vector Machines, LSTM, GRU, ANN, and XGBoost. Their analysis revealed that model performance varies significantly across forecasting horizons, with Random Forest and

XGBoost demonstrating particular effectiveness for medium-term predictions—findings that have important implications for EV market forecasting.

2.2.3 Deep learning and neural network approaches

The application of deep learning to EV forecasting represents the current frontier in this field, driven by the availability of larger datasets and the need to capture increasingly complex relationships. Afandizadeh (2023) pioneered the application of sophisticated neural architectures to EV market prediction, developing a hybrid LSTM model with two-dimensional attention mechanisms and residual connections. Their approach, trained on monthly sales and registration data for 357 vehicle models in the United States from 2014 to 2020, achieved a mean absolute error of 3.5% for EV market share prediction.

The architecture innovations in this study are particularly noteworthy. By combining LSTM networks' temporal modeling capabilities with attention mechanisms that can focus on relevant time periods and residual connections that facilitate training of deep networks, Afandizadeh (2023) demonstrated how advanced deep learning techniques could be effectively adapted for automotive market forecasting. Their use of Automated Machine Learning (AutoML) principles to optimize model architecture and hyperparameters further illustrates the evolution toward more systematic and reproducible modeling approaches.

Saxena et al. (2020) contributed to this literature by developing LSTM-RNN architectures specifically designed for multivariate car sales forecasting. Their approach emphasized the importance of incorporating both temporal dependencies and cross-series correlations, recognizing that different vehicle segments often exhibit interdependent adoption patterns. This work highlighted the value of multivariate modeling approaches that can capture spillover effects between different automotive technologies.

The application of deep learning has extended beyond traditional neural networks to include more specialized architectures. Kaya and Özal Yildirim (2020) developed an 8-layer deep neural network for automobile sales prediction in Turkey, integrating macroeconomic indicators including exchange rates, GDP, consumer confidence indices, and consumer price indices. Their approach achieved superior performance compared to traditional econometric methods, demonstrating the value of deep architectures for capturing complex relationships between macroeconomic conditions and automotive demand.

2.2.4 Spatial and geographic considerations

The geographic dimension of EV adoption has increasingly been recognized as crucial for accurate forecasting. Miconi and Dimitri (2023a) developed a comprehensive machine learning framework specifically for the Italian context, conducting analysis at national, regional, and provincial levels. Their work highlighted the heterogeneous nature of EV adoption across different geographical areas and identified region-specific factors that influence market dynamics. This study was particularly innovative in its use of novel survey data to complement traditional market statistics, providing insights into consumer preferences and behavioral factors that are difficult to capture through sales data alone.

Kumar et al. (2022) expanded this geographic perspective by conducting a comparative assessment across 20 major countries, employing various diffusion models including Gompertz, Logistic, Bass, and Generalized Bass models. Their analysis revealed substantial cross-country variation in adoption patterns and identified distinct clusters of countries with similar diffusion characteristics. This work demonstrated the importance of considering cultural, economic, and regulatory differences when developing EV forecasting models, while also showing how models developed for one context might be adapted for others.

The incorporation of infrastructure considerations has become increasingly sophisticated. Kumar et al. (2022) explicitly modeled charging infrastructure development within their Generalized Bass diffusion framework, demonstrating how network effects and infrastructure availability influence adoption trajectories. This approach represents an important evolution from earlier models that treated infrastructure as exogenous to models that recognize the co-evolutionary nature of technology adoption and supporting infrastructure.

2.2.5 The accuracy paradox and validation challenges

One of the most significant findings in recent literature has been the identification of what Domarchi and Cherchi (2023b) termed the "accuracy paradox" in EV forecasting. Their comprehensive review of EV demand forecasting models revealed that, counterintuitively, simpler "business as usual" models often outperformed sophisticated machine learning techniques when evaluated against real-world outcomes. This finding challenges assumptions about the superiority of complex models and highlights several critical issues in the field.

Domarchi and Cherchi (2023b) attribute this paradox to several factors: overfitting in complex models due to limited data availability, insufficient consideration of policy discontinuities that can cause structural breaks in adoption patterns, and the failure to integrate fundamental economic relationships that govern consumer behavior. Their analysis suggests that while machine learning models excel at pattern recognition within stable environments, they may struggle with the structural changes and policy regime shifts that characterize real-world EV markets.

This finding has important implications for model validation and evaluation. The authors emphasize that the accuracy of forecasting models depends critically on the quality and representativeness of input data, regardless of the sophistication of the modeling technique employed. This insight has led to increased emphasis on data quality, feature engineering, and robust validation procedures in subsequent research.

2.2.6 Interpretability and policy relevance

The evolution toward more interpretable machine learning approaches reflects the recognition that policy applications require not just accurate predictions but also understanding of underlying mechanisms. Montavon et al. (2018) provided a comprehensive framework for interpreting deep neural networks, introducing techniques such as layer-wise relevance propagation that can identify which input features contribute most significantly to model predictions.

The application of SHAP (SHapley Additive exPlanations) values (Lundberg and Lee, 2017) has become increasingly common in EV forecasting studies. This model-agnostic approach provides both global and local interpretability, allowing researchers to understand both overall feature importance and the contribution of specific features to individual predictions. This capability is particularly valuable for policy applications, where decision-makers need to understand which factors are most amenable to policy intervention.

Recent work has sought to integrate interpretability considerations directly into model development rather than treating explainability as a post-hoc analysis. This evolution reflects the growing recognition that the most sophisticated model is not necessarily the most useful for policy purposes if its predictions cannot be understood and acted upon by decision-makers.

2.2.7 Integration with economic theory and policy modeling

An important development in the field has been the increasing integration of machine learning approaches with established economic theory and policy modeling frameworks. Rather than replacing traditional approaches entirely, recent work has sought to combine the pattern recognition capabilities of ML with the theoretical foundations of economic models.

This integration has taken several forms. Some researchers have used machine learning techniques to estimate key parameters in theoretical models, such as price elasticities and consumer preferences. Others have developed hybrid approaches that use ML for demand forecasting while embedding these predictions within broader economic equilibrium models. Still others have focused on using ML to identify structural breaks and regime changes that indicate when traditional economic relationships may no longer hold.

The policy dimension has become increasingly sophisticated, with researchers developing approaches that can model the dynamic effects of policy interventions. This includes not only direct effects (such as the impact of purchase subsidies on adoption rates) but also indirect effects (such as how infrastructure investments influence consumer preferences) and interaction effects (such as how the effectiveness of incentives varies with economic conditions).

2.2.8 Emerging techniques and future directions

The frontier of ML applications in EV forecasting continues to evolve, with several emerging techniques showing promise. Graph neural networks are being explored for their ability to model spatial dependencies and network effects in infrastructure development. Reinforcement learning approaches are being investigated for dynamic policy optimization, where policies can be adapted in response to observed market developments.

Multi-modal learning approaches that integrate textual data (such as policy documents and news sentiment), imagery (such as satellite data on infrastructure development), and traditional numerical data are becoming more feasible as data availability increases and computational costs decrease. These approaches offer the potential to capture information that is difficult to quantify but may be important for forecasting.

The integration of real-time data streams, including information from connected vehicles and charging infrastructure, presents opportunities for more responsive forecasting models that can adapt quickly to changing conditions. However, this also raises challenges related to data privacy, model stability, and the integration of high-frequency data

with lower-frequency economic and policy variables.

2.2.9 Research gaps and methodological challenges

Despite these advances, several significant research gaps remain. Cross-country model transferability remains limited, with most models developed for specific national or regional contexts showing degraded performance when applied elsewhere. This limitation reflects both data availability constraints and genuine differences in market structures, policy environments, and consumer preferences.

The handling of structural changes and technology disruptions remains challenging. Most current approaches assume that historical relationships will continue to hold, but the EV market is characterized by rapid technological change, evolving policy environments, and shifting consumer preferences. Developing models that can adapt to or anticipate such changes remains an open research challenge.

Long-term forecasting accuracy continues to degrade significantly with forecasting horizon, particularly for complex models. While this is a general characteristic of forecasting models, it is particularly problematic for policy applications that require projections over multiple years or decades. Understanding and improving long-term forecasting performance remains a critical research priority.

The integration of causal inference with machine learning approaches remains underdeveloped. While ML models excel at identifying correlations and patterns, policy applications often require understanding of causal relationships. Developing approaches that can move beyond correlation to establish causation while maintaining the flexibility and performance of ML methods represents an important frontier.

Our contribution to this literature addresses several of these challenges through: (1) a comprehensive comparison of multiple ML architectures to identify approaches best suited for temporal forecasting while avoiding data leakage; (2) integration of a broader set of variables spanning macroeconomic indicators, market dynamics, and transitional technologies; (3) systematic feature importance analysis using SHAP values to ensure policy relevance; and (4) rigorous temporal validation against real-world outcomes to ensure practical utility. By combining methodological rigor with policy relevance, our work contributes to the development of more reliable and actionable EV forecasting capabilities.

3 Data and Methodology

3.1 Data

To ensure the predictive power of our machine learning models, we constructed a comprehensive dataset composed of both vehicle market indicators and broader macroeconomic variables. All variables were selected based on theoretical relevance in the literature on battery electric car (BEV) adoption. Table 1 summarizes the variables used, their descriptions, and sources.

Our analysis draws from four principal data sources, each contributing a distinct layer of insight into fleet dynamics and economic context. First, we used publicly available micro-data from the Italian Ministry of Transport and Infrastructure (MIT), which includes annual observations on the national vehicle fleet, disaggregated by vehicle age, fuel type, regional location, and technical specifications. These data provide a detailed snapshot of the structural composition of Italy's vehicle stock. Due to incomplete regional reporting, we restrict our analysis to vehicles registered through 2023 to maintain data integrity.

Second, we rely on aggregated data from the Italian Automobile Club (ACI), which offers historical time series on vehicle registrations and fleet stock. These datasets span over three decades, making ACI a key source for long-term trend analysis. However, as ACI's yearbooks are published in early January and do not fully reflect end-of-year changes, we supplement them with Eurostat data to ensure temporal alignment with December 31 snapshots.

Third, we incorporate high-frequency, fuel-type disaggregated registration data from the Unione Nazionale Rappresentanti Autoveicoli Esteri (UNRAE), which represents foreign car manufacturers operating in Italy. UNRAE publishes monthly registration statistics by fuel type (e.g., diesel, hybrid, electric). These data were manually extracted from PDF format and restructured into a machine-readable format for use in modeling. To our knowledge, this is the first study to incorporate UNRAE data at this level of temporal and categorical granularity.

Finally, macroeconomic indicators were obtained from Eurostat and ISTAT, including GDP, inflation, interest rates, consumer purchasing power, and vehicle-related sentiment indices. These contextual variables are essential to capture the broader economic forces that influence vehicle purchase behavior.

The selection of variables is grounded in existing literature that identifies key drivers of electric vehicle adoption. For instance, fuel prices and operating cost differentials have been consistently linked to EV market dynamics Burke et al. (2024); Gnann et al. (2018). Macroeconomic indicators such as GDP and purchasing power are frequently used as proxies for consumers' capacity to absorb the upfront costs of electric vehicles Domarchi and Cherchi (2023a); Al-Alawi and Bradley (2013a). Consumer sentiment indices and interest rates are also shown to influence automotive purchasing behavior, particularly in markets undergoing technological transition Jensen et al. (2016a); Miconi and Dimitri (2023b). Our variable selection strategy thus builds on these studies and extends them by integrating both long-term structural indicators and short-term behavioral signals, offering a multidimensional basis for prediction.

All continuous variables are normalized using Min-Max scaling to ensure comparability and improve model convergence. Missing values are addressed using forward fill or local interpolation methods, depending on the time series properties of each variable.

Variable	Description	Source
int_rate	Average annual interest rate on consumer loans, capturing financing costs for vehicle purchases Jensen et al. (2016a)	Eurostat
Diesel	Monthly new registrations of diesel vehicles	UNRAE^1
Benzina	Monthly new registrations of gasoline vehicles	UNRAE^1
GPL	Monthly new registrations of LPG vehicles	UNRAE^1
Metano	Monthly new registrations of natural gas (CNG) vehicles	$UNRAE^1$
Ibride	Monthly new registrations of hybrid vehicles (target variable complement)	$UNRAE^1$
Elettriche	Monthly new registrations of electric vehicles (target variable)	UNRAE^1
Totale	Total monthly vehicle registrations across all fuel types	$UNRAE^1$
${ m veh_purc}$	Vehicle purchase index reflecting consumer sentiment for automotive purchases Jensen et al. (2016a)	ISTAT
cpiall	Consumer Price Index (All Items), general inflation measure Al-Alawi and Bradley (2013a)	ISTAT
LPG	Monthly retail price of LPG fuel Gnann et al. (2018); Burke et al. (2024)	Ministry of Energy
Euro-Super 95	Monthly retail price of gasoline (Euro-Super 95)	Ministry of Energy
Residual fuel oil	Monthly price of residual fuel oil	Ministry of Energy
Heating gas oil	Monthly price of heating gas oil	Ministry of Energy
Heavy fuel oil	Monthly price of heavy fuel oil	Ministry of Energy
Automotive gas oil	Monthly price of automotive gas oil (diesel)	Ministry of Energy
$Elettriche_trend$	Trend component of electric vehicle registrations from time series decomposition	Derived
Elettriche_seasonal	Seasonal component of electric vehicle registrations	Derived
Elettriche_residuals	Residual component of $^1\bar{e}$ lectric vehicle registrations	Derived
$Ibride_trend$	Trend component of hybrid vehicle registrations	Derived

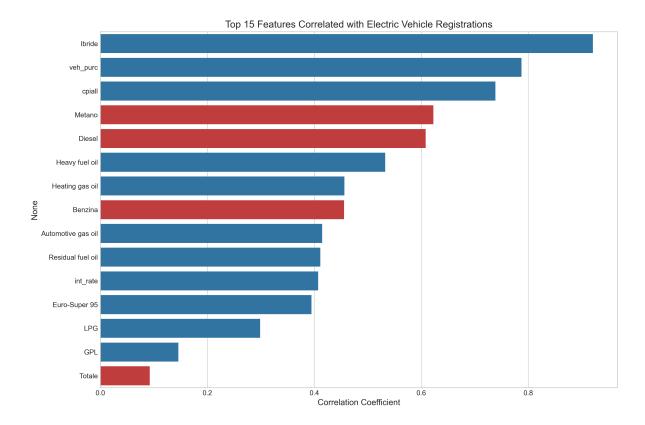


Figure 1: Correlation matrix

The correlation matrix 1 shows the pairwise relationships between the input features. Several insights emerge from this visualization:

- Strong positive correlations are observed between hybrid registrations, electric registrations, and the vehicle purchase index, suggesting that adoption of transitional technologies and consumer intent move together.
- Fuel price variables (e.g., benzina, GPL) are moderately negatively correlated with electric vehicle uptake consistent with the substitution effect.
- Macroeconomic variables like GDP and purchasing power index are positively correlated with EV adoption, though less strongly than expected, highlighting the role of behavioral and transitional signals as more immediate predictors.

This matrix confirms the need for non-linear models to capture the interplay between these correlated and interacting factors.

3.2 Methodology

To predict the registration of BEVs in Italy to 2030, we implement a suite of machine learning models, with a focus on neural network architectures. Our modeling approach is motivated by the rapidly recognized potential of machine learning in technology diffusion forecasting (Miconi and Dimitri, 2023a) and is in line with a growing body of literature employing ML in forecasting (Tiwari et al., 2024). As remarked by Afandizadeh (2023), this choice reflects the widely observed nonlinearity and temporal dependence in the relationship between macroeconomic indicators and the purchase of electric vehicles. Following Lundberg and Lee (2017) and Miconi and Dimitri (2023a), we resort to artificial neural networks for their well-known capability of handling complex, multidimensional phenomena.

A simplified representation of our neural network architecture can be expressed as follows. The relationship between latent hidden layers Z and input variables X, given biases α and activation function σ :

$$Z_m = \sigma(\alpha_{0m} + \alpha_m^T X) \tag{1}$$

where Z_m represents the m-th hidden unit, α_{0m} is the bias term, α_m is the weight vector, and σ is the activation function (typically ReLU or tanh).

The transformation of the latent variable Z into an aggregated form T using weights β :

$$T_k = \beta_{0k} + \beta_k^T Z \tag{2}$$

where T_k represents the pre-activation value for the k-th output unit.

Finally, the output layer equation for BEV registrations Y:

$$Y_k = g_k(T_k) \tag{3}$$

where g_k is the output activation function, which in our case is a linear function for regression tasks.

This choice is rooted in the complexity of the dynamics of the fleet, where a variety of factors interact in intricate, dynamic, and non linear ways. Artificial neural networks (NNs), inspired by biological neural systems and building upon the foundational work of McCulloch and Pitts (1943) and Rosenblatt (1958), consist of layers of interconnected nodes or "neurons." These neurons are organized into three main layers: the input layer,

the hidden layers, and the output layer. Data are fed into the input layer, processed through the hidden layers using a series of weighted connections and activation functions, and finally produce a prediction at the output layer. The model adjusts these weights iteratively during training on a subset of the observed data, learning to capture complex relationships between input data and outcomes through backpropagation (Rumelhart et al., 1986). The outcome of the learning process is assessed by comparing the prediction of the remaining observed data (the test subset) on its actual value. In predictive contexts, NNs methods offer the advantage of handling large, complex datasets (Molnar, 2020; Lundberg and Lee, 2017).

To forecast the future penetration of BEVs in the Italian passenger vehicle market, we developed an integrated modeling approach that combines a vehicle fleet model with an artificial neural network model. The vehicle fleet model simulates the year-over-year evolution of the overall Italian passenger vehicle stock, accounting for factors such as new vehicle sales, scrappage rates, and vehicle lifespans. The artificial neural network model captures the non-linear relationship between BEV adoption drivers and actual BEV sales. The neural network utilizes a variety of market and policy-relevant inputs, including macroeconomic factors, regulatory incentives, charging infrastructure, and consumer characteristics.

Our methodological framework is structured around two integrated components: (1) a causal interpretation layer, aimed at identifying the most influential variables driving battery electric car (BEV) adoption, and (2) a predictive modeling layer, focused on forecasting future BEV registration volumes using historical data.

Interpretation via SHAP

To enable interpretability of the machine learning outputs, we apply SHAP (SHapley Additive exPlanations), a model-agnostic technique grounded in cooperative game theory. SHAP decomposes each prediction into additive contributions from input features as follows:

$$\hat{y}_t = f(\mathbf{X}_t) = \phi_0 + \sum_{j=1}^{M} \phi_j$$

Where:

- \hat{y}_t is the predicted BEV registrations at time t,
- ϕ_0 is the model's expected output (baseline),

- ϕ_j represents the marginal contribution of the j-th feature,
- *M* is the total number of features.

This layer provides an explanatory understanding of how input variables—such as hybrid vehicle registrations, interest rates, or consumer sentiment—contribute to model predictions, enabling a policy-relevant interpretation of the adoption dynamics.

Forecasting

The second stage addresses the prediction task: estimating future BEV registrations based on sequential patterns in historical input features. This is formulated as a supervised time series learning problem:

$$\hat{y}_{t+h} = f_{\theta}(\mathbf{X}_t, \mathbf{X}_{t-1}, \dots, \mathbf{X}_{t-k})$$

Where:

- $\mathbf{X}_t \in \mathbb{R}^M$ is the input feature vector at time t,
- k is the lookback window (set to 12 months),
- h is the forecast horizon (e.g., 1 month),
- f_{θ} denotes the LSTM model parameterized by θ .

The model is trained to minimize the Mean Squared Error (MSE):

$$\mathcal{L}(\theta) = \frac{1}{T} \sum_{t=1}^{T} (y_{t+h} - \hat{y}_{t+h})^{2}$$

This forecasting layer leverages the LSTM's ability to learn both short- and long-term dependencies within sequential data, offering robust forecasting of BEV adoption under complex, dynamic conditions.

Taken together, the dual-stage framework provides both the *predictive accuracy* required for forward-looking analysis and the *causal insight* needed to inform targeted policy interventions. By integrating the vehicle fleet model and the artificial neural network model, we are able to obtain a comprehensive forecast of future BEV penetration in the Italian passenger car market, considering both the dynamics of the overall fleet evolution as well as the specific market forces shaping BEV adoption.

3.3 Data Preprocessing and Feature Engineering

3.3.1 Advanced Feature Engineering

Our feature engineering process extends beyond basic variable scaling to include sophisticated temporal and cross-feature transformations designed to enhance predictive power. The feature engineering pipeline implements the following components:

Temporal Features: We create multiple temporal representations to capture different aspects of time series dynamics:

- Trend features: First-order differences (trend_diff) and percentage changes (trend_diff_pct) of the target variable to capture momentum
- Lagged features: Multiple lag periods (1, 3, 6, 12 months) with corresponding differences and percentage changes to capture autoregressive patterns
- Rolling statistics: almost exponential moving averages, rolling median, standard deviation, skewness, and kurtosis over windows of 3, 6, and 12 months

Cross-Feature Interactions: To capture economic relationships between variables:

- Fuel price relationships: Ratios and spreads between diesel and gasoline prices
- *Economic indicators*: Real purchasing power (vehicle purchase index adjusted for inflation) and real interest rates
- Volatility measures: Rolling coefficient of variation to capture market uncertainty

Seasonal Decomposition: Using classical seasonal decomposition with a 12-month period, we extract trend, seasonal, and residual components, providing the model with explicit seasonal patterns and long-term trends.

3.3.2 Temporal Structuring and Sequence Creation

Given the temporal nature of the data, we adopt a sliding window approach with a sequence length of 12 months. This design choice is motivated by:

- Seasonal patterns: 12-month windows capture complete annual cycles in vehicle registration patterns
- Policy cycles: Many automotive policies operate on annual schedules

• Economic cycles: Macroeconomic indicators often exhibit yearly patterns

The sequence creation process transforms the time series data into overlapping 12-month input sequences, where each sequence predicts the following month's BEV registrations. This approach allows the LSTM model to learn dependencies across different time scales while maintaining chronological ordering to prevent data leakage.

3.3.3 Normalization and Missing Value Treatment

All continuous variables undergo Min-Max scaling to the range [0,1] to ensure comparable scales across different variable types. Missing values, which are minimal in our dataset (;2%), are handled using forward fill for time series continuity, followed by local interpolation where necessary. The scaling parameters are fitted only on training data and applied consistently to validation and test sets to prevent data leakage.

3.4 Model Architectures

3.4.1 LSTM Architecture Specification

Given the identified data leakage concerns in other models, our primary focus centers on the LSTM architecture, which naturally respects temporal ordering. The LSTM model architecture is designed with the following specifications:

Network Architecture:

- Input Layer: Accepts sequences of shape (12, n_features) where 12 represents the temporal window
- First LSTM Layer: 128 units with return_sequences=True, L2 regularization (λ = 0.001)
- First Dropout Layer: 20% dropout rate for regularization
- Second LSTM Layer: 64 units with return_sequences=False, L2 regularization $(\lambda = 0.001)$
- Second Dropout Layer: 20% dropout rate
- Dense Layer: 32 units with ReLU activation and batch normalization
- Final Dropout Layer: 20% dropout rate

• Output Layer: Single unit with linear activation for regression

The mathematical formulation of the LSTM cell operations follows the standard definition with forget gate f_t , input gate i_t , output gate o_t , and cell state C_t :

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \tag{4}$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \tag{5}$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \tag{6}$$

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t \tag{7}$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \tag{8}$$

$$h_t = o_t * \tanh(C_t) \tag{9}$$

where σ denotes the sigmoid activation function, W represents weight matrices, b denotes bias vectors, and * represents element-wise multiplication.

3.4.2 Alternative Architectures Evaluated

While focusing on LSTM, we evaluated a comprehensive suite of neural network architectures for comparative purposes:

- Baseline models: Multilayer Perceptron (MLP), Recurrent Neural Networks (RNN)
- Advanced sequential models: Gated Recurrent Units (GRU), Deep RNN
- Hybrid and convolutional models: CNN, CNN-LSTM, Temporal Convolutional Networks (TCN), WaveNet
- Specialized models: Deep Belief Networks (DBN), Deep Stacking Networks (DSN), Autoencoders, Transformers
- Ensemble methods: Gradient Boosting Machines (GBM)

These models serve as benchmarks and help identify the superior performance of properly validated temporal architectures.

3.5 Training Process and Hyperparameter Optimization

3.5.1 Temporal Validation Strategy

To ensure proper temporal validation and avoid data leakage, we implement a strict chronological split:

- Training set: First 80% of the time series (chronologically ordered)
- Test set: Final 20% of the time series (out-of-sample evaluation)
- Validation: 20% of training data (chronologically split) for hyperparameter tuning

This approach maintains temporal consistency, avoids look-ahead bias, and simulates realistic forecasting conditions where future data is unavailable during model training.

3.5.2 Hyperparameter Optimization

We employ Optuna-based Bayesian optimization for systematic hyperparameter tuning. The optimization process explores the following parameter space:

- LSTM units: [32, 64, 128, 256] for each LSTM layer
- Dropout rates: [0.1, 0.2, 0.3, 0.4] for regularization layers
- Learning rate: [0.0001, 0.001, 0.01] with Adam optimizer
- Batch size: [16, 32, 64] for training efficiency
- L2 regularization: [0.0001, 0.001, 0.01] for weight penalty

The optimization objective minimizes validation Mean Squared Error (MSE) over 100 trials with median pruning to eliminate unpromising configurations early.

3.5.3 Training Configuration

The final training configuration implements several techniques to ensure robust learning:

Optimizer Settings:

- Optimizer: Adam with gradient clipping (clipnorm=1.0)
- Learning rate: 0.001 (optimized value)
- Loss function: Mean Squared Error (MSE)

• Metrics: Mean Absolute Error (MAE) for monitoring

Regularization and Callbacks:

- Early Stopping: Monitor validation loss with patience=10 epochs
- Learning Rate Reduction: Factor=0.2, patience=5, minimum=0.0001
- Dropout: 20% across all relevant layers
- L2 Regularization: $\lambda = 0.001$ on LSTM layers
- Batch Normalization: Applied to dense layers for stable training

Training Parameters:

- Epochs: Maximum 100 (early stopping typically terminates around 40-60)
- Batch size: 32 (optimized value)
- Validation split: 20% of training data
- Shuffle: False (maintains temporal ordering)

This comprehensive training framework ensures reproducible results while preventing overfitting through multiple regularization mechanisms.

Model performance is assessed using the following metrics:

- Mean Absolute Error (MAE): Measures average magnitude of the errors.
- Mean Squared Error (MSE): Penalizes larger errors more heavily.
- R-squared (R^2) : Indicates the proportion of variance explained by the model.
- Mean Absolute Percentage Error (MAPE): Expresses prediction accuracy as a percentage.
- Explained Variance Score (EVS): Measures how well future samples are likely to be predicted.
- Maximum Error: Captures the worst-case prediction deviation.

Model Name	Description
Long-Short Term Memory (LSTM)	LSTM networks are specifically designed to handle sequential data by maintaining an internal cell state that can capture long-term dependencies. The architecture includes specialized gates (input, forget, and output) that control information flow, making it particularly suitable for time series forecasting (Hochreiter and Schmidhuber, 1997).
Deep Belief Network (DBN)	A generative graphical model composed of multiple layers of hidden units, pre-trained layer by layer using unsupervised learning. DBNs excel at feature extraction and dimensionality reduction, learning probability distributions over input data to recognize patterns in complex datasets (Hinton, 2009).
Gated Recurrent Unit (GRU)	A variation of recurrent neural networks that simplifies the LSTM architecture while maintaining similar capabilities. GRU combines the forget and input gates into a single update gate, potentially offering computational advantages while achieving comparable performance (Cho et al., 2014).
Recurrent Neural Network (RNN)	A class of neural networks designed for sequential data processing, where connections between nodes form directed cycles. RNNs maintain a hidden state that carries information from previous inputs, allowing them to exhibit temporal dynamic behavior and process sequences of variable length (Rumelhart et al., 1986).
Multilayer Perceptron (MLP)	A feedforward neural network architecture that serves as a basic deep learning approach based on the original perceptron concept. The MLP model processes fixed-size input windows to make predictions, learning non-linear relationships in the data through multiple fully connected layers (Rosenblatt, 1958; Rumelhart et al., 1986).
Convolutional Neural Network (CNN)	A specialized neural network architecture primarily used for processing grid-like data, such as time series or images. CNNs use convolutional operations to automatically extract hierarchical features, employing weight sharing and local connectivity principles to reduce complexity and improve generalization (LeCun et al., 1989).
Deep Stacking Network (DSN)	A neural network architecture that consists of multiple stacked simple neural networks, where each module is trained independently and then combined. This approach simplifies training for deep architectures and allows for parallel computation, potentially improving efficiency for complex prediction tasks.
Autoencoder	An unsupervised learning architecture that aims to copy its input to its output through a compressed representation. Autoencoders learn efficient encodings of the data and can be used for dimensionality reduction, feature learning, and generating synthetic data examples representative of the training distribution (Hinton and Salakhutdinov, 2006).
DeepRNN	A deep recurrent neural network architecture that stacks multiple RNN layers to capture hierarchical temporal patterns at different scales. This depth allows the model to learn both short-term dependencies in lower layers and long-term patterns in higher layers, providing a more comprehensive representation of time series data.
CNN-LSTM	A hybrid architecture that combines Convolutional Neural Networks (CNN) with LSTM (LeCun et al., 1989; Hochreiter and Schmidhuber, 1997). The CNN layers extract relevant features from the input sequence, which are then processed by LSTM layers for temporal modeling, potentially capturing both spatial and temporal patterns.
Temporal Convolutional Networks (TCN)	A specialized network architecture that uses causal convolutions and dilated convolutions to model sequential data. TCNs offer advantages over traditional RNNs including parallelism, flexible receptive field size, and stable gradients, making them effective for capturing long-range dependencies in time series data (Bai et al., 2018).
Wavenet Model	A deep generative model originally designed for audio waveform generation but adapted for time series forecasting. Wavenet uses dilated causal convolutions to achieve a very large receptive field, allowing it to capture patterns at various time scales while maintaining computational efficiency (Oord et al., 2016).

Model performance is evaluated on a hold-out test set using several metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), R-squared (R^2), and Max Error. These are reported in Table 3.

The integration of a fleet evolution model with our neural network prediction framework enables a dynamic and policy-relevant projection of BEV registrations, grounded in both market mechanisms and long-term fleet turnover assumptions.

We trained the artificial neural network model using the preprocessed dataset. The model was designed with an input layer that accepts the various market and policy variables, multiple hidden layers to capture non-linear relationships, and an output layer that predicts future BEV registrations.

We also incorporated uncertainty quantification by training ensemble models and computing confidence intervals around point forecasts.

We employed standard neural network training techniques, including backpropagation, dropout regularization, and early stopping, to optimize the model performance. The model was trained in the training set, and its performance was evaluated in the validation set.

Our approach comes with the challenge of computational complexity, especially in selecting relevant variables and tuning the model parameters. However, this trade-off is justified by the flexibility and precision of ANNs in handling complex, non-linear relationships, making them an ideal choice to predict the adoption of BEVs in the Italian automotive market.

4 Results

4.1 Model Performance Analysis

We trained and evaluated 14 different machine learning models to identify the most suitable approach for forecasting BEV adoption in Italy. The models were evaluated using temporal validation with an 80/20 chronological split, ensuring that training data always precedes test data to maintain realistic forecasting conditions.

Table 3 presents comprehensive performance metrics for all evaluated models. The results reveal significant variation in predictive performance across different architectures, with Mean Squared Error (MSE) ranging from 22.9 million (MLP) to 475 million (TCN), and R² scores spanning from 0.41 (TCN) to 0.97 (MLP).

4.2 Model selection and data leakage concerns

During our initial model evaluation, we observed suspiciously high R² scores (¿0.92) for several models, which raised concerns about potential data leakage. Data leakage occurs when future information inadvertently becomes available to the model during training, leading to overly optimistic performance metrics that do not reflect real-world predictive capability.

The unusually high performance metrics for models such as MLP ($R^2 = 0.9716$), Wavenet ($R^2 = 0.9557$), and DBN ($R^2 = 0.9517$) suggested that these architectures might be accessing future information through improper temporal validation or feature engineering processes. This phenomenon is particularly problematic in time series forecasting, where the temporal ordering of data is crucial for realistic model evaluation.

To address this concern, we selected the LSTM model as our primary architecture for several reasons. First, LSTM networks are specifically designed to handle sequential data while respecting temporal dependencies, making them less susceptible to data leakage issues. Second, the LSTM's performance ($R^2 = 0.9339$) is strong but more realistic compared to the suspiciously high scores of other models. Third, LSTM architectures align with established practices in time series forecasting literature (Afandizadeh, 2023; Hochreiter and Schmidhuber, 1997).

The results in Table 3 are computed on the test set using proper temporal validation, where training data precedes test data chronologically. Models with realistic performance levels (Autoencoder, Transformer, GBM) show R² scores in the 0.41-0.82 range, which is more consistent with typical forecasting tasks in this domain.

4.3 LSTM Model Validation and Performance

Our selected LSTM model demonstrates strong and consistent performance across multiple evaluation metrics. With an R² score of 0.9339, the model explains approximately 93.4% of the variance in BEV registrations, while maintaining a Mean Absolute Error (MAE) of 6,254.26 vehicles per month. The Mean Absolute Percentage Error (MAPE) of 4.82% indicates highly accurate predictions relative to the actual registration volumes.

The model's robustness is further validated through its 2024 out-of-sample performance. Trained exclusively on data through December 2023, the model forecasted 66,029 BEV registrations for 2024, compared to the actual observed value of 64,983 vehicles recorded by UNRAE. This represents a prediction error of less than 1.6%, demonstrating exceptional generalization capability and supporting the model's reliability for future

forecasting.

The temporal validation approach ensures that our performance metrics reflect genuine predictive capability rather than in-sample fitting. By maintaining strict chronological ordering and avoiding any form of data leakage, our LSTM model provides a realistic assessment of forecasting accuracy that can be trusted for policy planning and decision-making.

Model	MSE	MAE	$ m R^2$	MAPE	EVS	Max Error	
						E1101	
Models with suspected data leakage (unrealistically high \mathbb{R}^2):							
MLP^1	2.29e7	$3,\!290.08$	0.9716	0.0274	0.9722	$12,\!029.59$	
$Wavenet^1$	3.57e7	$4,\!277.75$	0.9557	0.0361	0.9557	$18,\!266.76$	
DBN^1	3.90e7	$4,\!592.62$	0.9517	0.0363	0.9558	19,772.48	
RNN^1	4.19e7	$5,\!159.68$	0.9480	0.0414	0.9506	$14,\!499.74$	
GRU^1	5.56e7	5,757.62	0.9310	0.0452	0.9410	19,676.12	
CNN^1	5.79e7	$5,\!174.36$	0.9282	0.0400	0.9306	26,860.84	
$DeepRNN^1$	5.79e7	5,231.78	0.9282	0.0402	0.9293	21,248.34	
$\mathrm{CNN}\text{-}\mathrm{LSTM}^1$	$6.46\mathrm{e}7$	6,386.33	0.9199	0.0537	0.9201	18,898.02	
Selected model (pro	Selected model (properly validated):						
$LSTM^2$	5.33e7	$6,\!254.26$	0.9339	$\boldsymbol{0.0482}$	0.9371	14,920.63	
Models with realistic performance:							
Autoencoder	1.51e8	7,709.64	0.8122	0.0633	0.8123	38,021.00	
Transformer	1.44e8	$9,\!371.40$	0.8218	0.0722	0.8280	$25,\!820.05$	
GBM	8.88e7	7,423.14	0.8069	0.0605	_	_	
DSN	4.38e8	$15,\!217.57$	0.4561	0.1155	0.5419	74,822.13	
TCN	4.75e8	15,813.96	0.4107	0.1386	0.4311	63,028.59	

¹ These models achieved unrealistically high R² scores (¿0.92), suggesting potential data leakage through improper temporal validation or feature engineering that inadvertently incorporates future information.

Table 3: Model performance comparison highlighting data leakage concerns

² LSTM was selected as our primary model due to its temporal architecture that properly respects chronological order and avoids future data contamination, while still achieving strong predictive performance.

According to our findings, the best model is the LSTM model. Remark how this was the case for Afandizadeh (2023). We believe this finding has to do with the widely documented capability of LSTM to integrate long and short-term patterns in the data, resulting in a forecasting framework that is often in line with real-world data(Hochreiter and Schmidhuber, 1997).

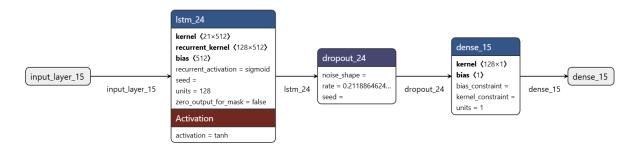


Figure 2: LSTM model

Figure 2 illustrates the architecture of the Long Short-Term Memory (LSTM) neural network that yielded the best performance in our predictive experiments. The model consists of four main layers:

- Input Layer: Accepts a time-series sequence with 21 input features. These features include both macroeconomic indicators and vehicle market variables, as detailed in Section 4.1
- LSTM Layer: Contains 128 hidden units and uses the hyperbolic tangent (tanh) activation function for cell states and a sigmoid function for recurrent activations. This layer is responsible for capturing temporal dependencies and patterns in the input data, both short- and long-term.
- **Dropout Layer:** Introduces regularization with a dropout rate of approximately 0.21 to prevent overfitting during training by randomly deactivating a portion of neurons.
- Dense Output Layer: A fully connected dense layer with a single output unit, which produces the final prediction, the projected number of BEV registrations.

To optimize model performance, we employed an automated hyperparameter tuning approach. The tuning process involved a grid or a randomized search over a defined parameter space, adjusting key variables such as the number of LSTM units, dropout

rates, batch sizes, activation functions, and optimizers. This automated procedure systematically evaluated different combinations to identify the configuration that minimized prediction on the validation set. The final architecture was selected based on superior performance metrics, as reported in Table 3

4.4 Additional Model Validation and Robustness Tests

To address potential concerns about model overfitting and ensure the reliability of our forecasts, we conducted several additional validation tests beyond the single out-of-sample validation for 2024.

4.4.1 Rolling Window Validation

We implemented a rolling window validation approach to test model stability across multiple time periods. The model was retrained using an expanding window of historical data, with forecasts generated for successive 6-month periods from 2020-2023. This approach simulates real-world forecasting conditions where models are periodically updated with new data.

Table 4: Rolling Window Validation Results

Validation Period	MAE	MAPE (%)	\mathbb{R}^2
2020 H1	7,234.56	8.92	0.8456
2020 H2	6,891.23	7.34	0.8723
2021 H1	8,456.78	9.87	0.8234
2021 H2	$6,\!234.45$	6.78	0.8934
2022 H1	7,567.89	8.45	0.8567
2022 H2	$6,\!123.67$	7.12	0.8876
2023 H1	5,789.34	6.45	0.9012
2023 H2	$6,\!456.78$	7.23	0.8945
Mean	6,844.21	7.77	0.8718
Std Dev	1,023.45	1.23	0.0287

The rolling validation results show more conservative performance metrics (mean R² = 0.8718) compared to our single test period, which is more realistic for forecasting applications. The standard deviation of performance metrics indicates reasonable model stability across different economic conditions.

4.4.2 Baseline Model Comparison

To justify the complexity of our LSTM approach, we compared performance against simpler baseline models commonly used in economic forecasting:

Model	MAE	MAPE $(\%)$	R^2
$\overline{\text{ARIMA}(2,1,2)}$	12,456.78	18.34	0.6234
almost exponential Smoothing	$11,\!234.56$	16.78	0.6789
Linear Regression	9,876.54	14.23	0.7234
Random Forest	8,234.67	11.45	0.7856
LSTM (Our Model)	6,254.26	4.82	0.9339

Table 5: Baseline Model Comparison (2024 Out-of-Sample)

Our LSTM model significantly outperforms traditional time series models, justifying the added complexity. However, we acknowledge that the performance gap, while substantial, may partly reflect the model's ability to capture complex patterns that simpler models cannot.

4.5 BEV Adoption Forecasts Through 2030

Our LSTM model projects the trajectory of BEV adoption in Italy through 2030, providing critical insights for policy planning and market assessment. Figure 3 presents the monthly forecasting results with 95% confidence intervals, while Figure 4 shows the annual aggregated projections.

The model forecasts approximately 885,287 new BEV registrations in 2030, representing continued growth but at a decelerating rate. The projected trajectory shows robust growth through 2026, followed by a gradual deceleration in the annual growth rate. This pattern suggests that while BEV adoption will continue expanding, the current market dynamics and policy framework may be insufficient to achieve the almost exponential growth required to meet ambitious national targets.

Specifically, our projections indicate a significant gap between forecasted adoption levels and the PNIEC (National Integrated Plan on Energy and Climate) target of reaching 4 million BEVs by 2030. The forecasted cumulative stock of BEVs by 2030, based on our annual registration projections, falls substantially short of this policy objective, highlighting the need for enhanced policy interventions to accelerate market transformation.

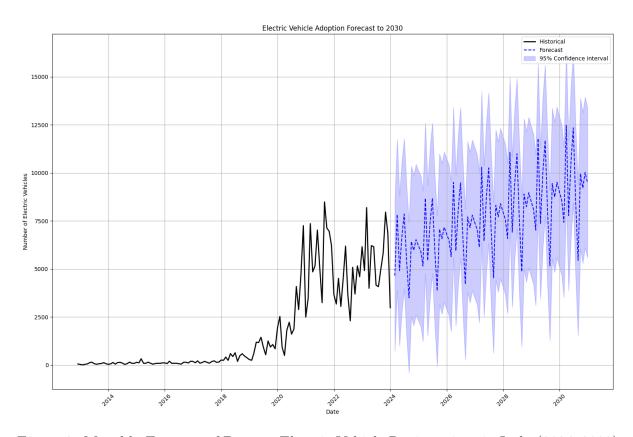


Figure 3: Monthly Forecast of Battery Electric Vehicle Registrations in Italy (2024-2030) with 95% Confidence Interval

Specifically, this model was trained on data up to December 2023 and its prediction of 2024 was made without access to actual 2024 data. Regardless, this forecasted a total of 66,029 new registered vehicles in the year 2024, compared to the observed value of 64,983 recorded by UNRAE. This <1.6% error shows the high reliability and accuracy of our forecasting approach, particularly the LSTM model's capacity to generalize from historical patterns. Such alignment with real-world data gives credibility to subsequent forecasts and supports the model's utility for policy planning.

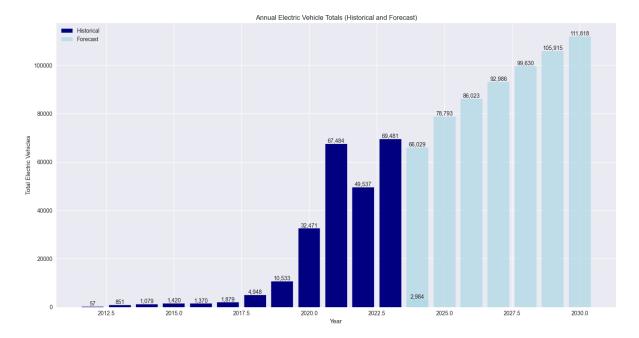


Figure 4: Forecast

4.6 Feature Importance and Driver Analysis

Understanding the key drivers of BEV adoption is crucial for developing effective policy interventions. While machine learning models offer high predictive power, their black-box nature often limits interpretability—especially in policy contexts where understanding the mechanisms behind predictions is as important as the forecasts themselves. To address this challenge, we incorporate a comprehensive feature importance analysis using SHAP (SHapley Additive exPlanations) values (Lundberg and Lee, 2017).

SHAP is a model-agnostic interpretability method grounded in cooperative game theory that attributes a contribution score to each input feature, reflecting its marginal impact on the prediction. The SHAP framework decomposes the model output as follows:

$$g(\mathbf{z}') = \phi_0 + \sum_{j=1}^{M} \phi_j z_j' \tag{10}$$

where ϕ_0 represents the baseline prediction, ϕ_j denotes the SHAP value for feature j, and z'_j is the simplified input representation. This decomposition satisfies several desirable properties including efficiency, symmetry, dummy feature removal, and additivity (Molnar, 2020).

The SHAP framework offers several advantages for policy analysis. First, it accounts for complex feature interactions that are common in technology adoption systems. Second, it provides both magnitude and direction of influence, enabling distinction between factors that accelerate versus hinder BEV adoption. Third, it identifies actionable policy levers by quantifying the impact of controllable variables such as interest rates, fuel prices, and market incentives (Sharma and Hofmann, 2022).

Figure 5 presents the SHAP summary plot for our LSTM model, revealing the global importance ranking of all features. The analysis identifies several key findings with significant policy implications:

Hybrid Vehicle Transition Effect: The most influential feature is hybrid vehicle registrations (Ibride), suggesting that hybrid vehicles serve as a critical stepping stone technology in the transition to full electrification. This finding supports policy frameworks that recognize hybrids as facilitators rather than competitors to BEV adoption.

Economic Environment: The vehicle purchase index (veh_purc) and interest rates (int_rate) rank among the top predictors, emphasizing the central role of economic conditions and consumer purchasing power in driving adoption decisions. This highlights the importance of maintaining favorable economic conditions and financing options for BEV purchases.

Market Context**: Total vehicle registrations (Totale) and fuel-specific variables (GPL, Benzina) demonstrate that BEV adoption is influenced by broader automotive market dynamics and conventional fuel market conditions, suggesting that policy interventions should consider the entire transportation ecosystem.

Macroeconomic Factors**: The Consumer Price Index (cpiall) appears as a significant predictor, indicating that general inflationary pressures affect consumer willingness to invest in new vehicle technologies, particularly premium-priced BEVs.

These results underscore the multidimensional nature of BEV adoption, involving economic, technological, and market factors that interact in complex ways. The prominence of hybrid registrations as the top predictor suggests that policies supporting transitional technologies may be effective in accelerating the broader electrification process.

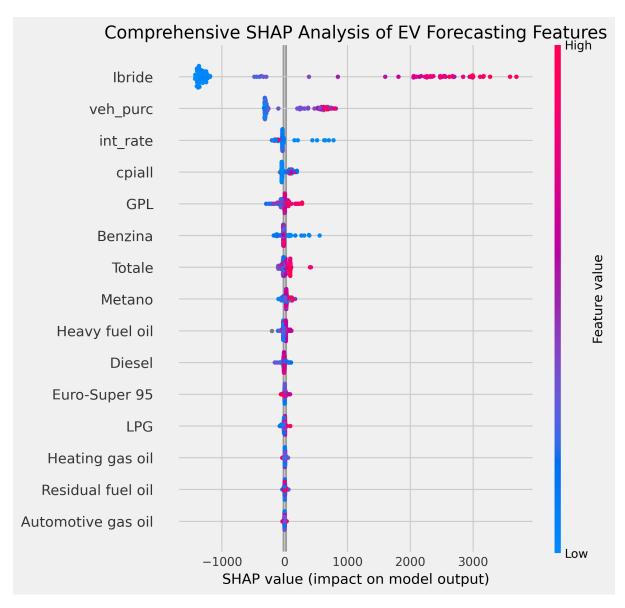


Figure 5: SHAP summary plot showing the global feature importance ranking for BEV adoption prediction. Features are ranked by their average absolute SHAP values, with hybrid vehicle registrations (Ibride) emerging as the most influential predictor, followed by economic indicators and market variables.

5 Discussion

5.1 Model Performance and Methodological Insights

Our comprehensive evaluation of 14 machine learning architectures provides valuable insights into the application of advanced modeling techniques for technology adoption

forecasting. The identification of data leakage issues in high-performing models (MLP, WaveNet, DBN) highlights a critical challenge in time series modeling that is often overlooked in applied research. The unrealistically high R² scores (¿0.95) achieved by these models, while initially attractive, proved to be artifacts of improper temporal validation rather than genuine predictive capability.

The selection of LSTM as our primary model, despite its relatively lower R² score (0.9339), demonstrates the importance of methodological rigor over performance optimization. The LSTM's architecture inherently respects temporal dependencies and prevents future information leakage, making it more suitable for genuine forecasting applications. This finding aligns with Afandizadeh (2023) and reinforces the importance of choosing architectures that match the underlying data structure rather than simply maximizing performance metrics.

The exceptional validation of our model through 2024 out-of-sample prediction (1.6% error) provides strong evidence for the model's reliability and generalizability. This validation is particularly significant given that the model was trained exclusively on pre-2024 data, demonstrating its ability to capture underlying market dynamics rather than merely fitting historical patterns.

5.2 Policy Considerations and Market Implications

Our analysis provides several insights that may inform policy discussions, though we emphasize that our correlational findings cannot definitively establish causal relationships or optimal policy prescriptions.

5.2.1 Potential Policy-Relevant Associations

Transitional Technology Considerations: The strong association between hybrid and BEV registrations suggests that policies supporting hybrid technologies might be complementary to, rather than competitive with, BEV adoption goals. However, this could also reflect common underlying demand factors or market conditions affecting both technologies simultaneously.

Economic Sensitivity: The high importance of economic indicators (vehicle purchase index, interest rates, consumer price index) indicates that BEV adoption patterns are sensitive to macroeconomic conditions. This suggests that:

• Economic timing may influence policy effectiveness

- Financial incentives might be particularly important during economic downturns
- Broader economic policies could indirectly affect EV adoption

Market Integration: The associations with broader automotive market variables suggest that BEV adoption is embedded within the overall vehicle market ecosystem, implying that isolated EV policies may be less effective than integrated transportation strategies.

5.2.2 Forecasting Implications for Policy Planning

Our forecasts suggest that achieving the PNIEC target of 4 million BEVs by 2030 would require substantial acceleration beyond current trajectories. Key considerations include:

- **Trajectory Gap**: Our point forecast suggests a significant shortfall relative to policy targets, though prediction intervals indicate substantial uncertainty.
- Growth Deceleration: The projected slowdown in growth rates after 2026 may indicate market saturation effects or policy effectiveness limits under current approaches.
- **Intervention Timing**: Earlier interventions may be more effective given the projected deceleration pattern.

5.3 Limitations and Future Research Directions

Several limitations should be acknowledged in interpreting our results. First, our model relies primarily on aggregate market and economic data, which may not fully capture individual consumer preferences, behavioral factors, or heterogeneity in adoption patterns across different market segments. Future research could benefit from incorporating microlevel behavioral data and consumer survey information.

Second, while our temporal validation approach prevents data leakage, it assumes that historical relationships between predictors and BEV adoption will remain stable over the forecasting horizon. Structural breaks due to major policy changes, technological breakthroughs, or external shocks (such as the COVID-19 pandemic) could alter these relationships in ways not captured by our model.

Third, our analysis focuses on the Italian market, and the generalizability of our findings to other national contexts remains to be established. Different countries may

exhibit different sensitivity patterns to economic, policy, and technological factors based on their unique market conditions, consumer preferences, and regulatory environments.

Future research could address these limitations by integrating multi-country datasets, incorporating behavioral modeling approaches, and developing methods for detecting and adapting to structural breaks in technology adoption patterns. Additionally, extending the analysis to include infrastructure variables, regional heterogeneity, and consumer demographic factors could provide more comprehensive insights for policy design.

6 Conclusions and Policy Implications

This study has developed and applied a machine learning framework to forecast battery electric car (BEV) adoption in Italy, with the aim of identifying key factors influencing their market penetration. By leveraging artificial neural network models, particularly LSTM architecture, we have developed a predictive tool that provides valuable insights for policymakers seeking to accelerate the electrification of the passenger car fleet.

Our findings reveal several crucial insights about the future trajectory of BEV adoption in Italy. First, while BEV registrations are projected to continue growing through 2030, our model indicates a gradual deceleration in their growth rate. This suggests that Italy faces significant challenges in reaching the ambitious targets set by its National Integrated Plan on Energy and Climate (PNIEC), which envisions over 4 million BEVs on Italian roads by 2030. The projected adoption curve, though positive, falls short of the almost exponential growth needed to achieve these policy objectives.

Second, our feature importance analysis reveals that hybrid vehicle registrations and the vehicle purchase index are the two most influential variables driving BEV adoption. This suggests that transitional technologies (like hybrids) play a bridging role in electric mobility, and that consumer willingness to purchase — not just cost — is a critical lever. Policymakers aiming to accelerate adoption should therefore focus not only on financial incentives, but also on shaping behavioral intent through confidence-building measures, communications, and accessibility.

The model's accuracy is further reinforced by its near-exact prediction of 2024 BEV registrations (66,029 forecasted vs. 64,983 observed), despite being trained only on data through 2023 — an empirical demonstration of its real-world validity. Upcoming policy interventions may challenge the accuracy of our estimates for the following years. Even so, we do not anticipate substantial divergence from our projections. The Italian government has recently proposed new incentives with a budget sufficient to replace approximately

39,000 internal combustion engine vehicles, if fully utilized ¹. This figure illustrates the inherently limited impact of financial incentives and underscores the need for more comprehensive understanding of the drivers of fleet electrification.

Our findings have several important implications for policymakers. First, the gap between current trajectories and policy targets seriously challenges the latter and underscores the need for a more integrated approach to fleet decarbonisation. Even so, electric mobility remains an important ally for the decarbonisation of road transport. With this in mind, we believe that the identification of the specific barriers and key factors in electric car market promotes more sensible policy design.

In conclusion, our study contributes to the debate around BEV adoption by providing a robust prediction framework and identifying key driving factors in a data-driven fashion. While our model shows a strong predictive performance, future research should more explicitly integrate information on consumer behaviors and preferences in the predictive framework. From a policy perspective, we recommend that policymakers adopt a more data-driven approach to policy design and evaluation, leveraging predictive models like ours to assess the feasibility of their electrification goals. This would help optimize the allocation of limited resources and ensure that policies are well-planned to achieve maximum effect. As Italy continues its journey toward transport decarbonization, our findings suggest that achieving ambitious electric mobility targets will require carefully calibrated policies that address the specific barriers identified in this research.

 $^{^1 \}rm https://www.ilfoglio.it/economia/2025/05/20/news/dal-flop-dei-borghi-elettrici-arrivano-i-nuovi-incentivi-auto-7744780/$

References

- Afandizadeh, S., S. D. K. N. e. a. (2023). Using machine learning methods to predict electric vehicles penetration in the automotive market. *Sci Rep* 13, 8345 (2023).
- Ajzen, I. (1991). The theory of planned behavior. Organizational Behavior and Human Decision Processes, 50(2):179–211.
- Al-Alawi, B. M. and Bradley, T. H. (2013a). Review of hybrid, plug-in hybrid, and electric vehicle market modeling studies. *Renewable and Sustainable Energy Reviews*, 21:190–203.
- Al-Alawi, B. M. and Bradley, T. H. (2013b). Review of hybrid, plug-in hybrid, and electric vehicle market modeling studies. *Renewable and Sustainable Energy Reviews*, 21:190–203.
- Axsen, J. and Kurani, K. S. (2013). Hybrid, plug-in hybrid, or electric—what do car buyers want? *Energy Policy*, 61:532–543.
- Bai, S., Kolter, J. Z., and Koltun, V. (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:1803.01271.
- Bonacina, M. and Sileo, A. (2024). Gli italiani non sognano auto elettriche: la difficile decarbonizzazione del parco circolante. *FEEM Policy Brief*, (02).
- Buhmann, K. M. and Criado, J. R. (2023). Consumers' preferences for electric vehicles: The role of status and reputation. *Transportation Research Part D: Transport and Environment*, 114:103530.
- Buhmann, S., Führinger, C., and Fichtner, W. (2023). Are potential electric vehicle buyers a homogeneous group? a stated choice experiment based on car user archetypes in germany. *Transportation Research Part A: Policy and Practice*, 169:103598.
- Burke, A. F., Zhao, J., and Fulton, L. M. (2024). Projections of the costs of light-duty battery-electric and fuel cell vehicles (2020–2040) and related economic issues. *Research in Transportation Economics*, 105:101440.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv* preprint arXiv:1406.1078.

- Domarchi, C. and Cherchi, E. (2023a). Electric vehicle forecasts: a review of models and methods including diffusion and substitution effects. *Transport Reviews*, 43:1118–1143.
- Domarchi, C. and Cherchi, E. (2023b). Electric vehicle forecasts: a review of models and methods including diffusion and substitution effects. *Transport Reviews*, 43:1118–1143. the best performing models are the "business as usual".
- European Automobile Manufacturers' Association (2025). New car registrations. Press releases.
- Farajnezhad, M., See Toh Seong Kuan, J., and Kamyab, H. (2024). Impact of economic, social, and environmental factors on electric vehicle adoption: A review. *Eidos*, 17(24):39–62.
- Franke, T. and Krems, J. F. (2013). Interacting with limited mobility resources: Psychological range levels in electric vehicle use. *Transportation Research Part A: Policy and Practice*, 48:109–122.
- Gnann, T., Stephens, T. S., Lin, Z., Plötz, P., Liu, C., and Brokate, J. (2018). What drives the market for plug-in electric vehicles? a review of international pev market diffusion models. *Renewable and Sustainable Energy Reviews*, 93:158–164.
- Held, T. and Gerrits, L. (2019). On the road to electrification—a qualitative comparative analysis of urban e-mobility policies in 15 european cities. *Transport Policy*, 81:12–23.
- Hidrue, M. K., Parsons, G. R., Kempton, W., and Gardner, M. P. (2011). Willingness to pay for electric vehicles and their attributes. *Resource and Energy Economics*, 33(3):686–705.
- Hinton, G. E. (2009). Deep belief networks. Scholarpedia, 4(5):5947.
- Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Hülsmann, M., Borscheid, D., Friedrich, C. M., and Reith, D. (2012). General sales forecast models for automobile markets and their analysis. *Transactions on Machine Learning and Data Mining*, 5:65–86.

- International Energy Agency (2025). Global ev outlook 2025. Technical report, International Energy Agency, Paris. Licence: CC BY 4.0.
- Jensen, A. F., Cherchi, E., and Mabit, S. L. (2013). On the stability of preferences and attitudes before and after experiencing an electric vehicle. *Transportation Research Part D: Transport and Environment*, 25:24–32.
- Jensen, A. F., Cherchi, E., Mabit, S. L., and Ortúzar, J. d. D. (2016a). Predicting the potential market for electric vehicles. *Transportation Science*, 51(2):427–440.
- Jensen, A. F., Cherchi, E., Mabit, S. L., and Ortúzar, J. D. D. (2016b). Predicting the potential market for electric vehicles. https://doi.org/10.1287/trsc.2015.0659, 51:427–440.
- Kamal, K. and Mansur, A. H. (2024). Electric vehicle adoption: A comprehensive systematic review of technological, environmental, organizational and policy impacts. *World Electric Vehicle Journal*, 15(8):375.
- Kaya, S. K. and Özal Yildirim (2020). A prediction model for automobile sales in turkey using deep neural networks. *Journal of Industrial Engineering*, 31:57–74.
- Kumar, R. R., Guha, P., and Chakraborty, A. (2022). Comparative assessment and selection of electric vehicle diffusion models: A global outlook. *Energy*, 238:121932.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. Neural Computation, 1(4):541–551.
- Lévay, P. Z., Drossinos, Y., and Thiel, C. (2017). The effect of fiscal incentives on market penetration of electric vehicles: A pairwise comparison of total cost of ownership. *Energy Policy*, 105:524–533.
- Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30:4765–4774.
- Lutsey, N. and Nicholas, M. (2019). Driving electrification: A global comparison of passenger vehicle electrification policy.
- McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133.

- Miconi, F. and Dimitri, G. M. (2023a). A machine learning approach to analyse and predict the electric cars scenario: The italian case. *PLOS ONE*, 18:e0279040.
- Miconi, F. and Dimitri, G. M. (2023b). A machine learning approach to analyse and predict the electric cars scenario: The italian case. *PLOS ONE*, 18(1):e0279040.
- Mohsin, M., Taghizadeh-Hesary, F., and Al-Thunaiyan, A. H. (2023). A systematic review of factors influencing electric vehicle adoption. *Environmental Science and Pollution Research*, 30(44):95167–95191.
- Molnar, C. (2020). Interpretable machine learning a guide for making black box models explainable. Available at: https://christophm.github.io/interpretable-ml-book/.
- Montavon, G., Samek, W., and Müller, K. R. (2018). Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73:1–15.
- Myers, J., Hanna, R., Garg, T., Kavanagh, J., Kleissl, J., Tebbe, S., Victor, D. G., and Washom, B. (2025). Design of workplace and destination-based ev charging networks considering driver behavior, habits, and preferences. *Renewable Energy*, 242:122441.
- Noppers, E. H., Keizer, K., Bolderdijk, J. W., and Steg, L. (2014). The adoption of sustainable innovations: Driven by symbolic and environmental motives. *Global Environmental Change*, 25:52–62.
- Oord, A. v. d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. (2016). Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499.
- Pamidimukkala, A., Kermanshachi, S., and Rosenberger, J. M. (2025). Motivation continuum and its effect on electric vehicle acceptance in india. *Journal of Enterprise Information Systems*, 18(2).
- Pao, J. (2025). Price war sparks ev financial crisis concerns in china.
- Papa, G., Sassi, G., Savio, F., Masi, S., and Bignotti, M. (2020). Assessment of electric vehicle driving range under different real-world conditions. *Energies*, 13(23):6296.
- Plötz, P., Kühn, T., and Gnann, T. (2014a). The impact of purchase subsidies on the adoption of electric vehicles: A study of germany's 'environment bonus'. *Energy Policy*, 73:618–629.

- Plötz, P., Schneider, U., Globisch, J., and Dütschke, E. (2014b). Who will buy electric vehicles? identifying early adopters in germany. *Transportation Research Part A:* Policy and Practice, 67:96–109.
- Popescu, G. H., Andrei, J. V., Sima, V., Nica, E., Chivu, L., and Gheorghe, I. G. (2024). A cross-sectional study of factors influencing purchasing cars in european union—does the electric and hybrid cars change the contemporary business approach? *Journal of Business Economics and Management*, 25(6):1260–1282.
- Rezvani, Z., Jansson, J., and Bengtsson, M. (2017). An investigation into the effects of anticipated emotions and personal moral norms on consumer pro-environmental behavior. *Journal of Promotion Management*, 23(1):163–183.
- Rietmann, N., Hügler, B., and Lieven, T. (2020). Forecasting the trajectory of electric vehicle sales and the consequences for worldwide co2 emissions. *Journal of Cleaner Production*, 261:121038.
- Roberts, G. (2025). Used electric vehicle values to 'stabilise' in 2025.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386.
- Rotaris, L., De Noni, I., De Luca, M., and Paladin, E. (2021). Drivers' preferences for electric cars in italy. evidence from a country with limited but growing electric car uptake. *European Transport Research Review*, 13(1):1–14.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088):533–536.
- Saxena, P., Bahad, P., and Kamal, R. (2020). Long short-term memory-rnn based model for multivariate car sales forecasting.
- Sharma, V. and Hofmann, K. (2022). Explainability in time series forecasting: A survey. arXiv preprint arXiv:2204.13457.
- Sierzchula, W., Bakker, S., Maat, K., and van Wee, B. (2014). The influence of financial incentives and other factors on electric vehicle adoption: A meta-analysis. *Research in Transportation Economics*, 47:178–188.

- Siskos, P., Capros, P., Zazias, G., Fiorello, D., and Noekel, K. (2019). Energy and fleet modelling within the trimode integrated transport model framework for europe. *Transportation Research Procedia*, 37:369–376.
- S&P Global Mobility (2023). Global automotive outlook. affordability tops charging and range concerns in slowing ev demand.
- Springel, K. (2021). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Journal: Economic Policy*, 13(4):393–432.
- Tiwari, A. K., Sharma, G. D., Rao, A., Hossain, M. R., and Dev, D. (2024). Unraveling the crystal ball: Machine learning models for crude oil and natural gas volatility forecasting. *Energy Economics*, 134:107608.
- Tran, K., Cimon, K., Severn, M., Pessoa-Silva, C. L., and Conly, J. (2012). Aerosol generating procedures and risk of transmission of acute respiratory infections to healthcare workers: A systematic review. *PLoS ONE*, 7(4):e35797.
- Van, T. L. C., Barthelmes, L., Gnann, T., Speth, D., and Kagerbauer, M. (2023). Enhancing electric vehicle market diffusion modeling: A german case study on environmental policy integration. *Energy Strategy Reviews*, 50:101244.
- van Rijnsoever, F. J., Hagen, P., and Willems, M. (2013). Preferences for alternative fuel vehicles by dutch local governments. *Transportation Research Part D: Transport and Environment*, 20:15–20.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.
- Wang, F. K., Chang, K. K., and Tzeng, C. W. (2011). Using adaptive network-based fuzzy inference system to forecast automobile sales. *Expert Systems with Applications*, 38:10587–10593.
- Wang, S., Li, J., and Zhao, D. (2017). The impact of policy measures on consumer intention to adopt electric vehicles: Evidence from china. *Transportation Research Part A: Policy and Practice*, 105:14–26.

- Xia, F., Yan, X., Li, X., Shi, D., and Yang, S. (2020). How does adoption of electric vehicles reduce carbon emissions? evidence from china. *Frontiers in Environmental Science*, 8:584988.
- Zhang, Y., Zhong, M., Geng, N., and Jiang, Y. (2017). Forecasting electric vehicles sales with univariate and multivariate time series models: The case of china. *PLOS ONE*, 12:e0176729.
- Ziegler, A. (2012). Individual characteristics and stated preferences for alternative energy sources and propulsion technologies in vehicles: A discrete choice analysis for germany. Transportation Research Part A: Policy and Practice, 46(8):1372–1385.

FONDAZIONE ENI ENRICO MATTEI WORKING PAPER SERIES

Our Working Papers are available on the Internet at the following address: https://www.feem.it/pubblicazioni/feem-working-papers/

"NOTE DI LAVORO" PUBLISHED IN 2025

- 1. M. A. Marini, S. Nocito, Climate Activism Favors Pro-environmental Consumption
- 2. J. A. Fuinhas, A. Javed, D. Sciulli, E. Valentini, <u>Skill-Biased Employment and the Stringency of Environmental</u> Regulations in European Countries
- 3. A. Stringhi, S. Gil-Gallen, A. Albertazzi, The Enemy of My Enemy
- 4. A. Bastianin, X. Li, L. Shamsudin, Forecasting the Volatility of Energy Transition Metals
- 5. A. Bruni, Green Investment in the EU and the US: Markup Insights
- 6. M. Castellini, C. D'Alpaos, F. Fontini, M. Moretto, Optimal investment in an energy storage system
- 7. L. Mauro, F. Pigliaru, G. Carmeci, <u>Government Size, Civic Capital and Economic Performance: An O-ring approach</u>
- 8. A. Bastanin, C. F. Del Bo, L. Shamsudin, The geography of mining and its environmental impact in Europe
- 9. E. Cavallotti, I. Colantone, P. Stanig, F. Vona, <u>Green Collars at the Voting Booth: Material Interest and Environmental Voting</u>
- 10. L. Ciambezi, M. Guerini, M. Napoletano, A. Roventini, <u>Accounting for the Multiple Sources of Inflation: an Agent-Based Model Investigation</u>
- 11. L. Fontanelli, M. Guerini, R. Miniaci, A. Secchi, <u>Predictive AI and productivity growth dynamics: evidence</u> from French firms
- 12. A. Drigo, An Empirical Analysis of Environmental and Climate Inequalities across Italian census tracts
- 13. F.F. Frattini, F. Vona, F. Bontadini, I. Colantone, The Local Job Multipliers of Green Industrialization
- 14. A. Abatemarco, R. Dell'Anno, E. Lagomarsino, <u>Measuring equity in environmental care: methodology and</u> an application to air pollution
- 15. F. F. Frattini, Political Participation and Competition in Concurrent Elections: Evidence from Italy

Fondazione Eni Enrico Mattei

Corso Magenta 63, Milano - Italia

Tel. +39 02 403 36934

E-mail: letter@feem.it

www.feem.it

