University of Kent

Kent Academic Repository

Raeesi, Ramin and Laporte, Gilbert (2026) *The impact of energy estimation inaccuracy on the optimality of electric vehicle routing decisions.* Transportation Research Part B: Methodological, 211 . ISSN 0191-2615.

Downloaded from
https://kar.kent.ac.uk/114633/ The University of Kent's Academic Repository KAR

The version of record is available from
https://doi.org/10.1016/j.trb.2026.103536

This document version
Author's Accepted Manuscript

DOI for this version

Licence for this version
CC BY (Attribution)

Additional information
For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

Versions of research works

Versions of Record
If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts
If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in ***Title of Journal*** , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries
If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies).

The impact of energy estimation inaccuracy on the optimality of electric vehicle routing decisions

Ramin Raeesi*

Centre for Logistics and Sustainability Analytics (CeLSA), Kent Business School, University of Kent, Canterbury, CT2 7PE, United Kingdom,

r.raeesi@kent.ac.uk

*(*Corresponding Author)*

Gilbert Laporte

Department of Decision Sciences, HEC Montréal, Montreal, Canada H3T 2A7;

School of Management, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom,

gilbert.laporte@hec.ca

Abstract

Accurate estimation of on-road energy consumption is often assumed to be critical in Electric Vehicle Routing Problems with Time Windows (EVRPTW). Yet the actual impact of estimation inaccuracies on routing feasibility and optimality remains largely unexplored. This paper focuses on assessing and tackling this impact through realistic testing and the development of an intuitive solution methodology based on a bi-objective ε-constraint optimisation framework. The proposed methodology relies only on the basic consumption rate reported by the vehicle manufacturer and builds on new intuitive results that suggest that despite the general tendency towards overestimating methods to remain conservative and risk-averse, the incorporation of an underestimating method might be more beneficial in implicitly ensuring the on-road optimality of routing decisions. Our results show that ignoring instantaneous speed variations and acceleration/deceleration rates can lead to significant inaccuracies, but the impact of disregarding factors such as slope, weather, or variable rolling resistance individually and collectively is rather negligible. Regardless of the size of estimation inaccuracy, however, we find inaccuracies to have very little impact on route-level outcomes as these rarely translate into infeasibility or meaningful optimality loss. Our analyses of the Pareto-optimal solutions on the generated efficient frontier of the considered EVRPTW instances, on the other hand, indicate that while the optimal solution of an instance is highly likely to also be on-road optimal, with a marginal sacrifice in instance optimality, significant allowance for inaccuracy is achievable to suit a strictly risk-averse decision maker.

Keywords: Routing; electric commercial vehicles; energy consumption estimation; inaccuracy; robustness.

1. Introduction

Energy estimation measures have been increasingly incorporated into the recent variants of the vehicle routing problems (VRPs) to plan energy-efficient or emissions-aware routes for conventional and alternative fuel freight vehicles (Bektaş & Laporte, 2011; Raeesi & O'Sullivan, 2014; Basso et al., 2019; Pelletier et al., 2019; Raeesi & Zografos, 2019; Basso et al., 2021; Bruglieri et al., 2023; Xiong et al., 2024). In Electric VRPs with Time Windows (EVRPTWs), energy consumption estimation is deemed an integral part of routing with a significant impact on the practical optimality (if not feasibility) of the solutions due to the restraining driving range of Electric Commercial Vehicles (ECVs). As such, unlike early papers in the field that simply relied on a basic energy consumption (or battery degradation) rate to ensure energy feasibility (Felipe et al., 2014; Schneider et al., 2014; Bruglieri et al., 2015; Desaulniers et al., 2016; Keskin & Çatay, 2016; Froger et al., 2019; Raeesi & Zografos, 2020; Raeesi & Zografos, 2022), several studies have recently tried to adopt "more accurate" energy consumption models that consider the impact of various variables such as vehicle characteristics, speed and payload (Basso et al., 2019; Pelletier et al., 2019; Basso et al., 2021; Yu et al., 2021; Bruglieri et al., 2023; Kim & Chung, 2023; Longhitano; et al., 2023; Xiong et al., 2024). These models, however, make multiple simplifying assumptions (e.g., by using average speeds and ignoring acceleration/deceleration (A/D) rates) for tractability, and as such contain an inevitable and sometimes high level of inaccuracy. Even without these simplifying assumptions, energy/fuel consumption estimation is a very difficult task due to a number of unobtainable, uncontrollable or unpredictable variables such as weather conditions, traffic and congestion situations, vehicle A/D, and driver behaviour, alongside several other variables such as vehicle speed and road gradient which are difficult to properly model. In the case of ECVs, this is even more complicated as other factors such as regenerative braking, the battery state of charge (SoC) and the battery state of health (SoH) have an impact on battery discharge along a driving cycle.

The estimation inaccuracy resulting from the use of simplified energy estimation models that ignore important variables such as vehicle's instantaneous A/D has already been flagged in the literature on emissions minimising VRPs (see for example, Turkensteen (2017) and Raeesi and Zografos (2019)). However, the real impact and cost of energy estimation inaccuracy on routing decisions, especially when it comes to planning ECV routes which are considered particularly prone to energy estimation inaccuracies, have never been investigated and quantified. At the same time, the gain achieved by improving the accuracy of estimation at the cost of extra complexity and intractability of the resulting routing model is not known either. In fact, it is easy to show, both theoretically and empirically, that a model based on a very crude energy estimation model can potentially yield the same optimal solution found by incorporating a very intricate microscopic second-by-second energy consumption estimation model.

The aim of this paper is to uncover the real cost and impact of energy estimation inaccuracy on EVRPTWs and shed light on the importance and priority of different factors affecting energy consumption such as vehicle variable speed, payload, A/D rates, variable rolling resistance, road slopes, and weather conditions. The expected estimation inaccuracy resulting from ignoring these variables independently and collectively is quantified through extensive realistic testing, and the impact of the resulting inaccuracy on the optimality or feasibility of routing decisions is analysed. Following the assessment of the impact of energy consumption estimation inaccuracy on routing decisions, a new approach to robustness is proposed based on some intuitive results used within a bi-objective integer programming optimisation framework to mitigate the impact of inaccuracy. The proposed methodology eliminates the need to incorporate a complex energy estimation model in the hope of achieving enhanced accuracy, and unlike most robust optimisation methods that are too conservative and often hedge against the worst-case scenario, gives the decision maker a freedom of choice based on their confidence in the estimation accuracy. Our results show that with even a basic energy estimation model, it is possible to identify an optimal or near-optimal solution with a high level of robustness against energy estimation inaccuracy.

1.1. Contributions

The contribution of this paper is fivefold:

- A detailed microscopic energy consumption estimation model for battery electric vehicles encompassing a comprehensive set of variables and pieces of route information is adopted as a benchmark model to predict battery discharge and regenerative braking of a particular ECV over 2,450 paths with varying distances, travel times, traffic and speed profiles, road gradients, and other factors generated across Greater London in the United Kingdom.
- A comparative analysis is established by applying eight alternative and less accurate energy estimation models ignoring one or several variables to assess the inaccuracies caused and to evaluate the importance of the disregarded variables.
- Energy consumption estimates based on the nine different models considered in the comparison are used within a set of load-dependent EVRPTW instances generated in Greater London to assess and report the impact of estimation inaccuracies on the feasibility and optimality of the routing solutions.
- New intuitive results are proposed to show that despite a general tendency to lean on an overestimating model to increase robustness of results, working with an underestimating model (in some cases just the basic energy consumption rate reported by the manufacturer) is more likely to yield an on-road optimal solution within a given confidence interval.

- Building on the proposed results, we develop a new approach based on a bi-objective integer programming optimisation framework to generate multiple Pareto-optimal solutions on the full efficient frontier of the problem, where each solution value is guaranteed to be within a certain factor of the optimal solution value and with a maximum energy estimation inaccuracy allowance.

This paper is, therefore, presented in two main parts that are dedicated to "assessing" the impact of energy estimation inaccuracy and "mitigating" this impact. In Section 2, the most pertinent studies are first discussed, and Sections 3 and 4 are devoted to assessing and tackling the impact of inaccuracy, respectively. The paper is concluded in Section 5. It is worth mentioning that all the collected data and details of the reported results are available online at https://data.kent.ac.uk/id/eprint/603.

2. Related literature

With the advent of the 21st century and with the rising global consciousness on climate change, a new wave of VRP papers have emerged to incorporate environmental considerations through new variants such as the energy minimising VRP (Kara et al., 2007), the pollution routing problem (Bektaş & Laporte, 2011) and its extensions (Franceschetti et al., 2013; Demir et al., 2014a; Koç et al., 2014; Dabia et al., 2017; Eshtehadi et al., 2017; Raeesi & Zografos, 2019), the green VRP (Erdoğan & Miller-Hooks, 2012; Salimifard & Raeesi, 2014), the bi-fuel VRP (Raeesi & O'Sullivan, 2014) and the EVRPTW (Schneider et al., 2014; Raeesi & Zografos, 2020; Raeesi & Zografos, 2022). Over a period of less than 20 years, more than 20 review papers such as those of Demir et al. (2014b), Bektaş et al. (2019), Marrekchi et al. (2021), Moghdani et al. (2021), Kucukoglu et al. (2021) and Garside et al. (2024) have been published. Almost all the papers cited by these review papers comprise some form of energy, fuel or emissions estimation, and the adopted approaches range from very basic linear models (e.g., a general energy/fuel consumption rate or emissions factor per distance unit travelled), to very detailed, nonlinear and complex models incorporating several variables.

Fuel/energy consumption estimation models (which are also usually used as a proxy for CO_2e) can be broadly categorised into (i) factor models, (ii) macroscopic models, (iii) microscopic models, and (iv) mesoscopic models, which differ based on how vehicle activities and fuel consumptions are aggregated over time and space (Demir et al., 2014c; Chen et al., 2017). Factor models are simple fuel consumption methods (e.g., based on manufacturer reported driving range or miles per gallon (MPG)); macroscopic models use average aggregate network parameters such as average travel speed to estimate fuel consumption; microscopic models follow a physics-based approach and estimate the instantaneous vehicle fuel consumption and emission rates at the most granular level and using second-by-second vehicle trajectory data; mesoscopic models statistically predict fuel consumption of vehicles using large-scale real-world driving data (Demir et al., 2014c; Chen et al., 2017). Some of these models have been reviewed and empirically compared in Demir et al. (2014b) and Demir et al. (2011).

While a few macroscopic models have been used in the early research on VRPs with environmental considerations (see for example, Figliozzi (2010), Figliozzi (2011) and Qian and Eglese (2014)), it can be confidently argued that following its application within the seminal paper of Bektaş and Laporte (2011), the comprehensive modal emissions model (CMEM) (Barth et al., 2005; Barth & Boriboonsomsin, 2009) has been by far the most used energy and emissions estimation model in the relevant literature. This is mainly because CMEM incorporates some of the most important energy consumption related factors such as vehicle speed and payload, as well as vehicle-specific characteristics that can be treated as decision variables within the routing model. Most papers that use CMEM have, however, incorporated a simplified fixed-speed and zero A/D version of CMEM within their routing models. This essentially amounts to turning a microscopic model inappropriately into a macroscopic average speed model as CMEM relies on specifying the values of all input parameters and variables on a second-by-second basis for generating accurate and reliable estimates. This issue has been flagged by several authors. Turkensteen (2017) compares CMEM results under fixed speeds with those computed under realistic driving conditions and reports that fixed speed estimates can be up to 50% inaccurate and underestimated. Kancharla and Ramadurai (2018a) report a maximum underestimation of 40% when driving cycles are ignored and an average speed is used within CMEM. Raeesi and Zografos (2019) show that large estimation inaccuracies can be expected when truck A/D rates are ignored in estimating fuel consumption using CMEM. A methodology to address the lack of driving cycle A/D data is proposed in both Kancharla and Ramadurai (2018a) and Raeesi and Zografos (2019).

More accurate methods for including ECV's battery discharge estimates have emerged in the EVRPTW literature. Thus, Goeke and Schneider (2015) investigate the routing of a mixed fleet of electric and conventional vehicles where the electric power that the electric motor of the ECV needs to provide is calculated using the fixed-speed version of CMEM. The authors treat payload, road gradients and speed as decision variables. Basso et al. (2019) study an EVRP with an energy consumption estimation model that incorporates approximations for factors such as road elevation, speed profiles, powertrain efficiency, A/D, and braking. The underlying energy estimation model used in their paper also resembles the engine power demand module within CMEM. Rastani et al. (2019) investigate the effect of ambient temperature on energy consumption, and its resulting impact on the routing decisions within an EVRPTW. Longhitano; et al. (2023) propose a battery degradation aware routing model by taking the battery SoC and SoH into consideration. Kim and Chung (2023) also incorporate a nonlinear SoC-dependent discharging function into an EVRP. Xiong et al. (2024) consider the impact of drivetrain losses as a nonlinear function of velocity and incorporate them into CMEM. They also consider the influence of traffic congestion on vehicle speed.

More and more papers contributing to improved energy consumption (or battery degradation) estimation accuracy may keep appearing as research in this field develops. However, there will always be an unavoidable level of inaccuracy in any energy consumption estimation, no matter how complex and comprehensive the model is. An alternative (less proactive, but more realistic) approach in the face of this inherent uncertainty is to accept and incorporate it into EVRPs using stochastic or robust optimisation methods. Thus, Pelletier et al. (2019) propose a robust optimisation framework to take into account the energy consumption uncertainty and determine minimum cost delivery routes capable of providing guarantees that ECVs will not run out of charge during their routes. Soysal et al. (2020) develop a chance-constrained EVRP with pickups and deliveries under stochastic battery depletion assumptions. Basso et al. (2021) use a Bayesian machine learning approach for predicting the expected energy consumption within a time-dependent EVRP with chance constraints and partial recharging where the uncertainty in energy demand is addressed by planning charging within a confidence interval. In a later study, Basso et al. (2022) model the dynamic stochastic EVRP as a Markov decision process to incorporate stochastic energy consumption as well as stochastic customer requests. Jeong et al. (2024) propose a two-stage adaptive robust optimisation framework for the EVRP to minimise the worst-case energy consumption while guaranteeing battery level adequacy. Han et al. (2025) consider uncertainty in travel time and energy consumption and develop risk-resistant routing and recharging decisions within a two-stage adaptive robust optimisation framework. Spinelli et al. (2026) model a stochastic EVRP with a threshold recourse policy to ensure that EVs recharge when their battery charge level falls below a specified threshold.

The proposed approach in this paper differs from existing robustness-oriented EVRP formulations discussed above in both its modelling philosophy and data requirements. Classical robust and stochastic EVRP approaches, such as robust optimisation, chance-constrained programming, stochastic programming, and adaptive robust optimisation, explicitly model uncertainty in energy consumption through scenarios, probability distributions, uncertainty sets, or recourse decisions. For example, Pelletier et al. (2019) and Jeong et al. (2024) seek solutions that remain feasible under adverse energy-consumption realisations, while Soysal et al. (2020) and Basso et al. (2021) use probabilistic constraints to control the risk of battery depletion. Basso et al. (2022) and Spinelli et al. (2026) go further by incorporating sequential or recourse-based decisions under uncertainty. In contrast, the method proposed in this paper does not require scenario generation, distributional assumptions, confidence levels, or dynamic recourse policies. Instead, it starts from the manufacturer-reported basic consumption rate and introduces an endogenous inaccuracy allowance, which quantify how much underestimation can be tolerated while preserving route feasibility. The method therefore does not optimise against a predefined uncertainty representation; rather, it identifies routing solutions that are intrinsically tolerant to estimation error. This distinction also leads to a different practical contribution. Existing robustness-oriented models are primarily designed to protect feasibility under explicitly modelled uncertainty, often at the cost of higher modelling and computational complexity and stronger data requirements. By contrast, the proposed bi-objective ε-constraint framework is designed to measure and maximise tolerance to imperfect energy estimates using only limited and readily available input data. It is therefore particularly suitable when high-quality stochastic information is unavailable, but the decision-maker still wishes to identify routes that are less sensitive to consumption-model error. In this sense, the proposed approach should not be viewed as a substitute for stochastic or robust EVRP formulations, but as a complementary and more lightweight robustness-oriented methodology that captures estimation resilience rather than probabilistic or worst-case robustness.

In all, the two strands of literature reviewed above are based on the premise that inaccuracies in estimating energy consumption can render routing decisions in EVRPs either suboptimal or infeasible. In this paper, we seek to uncover the real impact of such inaccuracies on routing decisions and address four critical but largely overlooked questions: (i) What is the impact of ignoring some supposedly key factors on the level of energy estimation accuracy? (ii) What is the real cost of energy estimation inaccuracy on routing decisions? (iii) How tolerant are routing decisions to energy estimation inaccuracies? and (iv) How can the impact of estimation inaccuracy be best mitigated?

3. Assessing the impact of inaccuracy

The importance and impact of different factors and parameters on the level of energy estimation accuracy has been quantified in several ways in a dedicated stream of research that focuses on sensitivity analysis of energy demand estimation for battery electric vehicles (Yang et al., 2014; Asamer et al., 2016; Liu et al., 2017; Wang et al., 2017; Yi & Bauer, 2017; Iora & Tribioli, 2019; Donkers et al., 2020; Schäfers et al., 2024; Abro et al., 2026). While some of these studies focus on the impact of only one particular factor such as the ambient temperature (Wang et al., 2017; Iora & Tribioli, 2019) or road gradient (Yang et al., 2014; Liu et al., 2017), others present a more inclusive analysis. Asamer et al. (2016) conduct a sensitivity analysis for a longitudinal dynamics model by collecting recorded GPS tracks of an electric vehicle under various driving conditions. They quantify the impact of various parameters on accuracy and show that the uncertainty of vehicle efficiency and rolling friction coefficient have the highest impact. They also analyse the energy shares related to several model component including acceleration, air drag, rolling and grade resistance and auxiliary energy demand and show that, while some of these components make up a large share of the overall energy demand, the uncertainty of parameters related to these components does not affect the accuracy of energy demand estimation significantly. Donkers et al. (2020) combine a microscopic traffic model and an energy prediction model to estimate the effects of driving styles, weather, traffic and infrastructure on the energy consumption and show that built environment variables have a large effect on the energy consumption of battery electric vehicles. Yi and Bauer (2017) conduct a sensitivity analysis with respect to the effects of wind speed,

rolling resistance, parasitic power and temperature and show that the sensitivity of energy consumption with respect to these four environmental variables greatly vary with operating conditions of the vehicle. Abro et al. (2026) examine the influence of vehicle mass on electric vehicle energy consumption and show that the energy consumption grows by 0.029 Wh/kg/km.

Unlike the above contributions, in this paper, we are more focused on routing applications and the impact of energy estimation inaccuracy on routing decisions. As such, we tend to examine factors such as speed, A/D rates, vehicle load, and road slope that are often considered as potential exogenous or endogenous variables to the optimisation model. To this end, we present as the benchmark model a previously tested and tried energy estimation model and procedure to collect detailed journey data. This methodology is quite flexible and reliable and can generate accurate results that are close to the actual on-road vehicle battery discharge and regenerative braking. This is followed by the presentation of eight alternative "less accurate" models. Seven of these models are based on simplifications applied to the benchmark model by disregarding the impact of one or multiple variables (e.g., the ambient temperature and weather conditions, or the second-by-second variable speed), and the last one is simply a distance-dependent energy consumption rate derived from the manufacturer's advertised driving range for the considered ECV. The experiments setting and the collected data for 2,450 paths in Greater London are then fully discussed, and summary results from the experiments are presented and analysed.

3.1. The benchmark model

We use the TU/e Microscopic Energy Consumption PRediction tOol 0.1 (TU/e MECPRO 0.1) (Beckers et al., 2021) developed and implemented by Beckers et al. (2020), which itself builds on the methodology developed by Wang et al. (2018) as the "benchmark model". This methodology was developed specifically for a particular fully electric delivery van (Voltia eVan[1]) and incorporates granular data on driving behaviour, road topography information, weather conditions, traffic situation, vehicle parameters, powertrain losses and regenerative braking, and can be applied to any real-world driving path with a high level of accuracy. The model was verified against a number of different urban, rural, highway and hilly driving tests and the mean absolute percentage error between the on-road measured energy consumption value, and the prediction result of all trips was reported to be within 5% of the observed result (Wang et al., 2018). A schematic summary of the data preparation and prediction procedures within the TU/e MECPRO is illustrated in Figure 1.

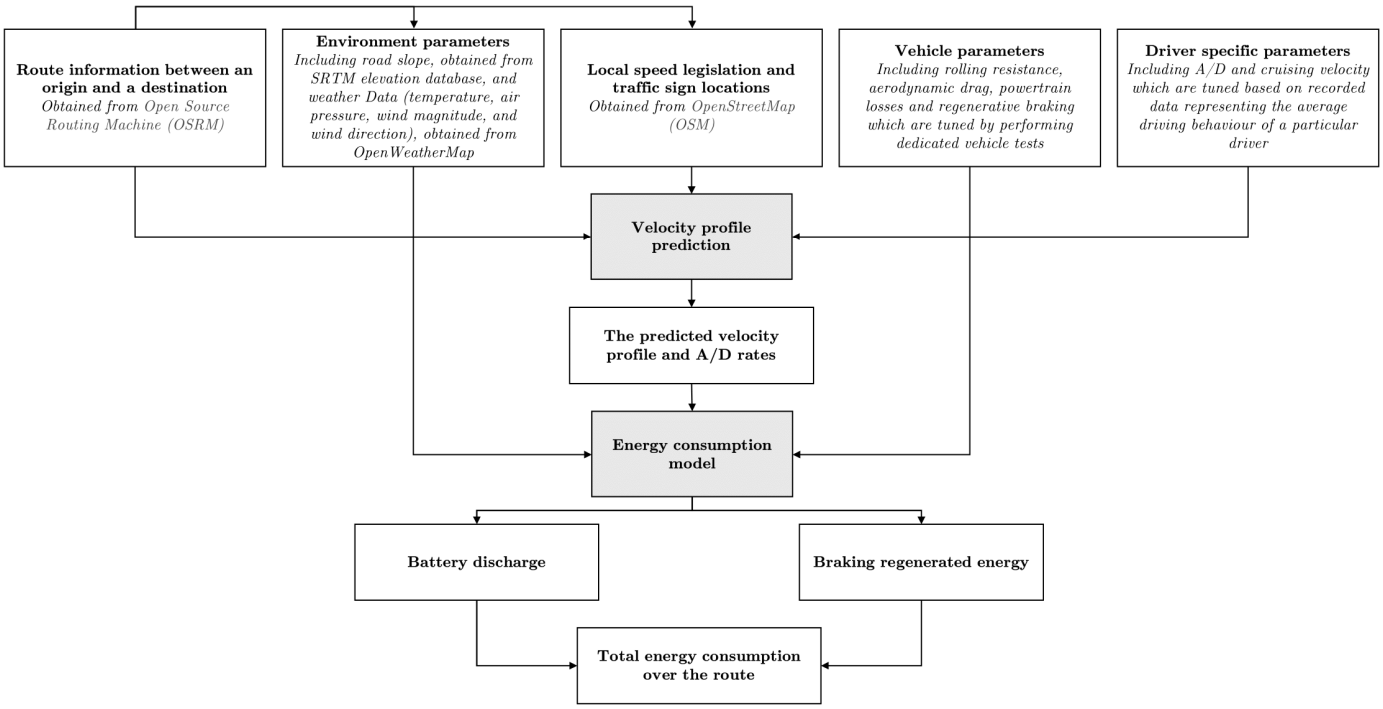


Figure 1 A schematic summary of the benchmark model

While the interested reader is referred to the original paper for a comprehensive description of the tool, we go through the working of the methodology using an illustrative example. To predict the energy consumption of the ECV, the procedure starts with extracting a shortest path between a given pair of origin-destination from the Open Source Routing Machine (OSRM)[2]. An Example Driving Route (EDR) between two points in Greater London is selected and shown in Figure 2. The considered EDR is around 34.35 km in length and takes around 57 minutes to traverse, based on the OSRM estimation obtained on 18/07/2024 at 13:15.

[1] https://www.voltia.com/
[2] http://map.project-osrm.org/



Figure 2 The EDR in Greater London

Once a route is calculated, the methodology determines several environment parameters relevant to the route, including the road slope and the weather data. The local road slope is determined from the Shuttle Radar Topography Mission (SRTM[3]) elevation database as a function of the travelled distance (see Figure 3).

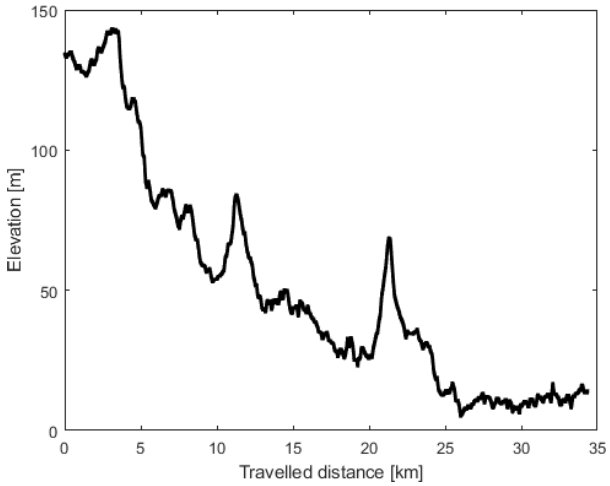


Figure 3 SRTM road elevation as a function of the travelled distance for the EDR

Up-to-date weather information comprising the temperature, air pressure, wind magnitude, and wind direction are obtained from OpenWeatherMap[4]. Air density is calculated using temperature and air pressure. The wind magnitude and direction are represented by indicating the wind direction with respect to the north, and together with the vehicle velocity vector (calculated from the vehicle heading) (see Figure 4), this vector determines the apparent wind vector.

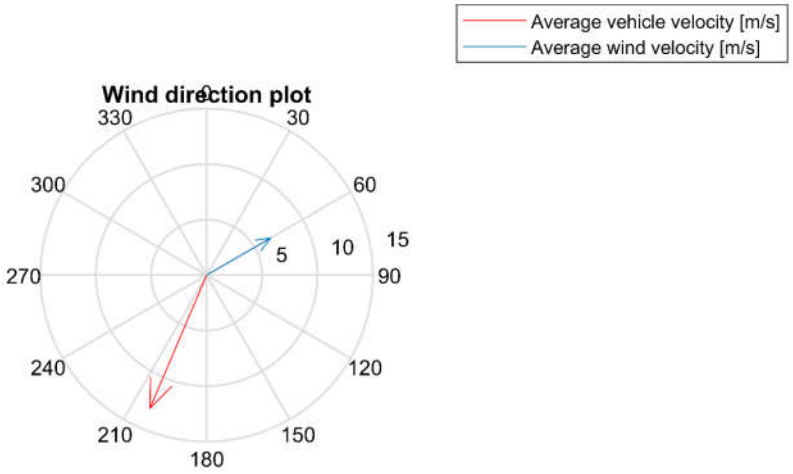


Figure 4 Wind direction plot based on data from OpenWeatherMap for the EDR

[3] https://eospso.nasa.gov/missions/shuttle-radar-topography-mission
[4] https://openweathermap.org/current

Data on local speed legislation and traffic sign locations are also obtained using OpenStreetMap Functions toolbox[5]. Figure 5 shows relevant data collected from OpenStreetMap for the EDR.

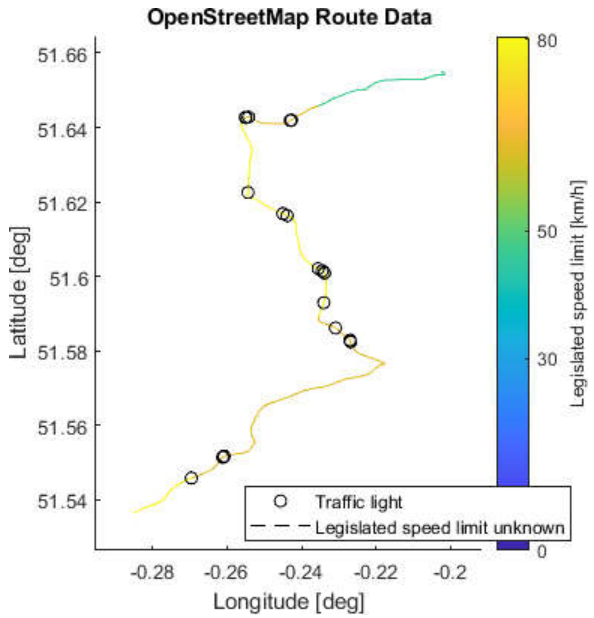


Figure 5 Legislated speed limit and traffic sign locations for the EDR based on OpenStreetMap

Route information and the collected data on local speed legislation and the traffic sign locations are then fed into a velocity predicting algorithm which is based on the methodology developed by Wang et al. (2018). In this methodology, the local speed legislation is assumed to dictate the maximum velocity along the route and further velocity constraining locations including traffic light locations and corners are also identified and applied. The vehicle is assumed to make a full stop at every encountered traffic light and to reduce velocity in a corner based on a function of the corner curvature (obtained by using the GPS-points and an assumed maximum lateral acceleration). The ultimate velocity profile (Figure 6) is then determined by including a limited vehicle A/D considering three driver specific parameters of acceleration, deceleration, and cruising velocity which are tuned based on recorded real-world data to represent the average driving behaviour of a particular driver.

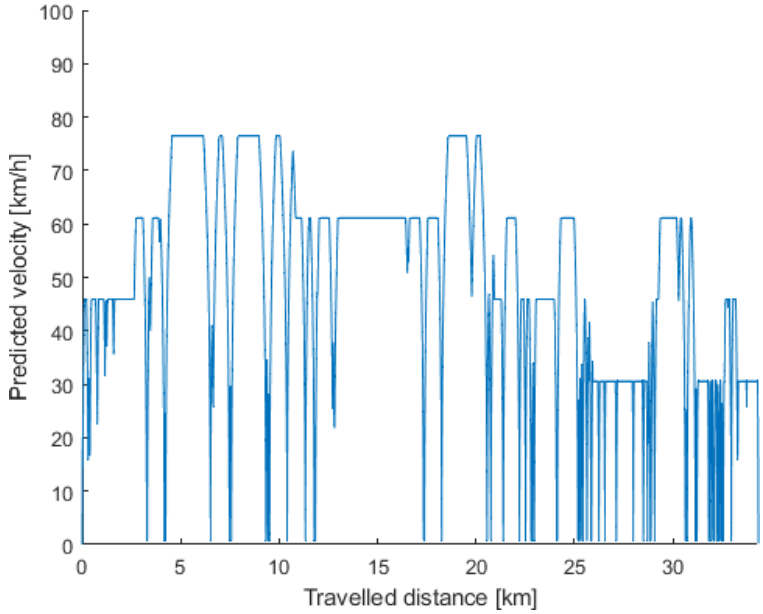


Figure 6 The predicted velocity profile along the EDR

The predicted velocity profile and A/D rates are then passed onto the physics-based microscopic energy consumption prediction model. The model uses the longitudinal dynamics of the vehicle (i.e., driving force, aerodynamic drag force, total rolling resistance force, and longitudinal gravity) to calculate the power required by the powertrain P_{pt} using the following expression:

$$P_{pt} = \left(m_{eff}\frac{dv}{dt} + f_r mg\cos(\alpha) + \frac{1}{2}\rho C_d A_f v_{wind,rel}^2 + mg\sin(\alpha)\right)v + P_{loss}(\omega_{wheel}, T_{wheel}). \quad (1)$$

[5] https://github.com/johnyf/openstreetmap

In this expression v is the vehicle velocity, m is the vehicle mass, f_r is the rolling resistance coefficient, g is the gravitational acceleration, C_d is the aerodynamic drag coefficient, A_f is the frontal area, ρ is the air density, α is the local road slope, $v_{wind,rel}$ is the relative wind velocity, and P_{loss} is powertrain losses as a function of wheel speed ω_{wheel} and wheel torque T_{wheel}. The total energy consumption for a trip is calculated by adding the auxiliary power and integrating over the trip time.

It must be mentioned that the tool does not rely on physical modelling alone and to acquire a knowledge of several vehicle-specific parameters, which are otherwise difficult to determine, it uses results of dedicated vehicle tests such as coast-down tests on different road surfaces (good asphalt, medium quality asphalt, and bad asphalt) and crosswind effects to determine both the rolling resistance coefficient and the aerodynamic drag coefficient, dynamometer tests to measure the lumped powertrain losses by operating the vehicle in steady-state at several combinations of velocities and torques, and coast-down tests on dynamometer to identify the regenerative braking characteristics of the vehicle.

The use of the benchmark model on the EDR to predict discharged and regenerated energy is shown in Figure 7.

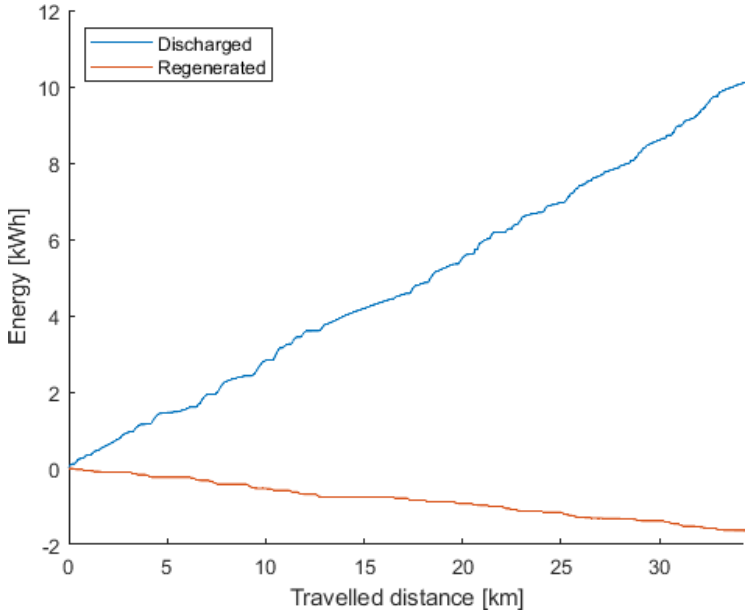


Figure 7 Predicted discharged and regenerated energy along the EDR using the bechmark model

As discussed earlier, there is no perfect tool to estimate energy consumption over a driving route, and while the adopted methodology works as a high-fidelity reference, we must note that the benchmark model itself is still an estimation model, and not a direct measurement of real-world energy consumption. Therefore, despite its sophistication, the benchmark model also contains inherent assumptions and simplifications (e.g., its velocity profile prediction algorithm, modelling of powertrain losses, regenerative braking characteristics) and any systematic error or bias within the benchmark model will propagate into the reported "inaccuracy" figures for all alternative models. For instance, if the benchmark model consistently underestimates consumption on hilly routes, a simplified model that ignores slope might appear more accurate than it truly is. That said, given the scale of the analysis and the need to simulate a large number of scenarios—many of which are impractical to test on-road—we consider the benchmark model a sufficiently reliable proxy for this study. This is due to its broad coverage of key factors and its empirically demonstrated ability to produce estimates close to real-world energy consumption. In Section 3.2, we introduce "less accurate" alternative methods which are essentially derived from ignoring certain factors within the benchmark model to assess their impact on accuracy of estimation.

3.2. The alternative methods

To measure the impact of inaccuracy on the optimality of EVRPTW solutions and also the importance of the various variables involved in the benchmark model, eight different alternatives are proposed. Seven of these are obtained by disregarding one or multiple factors from the benchmark model, and one is a very crude model based on the energy consumption rate per km travelled. These models are described below.

- **A1 - avgSpeedAllAcross**: As discussed above, the benchmark model benefits from an accurate instantaneous velocity profile and A/D rates prediction module in estimating energy consumption. However, in this alterative model, only an average travelling speed (40 km/h) across the network is considered to evaluate the impact of disregarding the actual on-road instantaneous variable speed and A/D rates.
- **A2 - noAcc**: Unlike the previous estimation method, here variable speed profiles are considered, but A/D rates are ignored by setting longitudinal vehicle acceleration on zero across all cycles, and hence, fully disregarding the acceleration force.

- **A3 - noVarFr**: In the benchmark model, the rolling resistance coefficient is assumed to vary as a function of road surface (good asphalt, medium quality asphalt, and bad asphalt) and is represented as a step function of vehicle velocity with the underlying assumption that the road quality is generally better on high-velocity roads such as motorways. In this alternative model, the effect of this variable rolling coefficient is ignored, and medium quality asphalt and velocity range is used across all road segments within a path.
- **A4 - noSlope**: This alternative model ignores the impact of road slope by assuming zero road angle for all road segments in the network.
- **A5 - noWeather**: In this model the impact of weather conditions is ignored by setting the relevant parameters on no wind and using the default air density.
- **A6 - noSlopeWeatherFr**: This model considers the collective impact of ignoring several variables comprising the road slope, weather conditions and variable rolling resistance altogether.
- **A7 - noVarSpeed**: This model is similar to the avgSpeedAllAcross model, but instead of using a priori determined average travelling speed across the network, here we use the average of the actual driving profile over the path as the speed value.
- **A8 - basicRate**: This is a crude and easily accessible model which is simply based on the manufacturer's reported average battery charge consumption rate per km travelled. This average value is almost always reported for all vehicles (MPG in the case of conventional vehicles and range in the case of EVs) and is readily available.

The consideration of these alternative models allows us to measure the impact of various factors on the energy estimation accuracy and the optimality of the routing decisions. In addition to these, we explore and report the impact of ignoring the regenerative braking in the considered ECV, and the vehicle payload.

Prior to discussing the experiments setting and the results, we describe the EVRPTW variant considered in this paper.

3.3. The load-dependent EVRPTW without en-route recharging

This section formally describes and presents a mathematical formulation for the selected problem variant. We start with some key assumptions:

- We assume that all vehicles are fully charged overnight at the starting depot and are not allowed to visit any intermediate Recharging Stations (RSs) along their delivery routes. We believe this is a reasonable assumption in urban contexts. Based on our observations and discussions with courier services in London, for example, en-route recharging is often disregarded as an option due to various operational challenges involved. Moreover, the range covered with one full charge is usually sufficient for urban deliveries. It is worth mentioning that this assumption is not restrictive for the proposed methodology discussed in Section 4, and relevant experiments and results are also presented for the case of en-route recharging.
- We model the problem on a complete graph constructed by the identification of all pairwise fastest paths between required nodes (i.e., the depot and the customers). As such there is a single path between every two customers, or between a customer and the depot. This is again a practical assumption as most courier services rely on navigation services such as Google Maps for their routing needs which essentially return the fastest path by default.
- Given the assumption above and the way energy consumption is calculated using the benchmark and the seven alternative models discussed earlier, we argue that all variables having an impact on energy consumption, except for vehicle load, can be captured off-line (i.e., outside the optimisation process) using the energy consumption estimation model, and fed into the problem. As such, a load-dependent variant of the problem is sufficient for the purpose of this paper. Time-dependency and the impact of time-varying speed are irrelevant here given how instantaneous speed and A/D profiles are constructed within the benchmark model based on the aforementioned velocity profile prediction module.
- Given the underlying longitudinal power requirement model presented in Expression (1), and following extensive tests with the benchmark model, we argue that energy consumption is approximately linear in vehicle payload when all other variables remain unchanged. Across 124,950 simulations covering 2,450 paths and 11 payload levels (0–1000 kg), linear regressions of energy consumption against payload yielded near-perfect fits, with average R^2 values exceeding 0.995 and typically above 0.999. The average RMSE was below 0.01 kWh, confirming that a linear approximation provides an excellent representation of payload-dependent energy consumption over the investigated operating range (further analytical and empirical evidence is provided in Appendix A). As such, for all considered models and over each path in the network we fit a linear model for energy consumption as a function of vehicle load with a dedicated y-intercept (α) and slope (β).

With these assumptions, the proposed EVRPTW without en-route recharging is defined on a complete, directed graph $G = (N, A)$, where N is the set of network nodes and A is the set of directed arcs representing the constructed paths between required nodes on the actual network. The set $N = \{D \cup C\}$ consists of the depot $D = \{0, n+1\}$, with $\{n+1\}$ being a dummy copy of $\{0\}$ (referred to as the final depot) and customer nodes $C = \{1,2, \dots, n\}$. Each customer $i \in C$ is associated with a certain demand q_i

to be delivered within its predetermined hard time window $[e_i, l_i]$, with service time s_i. The depot working hours, which define the planning horizon, are denoted by $T = [e_0, l_0]$. The set of directed arcs is defined as $A = \{(i,j) | i \in N_0, j \in N_\sigma, i \neq j\}$, where $N_0 = \{0\} \cup C$ and $N_\sigma = C \cup \{n+1\}$. Each arc $(i,j) \in A$ has a distance d_{ij}, a travel time t_{ij}, a y-intercept α_{ij} and a slope β_{ij} of the linear equation representing the load-dependent energy consumption function. There is a fleet of homogeneous ECVs that are all fully charged and located in the central depot. Each ECV has a maximum payload Q and a battery capacity B.

The mixed integer linear programming (MILP) formulation of the proposed EVRPTW variant works with the following decision variables:

- x_{ij}: binary variable equal to 1 iff arc $(i,j) \in A$ is traversed by an ECV.
- y_i: continuous variable denoting the time of arrival of an ECV at node $i \in N$.
- a_i: continuous variable denoting the service start time by an ECV at node $i \in N$.
- f_{ij}: continuous variable denoting the load carried by an ECV over arc $(i,j) \in A$.
- u_i: continuous variable denoting the residual battery charge level of an ECV on arrival at node $i \in N$.

The mathematical formulation of the problem is given by (2)–(11):

$$\text{Minimise} \sum_{(i,j)\in A} d_{ij} x_{ij} \tag{2}$$

subject to

$$\sum_{j \in N_\sigma} x_{ij} = 1, \qquad \forall i \in C \tag{3}$$

$$\sum_{i \in N_0} x_{ij} - \sum_{i \in N_\sigma} x_{ji} = 0, \qquad \forall j \in C \tag{4}$$

$$a_i = \max\{e_i, y_i\}, \qquad \forall i \in N \tag{5}$$

$$a_i + (t_{ij} + s_i) x_{ij} - l_0 (1 - x_{ij}) \leq y_j, \qquad \forall i \in N_0, j \in N_\sigma, i \neq j \tag{6}$$

$$e_i \leq a_i \leq l_i, \qquad \forall i \in N \tag{7}$$

$$\sum_{j \in N_0} f_{ji} - \sum_{j \in N_\sigma} f_{ij} = q_i, \qquad \forall i \in N_0 \tag{8}$$

$$q_j x_{ij} \leq f_{ij} \leq (Q - q_i) x_{ij}, \qquad \forall (i,j) \in A \tag{9}$$

$$0 \leq u_j \leq u_i - (\alpha_{ij} x_{ij} + \beta_{ij} f_{ij}) + B(1 - x_{ij}), \qquad \forall i \in N_0, j \in N_\sigma, i \neq j \tag{10}$$

$$0 \leq u_0 \leq B\,. \tag{11}$$

The objective function (2) seeks to minimise the total distance travelled by all ECVs. For experimental purposes discussed in the next section, we refer to (2) as the distance-based (DB) objective function, and further define the alternative energy-based (EB) objective in (12):

$$\text{Minimise} \sum_{(i,j)\in A} (\alpha_{ij} x_{ij} + \beta_{ij} f_{ij}) \tag{12}$$

Constraints (3) indicate that each customer must be visited exactly once by an ECV for delivery. Constraints (4) guarantee that if an ECV enters a customer node, it also exits this node. Constraints (5) define the service start time at each customer. Constraints (6) determine the arrival time of an ECV at each node by accounting for the service start time at the upstream node and its service time. Constraints (7) ensure that the service of a customer starts during its time window. Constraints (8) and (9) ensure demand fulfilment while guaranteeing that the capacity of the ECVs is not violated. Constraints (10) determine the battery charge level of ECVs after visiting a customer. Finally, constraints (11), in conjunction with (10), ensure that the battery charge level always remains within the ECV's battery capacity.

We discuss next how this formulation is used in assessing the impact of energy estimation inaccuracy on the optimality of the routing decisions.

3.4. Experiments setting

To gain practical insights into the size of inaccuracy caused by ignoring different variables, and the impact of the resulting inaccuracy on the optimality of routing decisions, a large set of realistic tests were conducted and are reported in this section.

To develop these instances, we have randomly generated a set of 50 nodes across Greater London (Figure 8). Following this, the pairwise fastest driving path between all pairs of nodes (a total of 2,450 paths) were identified, all environment parameters including the road slope and weather data were collected, and all local speed legislations and traffic sign locations were gathered as per the TU/e MECPRO procedure discussed in Section 3.1. These data along with the vehicle and driver related parameters were then fed into the velocity profile prediction module to generate instantaneous driving cycles along all paths. The generated paths cover a broad operating envelope, with distances ranging from approximately 1 km to 104 km and travel times from 1.5 min to 109 min. Across all paths, mean distance and duration are 37.7 km and 49.9 min, respectively, while substantial variability is observed in speed, acceleration/deceleration, terrain elevation gain and gradients, and weather-related conditions, ensuring that the evaluation is conducted under diverse and realistic operating scenarios. Appendix B provides a complete statistical analysis of various features of the generated paths.

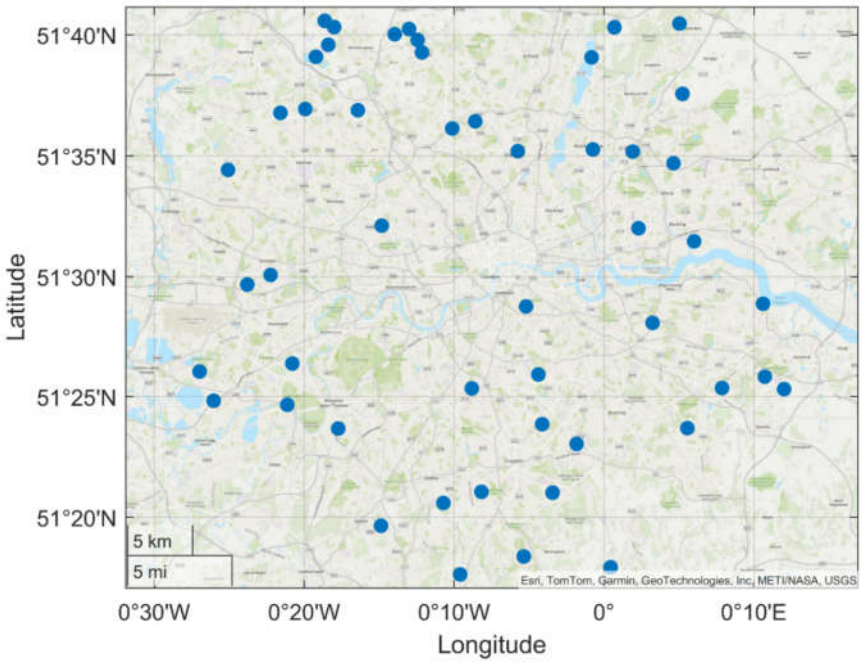


Figure 8 The 50 randomley generated nodes across Greater London

The generated driving cycles were fed into the benchmark model to estimate the energy consumed and regenerated using the regenerative braking considering different vehicle payloads. This was set as the basis for all comparisons with the eight alternative models. For the basicRate model, the average battery charge consumption rate per km distance travelled for Voltia eVan was set at 0.27 kWh/km, based on specifications from the manufacturer. It is also worth mentioning that Voltia eVans come in two battery options: 50 kWh or 75 kWh batteries. Both of these options are considered in our experiments.

For the routing tests, we have generated a set of 30 instances of sizes 5, 10, 15, 20, 25 and 30 (five of each) on the complete network constructed over the 50 nodes discussed earlier. For each instance, one of the nodes is randomly assigned as the depot and the remaining nodes were considered as customer nodes. The demand of each customer was generated randomly as an integer in the range of 1 kg to 100 kg. The time window of each customer was generated by first constructing nearest neighbour routes over the instances and then generating a time window around the arrival time τ at each node. The lower and upper bounds of the time windows were obtained by generating a random integer δ in the range $[e_0, l_0]$, and then using the expressions $\max(e_0, \lceil \tau - \delta/2 \rceil)$ and $\min(e_0, \lceil \tau + \delta/2 \rceil)$, respectively.

3.5. Experimental results

The first set of test results pertains to the comparative analysis of the calculated energy consumption over all the 2,450 paths using the benchmark model and the eight alternative methods described in Section 3.2. In order to summarise these results, we classify all paths into 10 groups based on their distance. Tables 1 and 2 show the summary results in two test settings of an empty and a fully loaded vehicle, respectively. In these tables, the numbers under 'Benchmark' column denote the average total energy consumed in kWh based on the benchmark method for each category, and columns under A1 to A8 report the average inaccuracy of each alternative estimation method. This was calculated for each path category and each alternative method using the expression $((EC'_c - EC^*_c)/EC^*_c) \times 100$, where EC'_c is the average estimated energy consumption over the paths in category c using an

alternative method, and EC_c^* is the average estimated energy consumption using the benchmark method for the paths in the same category. A positive inaccuracy indicates an overestimation, while a negative value indicates an underestimation.

As concerns the empty vehicle tests, Table 1 clearly shows that the largest inaccuracy occurs with alternative methods A1 (i.e., avgSpeedAllAcross), A7 (i.e., noVarSpeed) and A2 (i.e., noAcc), respectively. These values collectively indicate the importance of the incorporation of variable instantaneous speed profiles and A/D rates as also already confirmed by the extant literature (Turkensteen, 2017; Kancharla & Ramadurai, 2018b; Raeesi & Zografos, 2019). A8 (i.e., basicRate) is also, as expected, a rather inaccurate estimation as it is a very crude approach based on the generic figure reported by the vehicle manufacturer. It is, however, interesting to see that the other methods (i.e., A3 - noVarFr, A4 – noSlope, A5 – noWeather, and A6 – noSlopeWeatherFr) yield estimates very close to those of the benchmark model. Variable rolling resistance, weather conditions and road slopes, therefore, indicate limited value for a high-level context such as routing, and even the collective impact of ignoring all of these together is not significant.

Table 1
Summary results for comparison of different estimation methods based on the empty vehicle mode

Category (<km)	Benchmark (consumed kWh)	Average inaccuracy							
		A1	A2	A3	A4	A5	A6	A7	A8
10	1.87	65.0%	-26.6%	-0.9%	-0.8%	-0.7%	-2.5%	-30.6%	0.2%
20	4.20	62.4%	-26.9%	-0.8%	-0.4%	-0.8%	-2.0%	-32.7%	-3.6%
30	7.30	59.8%	-25.5%	-0.7%	-0.7%	-0.8%	-2.2%	-33.0%	-7.9%
40	10.14	56.9%	-23.9%	-0.6%	-0.4%	-1.0%	-2.0%	-33.0%	-8.6%
50	13.05	55.4%	-21.3%	-0.5%	-0.5%	-1.3%	-2.3%	-31.5%	-8.6%
60	16.45	44.6%	-14.1%	-0.1%	-0.5%	-1.2%	-1.8%	-27.4%	-13.2%
70	22.58	28.1%	-5.3%	0.4%	-0.2%	-1.4%	-1.3%	-20.4%	-23.9%
80	27.11	25.8%	-4.6%	0.4%	-0.2%	-1.7%	-1.5%	-21.1%	-25.6%
90	30.64	26.0%	-4.3%	0.4%	-0.4%	-1.5%	-1.5%	-19.9%	-25.8%
100	34.70	23.3%	-2.5%	0.5%	-0.2%	-1.9%	-1.6%	-15.3%	-27.4%
110	37.96	22.6%	-1.9%	0.5%	-0.4%	-2.0%	-1.9%	-11.9%	-28.5%
All	11.90	44.1%	-15.2%	-0.2%	-0.4%	-1.3%	-1.9%	-26.7%	-15.2%

Another interesting observation from Table 1, relates to the decreasing inaccuracy with the increase in the overall distance of the paths for most of the alternative methods, except for A8. This is particularly evident in the case of the variable speed and A/D rates ignoring methods, i.e., A1, A2 and A7. One explanation for this might be the fact that longer paths normally include faster speed road segments such as motorways, where speed variations are much less present, and the overall speed profiles are smoother. In the case of A8, the manufacturer's reported energy consumption expectedly tends to be a close estimate for an empty vehicle driven over a short distance (as this is normally how standardised test procedures are conducted), but this quickly starts to deteriorate as the distance of the path increases and inaccuracies are accumulated.

Table 2 provides almost similar results for the case of the fully loaded vehicle with some expected increase in the overall inaccuracy of some of the methods due to the added impact of the vehicle payload.

Table 2
Summary results for comparison of different estimation methods based on the full vehicle mode

Category (<km)	Benchmark (consumed kWh)	Average inaccuracy							
		A1	A2	A3	A4	A5	A6	A7	A8
10	2.44	55.4%	-34.03%	-1.0%	-1.0%	-0.5%	-2.5%	-36.8%	-23.0%
20	5.44	53.1%	-34.52%	-0.9%	-0.6%	-0.6%	-2.1%	-38.8%	-25.5%
30	9.32	52.1%	-32.92%	-0.7%	-0.8%	-0.6%	-2.2%	-38.6%	-27.9%
40	12.79	49.8%	-31.09%	-0.6%	-0.6%	-0.8%	-2.0%	-38.2%	-27.5%
50	16.26	48.9%	-28.33%	-0.5%	-0.7%	-1.0%	-2.3%	-36.5%	-26.7%
60	19.60	40.7%	-19.91%	-0.1%	-0.7%	-1.0%	-1.9%	-31.2%	-27.2%
70	25.18	27.2%	-8.15%	0.4%	-0.4%	-1.2%	-1.2%	-21.9%	-31.7%
80	30.05	25.0%	-7.19%	0.5%	-0.4%	-1.5%	-1.4%	-22.3%	-32.9%
90	33.91	25.3%	-6.76%	0.5%	-0.7%	-1.4%	-1.5%	-21.1%	-32.9%
100	37.88	23.0%	-4.10%	0.6%	-0.5%	-1.7%	-1.5%	-16.1%	-33.5%
110	41.25	22.7%	-3.17%	0.7%	-0.7%	-1.8%	-1.9%	-12.6%	-34.2%
All	14.22	40.5%	-21.16%	-0.2%	-0.6%	-1.1%	-1.9%	-30.8%	-29.1%

An overview of the impact of ignoring vehicle payload and vehicle regenerative braking in estimating the overall energy consumption is provided in Figure 9. The corresponding inaccuracy tends to be higher in shorter paths. A key explanation is again due to the lower speed variations observed in road segments within longer paths, as, for example, in the case of the regenerative braking, lower A/D rates and speed variations simply mean less braking.

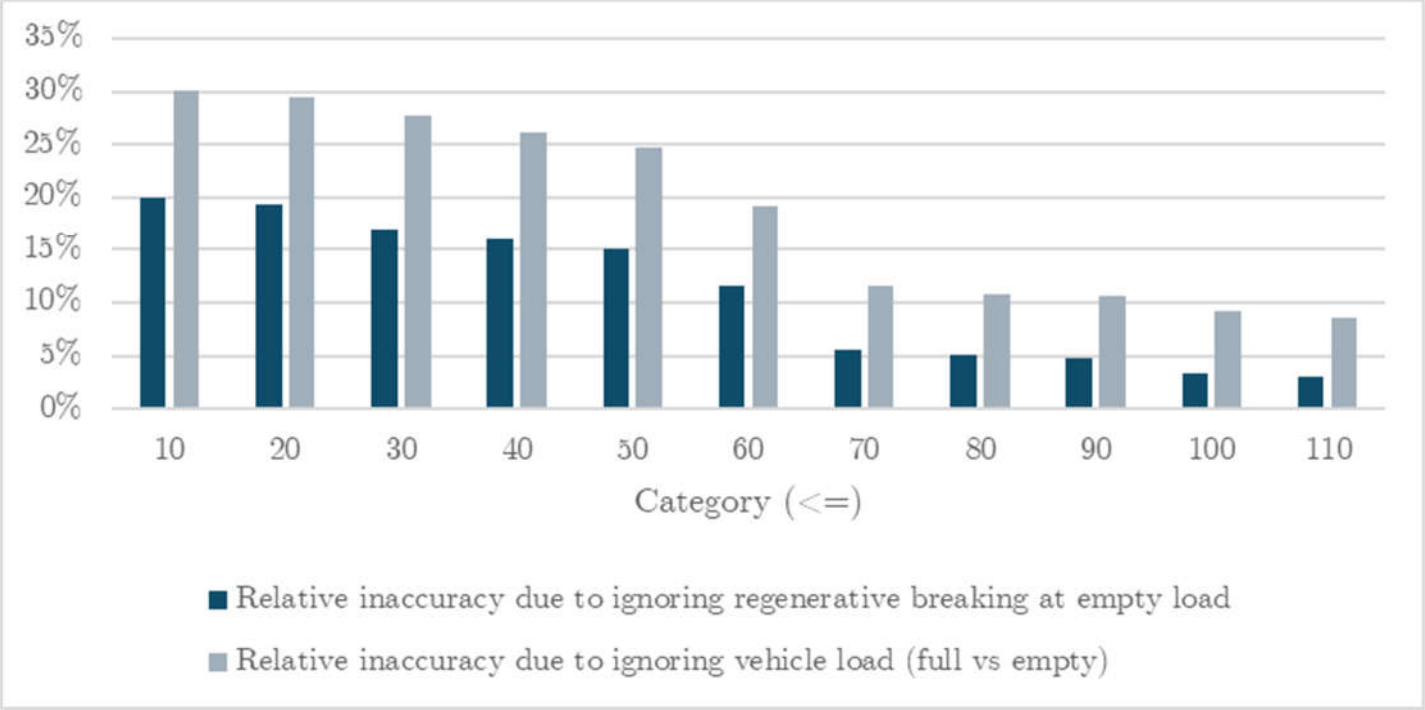


Figure 9 Impact of vehicle load and regenerative breaking

Following the analysis of the path-based results and assessing the impact of different variables on the estimation inaccuracy, we focus on the impact of the resulting inaccuracy on the optimality of routing decisions for EVRPTWs. To this end, all the 30 EVRPTW test instances discussed earlier were solved to optimality using the branch-and-bound solver of CPLEX 22.1.1 by assigning the distance-based (DB) and energy-based (EB) objectives to the mathematical model presented in (2)–(11), separately. To adapt and implement the proposed load-dependent EVRPTW mathematical model for each of the alternative estimation methods, the values of α_{ij}s and $\beta_{ij}s$ for each path $(i,j) \in A$ in the network and for each of the eight alternative methods were calculated and used. As a result, a total of 540 test instances were solved to optimality to establish a comprehensive comparative analysis.

The results are summarised into six categories of instances with 5, 10, 15, 20, 25 and 30 nodes (where each category contains five instances), and reported for DB and EB objectives in Tables 3 and 4, respectively. For each category we report the number of instances where no solution to the problem is found (NS), the number of instances where the solution found is infeasible with regard to the benchmark energy consumption model (IN), the average optimality gap for instances with a feasible solution (Opt Gap), and the average inaccuracy of the estimation method used against the benchmark model (Inaccuracy). Note that while the NS and IN figures both report infeasibility, NS denotes infeasible instances with respect to the alternative energy estimation model used (referred to as instance-infeasible hereafter), while IN denotes instances that are feasible with respect to the alternative estimation model (i.e., instance-feasible), but infeasible with respect to the benchmark model (referred to as benchmark-infeasible hereafter).

Table 3 shows test results for the DB objective function in two different vehicle battery capacity settings of 50 kWh and 75 kWh. Results against the NS rows in Table 3 indicate that over all instances only in two cases under A1 for the 50 kWh battery capacity there is no instance-feasible solution. In all other cases, an instance-feasible solution is found (regardless of whether these are benchmark-feasible or not which we will discuss shortly). This is mainly because A1 is a significantly overestimating model as was shown in Tables 1 and 2, and as can be seen in Table 3 against inaccuracy figures. As such this is expected, particularly in the case of a more constraining battery capacity.

Moving to the IN results in Table 3, it is interesting to see when A1 returns an instance-feasible solution, it is always a benchmark-feasible solution. However, the other seven models which always return an instance-feasible solution, in 16 out of 210 cases (around 8% of all cases) under the 50 kWh battery scenario return benchmark-infeasible solutions. When the battery size increases to 75 kWh, however, all tests return a benchmark-feasible solution over all estimation models.

The optimality gap results for the DB objective are calculated for each instance where a benchmark-feasible solution is found using the expression $((Obj' - Obj^*)/Obj^*) \times 100$ and then averaged for each category (Obj' is the optimal value returned by a model based on an alternative method, and Obj^* is the optimal value returned by the model based on the benchmark method). It is

interesting to see that except for the case of A1 which employs a significantly inaccurate energy estimation model, the average optimality gap for almost all other methods remains at 0%. Even in the case of A1 the optimality gap is not drastically high (9% on average).

Table 3
EVRPTW results for the distance-based experiments

Criterion	IG	DB_50 kWh								DB_75 kWh							
		A1	A2	A3	A4	A5	A6	A7	A8	A1	A2	A3	A4	A5	A6	A7	A8
	5	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	10	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
NS	20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	30	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	All	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	5	0	1	0	0	0	1	1	1	0	0	0	0	0	0	0	0
	10	0	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0
	15	0	1	0	0	0	1	1	1	0	0	0	0	0	0	0	0
IN	20	0	1	0	0	0	0	1	1	0	0	0	0	0	0	0	0
	25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	30	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	All	0	4	0	0	1	3	4	4	0	0	0	0	0	0	0	0
	5	7%	0%	0%	0%	0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	0%	0%
	10	14%	0%	0%	0%	0%	0%	0%	0%	2%	0%	0%	0%	0%	0%	0%	0%
	15	6%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Opt Gap	20	12%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	25	7%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	30	6%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	All	9%	0%	0%	0%	0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	0%	0%
	5	56%	-24%	-1%	0%	-1%	-2%	-34%	-9%	55%	-25%	-1%	0%	-1%	-2%	-34%	-9%
	10	64%	-29%	-1%	-1%	-1%	-2%	-34%	-4%	61%	-26%	-1%	-1%	-1%	-2%	-33%	-6%
	15	62%	-29%	-1%	0%	-1%	-2%	-35%	-6%	59%	-26%	-1%	0%	-1%	-2%	-33%	-6%
Inaccuracy	20	64%	-31%	-1%	-1%	-1%	-2%	-36%	-7%	63%	-30%	-1%	-1%	-1%	-2%	-35%	-6%
	25	64%	-31%	-1%	-1%	-1%	-3%	-35%	-7%	63%	-31%	-1%	-1%	-1%	-3%	-34%	-7%
	30	65%	-31%	-1%	-1%	-1%	-3%	-35%	-6%	64%	-31%	-1%	-1%	-1%	-3%	-34%	-7%
	All	63%	-29%	-1%	-1%	-1%	-2%	-35%	-7%	61%	-28%	-1%	-1%	-1%	-2%	-34%	-7%

Table 4 reports similar results for the case of the EB objective function. The main motivation behind these experiments is to measure the importance of energy consumption estimation accuracy when energy consumption itself is the objective function of interest. To clarify how the experiment results are reported here, let us denote by Obj^+ the actual energy consumption for the optimal solution returned by a model based on an alternative method if the benchmark method had been used. In other words, the optimal solution underpinning Obj' and Obj^+ is the exact same solution, but energy consumption over the routes within the solution has been calculated using two different methods: an alternative method and the benchmark method, respectively. Therefore, the reported optimality gaps and inaccuracy figures in Table 4 are based on the expressions $((Obj^+ - Obj^*)/Obj^*) \times 100$ and $((Obj' - Obj^*)/Obj^*) \times 100$, respectively.

While the figures in the NS rows of Table 4 remain consistent with those of Table 3 for the DB objective, those in the IN rows differ slightly as some of the previously benchmark-infeasible instances are now benchmark-feasible and vice versa. As far as the optimality gap is concerned, we observe only a negligible increase for some of the alternative estimation methods. Note that a 0% optimality gap against 'All' under a column with some reported optimality gaps is only due to rounding the results in this table. For example, the average optimality gap of A7 in the case of 50 kWh battery capacity over all instances is 0.39%, to be precise.

All in all, the results reported in this section provide multiple insights to address the key research questions investigated in this paper, and feed into the development of the intuitive methodology discussed in the next section to tackle energy estimation inaccuracy. As a key finding, we clearly see there is limited value in increasing the intricacy of an energy estimation model by incorporating hard-to-collect data such as road slope and weather conditions, which is the case for most routing applications. The impact of excluding these variables is not sufficiently important to jeopardise the feasibility of the prescribed solution or even compromise its optimality. In line with the extant literature, we also observe that ignoring instantaneous speed variability and A/D rates can lead to a significant estimation inaccuracy. Having said that, the impact of this inaccuracy on the feasibility and optimality of the resulting routing solution is marginal. For example, we observe that for a method with an average 34% underestimation of the actual energy consumption in the case of a 75 kWh battery capacity, all results match exactly those of the benchmark method with no infeasibility or optimality gap. Another interesting observation from the routing experiments (i.e., Tables 3 and 4)

corresponds to the overall inaccuracy of A8 (7% on average) which is the crudest method of all based on the manufacturer's reported consumption rate. This differs significantly from the reported performance in path-based experiments for A8, and is again a strong indication that a very simple approach to energy estimation might do the job excellently without the need to invest heavily in increasing intricacy and complexity of the considered energy estimation model.

Table 4
EVRPTW results for the energy-based experiments

Criterion	IG	EB_50 kWh								EB_75 kWh							
		A1	A2	A3	A4	A5	A6	A7	A8	A1	A2	A3	A4	A5	A6	A7	A8
	5	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	10	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	15	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
NS	20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	30	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	All	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	5	0	1	0	0	0	1	1	1	0	0	0	0	0	0	0	0
	10	0	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0
	15	0	0	0	0	0	0	0	1	0	0	0	1	1	1	0	0
IN	20	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0
	25	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	30	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	All	0	2	1	0	1	2	2	4	0	0	0	1	1	1	0	0
	5	8%	0%	0%	0%	0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	0%	2%
	10	13%	0%	0%	0%	0%	0%	0%	0%	2%	0%	0%	0%	0%	0%	0%	0%
	15	5%	2%	0%	0%	0%	0%	0%	0%	1%	2%	0%	0%	0%	0%	0%	1%
Opt Gap	20	12%	0%	0%	0%	0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	0%	1%
	25	7%	1%	0%	0%	0%	0%	1%	0%	0%	1%	0%	0%	0%	0%	1%	1%
	30	7%	1%	0%	0%	0%	0%	1%	0%	1%	1%	0%	0%	0%	0%	1%	2%
	All	9%	1%	0%	0%	0%	0%	0%	0%	1%	1%	0%	0%	0%	0%	0%	1%
	5	57%	-24%	-1%	0%	-1%	-2%	-33%	-9%	59%	-25%	-1%	0%	-1%	-2%	-33%	-9%
	10	63%	-29%	-1%	-1%	-1%	-2%	-34%	-4%	63%	-26%	-1%	-1%	-1%	-2%	-33%	-6%
	15	62%	-29%	-1%	0%	-1%	-2%	-34%	-6%	57%	-29%	-1%	0%	-1%	-2%	-34%	-5%
Inaccuracy	20	62%	-30%	-1%	-1%	-1%	-3%	-35%	-7%	61%	-30%	-1%	-1%	-1%	-2%	-35%	-6%
	25	64%	-32%	-1%	-1%	-1%	-3%	-36%	-7%	63%	-32%	-1%	-1%	-1%	-3%	-36%	-7%
	30	64%	-32%	-1%	-1%	-1%	-3%	-35%	-6%	62%	-32%	-1%	-1%	-1%	-3%	-35%	-7%
	All	62%	-29%	-1%	-1%	-1%	-2%	-35%	-7%	61%	-29%	-1%	-1%	-1%	-2%	-34%	-7%

The results reported in this section also unveil interesting theoretical and practical insights into a potential issue with the general conservative tendency in the literature to incline towards an overestimation of energy consumption rather than an underestimation. Our results indicate that there is only one overestimating approach (i.e., A1) which ironically tends to deliver the worst performance when compared to all other methods that are underestimating. We discuss this further in the Section 4 where we develop relevant intuitive results that underpin a simple but practical methodology to tackle the energy estimation inaccuracy in the EVRPTW.

4. Tackling the impact of inaccuracy

Section 3 focused on assessing the impact of energy consumption estimation inaccuracy on the optimality of EVRPTW solutions. In this section, we focus on the development of an intuitive, simple-to-implement and yet practical methodology to tackle this impact. It is based on a bi-objective ε-constraint optimisation framework and builds on intuitive results and findings from the previous section. It only uses the basic consumption rate of the vehicle as the energy estimation model, and finds multiple nondominated solutions along the efficient frontier of the problem instance. Each of these solutions are guaranteed to lie within a known factor of the optimal solution and have a certain robustness level against energy estimation inaccuracy.

We start this section by discussing the proposed methodology and its underpinning definitions and theoretical assumptions, and we then present the results of numerical experiments to assess its performance.

4.1. The bi-objective ε-constraint based methodology

The proposed methodology recognises the fact that accurately estimating the actual on-road ECV energy consumption is very difficult if not impossible, and any estimation method bears a certain level of unavoidable inaccuracy. Rather than incorporating an intricate and complex energy consumption estimation method, the proposed methodology applies the simplest of all alternative

methods discussed in Section 3, i.e., the generic energy consumption rate provided by the vehicle manufacturer, and shows that even with such a crude and readily available estimation method, and allowing for a high level of inaccuracy, there are ways to generate solutions that are likely to be on-road-optimal, or near-optimal within a known tolerance of the identified optimal solution value.

We start with some definitions and intuitive remarks and arguments that provide the underpinning basis for the proposed methodology:

Definition 1 *A(n) (in)feasible routing solution to an instance of the EVRPTW with an incorporated energy consumption estimation model is referred to as "an instance-(in)feasible solution". We use π_f to indicate an instance-feasible solution, and the set Π to refer to the full set of instance-feasible solutions.*

In this definition and all of the following definitions, (in)feasibility is computed with respect to vehicle battery charge level constraints.

Definition 2 *An instance of the EVRPTW with an incorporated energy consumption estimation model is referred to as "an infeasible instance" if no instance-feasible solution to the problem can be found (i.e., $\Pi = \emptyset$).*

Definition 3 *An optimal routing solution to an instance of the EVRPTW with an incorporated energy consumption estimation model is referred to as "an instance-optimal solution" (denoted by $\pi_o \in \Pi$).*

All reported optimal solutions to the instances in the previous section are therefore instance-optimal.

Definition 4 *A(n) (in)feasible on-road routing solution to an instance of the EVRPTW is referred to as "an on-road-(in)feasible solution". We use θ_f to refer to an on-road-feasible solution, and the set Θ to refer to the full set of on-road-feasible solutions.*

Definition 5 *An optimal on-road routing solution to an instance of the EVRPTW is referred to as an "on-road-optimal solution" (denoted by $\theta_o \in \Theta$).*

It is important to reiterate that as we assume the actual on-road energy consumption to be unknown, so is the actual on-road-optimal solution. In fact, it is even difficult to ensure on-road feasibility. As such, the proposed methodology tries to 'build confidence' in on-road optimality implicitly rather than explicitly as will be discussed shortly. For all of the following arguments we assume the problem at hand is on-road feasible and there exists at least one on-road-feasible solution (hence an optimal on-road solution) to the problem.

Owing to the observations derived from the experiments of the previous section, we argue that different methods of energy consumption estimation tend to 'underestimate' or 'overestimate' on average the actual on-road consumption. While any form of inaccuracy is unfavourable and may potentially lead to infeasibility or suboptimality, there is useful insight in the overall behaviour of an applied estimation method. Therefore, we put forth the following definition which is an important building block to the results that follow:

Definition 6 *An energy consumption estimation model that over(under)estimates the actual on-road energy consumption of the vehicle is referred to as "an over(under)estimating model".*

We must note that, in practice, it is unrealistic to expect an estimation model to strictly overestimate (or underestimate) energy demand in 100% of cases. However, as discussed in Appendix C, the basicRate model (i.e., the generic energy consumption rate reported by vehicle manufacturers) exhibits a strong overall underestimating tendency, which is good enough for inclusion within our proposed methodology. Across 26,950 comparisons involving 2,450 paths and 11 payload levels, the model underestimates the benchmark energy consumption in 92.3% of cases, with an average underestimation of 16.5%. The unlikely instances of overestimation are expected to be offset at route-level by more frequent and pronounced underestimations. In Appendix C, we also refer to a practical way to mitigate extreme cases of unexpected model overestimation.

Lemma 1 *Assume that Π^+ and Π^- denote the full set of instance-feasible solutions to an instance of the EVRPTW with incorporated overestimating and underestimating models, respectively. Then, it is guaranteed that $\Pi^+ \subseteq \Theta \subseteq \Pi^-$.*

Proof. Consider $\pi_f^+ \in \Pi^+$ as an instance-feasible solution to an instance of the EVRPTW with an incorporated overestimating model M^+, and an estimated energy consumption value of $M^+(\pi_f^+)$. Also, consider $\theta_f \in \Theta$ as an on-road-feasible solution to the same instance with actual on-road energy consumption value $M(\theta_f)$. We need to prove that: (i) $\pi_f^+ \in \Theta$, (ii) $\pi_f^+ \in \Pi^-$, and (iii) $\theta_f \in \Pi^-$. Assuming M^- is an underestimating model, it is guaranteed that $M^+(\pi_f^+) \geq M(\pi_f^+) \geq M^-(\pi_f^+)$, and $M(\theta_f) \geq M^-(\theta_f)$; as such π_f^+ is guaranteed to appear as a feasible solution in both Θ and Π^-, and θ_f is guaranteed to appear as a feasible solution in Π^-. □

Remark 1. *Under the worst-case scenario, incorporating an overestimating model into an instance of EVRPTW can lead to instance-infeasibility (i.e., $\Pi^+ = \emptyset$).*

To clarify this remark, assume there is only one on-road-feasible solution θ_f that consumes $M(\theta_f)$ units of energy, which also happens to be equal to the total capacity of the ECV battery B. Therefore, the incorporation of a model M^+ that overestimates the consumption of the very solution at $M^+(\theta_f)$, where $M^+(\theta_f) > M(\theta_f) = B$, makes this single solution infeasible, and hence leads to instance-infeasibility (i.e., $\Pi^+ = \emptyset$). Note that this is not in contradiction with Lemma 1, as the empty set ($\emptyset$) is a subset of any set.

Corollary 1. *All instance-feasible solutions (hence, also the instance-optimal solution) to an instance of the EVRPTW with an incorporated overestimating model are guaranteed to be on-road-feasible.*

This corollary is a direct implication of a part of Lemma 1, i.e., $\Pi^+ \subseteq \Theta$, and presents a simple and intuitive argument which constitutes the backbone to most conservative and robust optimisation approaches.

Remark 2. *Under the worst-case scenario, the instance-optimal solution to an instance of the EVRPTW with an incorporated underestimating model (π_o^-) is on-road-infeasible (i.e., $\pi_o^- \notin \Theta$).*

This is the case when $M^-(\pi_o^-) < M(\pi_o^-)$, and $M(\pi_o^-) > B$.

Corollary 2. *All on-road-feasible solutions (hence, also the on-road-optimal solution) are guaranteed to be instance-feasible for an instance of the EVRPTW with an incorporated underestimating model.*

This corollary is also a direct implication of a part of Lemma 1, i.e., $\Theta \subseteq \Pi^-$.

Up to this point, it seems enticing to adopt a conservative approach and incorporate an overestimating energy estimation method into the problem as π_o^+ is guaranteed to be on-road-feasible based on Corollary 1, but π_o^- can be on-road infeasible based on Remark 2. However, Theorem 1 establishes an important result that is the underpinning justification for the proposed methodology and the decision to incorporate an underestimating model of estimation:

Theorem 1. *An on-road-feasible instance-optimal solution to an instance of the EVRPTW with an incorporated underestimating model (i.e., $\pi_o^- \in \Theta$) is guaranteed to be on-road-optimal (i.e., $\pi_o^- \equiv \theta_o$).*

Proof. Let $Z(\theta_o)$ and $Z(\pi_o^-)$ be the objective function values of the on-road-optimal solution, θ_o, and the instance-optimal solution, π_o^-, respectively. The assumption $\pi_o^- \in \Theta$, and the on-road optimality of θ_o necessitates $Z(\theta_o) \leq Z(\pi_o^-)$. On the other hand, based on Lemma 1 and Corollary 2, we know that $\Theta \subseteq \Pi^-$, and $\theta_o \in \Pi^-$, and the instance-optimality of π_o^- necessitates $Z(\theta_o) \geq Z(\pi_o^-)$. Therefore, $Z(\theta_o) \leq Z(\pi_o^-) \wedge Z(\theta_o) \geq Z(\pi_o^-) \Longrightarrow Z(\theta_o) = Z(\pi_o^-)$. □

To further clarify this Theorem and understand how this may help develop a suitable solution methodology, assume that we incorporate the basicRate model into an instance of the EVRPTW which yields four instance feasible solutions r_1, r_2, r_3, and r_4 with total distances of 5, 6, 7, and 8, and total energy consumption values of 6, 5, 7, and 8, respectively; also, assume the vehicle battery capacity is 6, and the actual on-road energy consumption of these routes is 7, 6, 8 and 9. As such, the instance-optimal solution of the problem (i.e., r_1 in this case) is on-road infeasible; however, if we disregard r_1, then r_2 is an on-road-feasible, instance-optimal solution that happens to be also the on-road-optimal solution to the problem.

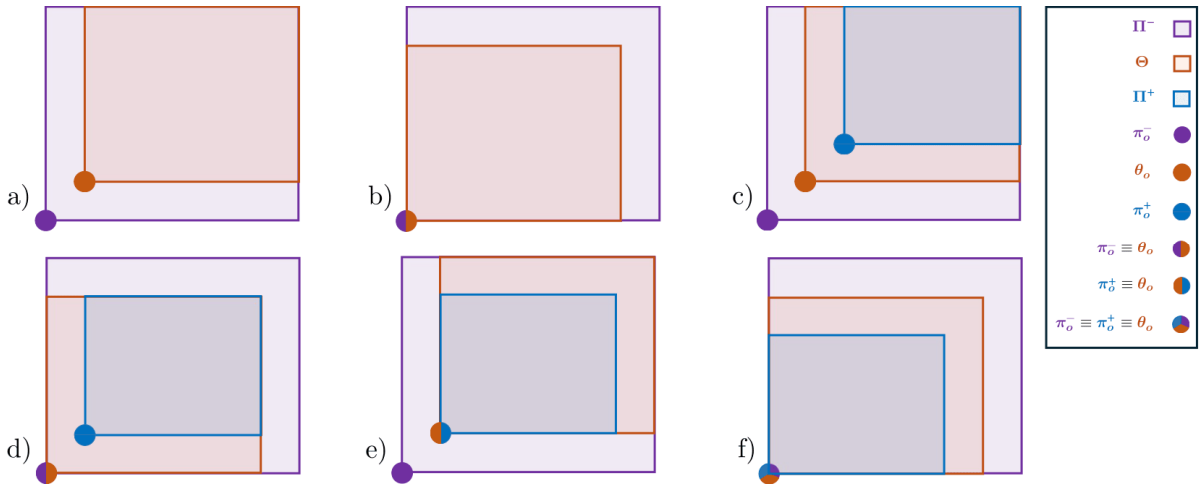


Figure 10 Set-relation diagrams of all possible scenarios of instance-optimality vs. on-road-optimality

A schematic set-relation diagram of all possibilities of the above-mentioned arguments is illustrated in Figure 10. In all cases, the $\Pi^+ \subseteq \Theta \subseteq \Pi^-$ relationship is implied by using nested rectangular feasible solution regions (represented by blue, orange and purple rectangles, respectively). For the case of the overestimating model, in two cases (i.e., Figure 10.a and b), $\Pi^+ = \emptyset$ (Remark 1), and in other cases the instance-optimal solution is either on-road-feasible (i.e., Figure 10.c and d) or on-road-optimal (i.e., Figure 10.e and f). For the case of the underestimating model, there are scenarios (i.e., Figure 10.a, c and e) where the instance-optimal

solution is on-road-infeasible (Remark 2), and scenarios (i.e., Figure 10.b, d and f) where the instance-optimal solution is on-road-feasible and hence on-road-optimal (Theorem 1).

Given the above discussion and results, the key in identifying the on-road-optimal solution lies in building a confidence measure of 'on-road feasibility' of the instance-optimal solution. One way to achieve this is to work out the maximum underestimation inaccuracy that leave the optimal solution unchanged. Clearly, if there is a sufficiently high level of inaccuracy allowance for a given instance-optimal solution, and the decision maker is confident that the incorporated estimation method is expected to perform significantly better than that, there is a high chance that the given instance-optimal solution is on-road-optimal based on Theorem 1. However, if the allowance for inaccuracy is low, there will be no fall-back solution, and the decision maker is left with a single solution that bears a high level of risk of on-road infeasibility. To address this situation, we propose a bi-objective optimisation framework in which we fully generate solutions on the efficient frontier of the problem such that each solution presents a certain level of inaccuracy allowance, and a certain distance from the instance-optimal solution. A simple example of this is given in Figure 11 for two different EVRPTW instances. In this figure, in the case of Example 1, the efficient frontier comprises a single optimal solution which remains optimal even with an energy underestimation inaccuracy as high as 100% (i.e., half of the actual on-road energy consumption). Therefore, it can be 'confidently' argued that the instance-optimal solution is expectedly on-road-feasible and hence on-road-optimal based on Theorem 1. In the case of Example 2, the efficient frontier comprises three Pareto-optimal solutions; the instance-optimal solution in this case remains optimal for an energy underestimation inaccuracy of as high as 40%, which is again quite a large allowance for inaccuracy; but if the decision maker is quite risk-averse, with small sacrifice in instance-optimality (less than 2%) the inaccuracy allowance increases twice as much to 80%.

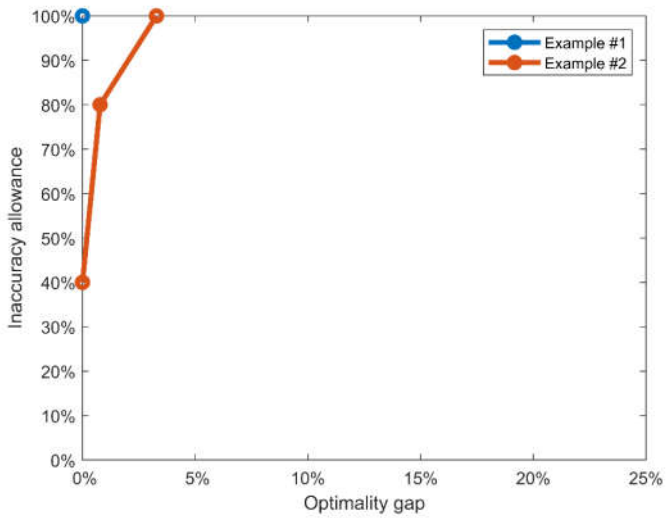


Figure 11 Efficient frontier examples

The core mathematical model that sits at the heart of the proposed ε-constraint based methodology to generate such efficient frontier for an instance of the EVRPTW is presented below:

$$z_1 = \sum_{(i,j)\in A} d_{ij}x_{ij} \tag{13}$$

$$z_2 = \Lambda \tag{14}$$

$$\text{Minimise } z_1, \text{Maximise } z_2 \tag{15}$$

subject to (3), (4), (5), (6), (7), (8), (9), (11) and

$$0 \le u_j \le u_i - \delta_{ij}\lambda_{ij} + B\big(1 - x_{ij}\big), \qquad \forall i \in N_0, j \in N_\sigma, i \ne j \tag{16}$$

$$\Lambda - 2\big(1 - x_{ij}\big) \le \lambda_{ij} \le \Lambda, \qquad \forall (i,j) \in A \tag{17}$$

$$\sum_{i\in N_0} \lambda_{ij} - \sum_{i\in N_\sigma} \lambda_{ji} = 0, \qquad \forall j \in C \tag{18}$$

$$\sum_{(i,j)\in A} d_{ij}x_{ij} \le \varpi \tag{19}$$

$$\Lambda \geq \emptyset, \tag{20}$$

where:

- Λ is a continuous decision variable in the range $[1,2]$ denoting the overall maximum inaccuracy allowance, with 1 providing no inaccuracy allowance, and 2 allowing maximum allowed inaccuracy of twice as much as the estimated figure, i.e., 100% inaccuracy,
- λ_{ij} is a continuous decision variable in the range $[1,2]$ denoting the maximum inaccuracy allowance over the arc $(i,j) \in A$, and
- δ_{ij} is the estimated load-independent energy consumption based on the basicRate model over the arc $(i,j) \in A$.

The two objective functions in (13) and (14) denote the total distance and the maximum inaccuracy allowed. As will be shortly discussed within the ε-constraint based methodology, in each iteration we assign one of the two objectives and different right-hand side values ϖ and ϕ to the inequalities (19) and (20), respectively, to guide the algorithm progression in search of Pareto-optimal solutions. Constraints (16) apply the inaccuracy multiplier λ_{ij} on each δ_{ij} over each arc $(i,j) \in A$. Note that while δ_{ij} denotes the nominal (deterministic) energy consumption obtained from the basicRate model, $\lambda_{ij} \in [1,2]$ scales this nominal value to account for innacuracy, such that $\lambda_{ij} = 1$ assumes zero buffer, and $\lambda_{ij} = 2$ allows for up to 100% deviation (i.e., consumption could be twice the estimated value). From a modelling perspective, λ_{ij} enables the optimisation to endogenously allocate robustness on each arc, and the term $\delta_{ij}\lambda_{ij}$ determines the maximum allowable energy consumption over the arc such that the vehicle battery capacity and the bounds on constraints (19) and (20) are not violated. While arcs may receive different multipliers, constraints (17) and (18) together ensure that a consistent inaccuracy level is applied across all arcs that are part of the optimal solution and the overall inaccuracy allowance is set on the smallest values of all λ_{ij}. Note that maximising Λ as per the objective function (14), ensures that the values of λ_{ij} are also maximised.

The pseudocode for the proposed ε-constraint based methodology is given in Algorithm 1. Here, we refer to the model presented in (13) to (20) as model "M", and the objective assignment is indicated by $z \leftarrow 1$ or $z \leftarrow 2$ depending on whether the objective function expression (13) or (14) has been used in the model, respectively. Also, M_1^* and M_2^* denote the values for the total distance and the maximum inaccuracy allowed as a result of optimising M, respectively.

Algorithm 1 The proposed ε-constraint based methodology

```
Input ε
Initialise P = {}, ϖ ← ∞, ϕ ← 1, and z ← 1;
while True do
    Solve M
    if M is infeasible then
        break;
    Else
        if z = 1 then
            z ← 2, ϖ ← M1*, and ϕ ← 1
        else
            P = {P, (M1* M2*)};
            z ← 1, ϖ ← ∞, and ϕ ← M2* + ε;
        end if
    end if
end while
Return P
```

The performance of the proposed methodology is discussed in Section 4.2 through a set of numerical tests, but we first provide a generalisation of the proposed results and the algorithm for the case of EVRPTWs with en-route recharging.

4.1.1. The case of the EVRPTW with en-route recharging

In this section, we discuss the implications of relaxing the initial assumption that all vehicles are fully charged overnight at the starting depot and no en-route visits to RSs are allowed.

Clearly, the inclusion of en-route recharging can change the solution space significantly; however, we argue that the key impact on the proposed results and the algorithm will be with respect to the time-windows feasibility. We recall that in all discussions presented in Section 4.1, we considered (in)feasibility with respect to vehicle battery charge level constraints only. This is because changing the estimation model used for energy consumption prediction does not change the travel times, and hence the arrival times. However, when en-route recharging is considered, the amount of time spent at an RS is a direct result of the amount of recharging needed, which is in turn a result of the energy consumption estimation on arcs leading to the RS. As such, arrival times

are affected by energy estimation, and so, time-windows violation can happen. In other words, an instance-feasible solution based on an underestimating energy demand model can be on-road feasible in terms of battery charge level, but violating customer time windows. To clarify this, we use a simple illustrative example shown in Figure 12. In this figure, node 0 is the depot, nodes 1 and 2 are customer nodes, and node 3 is an RS. A single number is used on each arc to denote distance, travel time and the estimated energy consumption (using an underestimating method). Suppose the ECV battery capacity is 18, and the inverse charging rate of the RS is 0.4 (i.e., fully recharging the battery from zero takes 7.2 time units at an RS). Also, assume that service time is zero and the time window on the depot, and both customer nodes is [0,24].

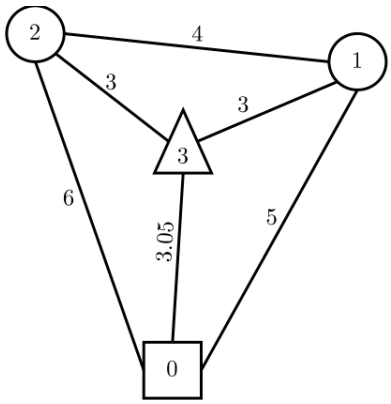


Figure 12 An illustrative example for the EVRPTW with en-route recharging

Considering the case of EVRPTW without en-route recharging, there are only two solutions on the efficient frontier of the problem using the proposed methodology corresponding to: (i) 0→1→2→0 (optimality gap = 0%, inaccuracy allowance = 20%), and (ii) 0→1→0, 0→2→0 (optimality gap = 46.67%, inaccuracy allowance = 50%). Any changes to the time windows in this example (as long as the problem remains time feasible), will not change the solutions on the efficient frontier.

Including the RS and extending the problem to the case of EVRPTW with en-route recharging, on the other hand, changes the solution space, and while the instance-optimal solution remains the same, i.e., 0→1→2→0 (optimality gap = 0%, inaccuracy allowance = 20%), there will be other interesting solutions on the efficient frontier such as 0→1→2→3→0 (optimality gap = 0.33%, inaccuracy allowance = 50%) and 0→1→3→2→0 (optimality gap = 13.33%, inaccuracy allowance = 100%). However, any changes to time windows can render some of the solutions on the efficient frontier infeasible or have an impact on their inaccuracy allowance. For example, if we change the time-window on customer 2 to [0,14], the inaccuracy allowance of the solution 0→1→3→2→0 drops from the initial 100% to 50%, and if we decrease the upper bound of the time window further to [0,11.5], the inaccuracy allowance drops to below 14% and the solution is no more Pareto-optimal and does not appear on the efficient frontier.

Given the discussion above, if we extend the assumption in the arguments and results presented in Section 4.1 to consider (in)feasibility with respect to both vehicle battery charge level and time windows constraints, all results stand relevant and correct (in the case of the considered simplified charging environment), and Theorem 1 can be extended as follows for the case of EVRPTW with en-route recharging:

Corollary 3. *If an instance-optimal solution to an instance of the EVRPTW with en-route recharging and an incorporated underestimating model is on-road-feasible with respect to time windows and battery charge level constraints, it is guaranteed that the same solution is also on-road-optimal.*

Proof. The proof used for Theorem 1 remains correct for this corollary. □

As regards the proposed solution methodology, the structure of the algorithm remains the same and we only need to replace the core model M presented in (13) to (20) with an extended version M' which accommodates en-route recharging. This can be simply achieved by applying few modifications and adding recharging-related constraints similar to those presented in Keskin and Çatay (2016).

The required bi-objective EVRPTW with en-route recharging (M') is defined on an extended graph $G' = (N', A')$, where $N' = \{D' \cup C \cup R\}$ includes a set of m identical RSs, $R = \{n+1, \dots, n+m\}$, the depot set $D = \{0, n+m+1\}$, where $\{n+m+1\}$ is the dummy copy of $\{0\}$ as the final depot, and the customers set C as before. All RSs are associated with the same inverse charging rate of g. We also define three new sets corresponding to $N'_r = C \cup R$, $N'_0 = \{0\} \cup C \cup R$ and $N'_\sigma = C \cup R \cup \{n+m+1\}$ in addition to the previously defined sets for ease of reference. Correspondingly, the set of directed arcs in G' is defined as $A' = \{(i,j) | i \in N'_0, j \in N'_\sigma, i \neq j\}$. A new continuous decision variable $Y_i \in [0, B]$ is also defined to denote the residual battery charge level of an ECV on departure from an RS $i \in R$. Then, M' can be presented as below:

$\text{Minimise } z_1, \text{Maximise } z_2$ (as per the objective functions presented in (13) and (14), respectively, where A is replaced with A' in z_1).

subject to (3), (4), (5), (6), (7), (8), (9), (11), (16), (17), (18), (19) and (20) where N_0 is replaced with N_0' in (4), (8) and (18), and with C in (16), N_σ is replaced with N_σ' in (3), (4), (6), (8), (16) and (18), C is replaced with N_r' in (4) and (18), N is replaced with $D' \cup C$ in (5) and (7), A is replaced with A' in (9), (17) and (19), and:

$$\sum_{j \in N_\sigma'} x_{ij} \leq 1, \qquad \forall i \in R \tag{21}$$

$$y_i + t_{ij}x_{ij} + g(Y_i - u_i) - (l_0 + gB)\big(1 - x_{ij}\big) \leq y_j, \qquad \forall i \in R, j \in N_\sigma', i \neq j \tag{22}$$

$$0 \leq u_j \leq Y_i - \delta_{ij}\lambda_{ij} + B\big(1 - x_{ij}\big), \qquad \forall i \in \{0\} \cup R, j \in N_\sigma', i \neq j \tag{23}$$

$$u_i \leq Y_i \leq B, \qquad \forall i \in \{0\} \cup R. \tag{24}$$

Constraints (21) allow each RS to be visited at most once. Note that, to allow multiple visits to the same RS, dummy copies can be added to R. Constraints (22) ensure the time feasibility of arcs leaving the RSs. Constraints (23) and (24) together keep track of the battery charge level following the ECV departure from an RS and the depot and make sure that it does not exceed its capacity.

We present numerical results on instances with en-route recharging in Section 4.2.1.

4.2. Analysing Pareto-optimal solutions on the efficient frontier

To analyse the performance of the proposed methodology, we have conducted a set of tests on the instances of sizes 5, 10, and 15 introduced in Section 3.4, for the distance-based objective function in two battery capacity settings of 50 kWh and 75 kWh (a total of 30 test instances). The results of all tests are grouped and graphically illustrated in Figure 13. The x-axis on the generated graphs shows the optimality gap (with respect to the instance-optimal solution), while the y-axis shows the maximum inaccuracy allowed.

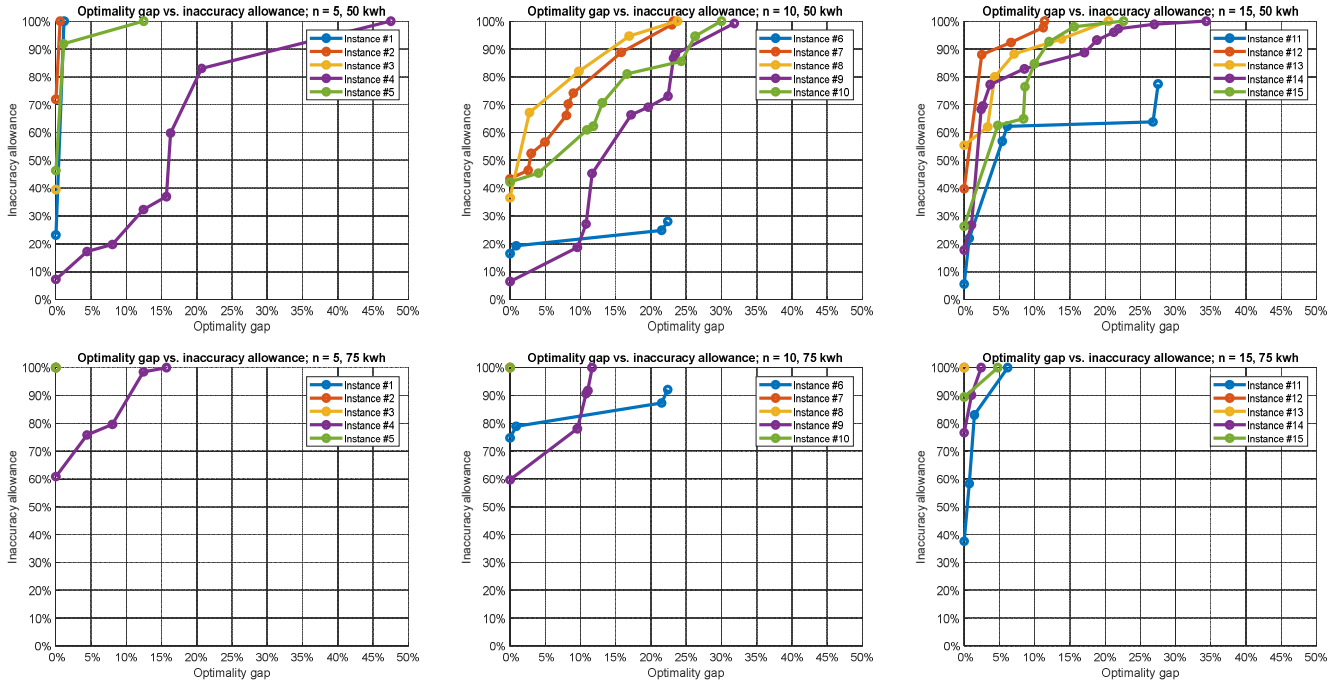


Figure 13 Optimality gap vs inaccuracy allowance of the Pareto-optimal solutions on the efficient frontier of test instances

Focusing first on the 50 kWh battery size setting, the most striking observation is that the instance-optimal solution in most cases already provides a significant allowance for inaccuracy: 32% on average and over all cases, and this is already significantly higher than the observed average inaccuracy of the basicRate model at 16.5% (see Appendix C). More importantly, the noticeable sharp rising shape of all graphs and the concentration of most points within the first 25 centiles of optimality gap (in fact, 37% of all points are within the first five centiles) suggest that with small sacrifice in instance-optimality in cases where we lack confidence in the solution due to its rather low allowance for inaccuracy, a very risk-averse but near optimal solution is obtainable. An extreme case of this is shown in instance #1, where with just about a 1% sacrifice in instance-optimality (from 146.57 km to 148.30 km), the inaccuracy allowance rises sharply from 23% to 100%. This being said, there are also stubborn instances where the inaccuracy

allowance is very resistant to an optimality gap increase. Instance #6 serves as a good example for this situation where increasing optimality gap to over 22%, only increases inaccuracy allowance from 16% to 28%.

It is also useful to analyse these results against those reported in Section 3, Table 3 for the case of the 50 kWh battery capacity. If we assume that the benchmark model is providing a close proxy for on-road energy consumption, we can already see that in three cases (one case under each group of 5, 10 and 15 nodes) the provided solution by incorporating A8 (i.e., the basicRate model) is infeasible, and in all other cases the returned solution is same as the optimal solution found using the benchmark model. Looking closer at the three infeasible solutions (which correspond to solutions to instances #4, #6 and #11), and cross-checking these with their corresponding plots in Figure 13, we can already see that the instance-optimal solutions to these instances provide a limited inaccuracy allowance of 7%, 16% and 5%, respectively, whereas the solutions that match with those of the benchmark model already appear on the efficient frontier with inaccuracy allowances of 17%, 25%, 22%. In all other instances except these three instances, the instance-optimal solution matches the 'on-road-optimal' solution.

Moving onto the case of 75 kWh battery size, we observe a significantly improved situation and a considerable decrease in the number of nondominated points on the efficient frontier. Not only does the instance-optimal solution always match the assumed 'on-road-optimal' solution (i.e., based on the benchmark estimation), but in several cases we observe that there is only one point on the generated efficient frontier of the instances. This suggests that the instance-optimal solutions for these instances are resistant to up to 100% underestimation inaccuracy.

To clarify the size of the achieved improvement by 25 kWh increase in battery size, we have generated two heat maps for each of the 50 kWh and 75 kWh battery size cases in Figure 14. The number of rows in these heat maps corresponds to the largest number of Pareto-optimal solutions observed considering all instances in each case (12 points on the efficient frontier of instance #14 in the case of 50 kWh battery capacity, and five points for instance #4 and #9 in the case of the 75 kWh battery capacity). Starting from the top, each row indicates the average optimality gap and inaccuracy allowance of the best solution (optimal solution in terms of total distance), second-best solution, third-best solution, and so on. The indicated values are based on the mean of values for all instances that have a corresponding Pareto-optimal point within their efficient frontier set. For example, only 14 instances in the case of 50 kWh battery capacity have a second-best solution, whereas there are only six such instances in the case of the 75 kWh battery capacity, and the calculated values are based on the mean of the optimality gap and inaccuracy allowance of the corresponding second-best Pareto-optimal solutions. As such, while solutions on the efficient frontier of a given instance are guaranteed to offer increased inaccuracy allowance with the increase in optimality gap, in the case of the provided heatmap for the 75 kWh battery capacity we have (0%, 86.6%) and (3.6%, 80.2%), which seems counter-intuitive at first sight. To clarify this better, assume there are two instances with the following efficient frontier sets: Instance #1: {(0%, 90%)} (containing one Pareto-optimal solution), Instance #2: {(0%, 50%), (2%, 60%)} (containing two Pareto-optimal solutions). Averaging (as done for the heatmaps in Figure 14) would then yield {(0%, 70%), (2%,60%)}.

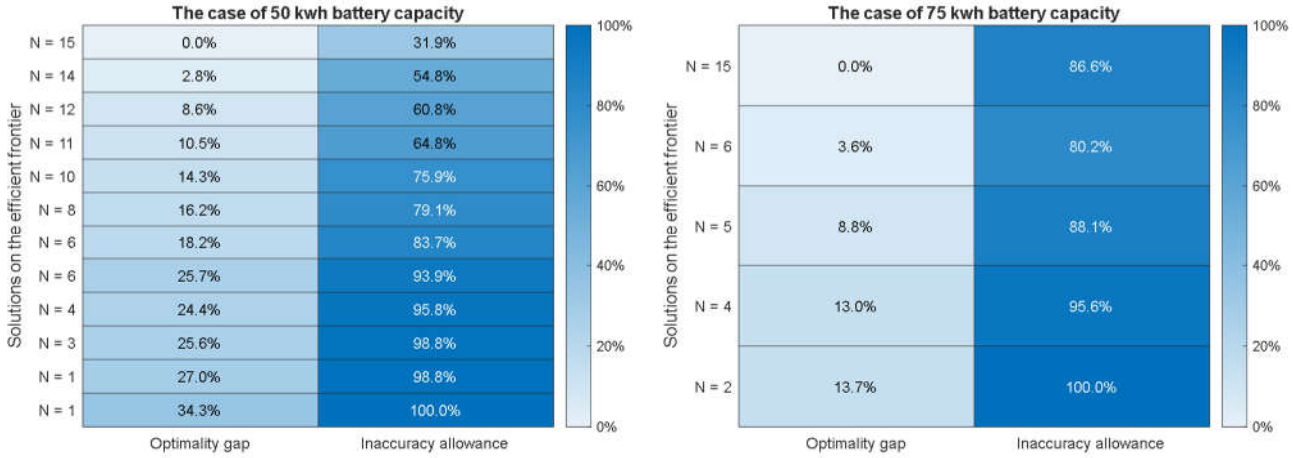


Figure 14 Heat maps summarising the efficient frontier of test instances in the case of 50 kWh and 75 kWh battery capacities

There are various clear take aways from these heat maps. For one thing, the instance-optimal solutions on the efficient frontier of instances in the case of 50 kWh battery capacity are tolerant against around 32% underestimation inaccuracy (which is, as discussed, already a quite good allowance for inaccuracy given the much lower expected underestimation inaccuracy of the basicRate method), whereas this is around 87% in the case of 75 kWh battery capacity, and this is a significantly high allowance for inaccuracy. Secondly, we can see that the optimality gap column in the case of the 75 kWh battery capacity is always in quite light colour as the solutions on the efficient frontier of instances in this case never need to go too far from the optimal solution to ensure complete confidence in on-road feasibility. It is also worth highlighting an interesting observation related to the case of 50 kWh battery capacity. This corresponds to the average optimality gap and inaccuracy allowance of the second-best solutions on the efficient frontier of relevant instances (i.e., 2.8% and 54.8%, respectively), which simply indicates that with a marginal sacrifice

in instance optimality, a significant allowance for inaccuracy is achievable, and second-best solutions overall seem to be an excellent option for a strictly risk-averse decision maker.

All in all, the results of this section complement those of Section 3 with a clear message that EVRPTW routing decisions seem to remain robust against a high level of estimation inaccuracy in most cases, and with only a small sacrifice in optimality, it is possible to confidently pick a solution that remains on-road-feasible and probably on-road-(near)optimal. As such, there is limited value in incorporating a highly intricate and detailed energy estimation model into routing decisions, as more often than not it is good enough to just use a crude model based on the manufacturer's reported consumption rate, which is readily available. More importantly, we observe and argue that with the ongoing improvements in the battery capacity of ECVs, energy estimation accuracy becomes less and less of a concern, at least in the case of urban last-mile deliveries. Having said that, we must also note that the results presented in this paper and the rather small optimality gaps reported as a result of using a simple energy estimation model in most cases can still be operationally significant in large-scale logistics operations where even small percentage increases in distance or energy consumption may have considerable economic or service-level implications. In such cases, the cost of error may well justify investment in incorporating a detailed and highly accurate energy estimation module.

4.2.1. Instances with en-route recharging

A set of 20 test instances with en-route recharging have been also generated and examined. For each one of these instances, first, we have randomly selected 13 nodes from the set of 50 nodes generated in Greater London, and assigned the nodes with shortest, second-shortest and third-shortest sum distance to all other nodes as the locations for RSs (with the first one being at the depot), and then created three sets of dummy copies from these nodes to allow each to be visited up to three times. Therefore, for each of these instances we have $n = 10$, and $m = 9$. The ECV battery capacity is also reduced to 30 kWh and we assume 50 kW RSs with an inverse charge rate of 1.2 (min/kWh). These settings reinforce the inclusion of Pareto-optimal solutions with en-route recharging on the efficient frontier. All other aspects of the instances such as the time windows and customer demands have been generated using the same procedure explained earlier for instances without en-route recharging.

Figure 15 illustrates the generated efficient frontier for each of the 20 instances considered and Figure 16 provides a summary analysis of the Pareto-optimal solutions on the efficient frontier of instances on a heatmap.

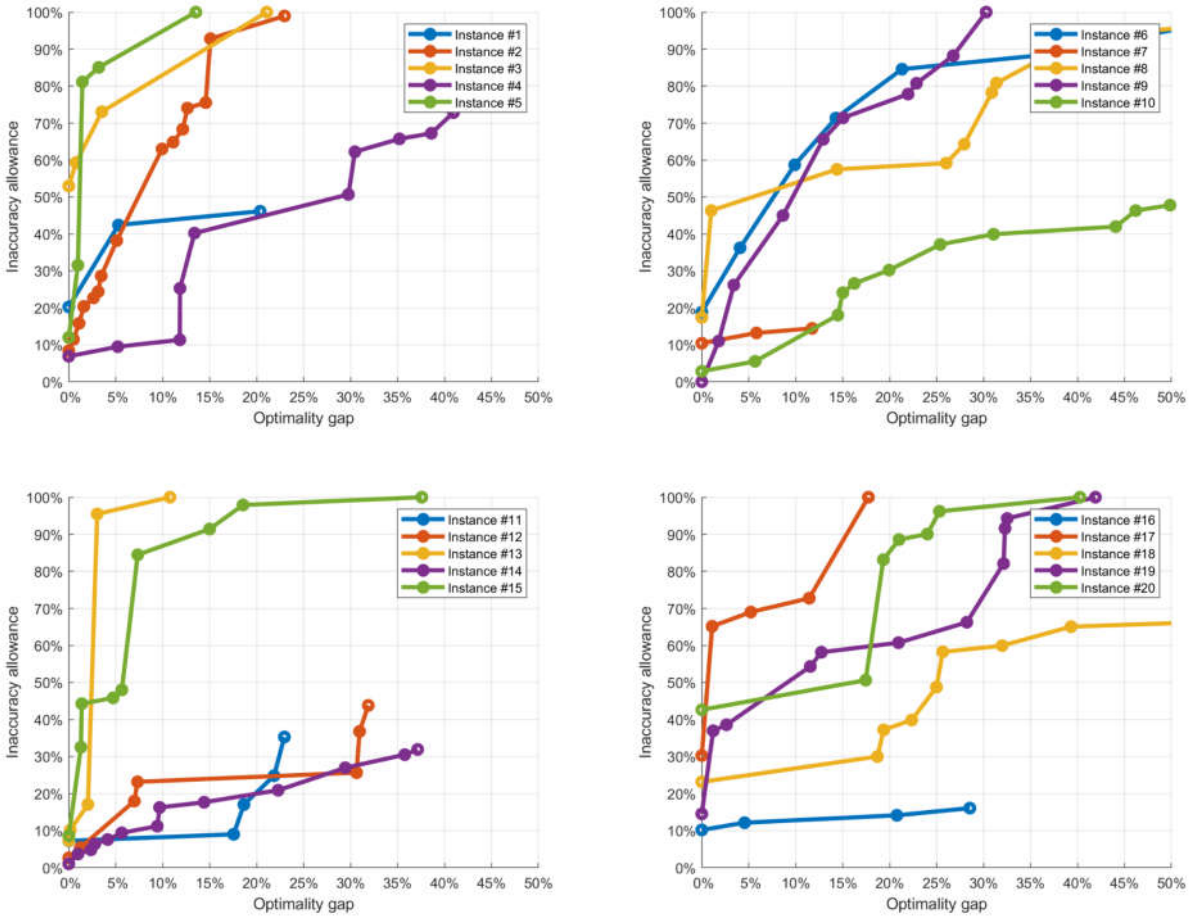


Figure 15 Pareto-optimal solutions on the efficient frontier of test instances with en-route charging

As expected, we observe a larger set of Pareto-optimal solutions on the efficient frontiers compared to the previous cases without en-route recharging. Solutions show to be more resistant against increases in optimality gap, and optimal solutions provide only

up to 15% inaccuracy allowance on average. While the second-best and third-best solutions provide some wiggle room (26% and 37%, on average), these would yield optimality gaps of up to 4.8% and 9.5% on average, respectively. This is mainly due to the more binding battery charge level constraint (i.e., 30 kWh) and the additional impact of time-windows constraints on the tolerance of solutions against inaccuracy which was earlier discussed.

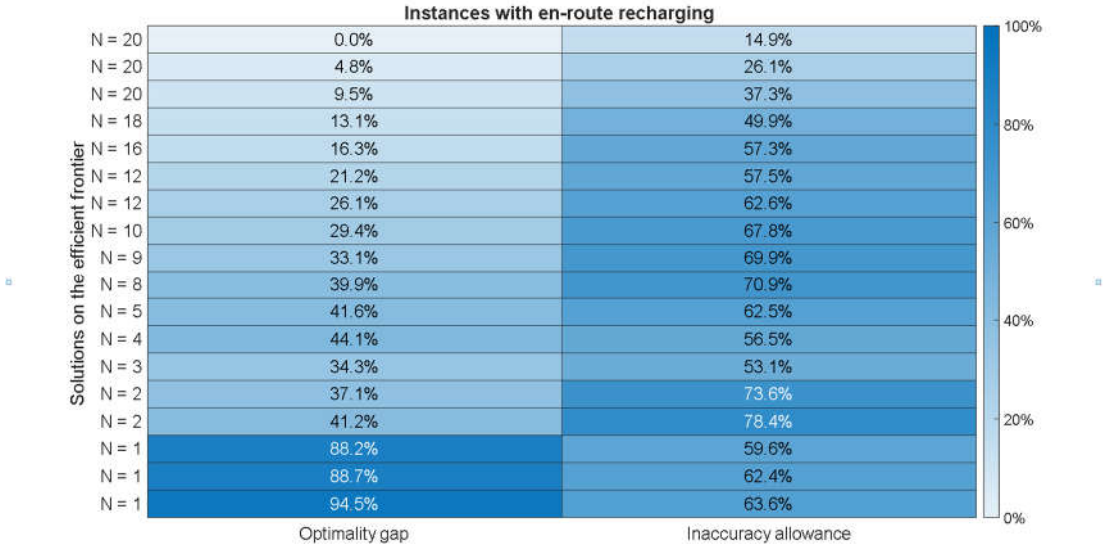


Figure 16 Heat map summarising the efficient frontier of test instances with en-route recharging

To visualise the structural differences of the solutions on the efficient frontier, we have illustrated the top three Pareto-optimal solutions (in terms of objective values) on the efficient frontier of instance 5 in Figure 17. Interestingly, the first two solutions in this figure indicate that a small change in the solution structure can lead to a significant jump in inaccuracy allowance (around 20%) while incurring very minimal increase in total cost (less than 1%). One of the routes in the instance-optimal solution is repeated on the second-best solution, and the other one has the exact same order of customer visits with a difference in locating the RS visit in the route.

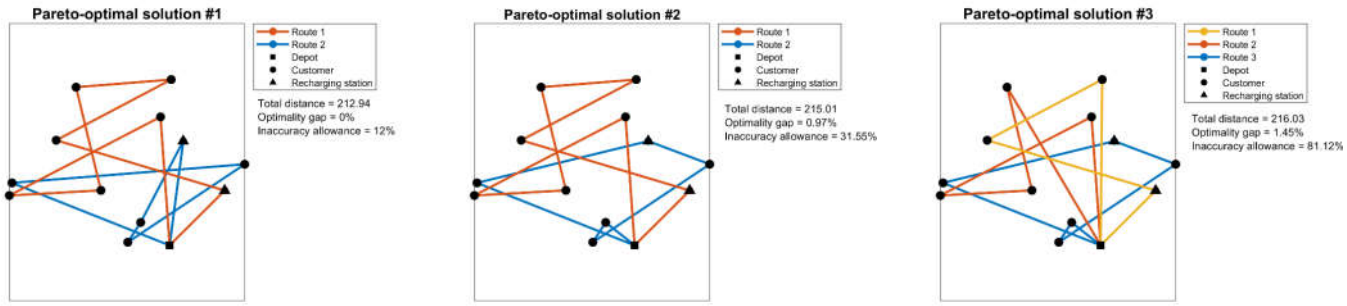


Figure 17 Designed ECV routes within the top three Pareto-optimal solutions on the efficient frontier of instance 5

As a final point, it must be noted that given the scope of the available data for this study, the instances used for the analysis of the case of the EVRPTW with en-route recharging have been generated within an urban setting by reducing the battery pack size. However, as we discussed earlier and backed by the results presented in Sections 3.5 and 4.2, within such context status quo of ECV battery technology and pack size provides sufficient range and flexibility for uninterrupted daily deliveries while relying solely on overnight recharging at the depot. The case of en-route recharging is more suited for extra-urban deliveries using ECVs, and the driving profile of such context can differ significantly (more motorway journeys and less city type stop-and-go profiles). This is particularly important and an interesting line of further research as there is evidence that unlike internal combustion engine vehicles, EVs thrive in stop–start conditions thanks to low-speed efficiency, regenerative braking and minimal idling losses. On motorways, however, aerodynamic drag grows rapidly at higher speeds, which has a large effect on energy consumption in EVs. In addition to this, the presented results on the case of the EVRPTW with en-route recharging and the arguments discussed in Section 4.1.1 hold in the case of the considered simplified charging environment with deterministic charger availability and charging behaviour. Factors such as charger congestion, nonlinear charging behaviour, charger availability uncertainty, battery degradation, and operational constraints associated with public charging infrastructure may influence routing decisions and could affect the extent to which the observed findings generalise to other settings.

5. Concluding remarks

Accurately estimating the actual on-road energy consumption of an ECV is a complex and difficult task. The added value in incorporating a detailed and intricate energy estimation module into routing models is not yet fully uncovered and the vast literature in the area is prevalently built on the premise that inaccuracy in energy consumption estimation leads to suboptimal or infeasible routing decisions. We have examined the real impact of energy estimation inaccuracy on routing decisions and proposed an intuitive but practical methodology to mitigate any such impact. A reliable benchmark method of energy estimation capable of generating close estimations of actual on-road energy consumption was employed, and a comprehensive comparative analysis through the application of nine alternative methods of estimation across 2,450 driving paths was constructed. The impact of ignoring various factors (such as vehicle load, speed profile and A/D rates, road slopes, weather conditions, regenerative braking, asphalt conditions, etc.) on the level of energy estimation inaccuracy and the optimality of routing decisions was examined. Finally, we have proposed a bi-objective ε-constraint based methodology to tackle the impact of estimation inaccuracy that is able to generate all Pareto-optimal solutions on the efficient frontier of the problem where each solution lies within a given factor of the instance-optimal solution and is robust against a given level of estimation inaccuracy.

Our results yield several useful technical and managerial insights:

- We observed that ignoring variable instantaneous speed profiles and A/D rates (which are widely used in most of the extant literature) can lead to significant estimation inaccuracies. Methods based on using an average speed across the network, and those that use a fixed speed value over each arc and ignore A/D rates, generate the worst estimations. This result is in line with findings from several relevant studies (Turkensteen, 2017; Kancharla & Ramadurai, 2018b; Raeesi & Zografos, 2019).
- Conversely, the impact of ignoring variables such as road slope, weather conditions and variable rolling resistance based on asphalt quality (whether disregarded individually or collectively) is insignificant, and methods that ignore these variables but incorporate variable speed profiles and A/D rates still yield very close estimations compared to those of the benchmark model. As such, we argue that given the difficulty of collecting data for these variables and the added complexity they introduce into routing models, their contribution appears to be of limited value in high-level applications such as routing.
- As concerns the main research question of the paper, we observe that overall, the impact of energy estimation inaccuracy on the feasibility and even optimality of routing decisions is quite limited. Across all the 540 tests conducted, in only two cases no feasible solution was found, and this occurred for the significantly overestimating method that uses an average driving speed of 40 km/h across the road network. In 16 cases, the identified solution is expected to be on-road infeasible (assuming the benchmark method generates figures that are closely aligned with the actual on-road figures) due to the incorporation of a significantly underestimating method. In all other cases, the average optimality gap remains negligible (nearly zero for the case of most methods, and 9% on average for the worst method with an average inaccuracy of 63%).
- We discussed various properties of overestimating versus underestimating methods of energy consumption prediction, and despite the prevalent tendency towards overestimating methods to remain conservative and risk-averse, we argue that incorporation of an underestimating method might be beneficial in ensuring on-road feasibility and hence on-road-optimality of routing solutions.
- We discussed results that help implicitly ensure the on-road optimality of a routing decision without relying on a complex energy estimation method.
- We showed that by just including the simplest energy consumption method of all (i.e., the basic consumption rate reported by the vehicle manufacturer) within the proposed bi-objective ε-constraint based methodology and generating the full efficient frontier of the Pareto-optimal solutions, it is possible to obtain optimal or near-optimal solutions that are proven to be significantly robust against a high level of estimation inaccuracy.
- Our analysis of the efficient frontier of the considered instances indicates that an instance-optimal solution is very likely to be an on-road-optimal solution; however, with a marginal sacrifice in instance optimality, a significant allowance for inaccuracy is achievable, and second-based solutions on the efficient frontiers overall seem to constitute an excellent option for a strictly risk-averse decision maker.

It must be mentioned again that some of the above conclusions (especially those on the impact of the factors on accuracy) are contingent on the benchmark model being a close proxy for reality. The on-road optimality discussions, however, are less affected as they are deduced by applying a methodology that does not make any assumption about the actual on-road energy consumption. While we may have inevitably used the benchmark model to generate a proxy for the actual on-road optimal solution (mainly for comparison purposes), any inherent inaccuracy within the model will not have a significant impact on these results.

All in all, we find that the inclusion of a highly detailed energy estimation model in routing optimisation yields only marginal benefits. In practice, a simplified approach relying on the manufacturer's reported consumption rate, which is readily available, is generally sufficient. Since there is already substantial empirical and academic evidence for the inaccuracy levels of such rates across a large variety of operating conditions and for many different types of EVs (see Appendix C), we believe including the basicRate in the

proposed methodology has the added advantage of addressing cases with heterogeneous fleets and extreme operating conditions, too. This is because these rates are universally available for virtually all EV models, and their underestimation error in different operating and weather conditions is well-studied and roughly known. As such, for the case of the EVRPTW with heterogeneous fleets, vehicle-specific consumption rates could be used without affecting the validity or structure of the proposed methodology. Also, as the proposed method generates a pool of solutions with different inaccuracy allowances rather than a single optimal solution, solutions can be simply filtered in extreme operating conditions based on the expected error scale. As a final point, we contend that, given the continuous improvements in ECV battery capacity, the importance of precise energy modelling is likely to diminish, particularly within the context of urban last-mile delivery.

Acknowledgement

The authors are grateful to the reviewers for their very constructive comments that have helped us to improve the presentation of this paper.

Appendices

Appendix A. Energy consumption as a linear function of vehicle mass

To investigate the dependence of energy consumption on vehicle payload when all other parameters remain unchanged, the benchmark model was executed across all the 2,450 paths generated in Greater London for payload levels ranging from 0 to 1000 kg in 20 kg increments (51 load levels per path). For each simulation, battery energy discharge, regenerated energy through braking, and net energy consumption were recorded. In total, 124,950 simulations were performed. For each path, three different linear regression models were fitted with payload as the independent variable, and battery energy discharge, regenerative braking and total energy consumption as the dependent variables, respectively, and the resulting R^2, adjusted R^2 and RMSE values of each regression model were calculated. Table A.1 groups all paths into 11 categories with increasing distances by increments of 10 km, and for each group shows the average for these values. The results clearly demonstrate a consistently high goodness-of-fit across all paths. In nearly all cases, R^2 and adjusted R^2 values approach unity, indicating that a first-order linear model explains almost all variance in energy consumption across the 0–1000 kg load range. The similarity between R^2 and adjusted R^2 further confirms that the linear specification is not overfitted and that higher-order terms would provide negligible explanatory improvement. RMSE values remain small relative to the magnitude of total path energy consumption, implying that residual deviations from linearity are practically insignificant.

Table A.1
Regression analysis results for the linearity between energy consumption and vehicle mass

	Discharge			Regenerative braking			Energy consumption		
Category (≤km)	R^2	Adjusted R^2	RMSE	R^2	Adjusted R^2	RMSE	R^2	Adjusted R^2	RMSE
10	0.99998	0.99998	0.00074	0.99241	0.99225	0.00117	0.99997	0.99997	0.00081
20	0.99998	0.99998	0.00177	0.99309	0.99295	0.00259	0.99998	0.99998	0.00165
30	0.99998	0.99998	0.00284	0.99379	0.99366	0.00382	0.99998	0.99998	0.00218
40	0.99998	0.99998	0.00383	0.99502	0.99492	0.00498	0.99999	0.99999	0.00271
50	0.99998	0.99998	0.00464	0.99560	0.99551	0.00626	0.99999	0.99999	0.00324
60	0.99997	0.99997	0.00494	0.99686	0.99679	0.00653	0.99998	0.99998	0.00431
70	0.99997	0.99997	0.00430	0.99850	0.99847	0.00453	0.99996	0.99996	0.00516
80	0.99998	0.99998	0.00410	0.99855	0.99853	0.00516	0.99997	0.99997	0.00496
90	0.99997	0.99997	0.00467	0.99866	0.99864	0.00534	0.99996	0.99996	0.00668
100	0.99996	0.99996	0.00594	0.99908	0.99906	0.00441	0.99995	0.99995	0.00705
110	0.99996	0.99996	0.00588	0.99912	0.99911	0.00456	0.99993	0.99993	0.00887
All	0.99998	0.99998	0.00343	0.99510	0.99500	0.00439	0.99998	0.99998	0.00302

The observed linearity can be further explained by the underlying longitudinal power requirement model presented in Expression (1) in Section 3.1. Within the model, inertial force ($m_{eff}\frac{dv}{dt}$), rolling resistance ($f_r mg\cos(\alpha)$), and gravitational force ($mg\sin(\alpha)$) components scale directly with vehicle mass (m). For a fixed velocity profile, fixed road gradient, and unchanged vehicle parameters, these components are strictly linear in mass, and since path energy consumption is obtained by integrating power over time, and knowing that integration preserves linearity, the slope of the linear equation captures mass-dependent inertial, rolling, and grade effects, and the intercept represents mass-independent contributions, primarily aerodynamic drag and baseline drivetrain losses. Although drivetrain losses $P_{loss}(\omega_{wheel}, T_{wheel})$ introduce nonlinear dependence on torque, and aerodynamic drag is independent of mass, these effects do not structurally invalidate an affine-linear relationship between energy and payload

over moderate load ranges. Therefore, from first principles, near-linearity in payload is theoretically expected when all other factors remain unchanged. Regenerated energy, on the other hand, is governed primarily by recoverable kinetic and potential energy, where both terms scale linearly with mass. While regenerative efficiency is not perfectly constant and may introduce mild nonlinearities, the regression results confirm that regenerated energy remains highly linear in payload across all paths. Consequently, the net energy consumption—defined as discharge plus regenerated energy—also exhibits near-linear behaviour, as it is the sum of two approximately linear functions.

Therefore, the strong empirical linearity observed across 124,950 simulations is consistent with the structural properties of the longitudinal vehicle dynamics model. While exact linearity cannot be claimed due to torque-dependent drivetrain efficiencies and regenerative system characteristics, deviations are statistically negligible within the investigated load range (0–1000 kg).

Appendix B. Diversity of the generated driving cycles

The dataset of the instantaneous driving cycles contains 2,450 paths generated within Greater London. For each path, we have calculated summary features including trip length and duration, driving dynamics (speed and acceleration statistics), terrain (elevation gain and gradients), wind-related conditions, and spatial extent (latitude/longitude range) to make sure the selected mix provides sufficient variability across these dimensions, and the required evaluation is performed under realistic conditions. Figure B.1 illustrates the spread of the paths over Greater London's spatial footprint.



Figure B.1 The coverage of the generated paths on the map of the Greater London

Descriptive statistics of the paths' features including the mean and standard deviation of their distance, duration, average speed, acceleration and gradient is presented in Table B.1 by grouping all paths into 11 categories with increasing distances by increments of 10 km.

Table B.1
Summary statistics of some of the paths' features

Category (≤km)	η	μ_d	σ_d	μ_t	σ_t	μ_s	σ_s	σ_a	max_g
10	183	7.01	2.32	12.58	5.44	10.04	4.58	0.63	0.021
20	402	15.12	4.66	26.70	8.30	9.89	5.30	0.66	0.026
30	442	25.09	5.74	43.63	12.68	10.35	5.93	0.66	0.029
40	471	34.61	5.58	57.75	16.09	10.78	6.40	0.66	0.033
50	373	44.51	5.77	69.50	17.21	11.51	6.61	0.64	0.036
60	190	53.30	5.08	66.48	19.31	14.69	7.64	0.58	0.037
70	99	64.17	5.89	54.74	12.32	20.34	8.87	0.46	0.042
80	128	75.28	6.20	61.12	6.75	20.74	9.39	0.46	0.043
90	72	84.91	5.36	67.99	8.70	21.13	9.12	0.44	0.047
100	63	94.12	5.36	67.53	5.16	23.34	8.27	0.37	0.045
110	27	101.36	4.13	69.08	5.25	24.59	7.37	0.33	0.053
All	2450	37.65	22.65	49.93	22.31	12.60	6.52	0.61	0.033

η: Number of paths, μ_d: mean distance (km), σ_d: Standard deviation of distance (km), μ_t: mean duration (min), σ_d: Standard deviation of duration (min), μ_s: mean speed (km/h), σ_s: Standard deviation of speed (km/h), σ_a: Standard deviation of acceleration (m/s^2), max_g: maximum gradient

As regards the variability of trip lengths (Figure B.2.A) and durations (Figure B.2.B), the dataset spans a broad range of trips, scaling from only 1.17 km all the way to 103.82 km, and from 1.48 min to 108.53 min in duration. As such the dataset includes very short trips (short urban hops), mid-length intra-city journeys, and long cross-city or orbital trips. Importantly, the middle 90% of trips still covers an order-of-magnitude spread, indicating variability is not driven only by a handful of outliers. Driving dynamics also show considerable diversity consistent with low-speed regimes (congested/urban), intermediate regimes (mixed traffic), and higher-speed regimes, and we observe a mean speed range of 4.9 km/h to 27.32 km/h (Figure B.2.C). The spread in the standard deviation of speed (Figure B.2.D) and acceleration (Figure B.2.F) provides evidence that the set includes both smoother-flowing cycles and more stop-go/aggressive cycles. It is worth mentioning that the acceleration mean (Figure B.2.E) of nearly zero for many cycles is not surprising and rather typical for a full trip where accelerations and decelerations balance. As a better discriminator of "aggressiveness" of the cycles, the standard deviation of acceleration is also presented (Figure B.2.F) which illustrates significant variability across the cycles.

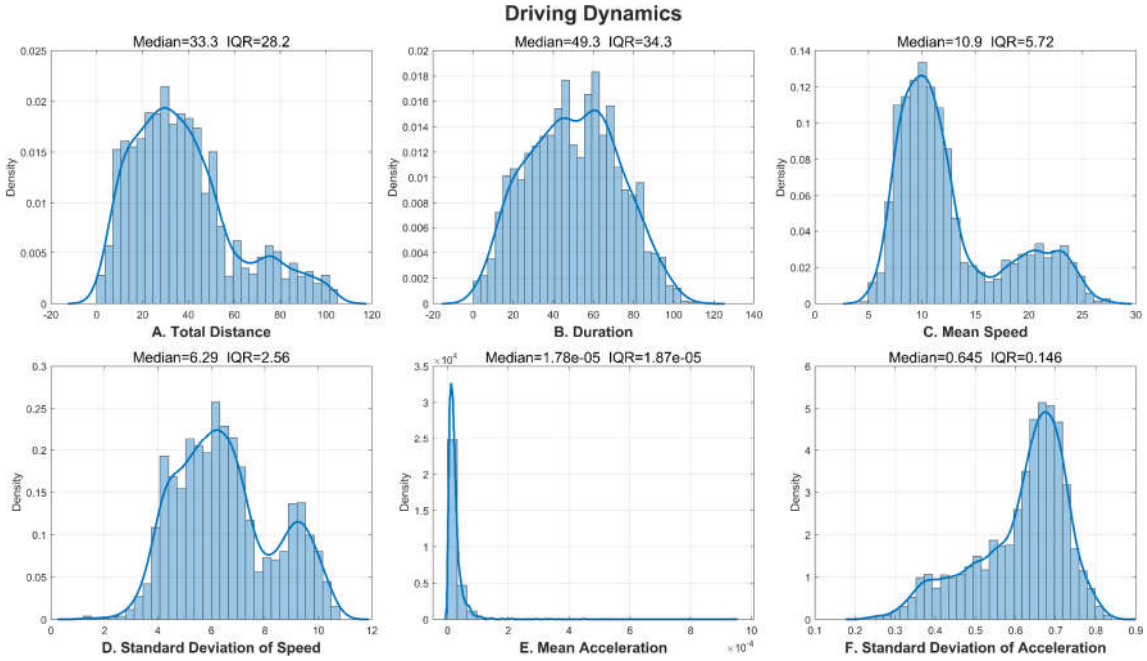


Figure B.2 Analysis of the driving dynamics of the generated paths

Terrain variability (elevation and gradients), on the other hand, is strong and not confined to a narrow band. The range of elevation gain (Figure B.3) is from 0 to 680.81, which suggests the dataset includes paths with essentially no climb as well as paths with substantial cumulative ascent, plus meaningful variation in steepness.

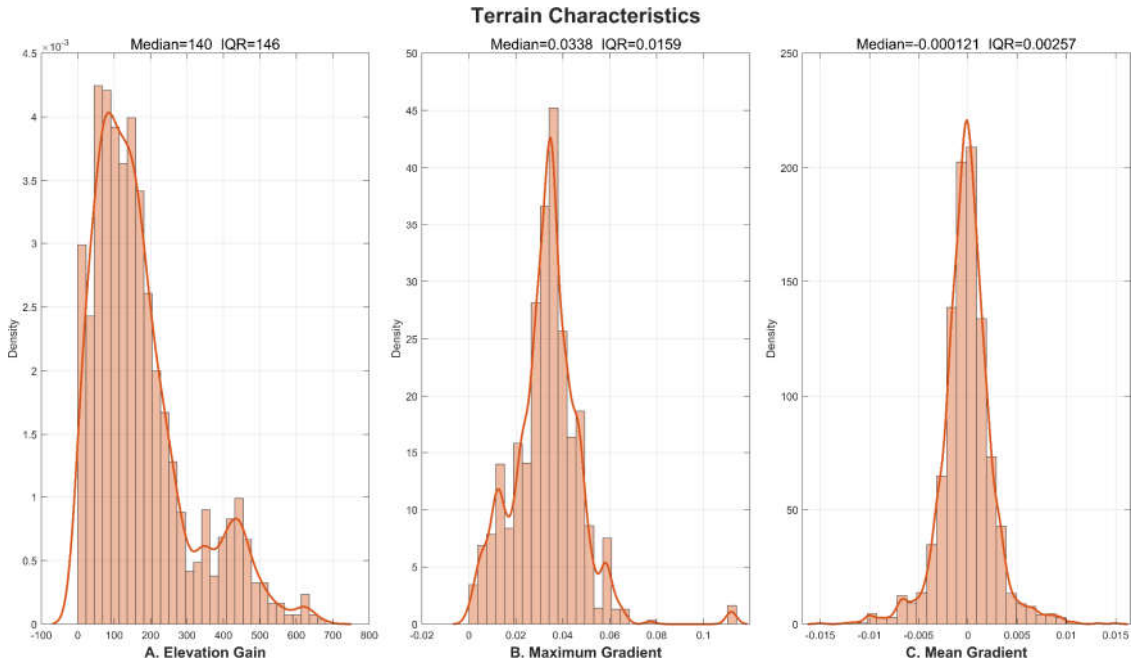


Figure B.3 Analysis of the terrain characteristics of the generated paths

Finally, weather and wind-related features show a wide spread in magnitude and directional variability (Figure B.4). The mean of wind strength spans a range from 4.67 m/s^2 to 27.34 m/s^2 and the wind angle standard deviation ranges from nearly 0 degree to 48.98 degree.

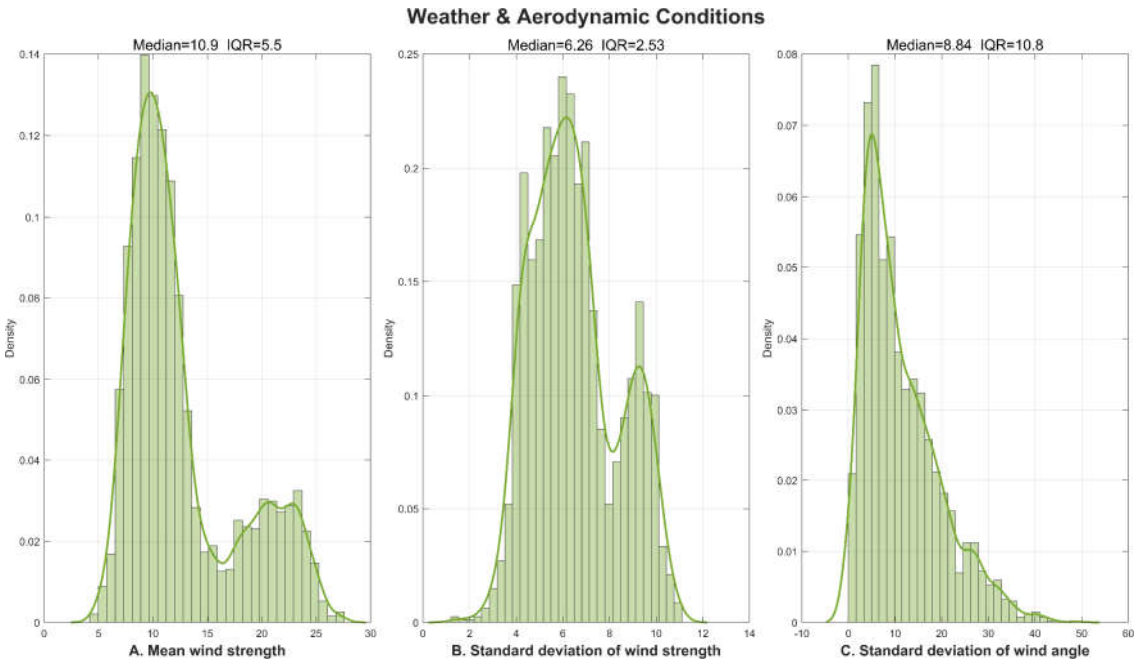


Figure B.4 Analysis of the weather and aerodynamic conditions over the generated paths

The above analysis examines variability of one feature at a time (distance, terrain, wind, etc.), but to further demonstrate true multidimensional variability, we have also conducted a Principal Component Analysis (PCA) on the standardised path-level feature set to evaluate multivariate diversity (Figure B.5). The first principal component explains 50.7% of the total variance, while the first five components explains 85.4% and the first eight components explains 97.1%. This indicates that the dataset spans multiple independent dimensions of variability rather than clustering around a narrow set of operating conditions.

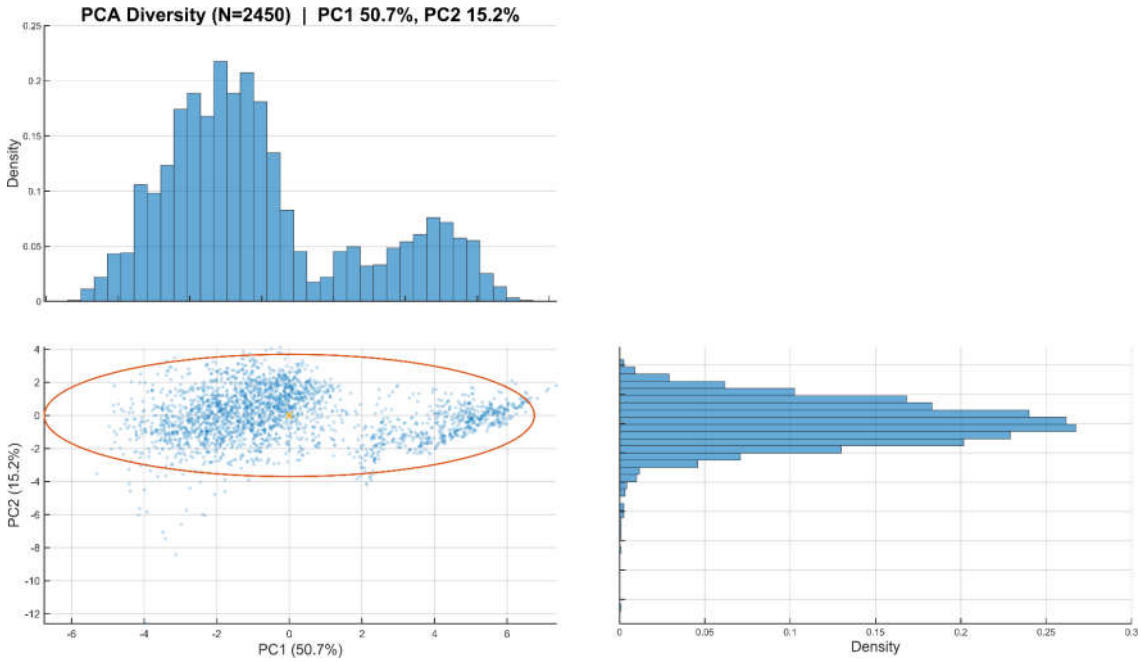


Figure B.5 PCA on the standardised path-level feature set to evaluate multivariate diversity

All in all, the presented analysis confirms that the 2,450 selected paths provide broad coverage of terrain, driving conditions, and environmental exposure within Greater London.

Appendix C. The estimation behaviour of the basicRate model

The proposed methodology in this paper relies on the incorporation of an underestimating model. In this section we show that the officially reported energy consumption of ECVs is underestimating the actual on-road energy consumption in most cases. For the cases of unexpected overestimation, we discuss the route-level impacts and a practical way to address them.

There is already substantial empirical evidence showing that the officially reported EV energy consumption (or range) often underestimates the actual energy demand observed in real-world driving. However, the magnitude of the gap is typically smaller than what was historically observed for internal combustion vehicles, and it varies with driving conditions and test protocols. Several studies comparing certification test cycles (e.g., WLTP, EPA) with real-world data have found systematic underestimation. Weiss et al. (2024) analyse European EV data and find that real-world energy consumption is significantly higher than certified values and the discrepancy is systematic across models. Van Gijlswijk et al. (2020) compare laboratory ratings with real-world usage and report that real-world EV energy consumption is 1–20% higher than lab-rated values, leading to an average 16% reduction in effective range relative to official figures. Evidence from independent real-world testing programs also consistently confirm this gap. For example, the Australian Automobile Association[6] testing found 5–23% lower range than laboratory claims for several EV models, and as range is inversely proportional to energy consumption, this implies that actual energy use per km is correspondingly higher. This discrepancy arises mainly from differences between laboratory test cycles and real driving conditions, as lab cycles occur under controlled conditions ignoring real-world driving factors such as weather, wind, terrain and driving behaviour.

In line with the above, we find very similar results when the benchmark model is used as the high fidelity point of reference. Based on the results of 26,950 comparisons (2,450 paths travelled with 11 different vehicle payload levels ranging from 0 to 1000 kg in 100 kg increments), we observe an average underestimation of 16.5%. Figure C.1 presents a direct comparison between the basicRate and the benchmark models showing clearly that in 92.3% of the cases the basicRate model is underestimating the predicted value by the benchmark model. In this figure, each point is an estimation by the basicRate model against the benchmark, and anything below the 45° line indicates $basicRate < benchmark$. The heavy saturation of the points under the line hence indicates the strong and statistically robust underestimating behaviour of the basicRate model, and it is clear that overestimation happens at very small energy consumption values and very close to the line. As such, it can be confidently argued that at route-level for EVRPTW applications, any potential overestimation is highly likely to be offset by other sources of dominant underestimations. It is also worth mentioning that the increase in underestimation gap in values over 25 kWh is due to the impact of vehicle load which is completely ignored by the basicRate model.

[6] https://realworld.org.au/

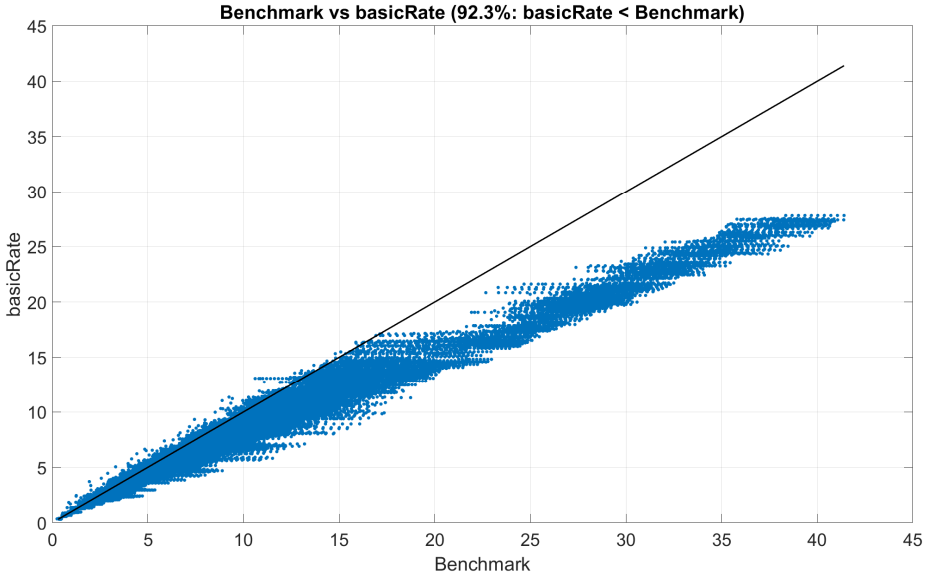


Figure C.1 Estimation behaviour of the basicRate model vs the benchmark

Figure C.2 further highlights the impact of vehicle load on the estimation behaviour of the basicRate model. As can be seen, with the increase in vehicle load the percentage of underestimation cases also increases and gets closer to 100%. Therefore, overestimation by the basicRate is more likely to happen when the vehicle runs nearly empty. At the EVRPTW routing level, the vehicle runs empty when it returns to the depot at the end of a route which represents only k arcs of the total $n + k - 1$ arcs in the solution (where k is the number of routes in the solution and n is the number of customer nodes). For example, within the optimal solution of a given EVRPTW with 20 customers and 3 routes, there are only 3 arcs out of the total 22 arcs where the vehicle runs empty. This again makes it unlikely for the overall route-level energy demand to be overestimated when the basicRate model is used.

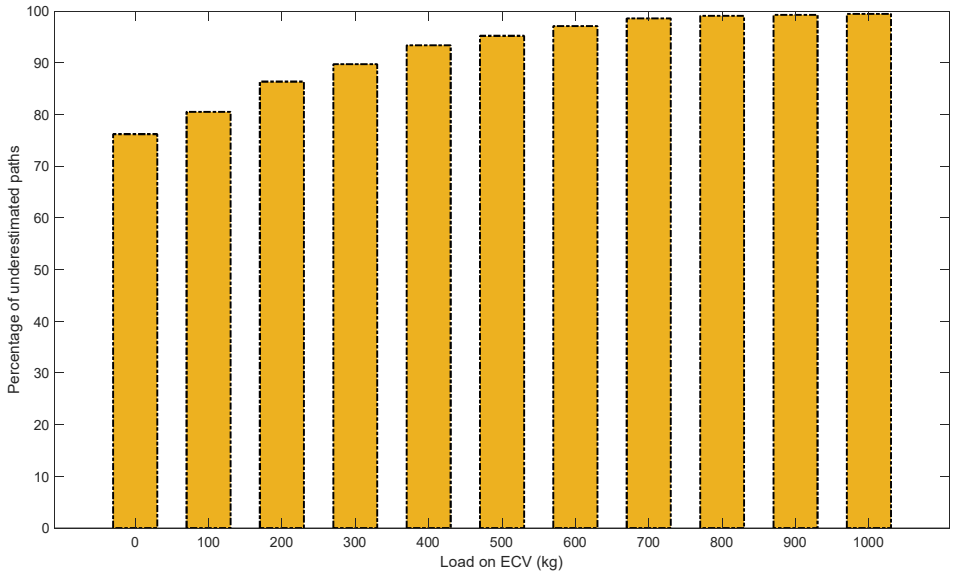


Figure C.2 Impact of vehicle load on the estimation behaviour of the basicRate model

Another set of analyses is presented in Figure C.3 to show the distribution of the difference $\Delta = benchmark - basicRate$. It is clear that the distribution is positively skewed, with the majority of values being strictly greater than zero. This confirms that underestimation is not only frequent but dominates the distribution. The magnitude of underestimation varies, with a long right tail, indicating that in some cases the discrepancy is substantial. It is also interesting to see that the existing few cases of negative Δ are mostly saturated around zero, suggesting that even if basicRate returns an overestimation in some cases, the difference is very small; another reason to expect the unlikely cases of overestimation will be compensated by more strongly underestimating cases within the overall route-level estimation outcome.

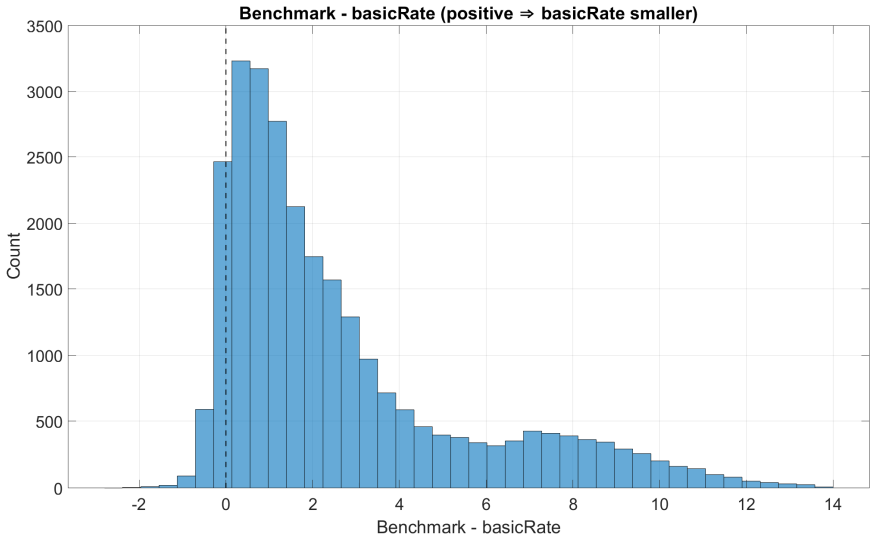


Figure C.3 Histogram of the basicRate under(over)estimation values vs the benchmark

Finally, Figure C.4 provides a comparison of mean over(under)estimation values at different vehicle payloads. The most striking feature of this figure is the negligible size of overestimation against underestimation in each case. Moreover, unlike the mean overestimation error, the mean underestimation error tends to increase as the vehicle payload increases.

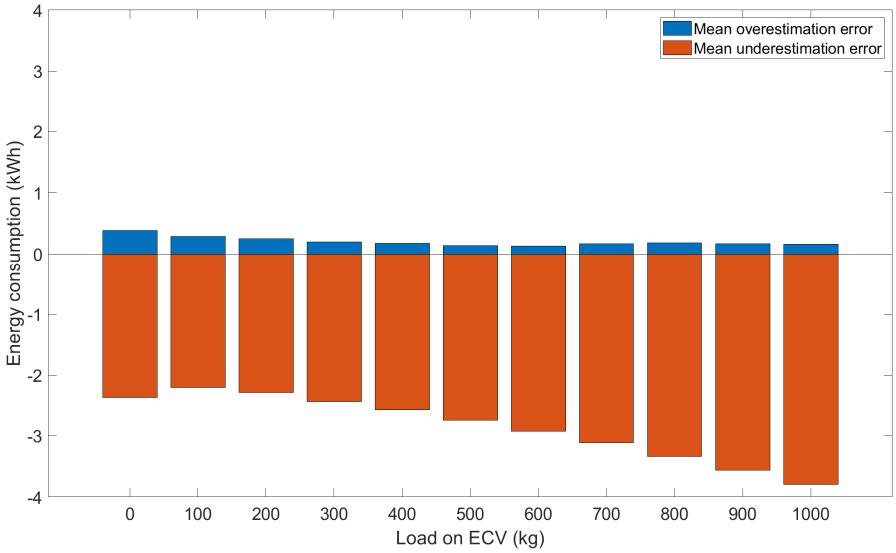


Figure C.4 Mean over(under)estimation error at different payload levels

All in all, overestimation by basicRate model is expected to be rare, and in most cases with no route-level impacts in the case of the EVRPTW. However, we present a theoretical worst-case in Remark C.1 and propose a practical mitigation approach next:

Remark C.1. *The worst-case scenario for the estimation behaviour of the basicRate model would be when the model is over-estimating energy consumption across all paths in the network despite the expected underestimating behaviour and as per Remark 1 for the worst-case scenario of overestimation, the efficient frontier is empty (i.e., no instance-feasible solution is generated).*

While this is a quite unlikely situation, we suggest that an empty efficient frontier must be dealt with as a potential sign of unexpected overestimation (assuming the problem at hand is on-road-feasible and there is at least one on-road-feasible solution for the problem). We suggest that a practical way to tackle such cases is to apply a small discount rate on all estimated energy consumption values for all paths in the network, re-solve the problem, and pick the instance-optimal solution regardless of its inaccuracy allowance. This is because an empty efficient frontier is essentially a sign of overestimation when we know the problem is on-road-feasible. If the selected discount rate is sufficiently small, based on Corollary 1, we know an emerged solution is guaranteed to be on-road feasible (and while not guaranteed, probably on-road (near)optimal). If no instance-feasible solution is found after applying the discount rate, this must be further increased, and the model must be re-solved. This procedure should continue until an instance-optimal solution is found.

References

Abro, M. A. J., Larik, A. S., Unar, M. A., Halepoto, I. A., Rajab, A., Al Reshan, M. S., Shaikh, A., & Alalhareth, M. (2026). Machine Learning Analysis of Vehicle Mass Influence on EV Energy Consumption Using Black Box Energy Consumption Model. *IEEE Access*, *14*, 12984–97.

Asamer, J., Graser, A., Heilmann, B., & Ruthmair, M. (2016). Sensitivity analysis for energy demand estimation of electric vehicles. *Transportation Research Part D: Transport and Environment*, *46*, 182–99.

Barth, M., & Boriboonsomsin, K. (2009). Energy and emissions impacts of a freeway-based dynamic eco-driving system. *Transportation Research Part D: Transport and Environment*, *14*(6), 400–10.

Barth, M., Younglove, T., & Scora, G. (2005). Development of a heavy-duty diesel modal emissions and fuel consumption model.

Basso, R., Kulcsár, B., Egardt, B., Lindroth, P., & Sanchez-Diaz, I. (2019). Energy consumption estimation integrated into the Electric Vehicle Routing Problem. *Transportation Research Part D: Transport and Environment*, *69*, 141–67.

Basso, R., Kulcsár, B., & Sanchez-Diaz, I. (2021). Electric vehicle routing problem with machine learning for energy prediction. *Transportation Research Part B: Methodological*, *145*, 24–55.

Basso, R., Kulcsár, B., Sanchez-Diaz, I., & Qu, X. (2022). Dynamic stochastic electric vehicle routing with safe reinforcement learning. *Transportation Research Part E: Logistics and Transportation Review*, *157*, 102496.

Beckers, C. J. J., Geraedts, T. A. G. H., Besselink, I. J. M., & Nijmeijer, H. (2021). *TU/e Microscopic Energy Consumption PRediction tOol 0.1 (TU/e MECPRO 0.1)* https://doi.org/10.4121/uuid:721b2ea6-2634-4dd5-a0c8-865a0aa41a99

Beckers, C. J. J., Paroha, M., Besselink, I. J. M., & Nijmeijer, H. (2020). *A microscopic energy consumption prediction tool for fully electric delivery vans*. 33rd Electric Vehicle Symposium (EVS33), Portland, Oregon, June 14-17.

Bektaş, T., Ehmke, J. F., Psaraftis, H. N., & Puchinger, J. (2019). The role of operational research in green freight transportation. *European Journal of Operational Research*, *274*(3), 807–23.

Bektaş, T., & Laporte, G. (2011). The pollution-routing problem. *Transportation Research Part B: Methodological*, *45*(8), 1232–50.

Bruglieri, M., Paolucci, M., & Pisacane, O. (2023). A matheuristic for the electric vehicle routing problem with time windows and a realistic energy consumption model. *Computers & Operations Research*, *157*, 106261.

Bruglieri, M., Pezzella, F., Pisacane, O., & Suraci, S. (2015). A variable neighborhood search branching for the electric vehicle routing problem with time windows. *Electronic Notes in Discrete Mathematics*, *47*, 221–8.

Chen, Y., Zhu, L., Gonder, J., Young, S., & Walkowicz, K. (2017). Data-driven fuel consumption estimation: A multivariate adaptive regression spline approach. *Transportation Research Part C: Emerging Technologies*, *83*, 134-45.

Dabia, S., Demir, E., & Van Woensel, T. (2017). An exact approach for a variant of the pollution-routing problem. *Transportation Science*, *51*(2), 607–28.

Demir, E., Bektaş, T., & Laporte, G. (2011). A comparative analysis of several vehicle emission models for road freight transportation. *Transportation Research Part D: Transport and Environment*, *16*(5), 347–57.

Demir, E., Bektaş, T., & Laporte, G. (2014a). The bi-objective Pollution-Routing Problem. *European Journal of Operational Research*, *232*(3), 464-78.

Demir, E., Bektaş, T., & Laporte, G. (2014b). A review of recent research on green road freight transportation. *European Journal of Operational Research*, *237*(3), 775–93.

Demir, E., Bektaş, T., & Laporte, G. (2014c). A review of recent research on green road freight transportation. *European Journal of Operational Research*, *237*(3), 775-93.

Desaulniers, G., Errico, F., Irnich, S., & Schneider, M. (2016). Exact algorithms for electric vehicle-routing problems with time windows. *Operations Research*, *64*(6), 1388–405.

Donkers, A., Yang, D., & Viktorović, M. (2020). Influence of driving style, infrastructure, weather and traffic on electric vehicle performance. *Transportation Research Part D: Transport and Environment*, *88*, 102569.

Erdoğan, S., & Miller-Hooks, E. (2012). A Green Vehicle Routing Problem. *Transportation Research Part E: Logistics and Transportation Review*, *48*(1), 100-14.

Eshtehadi, R., Fathian, M., & Demir, E. (2017). Robust solutions to the pollution-routing problem with demand and travel time uncertainty. *Transportation Research Part D: Transport and Environment*, *51*, 351–63.

Felipe, Á., Ortuño, M. T., Righini, G., & Tirado, G. (2014). A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges. *Transportation Research Part E: Logistics and Transportation Review*, *71*, 111–28.

Figliozzi, M. (2010). Vehicle routing problem for emissions minimization. *Transportation Research Record*, *2197*, 1–7.

Figliozzi, M. A. (2011). The impacts of congestion on time-definitive urban freight distribution networks CO2 emission levels: Results from a case study in Portland, Oregon. *Transportation Research Part C: Emerging Technologies*, *19*(5), 766–78.

Franceschetti, A., Honhon, D., Van Woensel, T., Bektaş, T., & Laporte, G. (2013). The time-dependent pollution-routing problem. *Transportation Research Part B: Methodological*, *56*, 265-93.

Froger, A., Mendoza, J. E., Jabali, O., & Laporte, G. (2019). Improved formulations and algorithmic components for the electric vehicle routing problem with nonlinear charging functions. *Computers & Operations Research*, *104*, 256–94.

Garside, A. K., Ahmad, R., & Muhtazaruddin, M. N. B. (2024). A recent review of solution approaches for green vehicle routing problem and its variants. *Operations Research Perspectives*, *12*, 100303.

Goeke, D., & Schneider, M. (2015). Routing a mixed fleet of electric and conventional vehicles. *European Journal of Operational Research*, *245*(1), 81–99.

Han, B., Yan, Y., Chi, T., & Park, Y. (2025). The electric vehicle routing problem with travel time and energy consumption uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, *202*, 104211.

Iora, P., & Tribioli, L. (2019). Effect of ambient temperature on electric vehicles' energy consumption and range: Model definition and sensitivity analysis based on nissan leaf data. *World Electric Vehicle Journal*, *10*(1), 2.

Jeong, J., Ghaddar, B., Zufferey, N., & Nathwani, J. (2024). Adaptive robust electric vehicle routing under energy consumption uncertainty. *Transportation Research Part C: Emerging Technologies*, *160*, 104529.

Kancharla, S. R., & Ramadurai, G. (2018a). Incorporating driving cycle based fuel consumption estimation in green vehicle routing problems. *Sustainable Cities and Society*, *40*, 214–21.

Kancharla, S. R., & Ramadurai, G. (2018b). Incorporating driving cycle based fuel consumption estimation in green vehicle routing problems. *Sustainable Cities and Society*, *40*, 214-21.

Kara, I., Kara, B. Y., & Kadri Yetis, M. (2007). Energy minimizing vehicle routing problem. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics),

Keskin, M., & Çatay, B. (2016). Partial recharge strategies for the electric vehicle routing problem with time windows. *Transportation Research Part C: Emerging Technologies*, *65*, 111–27.

Kim, Y. J., & Chung, B. D. (2023). Energy consumption optimization for the electric vehicle routing problem with state-of-charge-dependent discharging rates. *Journal of Cleaner Production*, *385*, 135703.

Koç, Ç., Bektaş, T., Jabali, O., & Laporte, G. (2014). The fleet size and mix pollution-routing problem. *Transportation Research Part B: Methodological*, *70*, 239-54.

Kucukoglu, I., Dewil, R., & Cattrysse, D. (2021). The electric vehicle routing problem and its variations: A literature review. *Computers & Industrial Engineering*, *161*, 107650.

Liu, K., Yamamoto, T., & Morikawa, T. (2017). Impact of road gradient on energy consumption of electric vehicles. *Transportation Research Part D: Transport and Environment*, *54*, 74–81.

Longhitano;, P. D., Bérenguer;, C., & Echard., B. (2023). *Battery degradation aware vehicle routing with speed and acceleration limitations.* MIMAR 2023 - 12th IMA International Conference on Modelling in Industrial Maintenance and Reliability Applications, Nottingham, United Kingdom.

Marrekchi, E., Besbes, W., Dhouib, D., & Demir, E. (2021). A review of recent advances in the operations research literature on the green routing problem and its variants. *Annals of Operations Research*, *304*(1-2), 529–74.

Moghdani, R., Salimifard, K., Demir, E., & Benyettou, A. (2021). The green vehicle routing problem: A systematic literature review. *Journal of Cleaner Production*, *279*, 123691.

Pelletier, S., Jabali, O., & Laporte, G. (2019). The electric vehicle routing problem with energy consumption uncertainty. *Transportation Research Part B: Methodological*, *126*, 225–55.

Qian, J., & Eglese, R. (2014). Finding least fuel emission paths in a network with time-varying speeds. *Networks*, *63*(1), 96–106.

Raeesi, R., & O'Sullivan, M. J. (2014). Eco-logistics: Environmental and economic implications of alternative fuel vehicle routing problem. *International Journal of Business Performance and Supply Chain Modelling*, *6*(3-4), 276–97.

Raeesi, R., & Zografos, K. G. (2019). The multi-objective Steiner pollution-routing problem on congested urban road networks. *Transportation Research Part B: Methodological*, *122*, 457–85.

Raeesi, R., & Zografos, K. G. (2020). The electric vehicle routing problem with time windows and synchronised mobile battery swapping. *Transportation Research Part B: Methodological*, *140*, 101–29.

Raeesi, R., & Zografos, K. G. (2022). Coordinated routing of electric commercial vehicles with intra-route recharging and en-route battery swapping. *European Journal of Operational Research*, *301*(1), 82–109.

Rastani, S., Yüksel, T., & Çatay, B. (2019). Effects of ambient temperature on the route planning of electric freight vehicles. *Transportation Research Part D: Transport and Environment*, *74*, 124–41.

Salimifard, K., & Raeesi, R. (2014). A green routing problem: Optimising CO2 emissions and costs from a bi-fuel vehicle fleet. *International Journal of Advanced Operations Management*, *6*(1), 27-57.

Schäfers, L., Iravanimanesh, S., Franke, K., Savelsberg, R., & Pischinger, S. (2024). Holistic sensitivity analysis for long-term energy demand prediction of battery electric vehicles. *Automotive Innovation*, *7*(4), 669–81.

Schneider, M., Stenger, A., & Goeke, D. (2014). The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, *48*(4), 500–20.

Soysal, M., Çimen, M., & Belbağ, S. (2020). Pickup and delivery with electric vehicles under stochastic battery depletion. *Computers & Industrial Engineering*, *146*, 106512.

Spinelli, A., Bezzi, D., Jabali, O., & Maggioni, F. (2026). A stochastic electric vehicle routing problem under uncertain energy consumption. *Transportation Research Part C: Emerging Technologies*, *183*, 105480.

Turkensteen, M. (2017). The accuracy of carbon emission and fuel consumption computations in green vehicle routing. *European Journal of Operational Research*, *262*(2), 647–59.

Van Gijlswijk, R., Paalvast, M., Ligterink, N. E., & Smokers, R. (2020). Real-world fuel consumption of passenger cars and light commercial vehicles. *TNO 2020 R11664*, 21–3.

Wang, J.-b., Liu, K., Yamamoto, T., & Morikawa, T. (2017). Improving estimation accuracy for electric vehicle energy consumption considering the effects of ambient temperature. *Energy Procedia*, *105*, 2904–9.

Wang, J., Besselink, I., & Nijmeijer, H. (2018). Battery electric vehicle energy consumption prediction for a trip based on route information. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, *232*(11), 1528–42.

Weiss, M., Winbush, T., Newman, A., & Helmers, E. (2024). Energy consumption of electric vehicles in Europe. *Sustainability*, *16*(17), 7529.

Xiong, H., Xu, Y., Yan, H., Guo, H., & Zhang, C. (2024). Optimizing electric vehicle routing under traffic congestion: A comprehensive energy consumption model considering drivetrain losses. *Computers & Operations Research*, *168*, 106710.

Yang, S., Li, M., Lin, Y., & Tang, T. (2014). Electric vehicle's electricity consumption on a road with different slope. *Physica A: Statistical Mechanics and its Applications*, *402*, 41–8.

Yi, Z., & Bauer, P. H. (2017). Effects of environmental factors on electric vehicle energy consumption: a sensitivity analysis. *IET Electrical Systems in Transportation*, *7*(1), 3–13.

Yu, V. F., Jodiawan, P., & Gunawan, A. (2021). An Adaptive Large Neighborhood Search for the green mixed fleet vehicle routing problem with realistic energy consumption and partial recharges. *Applied Soft Computing*, *105*, 107251.