Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Department of the Environment,
Transport, Energy and Communication DETEC

Swiss Federal Office of Energy SFOE
Energy Research and Cleantech

Final report from August 26, 2025

GAMES

Grid Aware Mobility and Energy Sharing

GRID AWARE MOBILITY AND ENERGY SHARING

Publisher:
Swiss Federal Office of Energy SFOE
Research Programme Grids
CH-3003 Bern
www.bfe.admin.ch
energieforschung@bfe.admin.ch

Co-financed by:
Austrian Research Promotion Agency (FFG)
Ministry of Energy of Israel

Subsidy recipients:
University of Applied Sciences and Arts of Southern Switzerland (SUPSI), CH-6928 Manno
Hive Power SA, CH-6928 Manno
Sun2wheel AG, CH-6012 Obernau
Azienda Elettrica di Massagno SA, CH-6900 Massagno

Consortium Partners (ERA-Net MICall20–EnerDigit):
e7 energy innovation & engineering, Vienna, Austria
Salzburg Research Forschungsgesellschaft GmbH, Salzburg, Austria
Reichman University, Herzliya, Israel

Authors:
Roberto Rocchetta, ISAAC-SUPSI roberto.rocchetta@supsi.ch
Lorenzo Nespoli, ISAAC-SUPSI, lorenzo.nespoli@supsi.ch
Jalomi Maayan Tardif, ISAAC-SUPSI jalomi.maayantardif@supsi.ch
Vasco Medici, ISAAC-SUPSI, vasco.medici@supsi.ch

SFOE project coordinators:
Dr. Michael Moser, michael.moser@bfe.admin.ch

SFOE contract number: SI/502361

Summary

As a part of the ERA-Net Smart Energy Systems Joint Call 2020 on digital transformation for green energy transition, the GAMES project supported the integration of EVs and shared electric mobility by providing fleet owners with decision-support tools for scheduling operations and maintaining their electrified carsharing fleets. These tools enable grid-aware planning, cost-efficient relocation strategies, and operational optimisation to enhance profitability and sustainability. A probabilistic modelling framework has been proposed to characterise the spatiotemporal variability affecting mobility service demand within a control area, with a formal statistical analysis of the mobility demand profile of a free-floating system in Tel Aviv (AutoTel) and a station-based system in Switzerland (Mobility). This analysis highlighted similarities and differences between the two systems, offering insights into the potential for developing transferable decision-support tools applicable to both. A probabilistic demand forecaster has been trained to predict and support the simulation of realistic demand patterns, including electricity demand and mobility services. This model forecasts the demand for mobility services in the coming hours, providing fleet operators with valuable insights for optimizing fleet operations, such as relocating idle vehicles and scheduling EV charging. Then, a novel decision-support framework based on Reinforcement Learning (RL) is introduced to optimize operations in free-floating electric vehicle carsharing systems (FFEVCS). This framework explores how user participation in relocation and rebalancing activities can be incentivised using dedicated mobility incentives. It integrates an economic dispatch model to incorporate grid information, enabling grid-aware charging decisions that optimize charging costs and mobility revenues. Within the RL framework, a centralized control agent assigns dynamic, neighborhood-specific discounts on minute tariffs and mobility credits to incentivize users to participate in relocation activities. This reduces reliance on costlier crew-based relocations. The incentives are adapted to various operational conditions and fleet states, supported by predictions from the probabilistic demand forecaster. A transportation mode choice model is also integrated within the simulation environment to realistically emulate the reservation process and test different operational policies and relocation strategies. The system is demonstrated on the free-floating carsharing model in Tel Aviv, using real-world data from AutoTel. Four RL agents are trained using advanced algorithms (PPO, TD3, DDPG, and SAC) designed for continuous action spaces, are compared against baseline strategies such as crew-based and mixed relocation. Results show that these agents achieve higher revenues and lower charging and relocation costs. To evaluate scalability, two observation spaces with increasing dimensionality are tested, incorporating both local and global observations.

Finally, the potential of the integrated platform is demonstrated through two digital cross-sector platforms that connect mobility and energy data within a common framework. One platform enables corporate EV fleets to participate in grid services by providing real-time vehicle data and integrating with market mechanisms, while the other offers fleet operators decision-support tools for grid-aware relocation optimization, enhancing efficiency and sustainability in fleet management. This project ultimately supports the integration of EVs and shared electric mobility by providing fleet owners with decision-support tools that optimize operations, reduce costs, and enhance profitability. Grid-informed planning and relocation strategies enable more efficient management, helping fleet operators achieve a techno-economic optimum while promoting sustainable mobility solutions.

Zusammenfassung

Im Rahmen der gemeinsamen Ausschreibung „ERA-Net Smart Energy Systems 2020“ zum Thema digitale Transformation für die Energiewende unterstützte das Projekt GAMES die Integration von Elektrofahrzeugen und gemeinschaftlicher Elektromobilität, indem es Flottenbesitzern Entscheidungshilfetools für die Einsatzplanung und Wartung ihrer elektrifizierten Carsharing-Flotten zur Verfügung stellte. Diese Tools ermöglichen eine netzorientierte Planung, kosteneffiziente Umverteilungsstrategien und eine Optimierung des Betriebs, um die Rentabilität und Nachhaltigkeit zu verbessern. Es wurde ein probabilistischer Modellierungsrahmen vorgeschlagen, um die räumlich-zeitliche Variabilität zu charakterisieren, die die Nachfrage nach Mobilitätsdiensten innerhalb eines Kontrollgebiets beeinflusst. Dazu wurde eine formale statistische Analyse des Mobilitätsnachfrageprofils eines Free-Floating-Systems in Tel Aviv (AutoTel) und eines stationsbasierten Systems in der Schweiz (Mobility) durchgeführt. Diese Analyse hob Ähnlichkeiten und Unterschiede zwischen den beiden Systemen hervor und lieferte Erkenntnisse über das Potenzial für die Entwicklung übertragbarer Entscheidungshilfetools, die für beide Systeme geeignet sind. Ein probabilistischer Bedarfsprognostiker wurde trainiert, um realistische Bedarfsmuster, einschließlich des Strombedarfs und der Mobilitätsdienstleistungen, vorherzusagen und deren Simulation zu unterstützen. Dieses Modell prognostiziert den Bedarf an Mobilitätsdienstleistungen für die kommenden Stunden und liefert Flottenbetreibern wertvolle Erkenntnisse für die Optimierung des Flottenbetriebs, wie z. B. die Verlegung von stillstehenden Fahrzeugen und die Planung der Ladung von Elektrofahrzeugen. Anschließend wird ein neuartiger Entscheidungshilferahmen auf Basis von Reinforcement Learning (RL) eingeführt, um den Betrieb von Free-Floating-Carsharing-Systemen für Elektrofahrzeuge (FFEVCS) zu optimieren. Dieser Rahmen untersucht, wie die Beteiligung der Nutzer an Umverteilungs- und Ausgleichsmaßnahmen durch spezielle Mobilitätsanreize gefördert werden kann. Er integriert ein wirtschaftliches Dispatch-Modell zur Einbeziehung von Netzinformationen und ermöglicht so netzorientierte Ladeentscheidungen, die die Ladekosten und Mobilitätserlöse optimieren. Innerhalb des RL-Rahmens weist ein zentraler Steuerungsagent dynamische, nachbarschaftsspezifische Rabatte auf Minutentarife und Mobilitätsgutschriften zu, um die Nutzer zur Teilnahme an Umstellungsmaßnahmen zu motivieren. Dadurch wird die Abhängigkeit von kostspieligeren Umstellungen durch Personal reduziert. Die Anreize werden an verschiedene Betriebsbedingungen und Flottenzustände angepasst, unterstützt durch Vorhersagen des probabilistischen Nachfrageprognosemodells. Ein Modell zur Wahl des Verkehrsmittels ist ebenfalls in die Simulationsumgebung integriert, um den Reservierungsprozess realistisch nachzubilden und verschiedene Betriebsrichtlinien und Umstellungsstrategien zu testen. Das System wird am Beispiel des Free-Floating-Carsharing-Modells in Tel Aviv unter Verwendung realer Daten von AutoTel demonstriert. Vier RL-Agenten werden mit fortschrittlichen Algorithmen (PPO, TD3, DDPG und SAC) trainiert, die für kontinuierliche Aktionsräume entwickelt wurden, und mit Basisstrategien wie crew-basierten und gemischten Umverteilungen verglichen. Die Ergebnisse zeigen, dass diese Agenten höhere Einnahmen und niedrigere Lade- und Umstellkosten erzielen. Um die Skalierbarkeit zu bewerten, werden zwei Beobachtungsräume mit zunehmender Dimensionalität getestet, die sowohl lokale als auch globale Beobachtungen einbeziehen.

Schließlich wird das Potenzial der integrierten Plattform anhand von zwei digitalen branchenübergreifenden Plattformen demonstriert, die Mobilitäts- und Energiedaten in einem gemeinsamen Rahmen verbinden. Eine Plattform ermöglicht es Unternehmensflotten von Elektrofahrzeugen, sich an Netzdiensten zu beteiligen, indem sie Echtzeit-Fahrzeugdaten bereitstellt und sich in Marktmechanismen integriert, während die andere Flottenbetreibern Entscheidungshilfetools für die netzorientierte Relokalisierungsoptimierung bietet, wodurch die Effizienz und Nachhaltigkeit im Flottenmanagement verbessert werden. Dieses Projekt unterstützt letztlich die Integration von Elektrofahrzeugen und gemeinschaftlicher Elektromobilität, indem es Flottenbesitzern Entscheidungshilfetools zur Verfügung stellt, die den Betrieb optimieren, Kosten senken und die Rentabilität steigern. Netzintegrierte Planungs- und Verlagerungsstrategien ermöglichen ein effizienteres Management und helfen Flottenbetreibern, ein techno-ökonomisches Optimum zu erreichen und gleichzeitig nachhaltige Mobilitätslösungen zu fördern.

Sommario

Nell'ambito del bando congiunto ERA-Net Smart Energy Systems 2020 sulla trasformazione digitale per la transizione verso l'energia verde, il progetto GAMES si è proposto di supportare l'integrazione dei veicoli elettrici (EV) e della mobilità condivisa, fornendo ai proprietari di flotte strumenti avanzati di supporto alle decisioni per la pianificazione delle operazioni e la manutenzione delle flotte elettrificate. Questi strumenti si concentrano sulla pianificazione informata dalla rete elettrica, sulle strategie di ricollocamento e sull'ottimizzazione delle operazioni, con l'obiettivo di migliorare la redditività, ridurre i costi di ricarica e gestione e raggiungere un equilibrio tecnico-economico. In primo luogo, il progetto propone un quadro di modellazione probabilistica per caratterizzare la variabilità spaziotemporale che influisce sulla domanda di servizi di mobilità all'interno di un'area di controllo. La fase iniziale include un'analisi formale delle curve di domanda di mobilità per due modelli tradizionali di carsharing: un sistema "free-floating" a Tel Aviv (AutoTel) e un sistema basato su stazioni in Svizzera (Mobility). Questa analisi mette in evidenza le somiglianze e le differenze tra i due sistemi, fornendo spunti per lo sviluppo di strumenti di supporto alle decisioni generalizzabili e applicabili a entrambi. Successivamente, viene addestrato un previsore probabilistico della domanda per minimizzare le discrepanze tra i dati di carico osservati e la domanda prevista. Questo modello è in grado di prevedere la domanda di servizi di mobilità nelle ore successive, offrendo agli operatori di flotte informazioni preziose per ottimizzare le operazioni, come la rilocazione dei veicoli inattivi e la programmazione della ricarica degli EV. In una seconda fase del progetto, viene introdotto un nuovo quadro basato sull'apprendimento per rinforzo, particolarmente adatto per ottimizzare le operazioni in sistemi complessi come quelli del carsharing elettrico di tipo free-floating. Questo quadro di ottimizzazione esplora come incentivare la partecipazione degli utenti alle attività di rilocazione e bilanciamento attraverso incentivi di mobilità dedicati. Viene inoltre integrato un modello di economic dispatch per sfruttare le informazioni della rete elettrica, consentendo decisioni di ricarica consapevoli che ottimizzano sia i costi di ricarica sia i ricavi derivanti dalla mobilità. All'interno del quadro RL, un agente di controllo centralizzato assegna sconti dinamici e specifici per il quartiere sulle tariffe a minuto e crediti di mobilità per incentivare gli utenti a partecipare alle attività di rilocazione. Questo riduce la dipendenza dalle rilocazioni basate su equipaggi, più costose. Gli incentivi sono adattati a varie condizioni operative e stati della flotta, supportati da previsioni del previsore probabilistico della domanda. Un modello di scelta del mezzo di trasporto è stato integrato nell'ambiente di simulazione per emulare realisticamente il processo di prenotazione e testare diverse politiche operative e strategie di rilocazione. Il sistema è dimostrato sul modello di carsharing "free-floating" a Tel Aviv, utilizzando dati reali di AutoTel. Quattro agenti di RL, addestrati con algoritmi (PPO, TD3, DDPG e SAC) progettati per spazi d'azione continui, sono confrontati con strategie di base come la rilocazione basata su equipaggi e quella mista. I risultati mostrano che questi agenti ottengono maggiori ricavi e minori costi di ricarica e rilocazione. Per valutare la scalabilità, vengono testati due spazi di osservazione di crescente dimensionalità, incorporando sia osservazioni locali che globali.
Infine, il potenziale della piattaforma integrata è dimostrato da due piattaforme digitali intersettoriali che collegano i dati sulla mobilità e sull'energia in un quadro comune. Una piattaforma consente alle flotte aziendali di veicoli elettrici di partecipare ai servizi di rete fornendo dati in tempo reale sui veicoli e integrandosi con i meccanismi di mercato, mentre l'altra offre agli operatori delle flotte strumenti di supporto alle decisioni per l'ottimizzazione della ricollocazione attenta ai bisogni della rete, migliorando l'efficienza e la sostenibilità della gestione delle flotte. Questo progetto supporta infine l'integrazione dei veicoli elettrici e della mobilità condivisa, fornendo ai proprietari di flotte strumenti di supporto alle decisioni che ottimizzano le operazioni, riducono i costi e migliorano la redditività. La pianificazione informata dalla rete elettrica e le strategie di rilocazione consentono una gestione più efficiente, aiutando gli operatori delle flotte a raggiungere un optimum tecnico-economico e promuovendo soluzioni di mobilità sostenibile.

Contents

List of abbreviations

B2B	Business to Business
B2C	Business to Customer
BMC	Binary Transpiration Mode Choice
CDF	Comulative Distribution Function
CS	Carsharing
DDPG	Deep Deterministic Policy Gradient
DSO	Distribution System Operator
ED	Economic Dispatch
EV	Electric Vehicle
FFCS	Free-Floating Carsharing
FFEVCS	Free-Floating Electric Vehicle Carsharing
KDE	Kernel Density Estimator
O&M	Operations and Management
ODI	Origin Destination Incentives
OWCS	One-Way Carsharing
PDF	Probability Density Function
PPO	Proximal Policy Optimization
PV	Photovoltaic
RES	Renewable Energy Surces
RL	Reinforcement Learning
SAC	Soft Actor-Critic
SBCS	Station-Based Carsharing
SOC	State-of-Charge
SVF	Shared Vehicle Fleet
TD3	Twin Delayed Deep Deterministic Policy Gradient
TSO	Transmission System Operator

1 Introduction

The increasing popularity of electrified mobility services, such as shared electric vehicle fleets, is rapidly changing the landscape of the transportation industry. As a result of this growing electrification, the transportation network and electrical power grid (once considered independent) are forming a stronger bond while posing new challenges for operators of both systems. Despite the many challenges, this higher connectivity and electrification also bring new opportunities. These opportunities include the potential for enhanced synergies between these two systems, which can, in turn, increase profitability and sustainability by introducing new revenue streams. The project GAMES, which stands for *Grid Aware Mobility and Energy Sharing*, aims to explore viable options for highly electrified carsharing fleets and power grid hosting these systems. Ultimately, the project seeks to explore promising directions to foster a higher electrification of carsharing fleets by prescribing advanced data analytic tools to support decision-making and policy-making, finally hoping to encourage a long-lasting collaboration between the transportation and electric sectors.

1.1 Context and motivation

This project was carried out under the *ERA -Net Smart Energy Systems & Mission Innovation Joint Call 2020 "Digital Transformation for Green Energy Transition (MICall20 – EnerDigit)"*[1] initiative.

The electrification of mobility will reduce fossil fuel consumption in the transport sector, along with the related nocive emissions and associated health hazards. In the context of the "*Roadmap for Electric Mobility 2022*" and the "*Digital Switzerland*" strategy, the Swiss Federal Office of Energy (SFOE) has sought to promote electric and shared modes of transport. Concurrent with the increasing availability and diversity of electric vehicles on the market and with governmental incentives to electrify transport (as a means to decarbonize this sector), many vehicle fleet owners are contemplating how to shift partially or entirely to EVs. However, this transition poses a challenge for fleet owners, especially as the required charging infrastructure can be expensive and complex to position when considering electric grid and parking space restrictions. Additionally, to truly reduce their environmental footprint by using RES to recharge the EVs, the vehicle fleet owners need to coordinate the charging events with electricity suppliers and distributors and invest in their renewable energy installations. Thus, the electrification of shared vehicle fleets (SVFs) heavily depends on coordination and collaboration with complementary partners which could be facilitated by creating cross-sector services made possible through digitalization. Analysing both qualitative and quantitative data from SVF providers and users will allow us to assess distinct shared mobility requirements and use cases.

GAMES evaluated cross-sector exchanges between SVFs and the electric grid, which could present viable business models that benefit both the mobility and energy sectors, as well as support more sustainable lifestyles from environmental, social, and economic perspectives. Hence, investigating novel opportunities for grid-aware operational and management of electrified carsharing fleets is also paramount in this context and the project GAMES tackled this pressing issue. Among the main aims of the project, is to unveil the potential for grid-awareness to support carsharing fleet operations, possibly increasing revenues, reducing charging costs and providing ancillary services to DSO and the grid. Finally, the project aims to foster sustainable development within the energy and transportation sectors. The principles that lay the ground for the GAMES project are that:

1. The development and implementation of innovative and sustainable business models involving asset-sharing mechanisms help minimize customer risk, provide wider and fair access to sustainable energy technologies, and offer economically valuable cross-sector services.
2. Cross-sector coupling can improve the efficiency and resilience of energy systems at various geographical scales, and these rely partially on partnerships established via digital monitoring and analysis, and a comprehensive mapping of collateral benefits.

[1] https://www.eranet-smartenergysystems.eu/Calls/EnerDigit_Calls_funding/Joint_Call_2020

3. Understanding users, their mobility needs and behavior, their motives and engagement patterns with the service and its providers should be at the center of any mobility planning.
4. Shared vehicle fleets can have the critical mass required to offer reliable grid services, can reduce the number of vehicles and lower GHG emissions, reduce the need for parking space by reducing private vehicle holdings, and at least partially complement public transport.

Seeking an attack on these challenges, GAMES as focused on three main demonstrative use cases, defined thanks to the joint effort of academic and industrial partners within the consortium. Each use case presents unique opportunities and challenges in terms of grid-awareness and grid-fleet interactions, reflecting distinct operational models, regional electricity grids, and user behaviours:

1. A free-floating carsharing system, operated by AutoTel in the city of Tel Aviv;
2. A station-based carsharing system, operated by Mobility, focusing on the Zurich area, Switzerland;
3. A small corporate fleet in Austria, Operated by the group Windkraft Simonsfeld.

Overall, GAMES endeavors to demonstrate the feasibility and benefits of integrating electrified carsharing fleets with grid operations, using the three diverse use cases as practical examples of how grid-aware management can contribute to sustainable, cost-effective, and resilient transportation and energy systems.

1.2 Project objectives

The main objectives of GAMES are as follows:

1. Road-map for the transition of shared vehicle fleets to electric vehicles.
2. Cross-sector digital platform interoperability proof-of-concept.
3. Enhance efficiency, sustainability, stability, and security of distribution networks.
4. Amplify electric and shared mobility appeal to a wider and more diverse public.
5. Quantify the economic value of the flexibility provided by EV fleets

The objective and outcomes of the GAMES project are therein organized in nine deliverables (D1 to D9), delivery month (M), and partner responsible for the delivery. The scope, leading team, timing, and current status are summarized as follows:

- **D1 - M12 - SUPSI** [Completed]: Annual Report to the joint call initiative and national funding agencies
- **D3 - M9 - SRFG** [Completed]: Data Management Plan & Ethics: management of non-personal and personal data and ethical aspects of behavioral nudging
- **D4 - M12 - E7** [Completed]: Policy brief - publicly available on the GAMES website games-innovation.net
- **D5 - M13 - SUPSI** [Completed]: Forecasting algorithm code in GitHub - publicly available on the Github repository: GAMES_public_GitHub
- **D2 - M24 - SUPSI** [Completed]: Annual Report to the joint call initiative and national funding agencies
- **D6 - M25 - IDC** [Completed]: Report on shared mobility user characteristics and transport mode decision making - publicly available on the GAMES website games-innovation.net

- **D7 - M25 - E7** [Completed]: Industry white-paper - publicly available on the GAMES website games-innovation.net
- **D8 - M25 - SUPSI** [Completed]: Report on digital cross-sector platform proof-of-concept - publicly available on the GAMES website games-innovation.net
- **D9 - M30 - SUPSI** [Completed]: Final reporting and abstract of main results for the joint call initiative and national funding agencies

Within the GAMES project, the Swiss partners have concentrated their efforts on key technical areas, outlined as follows:

1. **Probabilistic Forecasting and Mobility Simulation Models**: a comprehensive mobility simulation model was developed and advanced probabilistic forecasters for energy and mobility demand. These simulation tools enable the testing of control policies under realistic conditions, thereby minimizing the need for extensive and costly field testing. The forecasters predict demand trends and help quantify uncertainties associated with mobility and energy needs, identifying patterns in carsharing usage and variability across times and locations.
2. **Grid-Aware Optimal Relocation and Dispatch Scheduling**: This approach emphasizes the strategic relocation of electric vehicles to optimize revenue from mobility services while aligning with grid conditions. Optimal relocation can leverage excess production from distributed renewable sources, reducing charging costs, storing surplus power, and enhancing energy flexibility. Since relocation and economic dispatch are tightly interrelated, economic dispatch (ED) optimizes the charging and discharging schedules to control costs, ensure vehicle availability, and alleviate grid congestion. These ED aspects were further developed in collaboration with Austrian partners under
3. **Digital Cross-Sector Platform**: gathering critical information to explore the information exchange enabling shared EV fleets to offer flexible services to the energy and distribution sectors. As part of this initiative under, a web interface was developed for fleet operators to visualize and monitor the fleet's status and the energy grid in real-time . This platform allows operators to identify EVs needing relocation, track their status, and monitor mobility trends. Additionally, it integrates statistical forecasts on mobility demand, supporting data-driven decision-making for optimized operations and system efficiency. This proof-of-concept illustrates the potential for cross-sector interoperability, including data integration from the fleet, charging, and grid sources and forecasting models.

The following activities have been successfully completed by SUPSI:

- **Data Collection and Use Case Definition**: Carsharing data has been extensively collected and analyzed, resulting in the identification of three distinct use cases, each offering unique insights and opportunities for testing grid-aware mobility strategies, see section 2.2 for details.
- **Development of Forecasting Algorithms**: designed, trained, and rigorously validated a suite of forecasting algorithms capable of accurately predicting both mobility and energy demands. These algorithms provide a solid foundation for adaptive and grid-responsive fleet management, enabling real-time decision-making with heightened accuracy and reliability.
- **Macroscopic Traffic Modeling for Optimal Relocation**: A static and macroscopic traffic model was developed to simulate reservation patterns, offering a comprehensive view of fleet mobility dynamics. This model has been seamlessly embedded within the optimal relocation and mobility incentive framework proposed, contributing to more precise, responsive relocation strategies. Refer to section 2.2.3.
- **Probabilistic Global Mobility Demand Forecasting Tool**: created an advanced probabilistic forecasting tool that leverages sophisticated spatiotemporal models to project global mobility demand patterns. This tool enables more accurate anticipation of demand fluctuations, supporting optimal fleet distribution and improved service reliability.

- **Completion of Deliverable D5 and Open-Source Contributions**: The deliverable D5 has been successfully finalized by SUPSI, encompassing all developed codes, toolboxes, and the core forecasting algorithms. To support transparency and collaboration, these resources have been published in a dedicated GitHub repository: https://github.com/supsi-dacd-isaac/GAMES_public. This open-access repository ensures that the tools and findings from GAMES are readily accessible for further research and development within the industry.
- **Completion of Deliverable D8 and web-based deployment of visualisation dashboard**: The deliverable D8 has been successfully finalized by SUPSI. To support transparency and collaboration, the proof of concept has been deployed here: https://games.hivepower.tech/.

2 Approach, method, results, and discussion

In this section, we outline the key aspects of the modelling framework developed for the GAMES project. We will discuss forecasting models and simulation frameworks for grid-aware fleet optimization. Additionally, we will introduce simulation results, selected use cases and data provided by project partners that are essential for testing and validating the proposed methodology. The proposed framework seeks to investigate whether interactions between the power grid and carsharing fleet are possible and to what extent optimized operational policies informed by forecasting algorithms and trained by advanced machine learning methods could support synergic interactions. To this aim, this investigation explores the potential for grid-aware relocation, incentives, and dispatch policies.

2.1 Preliminaries

Carsharing (CS) systems are generally divided into different classes based on designated rules for trips and the allowed locations where vehicles can be picked up and returned [1]. Free-Floating Carsharing (FFCS) and Station-Based Carsharing (SBCS) systems define different models depending on the rental locations where cars can be found and delivered. Users of FFCS systems can rent vehicles anywhere within a designated service area, whereas SBCS requires users to collect and return vehicles only at specific locations, such as charging stations and parking lots [2]. Thus, by design, SBCS systems offer greater control for fleet managers, but lower flexibility for users. In contrast, FFCS systems enable users to be more versatile in their trips, leading to higher booking rates and generally shorter trips within a relatively small area. Return-Based Carsharing (RBCS), also known as round-trip CS [3], differs from One-Way CS (OWCS) in terms of the trip types allowed rather than the pick-up and delivery points. Return-based systems require the locations of departure and arrival to overlap, while one-way systems allow drop-offs at different locations. Mobility data sets generated by these systems can provide valuable insights into user behaviour and can differ significantly. For instance, estimating vehicle flows between geographical areas may be inherently more challenging for RBCS [4], due to a lack of geographical information on the trip's intermediate stops. On the other hand, free-floating systems and one-way mobility data offer valuable insights into vehicle usage, but the non-finite set of delivery points may require dedicated aggregation approaches (e.g., zoning, clustering) to facilitate analysis. Despite these challenges, leveraging geographical information on trip destinations can enhance our understanding of the geographical variability of mobility demand patterns. In particular, the efficacy of electrified mobility services in meeting transportation and sustainability requirements in urban areas relies on suitable tools to process and analyze this data, reliable information from the energy grid, and techniques to improve user participation, ultimately seeking optimized operational and maintenance policies.

2.2 Definition of carsharing use cases

Throughout the GAMES project, both real-world data and high-fidelity simulated data were provided and exchanged by academic and industry partners. Specifically, synthetic and raw carsharing mobility data were shared by the companies: Mobility (CH), AutoTel (IL), and Windkraft Simonsfeld (AT). These datasets include carsharing events and form the foundation for our investigation of grid-aware operational policies for carsharing systems. These datasets represent mobility patterns from different carsharing systems in three different countries: Switzerland, Israel, and Austria. This present an opportunity to investigate differences and similarities in the operations of different CS systems, potential for electrification and usefulness of grid-aware operations to enhance profit, efficacy, and sustainability of the system.

To comprehensively assess the impact and opportunities of highly electrified carsharing on different business models and system paradigms, three representative use cases were selected and defined. Each use case, represents a specific carsharing model and scenario, namely; a medium-size station-based carsharing system, a medium-size free-floating one-way systems, and a small-size corporate fleet. These use cases forms the cornerstone of our research efforts and are explored to better understand the complex dynamics (and potential for synergies) relating electrified carsharing fleet, the transportation system, and energy grids. Table 2 gives a brief summary of the available data for the selected use cases.

	Data		
State	Switzerland (CH)	Israel (IL)	Austria (AT)
Group	**Mobility**	**AutoTel**	**Windkraft Simonsfeld**
Available	✓	✓	✓
Trips	return	one-way	return
System	SBCS	FFCS	SBCS
Business model	B2B/B2C	B2C	B2B
	Use case		
Focus area	**Zurich**	**Tel-Aviv**	**The HQ in Ernstbrunn**
Domain	Municipality-level	Government offices	Corporate fleet
Transportation grid (G_{tr})	✓	✓	✗
Number of columns	12	13	9
Number of rows	256146	1033624	12084
Stations identifiers	230	n.a.	1 (the HQ)
Unique EV identifiers	495	620	20

Table 2: A summary of the gathered mobility data and the three use cases considered within the scope of the project.

A station-based system The first use case focuses on **station-based EV carsharing system** where vehicles are collected and returned to a limited number of designated stations.

- **Data Set Description:** A simulated dataset of mobility events is used for this case study and obtained via a pre-existing high-fidelity simulation model, see e.g., [5, 6], specifically trained to replicate mobility patterns in Zurich.The mobility data is thus synthetic and generated from pre-existing simulation tools. Example data include about 250 thousand simulated mobility events, 230 unique identifiers for stations and a fleet with approximately 500 vehicles, and 12 mobility features.
- **Geographic Area**: Zurich, CH.
- **Objective**: Evaluate the challenges and opportunities of grid-awareness, forecasters and optimal management of O&M strategies in station-based (and return-based) carsharing systems at the municipal level.

Figure 1 present a visualization example for the simulated Mobility data in SBCS system in Zurich. The red crosses represent user demands for trips over a simulated time step, providing insight into the spatial distribution of trip requests. The distribution of stations is represented by black squared markers, allocated over the street network G_{tr} depicted by blue solid lines. Because the system is station-based, the EVs are not visualized as their location coincides with the one of the stations.

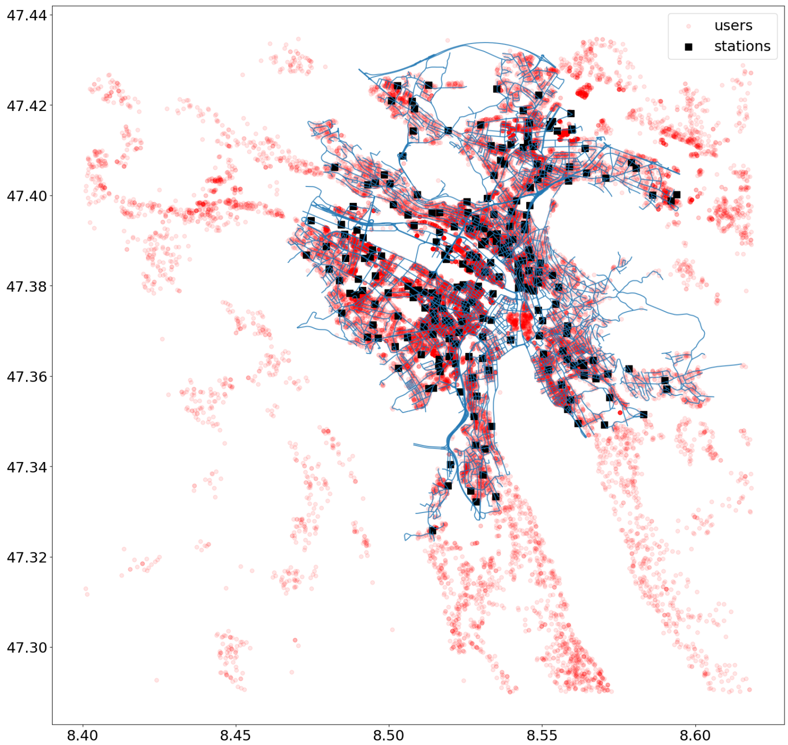

Figure 1: The distribution of carsharing users, i.e., their home location (red markers) ,and stations (black squared markers) over the street network $\mathcal{G}_{tr}$ of Zurich.

A free-floating system This use case focuses on **free-floating one-way EV carsharing** system in Tel Aviv, which operates without fixed stations, allowing vehicles to be picked up and dropped off anywhere within the service area. We aim to explore common carsharing usage patterns in this system, characterized by relatively short and frequent urban trips.

- **Data Set Description:**. The dataset is historical data collected over two years of operations. The data comprises 2 years of historical reservation events, including about one million of events and 620 unique EV identifiers, and 13 trip features. Table 3 gives an example of available data features, formats and types, which are expected by a mobility data set after pre-processing, filtering, and cleaning the raw data. Note that station data for the free-floating carsharing system was not made available.
- **Geographic Area**: Tel Aviv, IL.
- **Objective**: Evaluate the advantages and disadvantages of grid awareness, forecasters and optimal management of O&M strategies in free-floating one-way carsharing systems at the municipal level. Additionally, the case evaluates optimal policies for managing fleet relocation, mobility incentives, and economic dispatch of vehicles to balance operational efficiency and user satisfaction in a dynamic, urban carsharing system. This analysis will examine the role of grid-awareness in these systems, including shorter idle times, rapidly changing SOC, and unstructured pickup and delivery points. Specifically, we will evaluate the advantages and challenges of optimal management

policies—such as relocation, mobility incentives, and economic dispatch—applied to free-floating electric vehicle carsharing systems at the municipal level.

2.2.1 A corporate fleet

This use case focuses on a **small corporate EV fleet** managed by Windkraft Simonsfeld at their headquarters in Austria.

- **Data Set Description**: The dataset was gathered from a fleet of 20 electric vehicles and includes around 12,000 rows with 9 features. Only one station is considered (the headquarters).
- **Geographic Area**: Ernstbrunn, AT.
- **Objective**: Analyze the effects of optimizing charge and discharge strategies at the building level, with a focus on integrating electric vehicles with local renewable energy sources. This case provides insights into fleet management at a corporate level, exploring localized usage patterns, energy challenges, and how EVs can support grid services in such an environment.

Modelling grid-aware electric fleets A comprehensive evaluation of the voltage profile and disturbances in the electric grid relies on a deep understanding of the distribution grid's topology and the operational state and dynamics of the electrical grid. Similarly, optimizing EV fleet operational policies while assessing their impact on revenue streams from car booking services necessitates a thorough examination of the transportation network. This includes the development of models for estimating traffic congestion, characterizing EV idle times (related to fueling/charging duration), and employing forecasting tools to predict future trip patterns. In cases where the adoption of a large number of EVs by CS companies is widespread, several factors, such as power and mobility demand patterns and distribution grid stability could be significantly influenced. Thus, it is vital to recognize the complex and growing interplay between the power distribution grid and the transportation network. Yet, traditional single-network models may overlook these complex interactions, highlighting the necessity for a more integrated approach. Figure 2 illustrates this concept with a schematic and conceptual integrated modelling framework. It is imperative to incorporate grid-awareness considerations—such as electrical distribution grid structure, distributed renewable energy production, and the potential for load flow simulations—within frameworks for modeling transportation systems. Vice versa, models for analyzing power grids must adopt an integrated approach that considers electric vehicles and mobility behavior. In summary, a combined modelling approach is essential for accurately quantifying the collective impact of highly electrified mobility fleets on power distribution loads, voltage profiles, and the potential for disturbances within the power grid.

For the GAMES project, a comprehensive data analytics framework was developed to evaluate the potential and implications of optimized grid-aware operational and management strategies for electric carsharing fleets.

2.2.2 Energy grid

In the GAMES conceptual modeling framework, we consider a power grid model denoted as $\mathcal{G}_{el} = (\mathcal{N}_{el}, \mathcal{E}_{el}, \mathcal{W}_{el})$. This model includes a set of nodes $\mathcal{N}_{el}$, a set of edges/feeders/cables $\mathcal{E}_{el}$ with associated weights $\mathcal{W}_{el}$, allowing us to simulate the potential impact of electrified vehicle fleets on the energy systems. A complete power grid model would require a topological/connectivity structure and the definition of an admittance matrix for both medium voltage and low voltage distribution systems. However, identifying an accurate structure for the grid can be challenging due to missing or incomplete information. At SUPSI, researchers have developed graph-based methods to reconstruct the structure of the grid, as demonstrated in the work of F. Rosato [8]. These methods leverage publicly available geographical data from street maps, including building locations, sizes, and estimates of energy demand. Using this

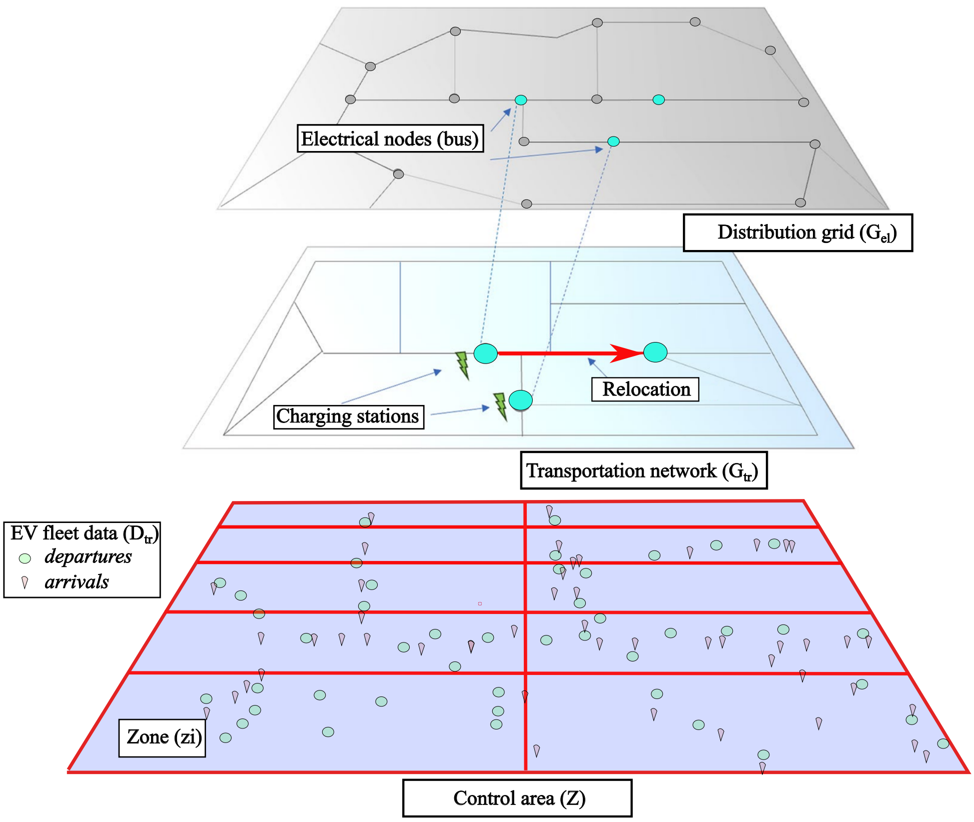

Figure 2: Conceptualization of a methodological framework for interconnected power-transportation grids. Figure adapted from [7].

heterogeneous dataset, a surrogate representation can be reconstructed and adopted as a plausible approximation of the true (unknown) distribution grid admittance. Alternatively, works like that of Corigliano et al. [9] offer methods to infer network structure and the corresponding admittance matrix. These methods are based on geographical information and provide sensible hypotheses on the distribution gird topological structure if this information is missing. The availability of actual electrical measurements from the underlying grid can complement these results, allowing for an orchestrated reconciliation that improves adherence to the real grid. It is important to emphasize that the accuracy of the reconstruction is entirely contingent upon the availability of data. As of the current project stage, energy measurements from the distribution grids are not yet accessible. However, in future developments, if metering data and grid measurements become available, these reconstruction tools will be employed to define a plausible grid $\mathcal{G}_{el}$. Note, that a realistic connectivity and admittance structure is only essential if accurate power flow analysis must be carried out, e.g., voltage profile estimation and load flow analysis. At this stage of the investigation, this level of detail was not strictly required and we therefore preferred to focus on nodal information and on aggregated power localized at charging stations and distributed generators. In other words, we did not account for the fine modeling of connectivity, voltage and flows in the distribution gird and feeders, but rather focused our aim on the nodes $\mathcal{N}el$ and their linkage with the transportation system, e.g., via the presence of charging stations. Additionally, preliminary analysis from our partner E7 demonstrates that the impact of electric vehicle fleets on feeders and power grid balance is moderate in magnitude if small and medium-sized fleets are considered. This observation justifies the decision to neglect a detailed topological structure at this stage. Refer to appendix 7.3 for more details on the economic dispatch modelling and analysis carried out by E7.

2.2.3 Street network and macroscopic traffic model

Several approaches are available in the literature and have been considered, including microscopic, macroscopic, and mesoscopic models as discussed by e.g., [10, 11]. Microscopic models simulate individual vehicle behaviour, offering fine-grained insights into traffic flow, congestion, and dynamic patterns at the street level. However, they can be computationally intensive. In contrast, macroscopic models focus on traffic density, speed, and flow in larger control zones, making them more efficient for network-level analysis. Mesoscopic models strike a balance by considering vehicle groups (platoons) and are well-suited for complex urban traffic scenarios. Recently, there has been growing interest in using probabilistic models and machine learning for traffic model calibration and prediction. For a comprehensive review, see the recent work [12].

SUPSI developed macroscopic traffic models; models that can provide critical information supporting planning in transportation engineering and management of carsharing fleets. A static macroscopic traffic flow model represented as a directed multi-weighted street graph is adopted here and analogously to the power grid, denoted as $\mathcal{G}tr = (\mathcal{N}_{tr}, \mathcal{E}_{tr}, \mathcal{W}tr)$. The network $\mathcal{G}tr$ represents the structure of the transportation system, where nodes in a set $\mathcal{N}_{tr}$ represent intersections (or stations and parking hubs) whilst the edges in $i, j \in \mathcal{E}_{tr}$ represent street segment between nodes $i \in \mathcal{N}_{tr}$ and $j \in \mathcal{N}_{tr}$. A set of weight vectors, $\mathcal{W}_{tr}$, defines key attributes for both nodes and edges in a transportation network, including segment length, expected travel time, maximum speed, and the geographical coordinates of nodes. Figure 3 presents two examples of road networks in Zurich, Switzerland, and Tel Aviv, Israel—two municipalities being examined within the GAMES project (refer to section 2.2). Each edge incorporates a vector of weights, capturing essential statistics and thus defining the static macroscopic model. Examples of features included within this model are the length of streets, their average travel time, and other important data for transportation modelling (for instance, the segments in red have a speed limit exceeding 50 km/h).

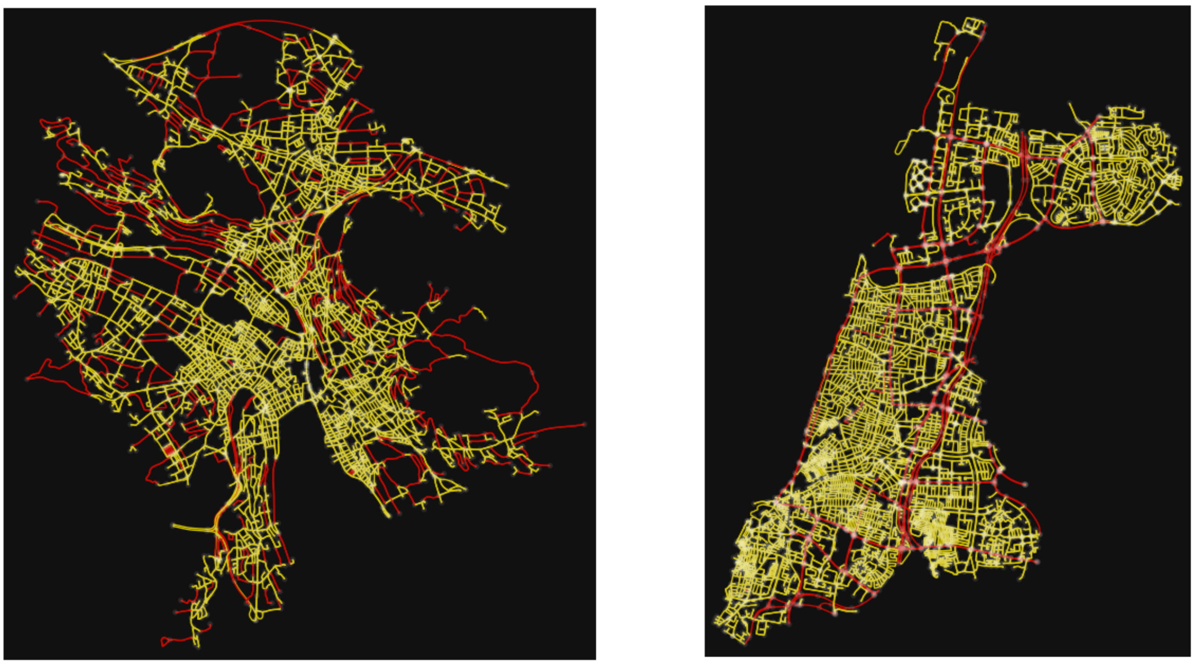

(a) The streets network of Zurich, Switzerland.

(b) The streets graph for Tel-Aviv, Israel.

Figure 3: The streets graph $\mathcal{G}_{tr}$ for two municipalities, the edged printed in dark red color show streets with higher average traveling speeds higher than 50 km/h.

2.2.4 Stations-assigned zones

The grids $\mathcal{G}tr$ and $\mathcal{G}el$ lay within a geographical area of interest denoted by $\mathcal{Z} \subset \mathbb{R}^2$, e.g., a set of latitudes and longitudes enclosing the nodes, reservation data (sources and the target destinations), stations, cars, and users of the carsharing services. For free-floating systems, characterized by more complex mobility patterns, it is helpful to partition $\mathcal{Z}$ in a set of smaller, non-overlapping, control zones $z_i, i = 1, ..., n_z$ such that $\bigcup z_i = \mathcal{Z}$ and $\bigcap z_i = \emptyset$. Note that this approach, often referred to as zoning, can be used to discretize the non-finite set of delivery and return points, and greatly simplifies the analysis of free-floating systems. Many approaches have been proposed in the literature, including uniform zoning, cluster-based zoning [13], and zoning based on Voronoi Cells [14]. In this work, we consider partitions of $\mathcal{Z}$ in non-overlapping station-assigned zones. Each zone is associated with a graph partition $\mathcal{G}_z \subset \mathcal{G}$, and a data set partition $\mathcal{D}_z \subset \mathcal{D}$, and one charging station. If charging station coordinates are not provided, a dedicated procedure to identify plausible stations coordinates and assigned zone is applied (see station allocation problem in the annex). Figure 4 illustrates in the Tel Aviv area an example of zoning and station allocation, where 50 and 10 stations are allocated in the left and right panels, respectively.

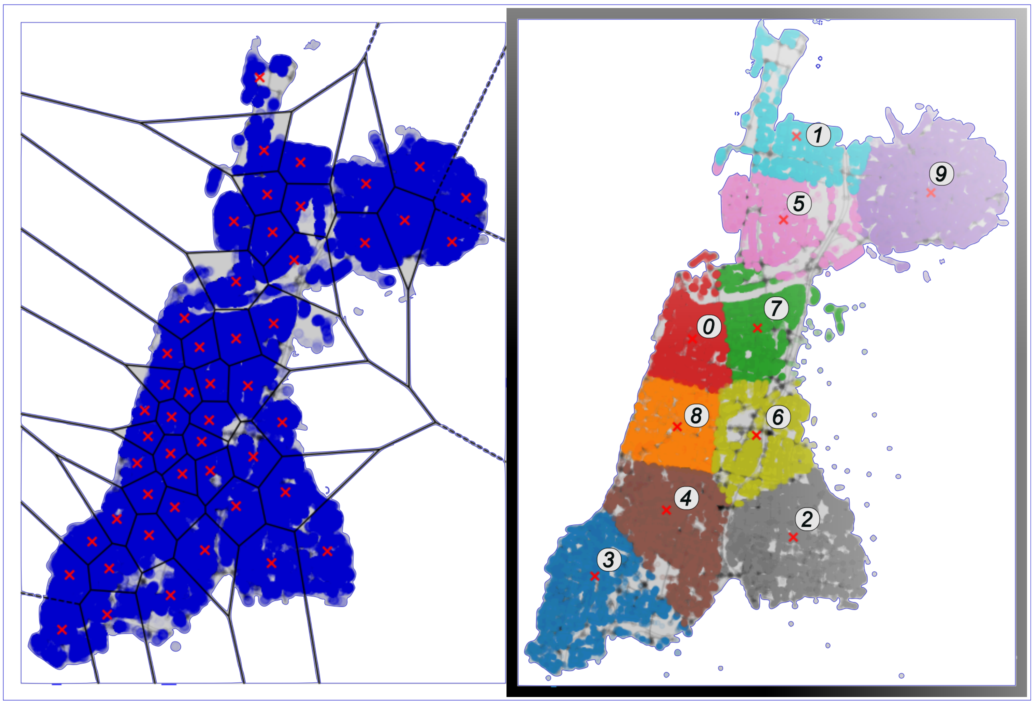

Figure 4: Two examples of $\mathcal{C}$ allocation: 50 station-assigned zones (on the left panel) and 10 station-assigned zones (on the right panel).

2.2.5 Formalize available data

We will now give a mathematical definition to the dataset used in this work. An electric fleet $\mathcal{E}$ and a set of charging stations $\mathcal{C}$ are the essential components that define a carsharing system. A combined evaluation of transportation system and the power grid require linkage of these two systems and the set $\mathcal{C}$ is thus essential.

Operational area and charging station data: The vehicles in the fleet can operate anywhere within a geographical area $\mathcal{Z} \subset \mathbb{R}^2$. For free-floating systems, the control area is divided into n_z disjoint zones, each assigned to a charging station $cs \in \mathcal{C}$. This greatly simplifies treatment and analysis. A charging station $cs \in \mathcal{C}$ has two important properties: a location $x_{cs} \in \mathcal{Z}$ (a coordinate), and a nominal power for the chargers P_{cs}^{BAT}. Other variables can be considered, including the number of chargers available, the voltage information of the energy grid, and the properties of the photovoltaic generators connected to the station, if available. In this project, if the coordinates x_{cs} are not provided, they are assigned using a station allocation Algorithm 2, described in detail in the Annex section.

Carsharing fleet data: The EV fleet is defined by a set $\mathcal{E}$, which contains features, parameters, and variables for each car in the fleet. Formally, we define the carsharing fleet via a set of states given by:

$$\mathcal{E} = \left\{x_{ev}, op_{ev}, SoC_{ev}, \tau_i^{idle}, \tau_i^{return}\right\}_{ev=1}^{n_{ev}}, \tag{1}$$

where n_{ev} is the fleet size and each car $ev \in \mathcal{E}$, has a set of coordinates, $x_{ev} \in \mathcal{Z}$, an operational state, $op_{ev} \in \{0, 1, 2\}$, a state-of-charge, $SoC_{ev} \in [0, 100]$, an idle time, $\tau_{ev}^{idle} \in \mathbb{R}^+$, and a time-to-return $\tau_{ev}^{return} \in \mathbb{R}^+$, both in minutes. The three operational states are given as follows:

$$op_{ev} = \begin{cases} 0 & \text{if } ev \text{ is unavailable} \\ 1 & \text{if } ev \text{ is available, not charging} \\ 2 & \text{if } ev \text{ is available, charging} \end{cases} \tag{2}$$

where $op_{ev} > 0$ if a vehicle is available and $op_{ev} = 0$ otherwise. The idle time is $\tau_{ev}^{idle} > 0$ if the EV has been returned and it is not in use. Hence, if the car is available, $\tau_{ev}^{idle} > 0$ and $\tau_i^{return} = 0$. Rental or relocation activities may change this state, leading to $\tau_{ev}^{idle} = 0$ and $op_{ev} = 0$. Additionally, EV-specific parameters include driving distance at full SoC, d_{ev}^{range}, rated charging power in kW, P^{rd}, brand and model type, battery capacity, E^{rd}, and charging efficiency, η.

Figure 5, taken from the demonstrative cross-sector platform developed (refer to section 2.8), illustrates an example of fleet state for the use case in Tel Aviv. In this example, we display a fleet with 45 electric vehicles and 10 charging stations allocated using Algorithm 2. The blue icons show the fixed locations of charging stations, and some are hosting charging EVs, i.e., $op_{ev} = 2$. The red car-shaped icons show the set of available vehicles $op_{ev} = 1$, while the unavailable vehicles ($op_{ev} = 0$) are gray. Some of the users will reserve the EVs to carry out their trip (displayed in green). An approach to simulate the user's requests and reservations based on historical data will be addressed next.

Mobility data: A data set of mobility events (carsharing events) is a collection of mobility events defined as $\mathcal{D}_{tr} = \{e_i\}_i^n$, where i^{th} event e_i defines a mobility event, a trip or a parking event. Each mobility event can be defined by a set of features such as the one presented in Table 3. An example of mobility event is defined by the following features:

$$e_i = (ev, x_s, x_e, t, \mathbb{1}s, \tau, l....)_i \tag{3}$$

where $ev \in \mathcal{E}$ is the unique identifier of a vehicle, $x_s \in \mathbb{R}^2$ represent the starting latitude and longitude of the event trip, $x_e \in \mathbb{R}^2$ is the end locations for the event (latitude and the longitude), t is the timestamps describing the occurrence of the event, e.g., $t = (t_b, t_s, t_e)$ the event booking time (for trips), trip start time and trip end time. The quantity τ is the duration of the event, and l defines the distance travelled. An indicator function $\mathbb{1}s$ is a Boolean indicator that defines the event type, that is, $\mathbb{1}s = 1$ for a departure event and $\mathbb{1}s = 0$ for a return event.

Energy data: Similarly to the data set of mobility events, we also consider a data set $\mathcal{D}_{el}$ that contains, when available, relevant information from the energy distribution grid, including load data, energy prices, demand forecasts, and other key metrics that influence the operation and management of electric vehicles within the carsharing system. This data set is crucial for understanding the interaction between mobility patterns and energy consumption, enabling more efficient economic dispatch and optimization of electric vehicle operations.

2.3 Results - a comparison between station-based and free-floating systems

We conducted a comparative analysis of statistical models for the station-based and free-floating carsharing use cases, using available (simulated) data pre-trained on raw Mobility data (CH, Zurich) and

Figure 5: Visualisation of simulated reservations in an episodic step and FFEVCS system state. The blue icons: the set of charging stations and plugged-in EVs. Red icons: available free-floating EVs. The grey and green icons present unavailable EVs and carsharing reservations, respectively.

Identifiers:	**Examples**	**Data type**	**Format**
reservation_no	$\{0, 1, 2, ..\}$	integer	list of indices
user_no	$\{0, 1, 2, ..\}$	integer	list of indices
vehicle_no	$\{0, 1, 2, ..\}$	integer	list of indices
station_no	$\{0, 1, 2, ..\}$	integer	list of indices
pv_no	$\{0, 1, 2, ..\}$	integer	list of indices
Coordinates/location:			
station_coordinates	(46.23, 6.921)	(lat,lon)	ISO 6709 List of coordinates
vehicle_coordinates	(46.312, 8.925)	(lat,lon)	ISO 6709 List of coordinates
user_coordinates	(46.312, 8.925)	(lat,lon)	ISO 6709 List of coordinates
county	CH	string	ISO 3166-2 country code
Vehicles parameters:			
brand_name	Tesla	string	brand
model_name	Model S 90	string	model
fuel_type	Electric, Hybrid,	string	fuel
vehicle_category	Luxury, Economy, Van	string	category
driving range	550	float	[km] at full capacity
fuel economy	0.00182	float	SoC % loss per km
charge_power	16.5	float	Peak charge power [kW]
battery_capacity	90	float	Battery capacity in [kWh]
Event data (trips, booking)			
reservation type	Normal, Relocation...	string	{Normal, Relocation, Service }
reservation state	Completed	string	{Completed, Cancelled, Pending}
system create date	28/01/2022 9:12:40	date time	mm/dd/yyyy hh:mm:ss
reservation from date	28/01/2022 11:00:00	date time	mm/dd/yyyy hh:mm:ss
reservation to date	28/01/2022 19:00:00	date time	mm/dd/yyyy hh:mm:ss
reservation coordinates	(46.2, 8.84)	(lat,lon)	ISO 6709
reservation from station	1	integer	in station_no
reservation to station	5	integer	in station_no
cancel date		date time	empty if not cancel
driven km	62.4	float	[km]
drive_time start	28/01/2022 11:35:00	date time	mm/dd/yyyy hh:mm:ss
drive_time end	28/01/2022 18:01:00	date time	mm/dd/yyyy hh:mm:ss
station_no start	1	integer	in station_no
station_no end	5	integer	in station_no
trip_start_coord	(46.2, 8.84)	(lat,lon)	ISO 6709
trip_start_coord	(46.31, 8.91)	(lat,lon)	ISO 6709
SoC start	88.4	float	SoC start [%]
SoC end	62.1	float	SoC end [%]
Revenue and costs data (Mobility):			
revenue_duration	3.5	float	monetary unit per hour [mu/h]
revenue_distance	2.1	float	monetary unit per distance [mu/km]
revenue_booking	5.0	float	monetary unit [mu/event]
relocation_cost	...	float	monetary unit [mu/h] or [mu/event]
service_cost	...	float	monetary unit per hour [mu/h]
revenue_other	...	float	o
Revenue and cost data (Energy):			
revenue_grid_services	...	float	...
cost_charging	7	float	charging cost [mu/%], [mu/event]

Table 3: The table present examples of data features, types, and format expected by the carsharing data sets.

the raw data provided by Autotel (IL, Tel Aviv). Figure 6 presents a summary table of statistical results for the probability distributions of booking and idle times, with free-floating (Autotel) data on the left-hand side (in red) and station-based system data on the right-hand side (in blue). This table provides insights into the characteristics of trip duration and idle times across these two distinct carsharing systems and confirms pre-existing results on similarities and differences between these two systems. Notably, in Tel Aviv the average idle time is of 242 minutes, while in Zurich it is significantly higher, 1014 minutes. Similarly, the mean duration of reservations is only 27 minutes in Tel Aviv, and over six hours in Zurich. Only 5% of bookings last less than one hour for the station-based system. These differences in the statistical properties of the reservations and idle times confirms operational differences between station-based and free-floating carsharing systems (as well as their target customer groups). Station-based systems, like Mobility, exhibit a slower dynamic, with cars picked up and returned at the same station, often used for longer trips spanning multiple days for leisure or business. In contrast, free-floating carsharing systems, such as Autotel in Tel Aviv, are designed for shorter, more frequent trips, with bookings starting and ending within the municipality's control area. Users are often offered contracts for a specified period, allowing them to pick up any vehicle and drop it off anywhere in the network. Consequently, idle times are significantly reduced in free-floating systems, with only 5% of drop-off events resulting in idle times exceeding 10 hours. In contrast, 15% of return events in station-based systems result in EVs remaining idle for at least one day, and 5% remain idle for two days or more.

	Autotel		Mobility	
	(free-float)	(free-float)	(station-based)	(station-based)
	Idle time [min]	trip [min]	Idle time [min]	trip [min]
μ (mean)	242.04	27.23	1014.28	337.69
σ (std)	308.34	41.2	719.24	~300
$q_{.05}$	10.0	8.0	30.0	60.0
$q_{.15}$	21.0	11.0	120.0	120.0
$q_{.85}$	536.0	34.0	1470.0 (~1d)	540.0
$q_{.95}$	950.0	60.0	2820.0 (~2d)	750.0

15% of the idle events have duration lower than 21 minutes

On 5% of occasions when a vehicle is dropped off, it will remain idle for more than 2 days

Figure 6: The table reports the statistics of the probability distributions of bookings and idle times for the FFCS (Autotel) and SBCS (Mobility) systems. We report the mean values, standard deviations and four quantiles for the lower and upper tails.

This analysis aggregates all events, regardless of return and pick-up times, resulting in an average idle time of 242 minutes, combining both longer night idle times and shorter daytime periods. For a clearer view of the systems' dynamics, further conditional analyses (e.g. by hour or location) are recommended. Figure 7 presents joint Kernel Density Estimates (KDE) for trip durations and idle times across the day for free-floating and station-based systems, revealing common usage patterns driven by day/night cycles, weekends, and return patterns. However, free-floating systems show greater dynamism, with shorter trips and lower idle times than station-based systems. This visual comparison highlights the distinct operational characteristics of each carsharing type.

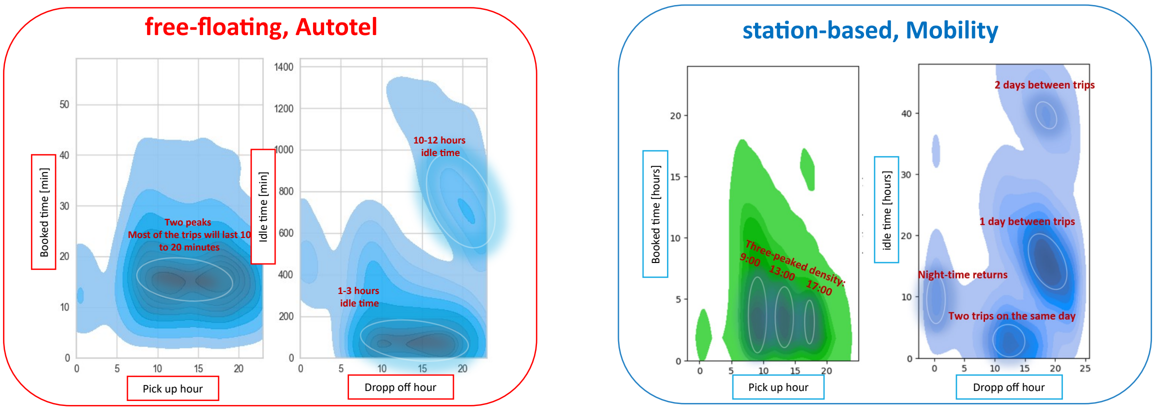

Figure 7: The figure shows joint KDE for trip durations and hour of the day and idle times vs. hour of the day for the free-floating and station-based systems. Note that the distribution is statistically similar due to common usage patterns. However, some differences can be observed in the shape and, in particular, on the scale of the different events. That is, free-floating systems are much more dynamic, characterized by short trips and shorter idle times.

2.4 Forecasting tools and models

Predicting future operational conditions for both the energy grid and transportation network is crucial for effective decision-making. Key indicators, such as aggregated and local power production, as well as the expected number of requests in specific municipalities or neighborhoods, provide valuable insights. Thus, SUPSI developed forecasting algorithms that are essential for estimating the availability of electric vehicles, which can generate potential revenues from grid-related services and enhance carsharing services.

In the following sections, we introduce parametric and non-parametric probabilistic forecasters designed to capture the variability of mobility features, such as demand, across both spatial and temporal dimensions. For simplicity, we define a **deterministic forecaster** for a target variable y as a function $y = f(x)$, which predicts the target based on input features x. Similarly, a **probabilistic forecaster** for spatiotemporal variables is represented as a conditional probability model $F(y \mid x, t)$, which accounts for both spatial coordinates x (and possibly other features) and time t.

2.4.1 Supervised learning

To establish a foundation for our forecasting tasks, we introduce a general framework for supervised learning. In this setup, a sample (x, y) is drawn from a dataset $\mathcal{D}$, where $x \in \mathcal{Z}$ represents the input variables (features), and $y \in \mathcal{Y}$ is the target (dependent variable). The objective is to train a model f that maps x to y, with the goal of minimizing prediction errors. The model f belongs to a class of functions $\mathcal{F}$, which could be, for example, linear models or more complex functions such as neural networks. Given a new input x, the model generates a prediction $\hat{y}$, aiming to approximate the true value y. The learning process involves selecting the best model $f^\star$ that generalizes well to unseen data by minimizing the error over the training set. This can be framed as an optimization problem:

$$f^\star = \arg\min_{f \in \mathcal{F}} \mathbb{E}[L(x, y)],$$

where $L(x, y)$ is a loss function that measures the error, and $\mathbb{E}[L(x, y)]$ is the expected loss across the dataset $\mathcal{D}$. This approach leads to an optimized forecaster that can make accurate predictions based on the provided input.

2.4.2 Probabilistic forecaster

A probabilistic forecaster goes beyond point estimates and incorporates uncertainty into the predictions. It does so by modelling the target variable y as a probability distribution. For instance, the conditional distribution of y given x can be represented as $F_{y|x}$. In this case, the forecaster provides not just a prediction but also a measure of uncertainty associated with that prediction. Probabilistic models can be parametric or non-parametric. Non-parametric models are often preferred as they do not make assumptions about the underlying data distribution, which is useful when the data generation process is not well understood. The goal is to learn a model $\hat{F}_{y|x}$ that approximates the true distribution $F_{y|x}$, thereby capturing both the expected value and the variance (uncertainty) of the target variable. **Non-parametric** and distribution-free models approximate the true distribution without making assumptions about family and parameters of the underlying target distribution. Conversely, a **parametric model** is defined based on fitting a set of parameters, denoted as θ, and represented as $\hat{F}_{y|x;\theta}$. This approach combines machine learning techniques with statistical methods to effectively model prediction uncertainty. It is worth noting that a parametric model can be distribution-free if the cumulative distribution function (CDF) does not belong to a standard distribution type.

2.4.3 Conditional Kernel density estimator (KDE)

Let $(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ be a set of coordinates, e.g., latitude and longitude, of specific mobility events. A mobility event can refer to, for instance, a departure or an arrival event in the transportation grid. A function $\hat{f} : \mathbf{x} \rightarrow [0, 1)$ can be applied to map the the coordinated $\mathbf{x}$ to a probability distribution, thus generating a probabilistic model that applies to the entire set of coordinates in the control area $\mathbf{x} \in \mathcal{X}$. A KDE is applied to obtain this function as, $\hat{f}(\mathbf{x}) = \frac{1}{n}\sum_{i=1}^{n} K_h(\mathbf{x}-\mathbf{x}_i)$, where $K_h()$ is a kernel function with its bandwidth, controlling smoothness of the distribution. This is, in essence, a stationary probabilistic forecaster for the events, e.g., departures and arrivals in our network, and can be directly estimated from data $\mathcal{D}_{tr}$. Note that there are different choices of kernel functions, such as the Gaussian kernel or the Epanechnikov kernel, and the bandwidth can be chosen using different methods, such as the rule-of-thumb or cross-validation, which will finally affect the resulting $\hat{f}$. To account for both time and space variability, a set of KDE models can be obtained by conditioning on time indicators. Let h denote the hour of the day (in 24-hour format). A set of KDE modelling departures and arrivals densities conditional to the hour h when the event occurred are given by:

$$\hat{f}_{\text{dep}}(\mathbf{x}|h) = \frac{1}{n_{\text{dep}(h)}} \sum_{i=1}^{n_{\text{dep,h}}(h)} K_h(\mathbf{x} - \mathbf{x}_{s,i}), \quad h = 1, .., 24 \tag{4}$$

$$\hat{f}_{\text{arr}}(\mathbf{x}|h) = \frac{1}{n_{\text{arr}(h)}} \sum_{i=1}^{n_{\text{arr,h}}(h)} K_h(\mathbf{x} - \mathbf{x}_{e,i}), \quad h = 1, .., 24 \tag{5}$$

where $n_{\text{dep,h}}, n_{\text{arr,h}}$ represent, respectively, the number of departures and arrivals in the control area and at hour h, $\mathbf{x}$ is a location in the transportation grid, $f_{\text{dep}}(\mathbf{x}|h), f_{\text{arr}}(\mathbf{x}|h)$ are estimated densities of departures and arrival events at location x and conditioned on the hour of occurrence h (hour of departure or hour of arrival), and $\mathbf{x}_{s,i}, \mathbf{x}_{e,i}$ represent the i-th starting and ending locations of mobility event.

2.5 Results - forecasting tools

In this section, we present the results obtained using the proposed forecasters on the project use cases. These results illustrate the effectiveness of these models in capturing spatiotemporal variability and provide insights into their performance across different scenarios.

Forecaster	MAE (#ev-a...	MAE (#ev-test)	R2(all)	R2(test)
RandomForestRegression	[3.17, 3.35]	[7.46, 7.94]	[0.981,0.978]	[0.912,0.902]
Mean hourly demand	11.50	-	0.784	-
AdaBoostRegressor	12.22	12.67	0.826	0.806
XGboost	[5.22, 5.5]	[7.786, 8.207]	[0.958, 0.951]	0.90, 0.89

Forecaster	nMA	RMSE (MW)
Holt-Winters (multiple models with double	0.058	6.83
Linear	0.032	3.29
Stack of 96 LightGBM models	**0.028**	**3.1**

Figure 8: Performance scores for the mobility demand forecaster (top) and the power demand forecaster (bottom table).

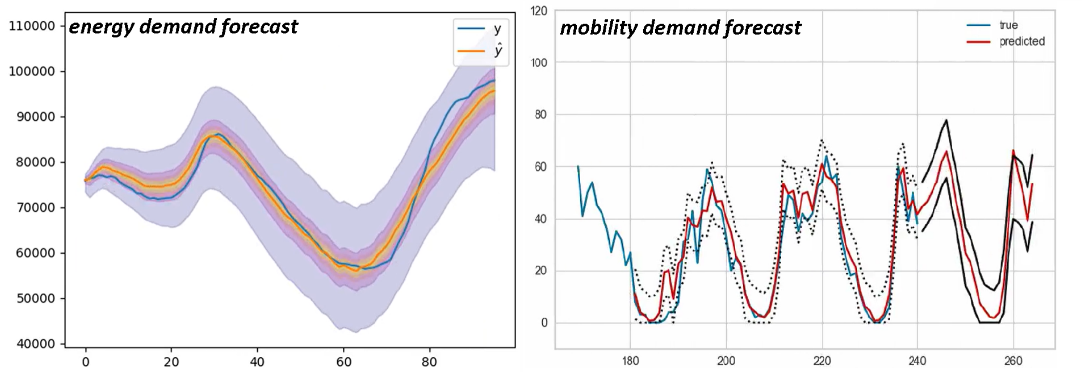

Figure 9: Left: Probabilistic forecast of aggregated power demand in Zurich over a prediction horizon from time t to $t+90$ time steps, including prediction intervals. Right: Probabilistic forecast of mobility demand (total departures) in Tel Aviv, showing observed demand (blue), predicted mean (red), and uncertainty bounds (dashed black for past predictions and solid black for 24-hour-ahead forecasts).

2.5.1 Forecasting aggregated mobility and power demand

We present two examples of probabilistic forecasters that capture the time variability of mobility and energy demand. These models predict future variations in aggregated factors and can be utilised to evaluate user-fleet interaction scenarios. Using the datasets $\mathcal{D}_{el}$ and $\mathcal{D}_{tr}$, respectively, energy and mobility data, we forecast aggregated power demand in Zurich as well as mobility demand in Tel Aviv.

For the **total energy demand** in Zurich, we assessed three predictors:

1. a Holt-Winters model with double seasonality de-trended by weather;
2. a simple linear prediction model;
3. a stacked ensemble of 96 LightGBM models.

For the **aggregated mobility demand** in Tel Aviv, we tested four models:

1. a Random Forest regression model;
2. a simple mean hourly demand model;
3. a AdaBoost regression;

4. an XGBoost regression.

Figure 8 summarises model performance, including metrics such as mean absolute error (MAE), root mean squared error (RMSE), and R^2 regression score. Figure 9 illustrates the results from the best-performing models, with predicted values in red and actual demand in blue, along with uncertainty bands that highlight potential error ranges.

2.5.2 Empirical distributions of mobility features: duration and occurrence

The distributions of mobility features, such as trip duration, idle time, and arrival and departure times, are directly estimated from the dataset, $\mathcal{D}_{tr}$. Non-parametric Kaplan-Meier estimators are used to derive the empirical CDFs for trip duration, $\hat{F}_{\tau}(x)$, and idle time, $\hat{F}_{\tau^{\text{idle}}}(x)$. These distributions are time-dependent; for instance, parking durations tend to be longer in the evening. Thus, time-conditional estimators are also computed. Note that correlations and dependencies between features are not accounted for at this stage. However, space-time dependencies will be accounted for in the next section. Furthermore, simulation-based analysis will later incorporate dependencies by embedding the empirical data directly within the optimisation framework.

The CDF of trip duration, conditional on the hour of arrival, is given by, $\hat{F}_{\tau^{\text{trip}}|h_s}$, and the CDF for idle time conditional on the hour of departure is expressed as, $\hat{F}_{\tau^{\text{idle}}|h_e}$. Similar methods can be used to estimate distributions for booking times and other mobility features. Parametric PDFs were initially used to model idle and booking durations. The data were segmented by hour, and conditional PDFs were estimated for each hour using the *distfit* Python package, a bootstrapping approach with 20 iterations, and a significance level of $\alpha = 0.01$. Outliers, defined as data points beyond the 99.5th percentile, were excluded. Additionally, uniform random noise (within a $\pm$ 30-minute range) was introduced to address measurement inaccuracies. We explored multiple PDF families, including Exponential, Weibull, Log-normal, and Exponential-Weibull, selecting the best fits based on likelihood minimization and information criteria like BIC and AIC to prevent overfitting. Figure 10 presents the best-fitting distributions for the carsharing departures in Tel Aviv (2021 data). The means, μ, and standard deviations, σ, for both the empirical and fitted PDFs are presented ($\hat{\sigma}$ for the fitted standard deviation). Most PDF effectively captured the first two moments of idle and booking times, though some, like the Fisk distribution for hour 17, showed high errors in σ. Additionally, the distribution shape for booking duration was often inaccurate, displaying multi-modality, as illustrated in Figure 11. Fridays and Saturdays have fewer departures in the morning (8:00-11:00). This is due to their status as rest days in Israel. While machine learning models for mobility demand significantly outperformed mean-trend models, none captured spatial demand variability. To address this, we construct a space-time probabilistic forecast model, presented in section 2.5.3.

	Idle time					Booking duration				
hour	μ	σ	$\hat{\mu}$	$\hat{\sigma}$	*best fit*	μ	σ	$\hat{\mu}$	$\hat{\sigma}$	*best fit*
1	15.93	12.63	15.77	14.96	*WeiMin*	5.47	6.26	5.49	10.66	*WeiMin*
2	12.39	8.85	14.25	13.50	*Fisk*	5.3	4.88	5.06	5.81	*WeiMin*
3	13.04	10.37	12.93	13.16	*WeiMin*	5.54	7.37	5.45	8.31	*ExpWeib*
4	12.67	12.85	12.85	16.10	*WeiMin*	5.63	6.03	5.61	5.89	*WeiMin*
5	11.32	11.63	11.43	16.25	*WeiMin*	5.21	7.44	4.97	5.58	*ExpWeib*
6	11.84	12.15	11.9	15.72	*WeiMin*	7.20	7.95	7.21	7.23	*Falife*
7	10.26	21.5	9.05	12.90	*Betapr*	9.18	6.67	9.17	7.68	*WeiMin*
8	10.62	22.51	10.79	17.65	*ExpWeib*	8.31	7.93	7.93	7.42	*Fisk*
9	11.82	156	10.43	20.48	*Fisk*	6.57	5.03	6.56	6.25	*WeiMin*
10	9.569	13.40	10.02	19.56	*WeiMin*	6.08	4.88	6.08	6.38	*WeiMin*
11	10.26	103.3	16.65	17.34	*Fisk*	5.75	5.23	5.83	6.10	*ExpWeib*
12	10.37	14.39	10.62	16.56	*WeiMin*	5.32	4.39	5.39	5.25	*ExpWeib*
13	27.90	238.7	11.53	17.80	*Betapr*	5.48	4.36	5.37	5.84	*WeiMin*
14	12.81	18.66	12.5	17.82	*Falife*	4.69	3.46	4.63	4.50	*WeiMin*
15	14.41	21.86	14.59	19.86	*Falife*	4.53	3.78	4.53	5.39	*WeiMin*
16	15.07	22.36	15.91	19.27	*Falife*	4.7	4.04	4.63	6.14	*WeiMin*
17	35.3	277.8	17.2	18.84	*Fisk*	4.81	4.50	4.92	6.58	*WeiMin*
18	20.12	25.19	18.77	19.57	*Fisk*	4.526	3.97	4.54	5.79	*WeiMin*
19	19.26	15.45	19.7	18.20	*Fisk*	3.83	3.41	4.09	5.65	*ExpWeib*
20	18.53	13.69	13.60	17.29	*Fisk*	3.43	3.04	4.70	3.51	*WeiMin*
21	17.82	12.36	19.79	17.19	*Fisk*	3.35	3.09	3.34	4.85	*WeiMin*
22	16.97	11.09	11.09	16.24	*Fisk*	3.91	8.81	3.83	5.61	*Betapr*
23	15.92	10.24	17.96	16.13	*Fisk*	4.44	4.67	4.48	7.35	*WeiMin*
24	16.42	10.64	16.85	15.98	*Expnweib*	4.22	6.29	4.19	6.01	*Expnweib*

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
1:00	20	23	25	26	38	31	26
2:00	9	9	10	12	24	19	9
3:00	3	4	4	5	12	10	5
4:00	1	2	2	3	6	6	2
5:00	2	1	1	2	4	5	2
6:00	5	6	5	5	6	3	6
7:00	19	19	18	17	14	8	17
8:00	56	57	56	55	34	10	52
9:00	71	72	73	69	53	17	63
10:00	67	67	66	65	63	33	64
11:00	57	55	55	59	71	48	53
12:00	54	53	56	52	79	55	48
13:00	54	56	56	57	81	59	52
14:00	56	58	59	59	78	60	53
15:00	59	59	62	60	72	57	58
16:00	72	72	74	73	68	60	73
17:00	71	69	72	72	64	69	74
18:00	76	79	76	79	61	70	81
19:00	86	86	85	85	56	72	86
20:00	81	83	84	87	50	74	78
21:00	71	75	77	87	48	70	68
22:00	52	55	57	68	55	51	53
23:00	43	45	43	49	59	45	38
24:00	35	39	37	45	44	42	32

Figure 10: Left - Mobility use case: best-fitting distribution and statistics for the idle time and booking duration conditional to the drop-off and pick-up hour. Right - Autotel use case: The average number of trip departures for a typical week in Tel Aviv (Autotel data - 2021)

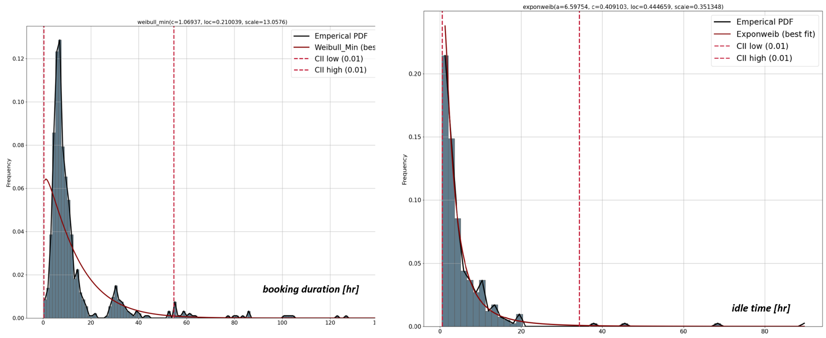

Figure 11: An example of multi-modality for the distribution of the duration of the booking events.

An example of non-parametric modelling is illustrated in Figure 12, where distribution of daily departures (red) and arrivals (blue) for the AutoTel system in Tel Aviv are obtained directly from available data for different hours and week days. The net outflow of vehicles (difference between the departures and arrivals) is also presented in green color on the right panels. This straightforward yet valuable model can be utilized for scenario-based analysis with limited assumptions on the data generating mechanism. The three panels in the top row of the figure present, from left to right, the variability in the number of departures (in red), the number of arrivals (in blue), and the net outflow (in green) for any day (without day distinctions). The lower panels condition the profile by day of the week, revealing different trends: a significant positive outflow peak occurs during weekday mornings (7:00-9:00), while Fridays and Saturdays show a more uniform distribution of net outflow throughout the day.

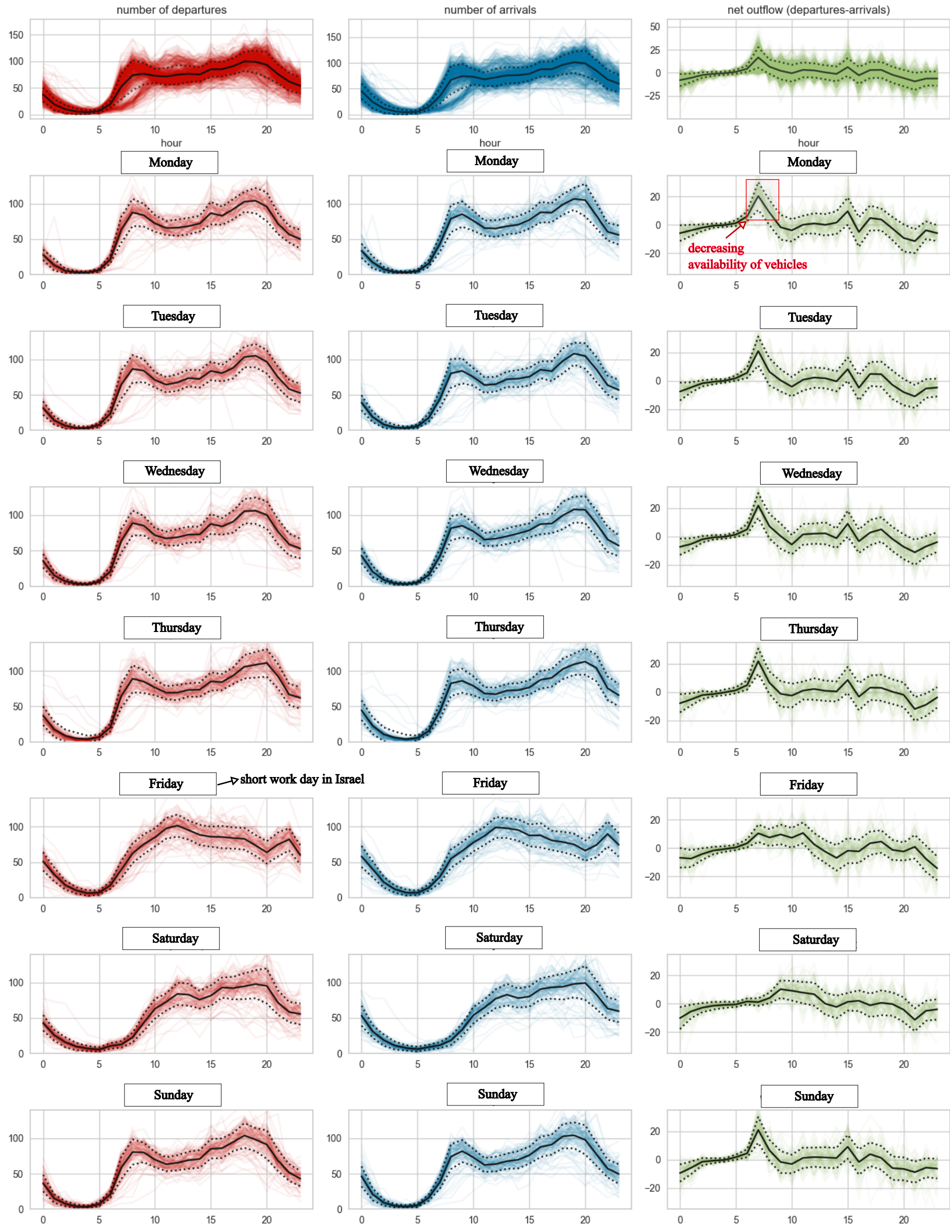

Figure 12: Daily and hourly profiles for the number of departures, arrivals, and the net outflow of Autotel vehicles (difference between the number of departures and arrivals) in Tel Aviv.

2.5.3 Probabilistic characterization of space-time variability

In this section, we introduce simple parametric and non-parametric probabilistic forecasters that can capture the variability of mobility demand across both space and time. For simplicity, these can be expressed as a conditional probability model $F(y|x,t)$ on both coordinates x and time t.

Non-parametric KDE of departures and arrivals: We applied the conditional KDE model presented in the previous section to characterize the spatiotemporal variability of departures and arrivals. This approach uncovers the complex interplay between geography and time, providing valuable insights for decision-making and future optimization. By conditioning on the hour of the day, or hour of the week, we capture interesting patterns in the probability density of departure and return events, also across the spatial dimensions. Figure 13 illustrates the distribution of Autotel's departures and arrivals at three distinct times: 7:00 AM, 1:00 PM, and 9:00 PM. The KDEs of return/arrival events are presented in red whilst the departures/pick-up events are in blue. These reveal interesting evolving mobility patterns, highlighting high-utilization and low-utilization areas (hot- and cold-spots) and variations in car-sharing activity. This visualization and models are fully data-driven and are essential for our non-parametric space-time probabilistic demand model, providing insights into usage patterns and enhancing our ability to meet the needs of Tel Aviv's residents.

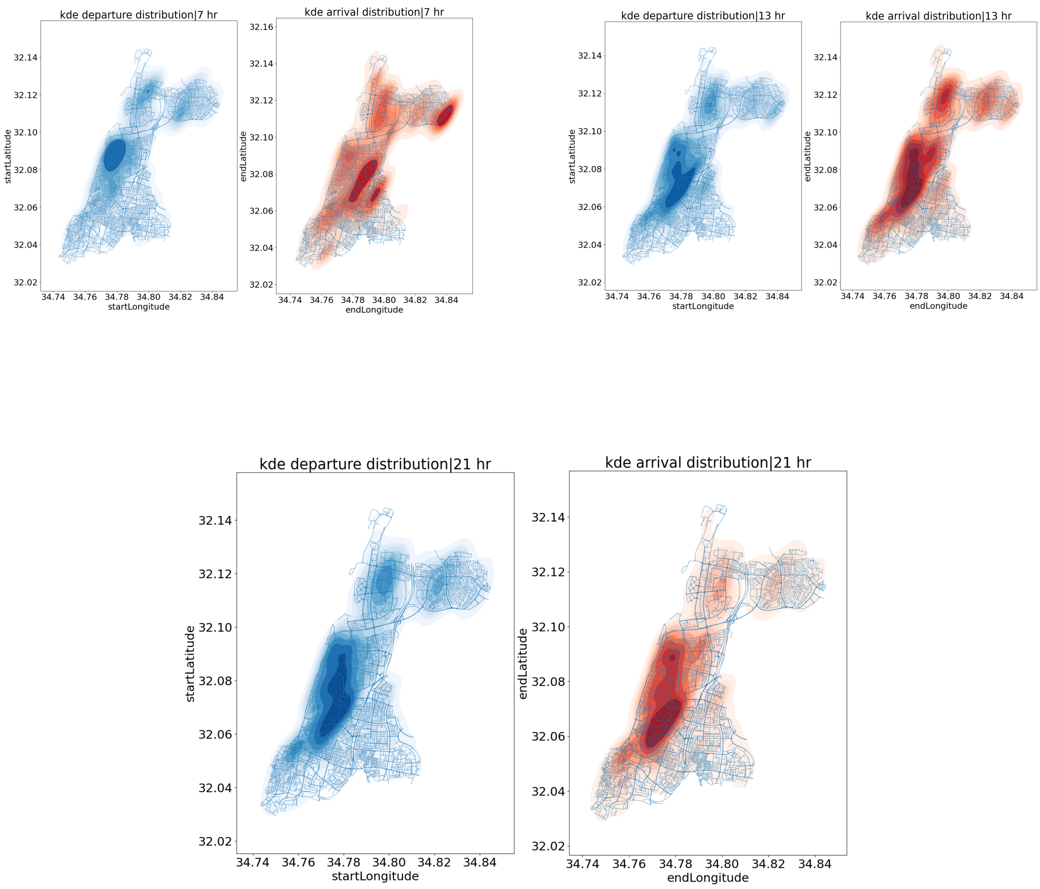

Figure 13: Kernel density estimator of trip departures and arrivals, $\hat{f}_{\text{dep}}(\mathbf{x}|h)$; $\hat{f}_{\text{arr}}(\mathbf{x}|h)$, at $h = 7:00$ AM, $h = 1:00$ PM, and $h = 9:00$ PM.

2.5.4 Scenario generation from probabilistic forecaster

A procedure is defined to generate realistic scenarios, such as mobility demand, for the probabilistic forecasters. This procedure is particularly useful for simulating diverse conditions, enabling the evaluation of forecasting models under various spatiotemporal patterns. Algorithm 1 presents the procedure to generate a simulated demand data set $\mathcal{D}_{\text{sim}}$ for a probabilistic demand forecaster and a mobility data set $\mathcal{D}_{tr}$. The scenario generation is obtained for a time interval $[t, t+\Delta_t]$. First, the number of requests is drawn as $n_t \sim f_{n|\psi_h}$, where $f_{n|\psi_h}$ is the PDF for the number of mobility requests parameterized by ψ_h, i.e., PDF from the week hours in the simulated step. Then, n_t events are generated from the conditional CDF of the mobility event $e_i \sim \hat{F}_{e|h}$, without repetition. The simulated data in $\mathcal{D}_{\text{sim}}$ is then sorted and arranged in chronological order, determining the time difference between consecutive events, δ_t. An empirical CDF is used to generate n_t conditional to the simulated hours, however, this approach can bias the sampling during night hours with few historical requests. To address this issue, random n_t are generated according to triangular mass functions with minimum, maximum, and mean parameters ψ_h, estimated from 168 partitions of $\mathcal{D}$ (representing hours of a week). The triangular family allows exploring a wide range of demand scenarios while avoiding bias due to lack of data and excessive exploitation of extreme cases (as for a uniform distribution).

The sampling procedure and forecasting model is depicted graphically in Figure 14. The spatiotemporal variability in the mobility demand is captured by this approach as illustrated, including the variability of reservations over a week (left panel) and space variability over the operational area in Tel Aviv (right bottom panel). Refer to Algorithm 1 for a detailed implementation. Unlike the events in $\mathcal{D}$, which are historical (realized) reservation events, an event in $\mathcal{D}_{\text{sim}}$ represents a user request for transportation from its location x_s to a destination x_e. The fleet state at time t will be determine if sample demand in $\mathcal{D}_{\text{sim}}$ can be satisfied by the carsharing fleet or by other means. A transportation mode choice model will be applied to determine which of the simulated requests can be accommodated by the carsharing fleet, see Section 7.5.

Algorithm 1 Mobility Demand Sampler

```
Input: D, Δ_t
Output: D_sim^sort
Function FORECASTER.sample(D, Δ_t) D_sim ← ∅ ;      // Initialize empty output set
for h ∈ Δ_t do
    t ← h Get week hour from h
    f_{n|ψ_h}, F̂_{e|h} ← (D, h) ;                   // Probability models
    n_t ∼ f_{n|ψ_h} ;                                // Sample number of trips
    for i = 1, ..., n_t do
        e_i ∼ F̂_{e|h} ;                              // Sample without repetitions
        D_sim ← D_sim ∪ {e_i} ;                      // Append demand scenario
D_sim^sort ← sort(D_sim, δ_t) ;                      // Order chronologically
return D_sim^sort ;                                  // Return the sorted demand
```

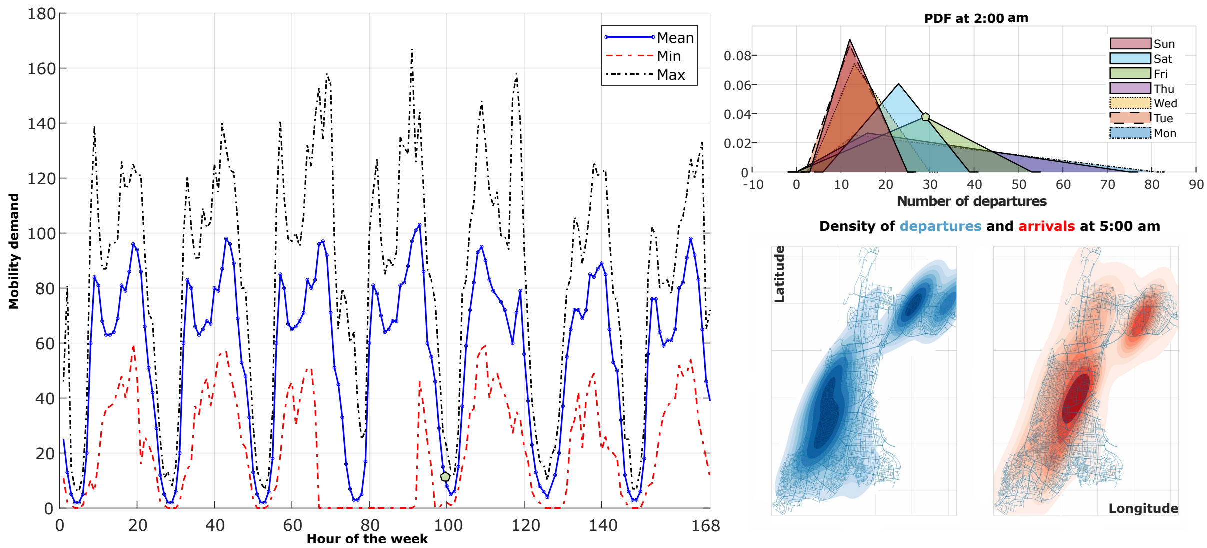

Figure 14: The probabilistic model of spatiotemporal variability of the mobility demand. Left: the parameters ψ_h for the number of departures for a typical week. Top right: triangular densities $f_{n|\psi_h}$ for each weekday at 2:00 AM. Bottom right: space variability depicted by Kernel densities of departures (blue) and arrivals (red) at 5:00 AM.

2.5.5 Validation of the simulated data set

A Monte Carlo simulation was conducted to assess the effectiveness of the sampling strategy and spatio-temporal forecast models. Figure 15 compares the KDE from simulated and raw data, specifically focusing on the joint distribution of booking events and idle times over 24 hours of a typical day. The panels on the left display the saturated KDE from the complete dataset in blue (left: trip duration; right: idle time), while the panels on the right show the simulated distribution from the prediction model in red (left: simulated trip duration; right: simulated idle time). Notably, the simulated and actual data distributions closely align in shape and scale, indicating strong agreement. In addition to visual inspection of the KDE, other performance indicators have also been considered, confirming that the simulation accurately captures the complex space-time variability within the CS system. This demonstrates that the simulation effectively captures the complex relationships between time-of-day and key carsharing system metrics, such as booking and idle times, providing a reliable basis for further investigations and policy assessments.

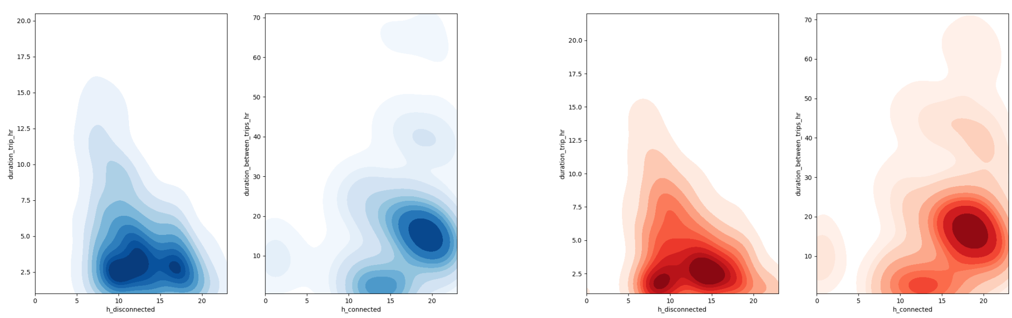

Figure 15: The joint probability density function (Kernel density estimator) of idle time duration $\hat{f}(\tau^{idle})$ and booking duration $\hat{f}(\tau^{book})$ conditional to the hour of the day $h = 1, ..., 24$ (on the x-axis). We present the KDE for the complete data set in blue (left) whilst in the red distributions display the simulated data generated by the proposed simulator sampling (right).

2.6 Optimization of carsharing policies for EV carsharing systems

Operational policies are necessary to manage various aspects of carsharing, including reservations management, customer support, or the scheduling of relocation and furling and charging activities. These policies can enhance system performance and revenues and are thus essential to gaining an edge against competitors and providing better services, e.g., by enhancing the quality of mobility services, reducing fuel costs and emissions, and boosting rental profits. These policies must be efficient and adaptable in controlling the fleet availability, and the time-space varying distribution of cars while ensuring engagement and satisfaction of the customers. Next, we will give a brief overview of important policies for electric carsharing systems, focusing on economic dispatch, relocation, and incentives.

2.6.1 An overview of decision-making policy for EV carsharing

Economic dispatch: methods are critical for optimizing the charging schedules of idle electric vehicles (see recent reviews like [15, 16] for further details). These methods leverage essential grid information, including expected energy demand, pricing structures, local renewable energy production, and the current state of the grid. By effectively balancing these variables, operators can minimize operational costs while enhancing service efficiency, ensuring that vehicles are available in the right place at the right time. ED methods are critical for optimizing the charging schedules of idle EVs. Data-driven methods like [17] can be used to support smart charging, effectively balancing operational costs. Examples include the following ED policies [15, 18]:

1. **Uncontrolled Charging**: This policy allows idle electric cars to be charged whenever plugged into a charging station, e.g., as soon as possible, without optimizing the charging times or rates.
2. **Smart G2V Charging**: This policy optimizes Grid-to-Vehicle (G2V) charging operations, determining the best times and rates for charging EVs based on electricity prices and grid conditions.
3. **Smart G2V and V2G**: This policy optimizes both G2V and Vehicle-to-Grid (V2G) operations, allowing charging, but also returning energy to the grid, when economically favourable, thus providing additional flexibility and revenue opportunities.

Uncontrolled charging and other ED policies heavily depend on the availability of idle EVs at charging stations and the remaining battery capacity. This reliance can conflict with the carsharing operator's goal of maximizing profits through increased vehicle utilization. The profitability of an ED policy is influenced by various operational strategies, including vehicle relocation. There exists an inherent tension: while parked vehicles contribute to energy-related revenues, vehicles in constant use require more frequent charging and maintenance. Although fuelling and idle times are inevitable, intelligent economic dispatch can optimize these periods for energy gains. Moreover, ED policies directly impact service quality by affecting the fleet's state of charge and overall availability.

Relocation Strategies: Especially for free-floating carsharing systems, high operational costs might be due to unpredictable user behaviour, leading to vehicle accumulation in low-demand areas. To address this issue, operators use relocation services to redistribute vehicles or manage EV charging, though these add to operational expenses and affect profitability [11, 19]. Fleet sizes are often minimized due to these costs, but effective relocation strategies, along with the ability of EVs to provide grid services while idle, could help alleviate this. Zhang et al. [20] reviewed optimal relocation approaches, classifying them by mobility demand models (deterministic, probabilistic, or fuzzy), charging strategies, and relocation policies (crew-led, user-based, or mixed). Other interesting modelling concepts for relocation include Markov Chain models [21], frameworks for rebalancing [22], approaches focusing on nighttime crew-led rebalancing operations [23], relocation strategies that exploit spatially-aware demand predictions [5]. Also, polynomial-time algorithms for user-based relocation are studied in [24], adaptive large neighbourhood search algorithm for rebalancing is investigated in [25] and [26], and meta-heuristic multi-agent strategy proposed for centralized management.

In GAMES, relocation strategies are classified into three categories following [20] and later tested on the selected free-floating use case:

1. **Crew-led Strategies**: Operators manage vehicle movements using simulations and optimization techniques, offering reliability but incurring higher costs due to staffing and transport needs.
2. **User-based Strategies**: Users are incentivized to relocate vehicles, lowering costs and environmental impact, though uncertainty in user behaviour and privacy concerns may limit their effectiveness.
3. **Hybrid Strategies**: These combine crew-led and user-based methods, incentivization users while scheduling crew-led relocations to address gaps, reducing costs and uncertainties.

Mobility and energy incentives for EV fleets: Mobility incentives play a pivotal role in encouraging user participation in relocation activities. Such incentives can take the form of adaptive pricing models or mobility credit schemes. For example, bonus minute incentives reward customers who retrieve vehicles from less desirable areas and return them to more central locations, thus helping to reduce idle times. Alternatively, reduced-rate incentives provide users with the flexibility to choose their return locations, although their impact on vehicle distribution may be less targeted compared to bonus minute incentives. Mobility credit schemes have been proposed to promote proactive user behaviour, and various strategies have been introduced for the design and management of vehicle-sharing systems. Incentives can be classified based on their influence on operational and maintenance policies, including:

- **EV-related Incentives:** EV incentives are typically categorized into three types [27]: (1) benefits on acquisition, (2) tax benefits on ownership, and (3) purchase incentives. For carsharing operators, 1)-3) and other commercial grants and incentives play a crucial role in making EVs more accessible and financially viable. Acquisition benefits can significantly reduce the upfront costs of purchasing EVs, making it easier for operators to expand their fleets. Tax benefits on ownership can lower the total cost of ownership over time, providing further financial relief. Lastly, purchase incentives, such as rebates or grants, can encourage operators to invest in more environmentally friendly options, aligning with sustainability goals.
- **Energy Grid-related Incentives:** Energy-related incentives can significantly influence the economic dispatch process by integrating factors such as tariff discounts and dynamic pricing. For example, [28] reviews flexibility and frequency services, presenting a classification of ancillary services from the EV perspective, including voltage regulation, valley-filling, and renewable power matching. By offering incentive rates that promote ancillary services and off-peak charging, carsharing operators can support the power grid (DSO and TSO), optimizing energy costs and enhancing grid stability.
- **Mobility and relocation Incentives:** These incentives are designed to enhance utilization rates and improve the efficiency of reservations. Tariff discounts for demand control, mobility credits, or surcharge adjustments can contribute to increased revenues and enhanced service offerings within carsharing systems, see e.g., [29, 30]. User engagement is essential in effective relocation strategies. Incentives can promote relocation activities by users, possibly leading to cost effective solutions. Example of related works include [29, 31, 32]. For instance, providing credits for using or relocating vehicles motivates users to actively participate, which can help reduce the necessity for crew-led relocation efforts.

2.6.2 Definition of a baseline for free-floating system

Baseline policy is proposed to evaluate the advantage and impact of optimized EV carsharing strategies. Two baselines and one random reference strategy are selected for this, and considered for sake of comparison. The policies are defined as follows

0) **Random:** This simple baseline scenario to test the effect of random incentives on the user participation and system revenues.

1) **Purely crew-led relocation:** This baseline scenario is defined by assigning null destination incentives which in our modelling framework stop user participation to relocation activities. Under this policy, no incentives are offered to users for relocating vehicles, and all relocation activities are managed exclusively by the carsharing staff.

2) **Hybrid and static relocation strategy:** This policy combines crew-led and user-based relocations by setting a positive incentive on user participation (so called destination incentive as defined in the next sections). By applying static (fixed) origin-destination incentives over designated zones, this mixed strategy also encourage users to pick up cars and participate in the relocation process by offering them mobility credits for parking vehicles at designated charging stations when certain conditions are met (e.g., low state-of-charge or proximity to a target location).

These baseline policies provide a base to evaluate effectiveness of different O&M strategies, including advanced solutions based on RL. By comparing purely crew-based relocations to hybrid strategy that involves user participation in relocation activities, the analysis aims to highlight the benefits and trade-offs associated with each approach. Also, for the ED policy, a baseline *charge as soon as possible* strategy is used to establish a reference case. This approach does not incorporate any smart charging techniques or grid-related optimizations to enhance efficiency or reduce costs.

2.6.3 A framework for the optimization of carsharing policies

The literature on optimal operations and maintenance policies for fossil fuel-based car fleets is extensive, with many contributions. However, most of the research focuses on traditional combustion vehicles, which pose different challenges compared to fully electrified fleets. For example, an optimization algorithm for on-demand trip-vehicle assignment in traditional taxi fleets was proposed in [33], using New York City taxi data. This approach demonstrated its ability to accommodate most demand while reducing waiting times and trip delays. Similarly, [34] explored a cooperative strategy for optimal vehicle relocation, focusing on areas with multiple carsharing companies. Tests conducted using data from Fuzhou, China, showed its effectiveness for traditional fuel-based vehicles, but the management of EV fleets was not addressed. Regarding electrified fleets, linear control policies for online vehicle relocation in shared mobility systems were investigated in [35], with a focus on data from public bike-sharing systems. [1] proposed a relocation model for FFCS systems involving both conventional and electric vehicles. However, only a few works have addressed both transportation and electrical services revenues within a unified framework, specifically for managing and optimizing EV fleets alongside relocation challenges. In [36], a RL framework was proposed to balance the distribution of electric vehicles across an operational area. For a broader survey on RL and optimal control, refer to [37], while data-efficient hierarchical RL approaches can be found in [38]. Additional insights into trajectory filtering and curriculum learning within RL can be explored in [39]. The joint effects of relocation and economic dispatch policies are investigated in the GAMES project, both in terms of generating new revenues for fleet operators (through higher utilization and reduced charging costs) and in providing energy flexibility for the grid. We define these two policies as follows:

- **Incentivised relocation policy** $\pi : \mathcal{S} \to \mathcal{A}$. This policy defines relocation actions by mapping the current carsharing system state ($\mathcal{S}$) to a set of relocation actions ($\mathcal{A}$). It determines which EVs should be relocated to nearby stations at specific times. Hybrid user-crew-led strategies are compared to pure crew-led approaches. In hybrid strategies, users are incentivized to participate in relocation, for instance, through mobility credits. The goal of this policy is to maximize factors such as revenue, booking rates, and user satisfaction, by efficiently managing vehicle distribution to meet demand while optimizing operational outcomes.
- **An economic dispatch policy** $\pi_{ED} : \mathcal{S} \to \mathcal{A}_{ED}$ defines a map from the current system state to a set of dispatch actions. This policy maps the current system state to a set of dispatch actions,

focusing on economic aspects related to energy, like managing the charging and discharging of EVs. Optimization of π_{ED} aims to minimize charging costs, ensure a minimum level of charge for future bookings, maximize revenues from ancillary services, or utilize other grid-related incentives. ED policies are designed to make economically sound decisions about the operation of the electric vehicle fleet, balancing cost-effectiveness with service quality.

It is important to note that relocation and economic dispatch policies are interdependent, as each influences and is influenced by the current state of the system. For example, relocating vehicles can affect their availability at charging stations, which in turn impacts the effectiveness of dispatch scheduling. Optimizing grid-related revenues through the economic dispatch policy may sometimes reduce transportation-related revenues generated by the relocation policy, and vice versa. To address this interaction, a joint optimization can be considered, seeking to maximize an overall revenue score that depends on both policies. This joint objective balances both energy-related costs and fleet management efficiency to improve overall system performance. The optimization problem can then be formulated as maximizing the expected revenue from both policies combined expressed as,

$$\max_{\pi,\pi_{ED}} \mathbb{E}[R(\pi,\pi_{ED})].$$

To simplify this complex optimization, it can be broken down into conditional sub-problems. For instance, given a fixed relocation policy, the economic dispatch policy can be optimized to minimize costs by adjusting charging and discharging schedules. Conversely, given a fixed economic dispatch policy, the relocation strategy can be adjusted to improve fleet utilization and reduce relocation costs. Through this coordinated approach, both policies can be optimized in a way that enhances overall performance, balancing economic efficiency with service quality. While a complete exploration of these interdependent policies is beyond the scope of the GAMES project, simulations are used to capture their competitive dynamics, providing valuable insights and potential directions for future research.

2.6.4 Formalize optimal relocation and carsharing policy problem

To attack this complex decision-making problem, we model the carsharing environment as a finite-horizon Markov Decision Process (MDP). This framework provides a structured approach for optimizing polices for decision-making in uncertain environments. The MDP for the carsharing system is defined as a tuple

$$(\mathcal{S},\mathcal{A},\mathbb{P},r,f_{s_0},T,\gamma,)\,,$$

where $\mathcal{S}$ is the state space for the carsharing system, $\mathcal{A}$ is the actions space (set of actions like incentives, relocation, and ED actions), $\mathbb{P}[s_{t+1}|s_t,a_t]$ are transition probabilities, $r:\mathcal{S}\times\mathcal{A}\times\mathcal{S}\to\mathbb{R}$ is a reward function, f_{s_0} is the probability distribution of the initial state, T is the finite episode horizon, and γ is a discount factor. In the MDP formulation, an initial state is randomly drawn from a stationary distribution $s_0\sim f_0$, and actions $a_t\in\mathcal{A}$ are selected at each step according to a policy π_θ. After applying a_t, a new state vector is generated as $s_{t+1}\sim\mathbb{P}[\cdot|s_t,a_t]$ and a reward score computed. The objective of the fleet operator, i.e., a centralized control agent, is to identify a good (possibly optimal) policy that maximizes the sum of discounted rewards $\sum_{t=0}^{T}\gamma^t r_t$.

Actions, incentive policy, and optimality: The policy π_θ defines *origin-destination incentive* (ODI) strategy and each action a_t at a time t are dynamic and spatially heterogeneous incentives to promote user participation in relocation activities and fleet utilization. Formally, a policy $\pi_\theta\in\Pi$ is probabilistic model mapping $\pi_\theta:\mathcal{S}\to\mathcal{A}$, and θ are model parameters. An action vector is obtained from the state of the system s_t as $a_t\sim\pi_\theta(s_t)=\mathbb{P}[a_t|s_t]$ and it is defined as follows:

$$a_t=(a_o,a_d)\ \in[-1,1]^{n_z+1}, \tag{6}$$

where a_o are zone-specific incentives on the origins of trips and a_d is a global destination incentive encouraging relocation of EVs by the users. While destination incentives are dynamically assigned and applied uniformly over the operative area, the origin incentives are both dynamic and zone-specific as they are applied heterogeneously over partitions of an operating area, i.e., over n_z stations-assigned zones. Relocation of EVs by users is encouraged for positive destination incentives and discouraged otherwise. Similarly, if $a_{o,z} > 0$, departures are incentivised from zone indeed by z. This incentive is obtained by applying a discount on the minute tariff or discourage by applying surcharges on the minute tariff for negative origin incentives. The optimal policy seeks to maximize the cumulative reward over time by adjusting policy parameters to maximize fleet efficiency and service quality. The goal is to optimize the expected cumulative value of rewards, representing an ideal balance between energy use, fleet availability, and demand satisfaction as follows:

$$V^{\star}_{\pi_\theta}(s_t) = \max_{\theta} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t r_t(s_{t+1}, a_t, s_t)\right], \quad \forall s_t \in \mathcal{S}. \tag{7}$$

where the state-value function is optimized thanks to the ODI policy π_θ. Problem (7) is analytically intractable because the stochastic transitions probabilities are unknown, which are typically unknown for most real-life applications. Deep RL models can overcome this tractability issue and learn an optimal policy without the need for an explicit model $\mathbb{P}[s_{t+1}|s_t, a_t]$. This is achieved by combining neural networks and cutting-edge RL algorithms, like Proximal Policy Optimization (PPO) [40], Soft Actor-Critic (SAC) [41], Deep Deterministic Policy Gradient (DDPG) [42], and Twin Delayed DDPG (TD3) [43]. Because RL algorithms require exploration of (random) actions for learning, a direct application to safety-critical systems is not advisable. On the other hand, model-based RL allows agents to be trained in a safe virtual environment before deployment, thus lowering the risk of catastrophic failures and economic losses. Hence, a simulation domain will be introduced to simulate the FFEVCS fleet using real data; see Section 7.5 for details.

Monitoring data and observations: A state vector $s_t \in \mathcal{S}$ is used by the policy network to prescribe incentive actions is defined as follows:

$$s_t = (s_t^{glob}, s_{t,1}^{loc}, \ldots, s_{t,n_z}^{loc}) \tag{8}$$

where s_t^{glob} are global states, e.g. including time of the day, average SoC and idle time of the fleet, and $s_{t,k}^{loc}$ are local observations on the station-assigned zones $k = 1, ..., n_z$. Note that s_t includes both global and local observations which are derived from the state variables in $\mathcal{E}$. To test the scalability of different policies, two state spaces are considered as follows:

Case 1: A small state space with 3 variables (hour of the day, day of the week and mean SOC of the fleet) with state vector defined by $s_t = s_t^{glob} = (t_{hd}, t_{dw}, \mu^{SoC})_t$. Note that policies learned using this small state space are expected to be periodic or cyclic to some extent.

Case 2: A larger state space $|\mathcal{S}| = 3n_z + 4$ with four global observations, same as Case 1 with the addition of $\mu^{idle,t}$, and three local states per area, $s_{t,k}^{loc} = (\mu_k^{SoC}, \mu_k^{idle}, n_{ev,k})_t$ with $k = 1, ..n_z$. These are zone-specific statistics for the fleet in the zones.

Note that a policy trained on **Case 1** will not exploit local information and will follow a periodic pattern (daily/weekly) influenced by the average fleet SoC. Differently, **Case 2** incorporates zone-specific information and provides a more detailed view of the fleet's state in different zones, allowing a better dynamic assignment of origin incentives and mobility credits for relocation. While incorporating vehicle-specific data can further enhance the policy, it significantly increases the dimensionality of $\mathcal{S}$ in proportion to fleet size n_{ev}, complicating the learning process. Multi-agent RL could potentially address this issue, but exploring these strategies is beyond the scope of this study and is left for future research.

Objective function - rewards, revenues and costs: A reward signal integrates various sources of revenues and costs, including gains from mobility services, cost for relocation activities and for charging. We define a global reward signal as follows:

$$r_t = (r_t^{min} + r_t^{res}) + (r_t^c + r_t^u) + r_t^{el} = r_t^{\text{tr}} + r_t^{rel} + r_t^{el}, \tag{9}$$

where $r_t^{\text{tr}} \in \mathbb{R}^+$ is the total reward achieved by the carsharing company for the reservations, $r_t^{rel} \in \mathbb{R}^-$ is the total cost of relocation activities, and $r_t^{el} \in \mathbb{R}^-$ defines the electricity cost for charging the fleet. The individual rewards (costs and revenues) all pertain directly to the carsharing company and are assumed to be received simultaneously during each step in a simulated episode. A fixed cost is assigned to each relocation event (personnel hours or contractual fees) and incentive-dependent credit is assigned to users for relocating the EVs. Likewise, mobility revenues are generated at time t due to reservations and increases for a higher number of bookings and duration.

The **reward for the number of reservation events** is defined as follows:

$$r_t^{res} = \rho^{res} \sum_{k=1}^{n_z} n_k^{res}(a_t) \tag{10}$$

where ρ^{res} is a fixed cost in monetary unit assigned to each reservation, i.e., [mu/reservation], and $n_k^{res}(a_t), k = 1, \cdots n_z$ are the number of reservations in the n_z zones at time t. The number of reservations depends on a_t and is obtained by simulating the environment for a simulated mobility demand $\mathcal{D}_{\text{sim}}$.

The **reward for the duration of reservations** is defined as follows:

$$r^{min} = \sum_{k}^{n_z} \sum_{i}^{n_k^{res}} \tau_{i,k}(\rho^{min} - a_{o,k}\rho^{min}) \tag{11}$$

where $\tau_{i,k}$ is the duration of the i-th reservation departing from zone k, ρ^{min} is a tariff per minute, and $a_{o,k}$ determines the incentive on departures in zone $k = 1, \cdots, n_z$. Time indices are omitted to simplify the notation.

The **relocation cost** includes crew activities and user credit. The definition is as follows:

$$r^{rel} = -\sum_{ev \in \mathcal{E}} (\rho^{crew} c_{ev}^{crew} + \rho^{user} c_{ev}^{user}), \tag{12}$$

where ρ^{crew} is a fixed tariff per crew relocation , c_{ev}^{crew} is an indicator function (true if ev is relocated by the crew), ρ^{user} is a dynamic incentive-dependent credit for user relocations, and c_{ev}^{user} is an indicator function (true if ev is relocated by a user). Hence, the total number of crew-led relocations, n^{crew}, and the total number of users-led relocations, n^{user}, at time t determine the relocation cost at time t. Albeit the dependency on a_d is omitted in the equation, ρ^{user} and n^{user} are ultimately affected by the destination incentives. The reader is reminded to the Annex, section 7.5.2, for a detailed description of the relocation model.

The last term in the reward signal is a **economic dispatch reward** given by:

$$r_t^{el} = -\sum_{ev \in \mathcal{E}} \rho_t^{el} E_{ev,t}^+(\pi_{ED}), \tag{13}$$

where ρ_t^{el} represents a dynamic price tariff for electricity per kilowatt-hour (kWh), and $E_{ev,t}^+$ denotes a positive increment in the energy level of an electric vehicle ($ev \in \mathcal{E}$), indicating charging. The energy increment is calculated based on a selected economic dispatch strategy, denoted as a charging policy π_{ED}. For simplicity, we consider a two-tier tariff structure for ρ_t^{el} (night/day) and apply a straightforward "charge as soon as possible" policy. Note that the optimization of π_{ED} is not directly addressed within the current RL framework. Instead, ED optimization is reserved for a post-processing phase, where mobility results obtained through an optimized relocation and incentive policy are examined.

2.6.5 A simulation environment

A simulation environment tailored to investigate the effects of user-based relocation and charging policies is proposed to optimize the ODI policy π. The environment captures relevant dynamics within an urban area, including user movements, carsharing requests, crew-led relocation, and activities related to EV charging and user-based relocation. The simulator can be used to evaluate, simulate, and optimize operational policies for EV carsharing systems. The simulator allows exploration of various realistic scenarios, including user demand for mobility, reservations, and charging patterns across different fleet configurations (sizes, car types, electrification levels). This tool is invaluable for optimizing carsharing fleet management strategies, including relocation and energy dispatch.

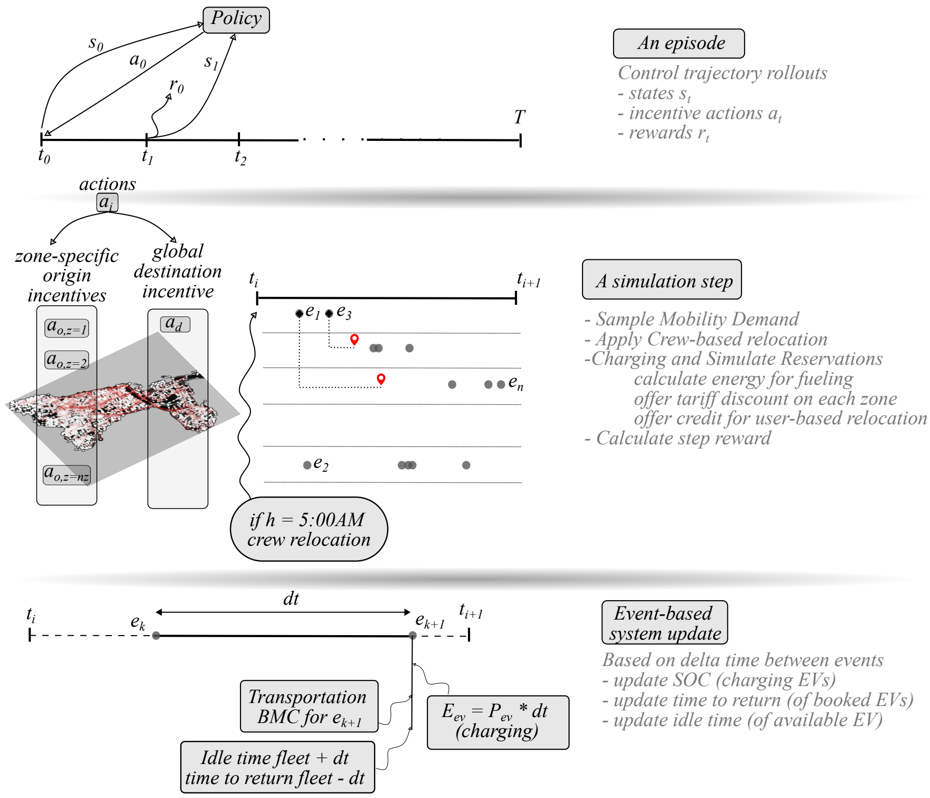

Figure 16: The episode, simulation step and event-based updating of the carsharing system state.

Algorithm 4 defines the core functionality of the environment, including initialization, reset, and episode simulator methods. Figure 16 offers a schematic summary of an episodic loop (top panel) and interaction with the learning agent (policy), a simulation of a time step (center panel), and discrete-event simulation and updating of the system state (bottom panel) Seeking a realistic simulation model, a custom gymnasium environment has been developed for simulating and evaluating different EV carsharing scenarios and test relocation policies. This environment initializes and loads essential data, including historical trip data, the street network information of the area under investigation, and the macroscopic traffic models. A vehicle fleet and charging stations is also defined in the initialization phase with size and distribution defined by a configuration file (see the algorithmic details in the Annex for a detailed description). Figure 5 illustrates an example for one simulated time step for the CS system in the Tel Aviv. For the use case in Zurich, Figure 1 visualizes the initial step of a simulated demand scenario, where black squares represent charging stations on the street network (indicated by blue lines). EV locations align with these stations, and red crosses illustrate user trip requests, highlighting the spatial distribution of demand. These visuals offer a glimpse into the simulation process and how the binary transportation mode choice function influences trip fulfilment in different scenarios.

Relocation models: Crew-led relocation, user-led relocation and hybrid relocations are applied within the simulation model. A relocation crew regularly undertake relocation by choosing vehicle based on a simple rule-based model which determines the EVs that must be moved to a charging station by the crews. In contrast to crew-based activities, user-based relocation encourages customers to park vehicles at designated charging stations. Similarly, to the transportation Binary Mode Choice (BMC) model in eq. (19), acceptance of the relocation offer is modelled and affected by the destination incentive (mobility credit). A relocation offer is proposed only if the SoC of the selected EV is below a minimum threshold. Hybrid relocation policies are obtained by combining these two strategies within the same control scenario, e.g., by modulating destination incentives for the user relocation (see detailed mathematical definition of the simulation and relocation models in the Annex, section 7.5).

Transportation mode choice model: Transportation Binary Mode Choice (BMC) models are a fundamental tool in analysing user behaviour and acceptance of mobility services. Building on previous contributions like [44, 45, 30], a Transportation BMC model is proposed to determine the acceptance of carsharing rides and rentals. Motivated by recent works on acceptable walking distances, see e.g. [45], the willingness to walk of users is modelled by an incentive-dependent maximum walking radius. Origin incentives, a_o, influence a user choice to pick up a car or not. A BMC model determines the user choice $mc = BMC(x_s, a)$ as a function of the starting location of a trip x_s and origin incentive. For $mc = 1$ a carsharing mode is selected, whereas if $mc = 0$ a different transportation mode is used, e.g., bus, bike, train, etc. Figure 17 illustrates how actions impacts willingness to walk at pickup and drop-off points in our model.

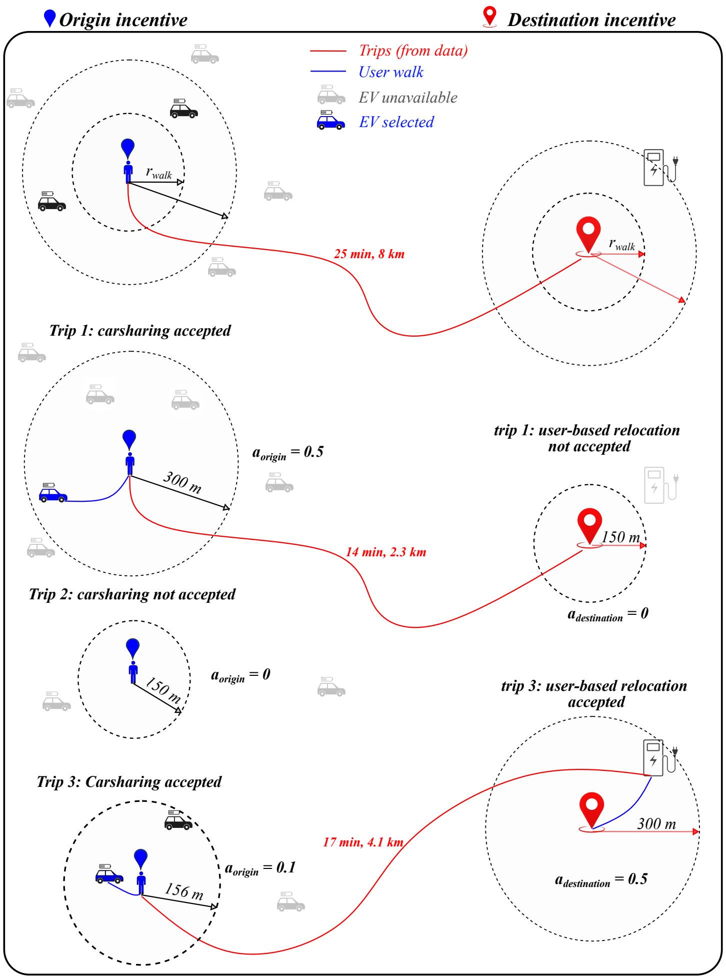

Figure 17: An illustrative example of user-based relocation and effect of origin and destination incentives on willingness to walk.

Apply ED policy and update SOC: The charging of EVs is managed by a battery storage module under the control of an economic dispatch policy, π_{ED}. This policy governs the charging strategy for EVs that are parked and connected to charging stations. A central battery model determines updates to each vehicle's SoC, considering factors like current battery levels, charging efficiency, and vehicle-specific parameters. Charging occurs only when the EV is parked and connected to the grid, with charging power set to the vehicle's rated capacity. The following updating strategy is implemented for plugged in EVs:

$$SoC_{ev}^{new}(t+\delta_t) = SoC_{ev}^{old}(t) + \frac{\pi_{ED}\eta_{ev}\delta_t}{E^{rd}}$$

which includes efficiency of the EV and rated energy capacity to define a new SOC after a charging step. On the other hand, when an EV is used for a trip, its SoC decreases based on the distance traveled as follows:

$$\mathrm{SoC}_{\mathrm{ev}}^{\mathrm{new}}(t+\tau) = \mathrm{SoC}_{\mathrm{ev}}^{\mathrm{old}}(t) - \frac{100 \cdot l}{d_{\mathrm{ev}}^{\mathrm{range}}}$$

Further mathematical details on SoC calculations and the implementation of the ED policy are provided in the Annex.

Discrete-event simulation: A numerical procedure (defined by Algorithm 3 in the Annex), defines a discrete-event simulator of reservations, energy consumed by the EV fleet, and relocation activities at each time step. The simulator starts by initializing reservations, reward function and polices. Crew-based relocation activities are carried out at a predefined time each day, for example, at $h_{\text{crew}} = $ 5:00 AM, and the relocation costs estimated from the total number of relocations by the crew and users. The model processes all the event $e \in \mathcal{D}_{\text{sim}}$ in chronological order, updating $\tau_{ev}^{\text{return}}$ and τ_{ev}^{idle} if time difference between consecutive events is $\delta_t > 0$. The energy demanded by grid-connected vehicles is obtained based on the time delta δ_t, and power injections, and EV states of booked and returned cars updated accordingly. Note that by using δ_t to update the state, the simulator enjoys a significant speed up. That is, the simulator avoids unnecessary operations for simultaneous events when $\delta_t = 0$. A binary transportation mode choice model determines reservations based on the origin incentives a_o and the fleet state, see equation (19). If carsharing is rejected, the model proceeds to the next event. If a carsharing trip is accepted, the event e is added to the reservation set $\mathcal{D}_{\text{res}}$, the nearest among the suitable vehicle is selected, its operational state $op_{ev} = 0$ is set to unavailable, its state of charge adjusted according to the SoC model (Section 7.5.3) and car moved to the new location x_e. Depending on the vehicle's state, users may be offered mobility credits (based on the action a_d) to relocate the car to the nearest charging station at the end of the trip, thus influencing target locations and the relocation cost r_t^u. The procedure concludes after processing all events in $\mathcal{D}_{\text{sim}}$, outputting the total relocation cost r_t^{rel}, simulated reservations $\mathcal{D}_{\text{res}}$, and charging costs r_t^{el}.

2.6.6 Software implementation and hardware specifications

To define the macroscopic traffic model, a street model $\mathcal{G}_{tr}$ is loaded with the numerical framework by sending a query to the python *osmnx* toolbox, while the simulation environment is implemented within the *Gymnasium* framework. The library *stable-baseline3* is adopted to train the RL agents as it offers a set of reliable implementations of RL algorithms in PyTorch. The simulation and training of the RL agents and forecasting models are trained on a standard machine installing a *12th Gen Intel(R) Core(TM) i7-1265U 1.80 GHz* processors and *32.0 GB RAM*. A public repository on GitHub, accessible through (GAMES-GitHub-repository), was established in order to facilitate the sharing and advancement of data analytic tools, forecasting models, and simulation models and results. This repository can serve as a centralized platform for researchers and practitioners to access various tools developed as part of the GAMES project. The repository includes comprehensive documentation, example codes, and data sets that allow users to replicate analyses and implement their own modifications. Figure 18 illustrates the main page of the GitHub repository, showcasing the GitHub project. This initiative not only enhances collaboration within the research community but also encourages transparency in methodologies and findings. By making these tools publicly available, SUPSI aims to promote further research and development in the field of electric vehicle carsharing systems and related analytical techniques.

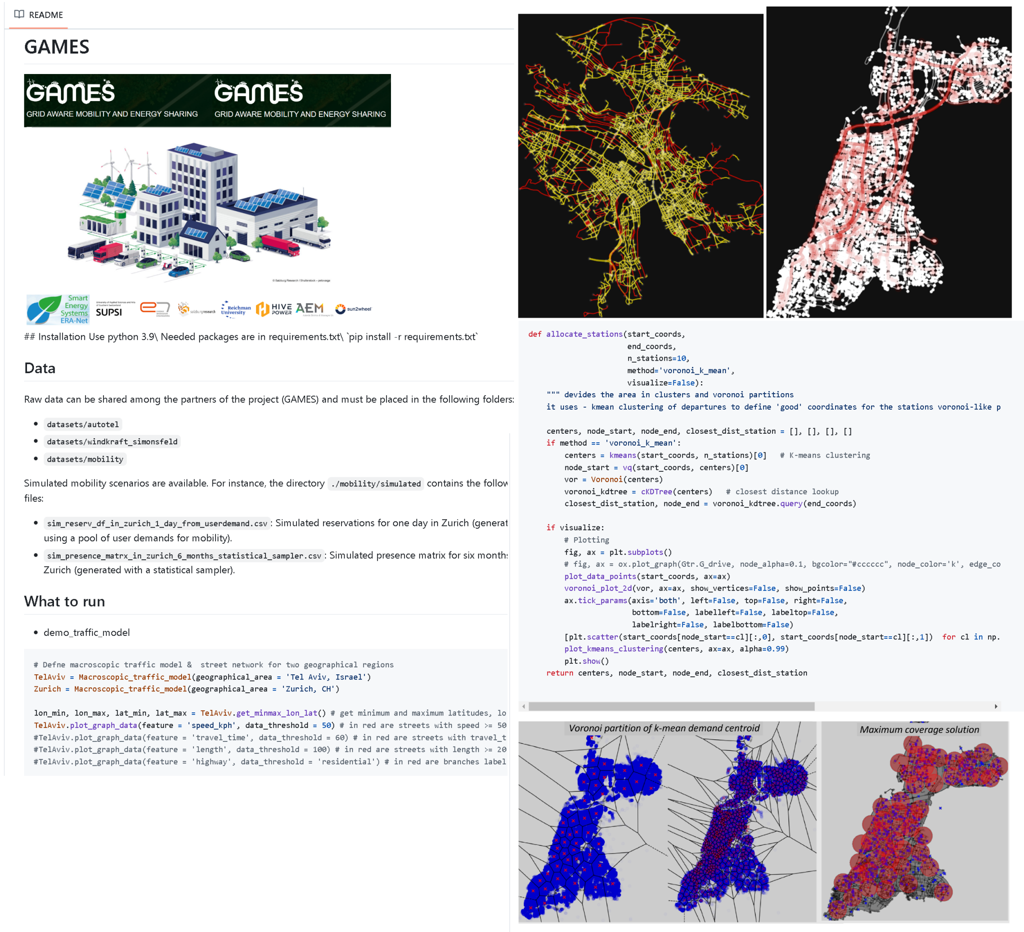

Figure 18: The GitHub repository for the GAMES project.

2.7 Results - optimization of carsharing policies

RL policies and training parameters: We investigate the capability of the four centralized RL strategies in learning dynamic incentives π_θ for effective FFEVCS relocation and management. Table 4 reports the parameters for the four considered learning strategies, i.e., SAC, TD3, PPO, and DGPP. The RL parameters include the learning rate, batch size, activation function, replay buffer (if used), and architectures of the policy defining π. We found that these default (dense fully connected feed-forward) neural network architectures, with the number of hidden nodes given in Table 4, achieve reasonably good results. Hyperparameter optimization could be considered to further improve the efficacy of the agents; however, this is not considered in this work to avoid over-complicating the study.

2.7.1 Simulated weekly profiles and potential rewards in energy and transportation

Crew-based relocation without incentives for users participation: First, we showcase the capability of the simulation environment and investigate its ability to capture relevant mobility patterns. The goal is to assess the model's capability to accurately replicate weekly and daily trends for the Free-Floating CS system and evaluate trends in performance scores potentially related to revenue sources from energy- and transportation-related services. These metrics are pivotal in understanding the potential for grid-related revenues, e.g., DSO flexibility provisioning, smart charging economic dispatch, etc. For this, a Monte Carlo analysis is performed by applying a baseline policy without incentives, i.e., $a_{o,k},\ k = 1, ..., 10$ and $a_d = 0$ and applied on the environment using 2021 data from the CS system in Tel Aviv (Autotel). Fifty independent episodes with horizon $T = 720$ and $\Delta_t = 1$ are considered (30 days with steps of

Reinforcement Learning Parameters				
	SAC	TD3	PPO	DGPP
Hidden Architecture	[256, 256]	[256, 256]	[64, 64]	[256, 256]
Activation	ReLU	ReLU	ReLU	ReLU
Optimizer	Adam	Adam	Adam	Adam
Learning Rate	0.0003	0.001	0.0003	0.001
Target Update Steps	T	T	2048	T
Batch Size	256	100	64	100
Discount Factor (γ)	0.99	0.99	0.99	0.99
Replay Buffer Size	10^6	10^6	-	10^6
Update Coefficient	0.005	0.005	-	0.005
a_t Noise	-	$\mathcal{N}(0, 0.2)$	-	$\mathcal{N}(0, 0.1)$
Clip Range (ϵ)	-	-	0.2	-
GAE (λ)	-	-	0.95	-

Table 4: Reinforcement learning configuration parameters. See default parameters in *stable-baseline3* for further details.

1 hour). Because credit is not assigned for user relocation, only crew-led relocation is carried out. The average computational cost per episode (with $T = 720$) was about 35 seconds. Note that the computational cost can increase to about 2 minutes per episode, as it depends on the computational cost of the policy and the average number of reservations per simulated step. When the episode runs, the environment is reset to its initial state, e.g., set $t = 0$ Sunday midnight. At $t = 0$, vehicles are assumed to be plugged in at the charging stations, and their SoC is randomised between 70% and 100% following a uniform distribution. This is done to explore various operational trajectories by applying a random noise to the initial state of the system. We assume crew-based relocation is carried out once a day at 5:00 AM, and the electricity price per kW follows a bi-level tariff (day/night). Conversely, the cost for crew-led relocation occurs at the designated hour h_{crew} at 5:00 AM, depending on the number of cars to be relocated. Figure 19 shows weekly trends resulting from this analysis and key mobility scores that can be gathered from the fleet. Black lines represent aggregated scores over the hours of a week (168 steps). The green curves in the top panels show electricity costs and relocation costs. As expected, crew relocation costs show regular spikes at 5:00 AM on all days, whilst user relocation costs are zero by definition. This is due to a lack of participation in the relocation activities when $a_d = 0$. Interestingly, because of the designated relocation hour h_{crew} and EVs charging pattern follows a charge as soon as possible policy π_{CS}, this yields a recurrent peak in the energy consumption, occurring daily in the first half of every morning between 7:00 AM to 10:00 AM (see top right panel). This is due to a combined effect of crew-based relocation at 5:00 AM and high electricity prices during daytime hours (the highest tariff of 0.84 Shekel/kWh from 7:00 AM). The central panels in the figure illustrate the average idle time, time to return, and availability of carsharing EVs, i.e., a percentage of idle cars. The bottom panels show the number of reservation events, the mean SoC of the fleet, and the total duration of the reservations at each step. This pattern aligns with historical data, where many reservations are during the day and last less than 30 minutes. The average availability of vehicles generally exceeds 90%, with a large night-day fluctuation, and the average time to return is less than 10 minutes for the whole fleet. This further confirms existing results on the potential for carsharing for V2G operations and ancillary services to the distribution system operator. The number of reservations fluctuates significantly between night and day hours, i.e., about 5-10 reservations per step (hour) and only 0-3 during night hours. This simulation demonstrates the ability to model realistic mobility patterns and operational dynamics for a free-floating EV system in an urban setting like Tel Aviv.

Hybrid relocation with users participation: We repeated the simulation study by applying relocation and mobility incentives and focusing on transportation-related performance scores while applying incentives to boost relocation and fleet utilization. Following this simulation, we collected the average number of idle cars, idle time in minutes, and mean remaining trip durations at each step. Notably, the fleet utilization rate increased substantially compared to the case without user participation and incentives. Also, it is important to highlight that the availability of parked vehicles never dropped below 70 per cent at any given hour throughout the simulations. This high availability of parked EVs (albeit for a possibly short time, at least 70 cars are parked at any given hour) could be exploited, strategically, for DSO-related services and smart charging in between reservations. This assuming that a high percentage of the unused EVs are gird connected. Additionally, we tracked metrics related directly to transportation services, such as the total number of bookings within the last hour, distance traveled in km, and the total duration of reservations in minutes, which are instrumental in assessing revenues generated from transportation services. The results of this simulation study are presented in Figure 20, where the x-axis represents the step ID (hour of the week), the red curves are grid-related scores for the simulated episodes, and the light blue envelopes are scores related to transportation services. The solid black lines delineate the mean trends across all episodes. Notably, there is a pronounced daily trend in the availability of parked EVs, with varying degrees of variability observed on different weekdays (Sunday and Friday as expected for Tel Aviv). It's worth highlighting that the mean idle time on Monday is notably high, exceeding 500 minutes. This is due to the setup of the simulated episode, which commences at 00:00 on Sunday and already assumes an average idle time of 300 minutes for the EVs.

This first comparative study provides valuable insights into the dynamics of EV fleet performance in Tel Aviv, shedding light on both grid-related and transportation-related revenue aspects. These findings serve as a foundational basis for further analysis and optimization strategies in the context of electric vehicle mobility services, with the high availability of parked vehicles presenting a unique opportunity for enhancing DSO-related services and smart charging strategies between reservations.

2.7.2 Effect of origin incentives on relocation

A comparison between a purely crew-led relocation strategy and a mixed policy (user-based together with crew-led activities) is proposed. Similarly to the previous section, MC evaluation of the policies is carried out over 50 independent episodes with horizon $T = 720$ and $\Delta_t = 1$. The purely crew-led relocation strategy is considered by setting $a_t = (0, \cdots, 0)$. Mixed relocation strategies are investigated for non-zero destination incentives, $a_t = (0, \cdots, 0, a_d)$, where various values $a_d > 0$ are investigated in the range [0,1]. Crew-based relocation is carried out regularly in all cases once a day but for a zero credit, $a_d = 0$, the user will not accept to carry our relocation. The results of this analysis are reported in Table 5, and the expected sum of crew relocation cost, $\mathbb{E}[\sum_t^T r_t^c]$, users relocation cost, $\mathbb{E}[\sum_t^T r_t^u]$, charging costs, $\mathbb{E}[\sum_t^T r_t^{el}]$, and reservation revenues $\mathbb{E}[\sum_t^T r_t^{tr}]$ are presented with a factor of 10^3 Shekel. The coefficients of variation (CV) are also reported showing the score variability across episodes.

	Crew Reloc. Cost	**User Reloc. Costs**	**Fueling Cost**	**Transport Revenue**
a_d	Mean (CV)	Mean (CV)	Mean (CV)	Mean (CV)
0.0	-13.4 (7.87%)	-0.00 (0.00%)	-1.55 (6.26%)	105.40 (3.10%)
0.1	-13.3 (7.21%)	-0.02 (27.26%)	-1.57 (6.39%)	106.55 (3.79%)
0.2	-13.3 (8.09%)	-0.06 (22.50%)	-1.61 (6.87%)	106.59 (3.82%)
0.3	-12.4 (9.04%)	-0.13 (23.37%)	-1.60 (6.10%)	106.66 (3.87%)
0.4	-11.8 (8.80%)	-0.22 (16.81%)	-1.59 (6.66%)	105.64 (3.90%)
0.5	-10.5 (9.81%)	-0.36 (16.60%)	-1.66 (7.16%)	105.52 (4.14%)
0.6	-9.39 (12.24%)	-0.56 (15.15%)	-1.63 (5.81%)	105.07 (3.81%)
0.7	-8.15 (13.64%)	-0.81 (12.17%)	-1.67 (5.36%)	105.34 (3.26%)
0.8	-8.26 (17.30%)	-0.91 (13.90%)	-1.64 (6.50%)	104.60 (3.37%)
0.9	-7.97 (16.73%)	-1.04 (13.81%)	-1.68 (5.39%)	104.78 (3.67%)
1.0	-7.98 (13.83%)	-1.17 (12.30%)	-1.70 (5.98%)	105.88 (4.12%)

Table 5: Effect of destination incentives on the expected revenues and costs (and their coefficients of variation) of the FFEVCS system. Results are obtained for 50 episodes and rewards reported with a factor of 10^3 Shekel.

The user relocation cost is zero when $a_d = 0$, i.e., a purely crew-led relocation policy is applied. As the destination incentive a_d increases, the crew relocation cost decreases whilst the cost of relocation by the users increases, i.e., these costs shift from the personnel to users. The transportation revenues remain stable at approximately 105'000 Shekel/month, hence not showing an evident dependency on

a_d. On the other hand, the charging costs positively correlate to a_d, increasing from about 1550 to 1700 Shekel/month for $a_d = 1$. This trend is probably due to longer charging times when user-led relocation becomes predominant. Specifically, users will occasionally connect EVs late in the evening, and most cars will have all night to complete the charging. Conversely, personnel are assumed to relocate cars only once daily at 5:00 AM, leaving just a few hours for charging the EV before rush hour.

Complementary to the previous analysis, the combined effect of origin and destination incentives is investigated, i.e., mixed crew-based and user-based relocation strategy is considered together with a 50% discount on the minute tariff to boost departures. Twenty simulated episodes with finite horizon $T = 2400$ (100 days) are considered and average key performance scores (per-step and per-episode) are reported in Table 6. The results show that the mixed relocation strategy combined with origin incentives improves the expected revenues, especially compared to the baseline without origin incentives. Because of the origin incentives, vehicles are more often booked, yielding higher transportation rewards and a reduction in the average idle time, respectively, 105,000 to 141,000 Shekel/month, and from 1139 to only 414 minutes. However, we have observed higher charging and relocation costs due to increased fleet utilization. While operational costs have increased, the rise in transportation revenues is greater but yet sub-optimal. To further maximize profits, a mixed and adaptive strategy informed by RL agents will be presented in the next section.

Strategy	Only Crew	Mixed	Mixed
Destination a_d	0	0.5	0.5
Origin $a_o = (a_1, ..., a_{10})$	$(0, ..., 0)$	$(0, ..., 0)$	$(0.5, ..., 0.5)$
Mean score (and σ) per-step			
Reservations	3.7 (2.8)	3.7 (2.9)	10.2 (6.4)
Distance [km]	23 (18.9)	23 (18.7)	61 (41.0)
Time driven [min]	85 (79.0)	84 (78.5)	232 (173.0)
Availability [%]	96 (3.0)	97 (3.2)	91 (6.6)
Time to return [min]	0.7 (1.0)	0.7 (1.0)	1.8 (2.0)
Idle time [min]	1122 (254.2)	1139 (262.5)	414 (134.0)
Fleet SoC [%]	67 (12.9)	69 (11.8)	52 (11.6)
$\mathbb{E}[\sum_t^T r_t]$ **(and σ) per-episode**			
Number of reservations	2.7 (0.1)	2.7 (0.1)	7.3 (0.2)
Revenue for duration	104.8 (4.5)	103.4 (3.8)	141.8 (3.6)
Cost of Charging	-1.6 (0.1)	-1.7 (0.1)	-4.9 (0.2)
Cost for Crew	-13.2 (1.0)	-10.4 (0.9)	-27.5 (1.6)
Credit to Users	0 (0)	-0.4 (0.0)	-2.4 (0.2)

Table 6: Average step metrics for the baseline incentive policies. Mean values and standard deviations of the episode rewards (in 10^3 Shekels) computed for 20 episodes of 30 days each are also presented.

2.7.3 Optimized RL agents and comparison with baselines

The reinforcement learning agents have been trained using SAC, TD3, PPO, and DDPG algorithms, with a maximum training size of 500,000 steps. To assess scalability, two state spaces of different sizes have been employed: case 1 with $|\mathcal{S}| = 3$ and case 2 with $|\mathcal{S}| = 34$, utilizing both global and local observations. This setup yields a total of eight agents, whose performances are compared against various baselines, including a random policy with $a_{t,i} \sim \mathcal{U}(0, 0.7)$, a purely crew-based relocation strategy, and mixed relocation strategies with static incentives ($a_t = (k, ..., k)$, where $k \in [0.1, 1.0]$). The agents are tested and validated over a 12-week horizon ($T = 2160$) to estimate expected cumulative rewards in thousand Shekels, averaged over 20 independent episodes. This involves running the algorithm twenty times and averaging the sum of revenues for each episode. The results of this comparison are summarized in Table 7 whilst Figure 21 shows the learning curves for the eight agents. Normalized rewards on the y-axis versus computational time (x-axis). In the extreme scenario of $a_t = (1.0, \ldots, 1.0)$, the fleet operator maximizes fleet utilization with an expected 10,920 reservations over the simulated month and

31,680 over 12 weeks. However, offering a full waiver of the minute tariff leads to a total loss in revenue due to null minute tariff revenues coupled with substantial relocation and charging costs. In contrast, the zero incentive policy $a_t = (0, \ldots, 0)$ yields better revenue, averaging 7,470 reservations for the 12-week episode, but results in lower fleet utilization and does not take advantage of dynamic incentives for user-based relocation and rebalancing. Actions set to $k = 0.5$ provided the best results among the considered baselines, though only slightly better than a random positive incentive policy. Note that charging costs and relocation costs increase with k (for higher mobility incentives). This is due to greater fleet utilization and energy demanded for charging the fleet. In all cases, the SAC, TD3, and PPO agents demonstrate significant improvements over the considered baselines, achieving higher revenues (number and duration of the reservations) and reducing relocation and charging costs. On the other hand, the DDPG agent for $|\mathcal{S}| = 3$ achieves a lower performance, possibly due to instabilities during the training phase. However, this issue was not observed for the case $|\mathcal{S}| = 34$. As a figure of merit, let us focus on the TD3 policy with $|\mathcal{S}| = 3$. The carsharing operator, if the TD3 agent is used, is achieving an average revenue of 339 k Shekel (about 113 k Shekel/month). This is much higher than the baseline case, where only crew-led relocation is considered (242 k Shekel or about 80 k Shekel/month) and the best-performing baseline (277k Shekel or 92 k Shekel/month). By including a larger number of observations, $|\mathcal{S}| = 34$, the average revenue grows to about 349 k Shekel, i.e., an additional 3k Shekel/month. Hence, the agents trained with the larger state space achieve higher scores compared to the $|\mathcal{S}| = 3$ counterparts. Hence, incorporating more observations and local information in the state space led to better policies. However, this improvement in perspective of the additional computational costs may be considered small. This suggests that while adding more states may capture comprehensively relevant system dynamics, the complexity introduced may not always translate into significant performance gains. Hence, it is important to find a balance between the complexity of the state space and the performance of the agents. Overall, the proposed RL framework and trained agents have proven proficient in reducing costs and boosting revenues, demonstrating their utility in promoting user-led relocation activities. The RL policies, particularly those trained with SAC and TD3 algorithms, significantly reduced relocation costs while maintaining high revenue from transportation services. This indicates the agents' ability to incentivize user behaviours that lead to effective fleet utilization. The simulated results underscore the effectiveness of RL-based approaches in addressing the complex optimization challenges in FFEVCS systems.

2.7.4 Distribution of actions across zones and time

Figure 22 illustrates the distribution of incentive values in a_t selected by the SAC agent over the simulated episodes, conditioned on the hour of the day, i.e., $f(a_{t,k}|t = h)$ with $h = 1, .., 24$ and $k = 1, .., n_z + 1$. The first 10 panels represent the origin discounts in the action vector a_t, while the last panel (bottom right) shows the global destination incentive. The blue lines indicate the control trajectories and actions taken each hour, thus graphically highlight the overall distribution of actions over a typical day. The red lines show the mean values plus and minus one standard deviation for the aggregated actions at that hour, providing a clear summary statistics and visualization of the policy dynamics and space variability. Notably, some incentives exhibit a clear temporal pattern. For instance, high incentives for departures are generally applied in the afternoon, especially in area 6, where from 14:00 to 20:00, the incentives are close to 0.8 with very little variance. In contrast, applying high discounts on departures from area 4 does not significantly boost fleet utilization, particularly in the evening hours when the distribution of actions is mostly flat and close to zero incentives. Origin incentives from area 4 in the early morning hours, between 4:00 and 7:00 AM, appears to be preferred by the SAC agent. Similarly, the distribution of incentives in area 3 shows a less clear temporal pattern and much higher variance, suggesting a more exploratory and less consistent approach in this zone. This demonstrates that the SAC agent may be adapting its strategy based on varying demand patterns and operational constraints across different zones and times of the day. An optimal policy must therefore account for both temporal and spatial variability and uncertainty that characterize mobility demand, energy needs, and the need for relocation activities of grid-couple FFEV carsharing systems.

	Expected Revenues/Costs						
	All	**Mobility**	**Minutes**	**N. Res.**	**Charging**	**Rel. (Crew)**	**Rel. (Users)**
Baselines							
Random	273.93	370.18	354.57	15.60	-11.07	-81.76	-3.41
$k = 0$	242.43	294.27	286.81	7.47	-5.14	-46.70	0.00
$k = 0.1$	229.14	282.21	274.33	7.88	-5.45	-47.50	-0.11
$k = 0.2$	240.09	300.92	291.46	9.46	-6.58	-53.86	-0.39
$k = 0.3$	251.99	326.31	314.64	11.67	-8.17	-65.16	-0.99
$k = 0.4$	272.15	373.85	358.33	15.52	-11.03	-87.52	-3.16
$k = 0.5$	277.14	406.14	386.15	19.99	-14.38	-106.92	-7.71
$k = 0.6$	261.61	423.33	397.51	25.82	-18.76	-125.06	-17.90
$k = 0.7$	196.37	397.89	366.10	31.79	-23.30	-138.72	-39.50
$k = 1.0$	-186.40	31.68	0.00	31.68	-23.24	-137.58	-57.26
RL policy π_θ	**Case 1,** $\lvert\mathcal{S}\rvert = 3$						
SAC	334.77	404.23	387.03	17.20	-12.37	-48.08	-9.01
TD3	338.91	407.44	388.36	19.08	-13.73	-42.62	-12.18
PPO	293.07	347.96	334.51	13.45	-9.50	-37.26	-8.12
DDPG	140.53	198.17	192.19	5.98	-4.28	-51.22	-2.13
RL policy π_θ	**Case 2,** $\lvert\mathcal{S}\rvert = 34$						
SAC	344.84	430.39	409.33	21.06	-15.25	-56.74	-13.56
TD3	349.69	434.11	412.29	21.82	-15.84	-54.40	-14.18
PPO	340.49	399.90	382.40	17.50	-12.55	-36.78	-10.08
DDPG	338.85	420.17	398.59	21.57	-15.55	-51.30	-14.47

Table 7: Summary of the efficacy of RL periodic policies and random and static baseline incentives with $a_{j,t}k\ \forall j$. Expected episodic sum of rewards (in 10^3 Shekel) are presented for the transportation services, charging costs and relocation costs and estimated over 20 independent episodes over a control horizons of 12 weeks and for the two state space cases.

2.7.5 Comparison between the weekly minute tariffs of TD3 and SAC

Weekly trends resulting from the SAC and TD3 are illustrated in Figure 23 using line charts, to provide a view of the temporal and spatial adaptations of origin incentives over simulated weeks. Figure 23 displays the weekly distribution of discounted minute tariffs determined by the SAC agent (blue) and the TD3 agent (yellow) across the ten zones, each represented as a separate panel. The x-axes represent the hours of a simulated week, while the y-axes indicate the Shekel/minute tariffs discounted by the agents. The mean values of the minute tariffs are shown, alongside their variability (represented as the band spanning the 10th and 90th percentiles). Additionally, the absolute difference between the mean values of the two tariffs is plotted as a dashed green line.

The numerical results indicate that both SAC and TD3 perform comparably well regarding the expected reward maximization, demonstrating their capacity to tackle this optimization problem effectively, see Table 7. Some similarities can be observed in the average discounted Shekel/minute rate, for instance, lower origin incentive (higher minute tariffs) in Area 7-9 during the morning between Monday and Wednesday. However, notable differences in the weekly incentive patterns emerge between the two agents. These differences may stem from the inherent complexity of the problem, where multiple local optima exist within the policy space. The observed discrepancies suggest these two agents use distinct strategies while achieving similarly high revenues.

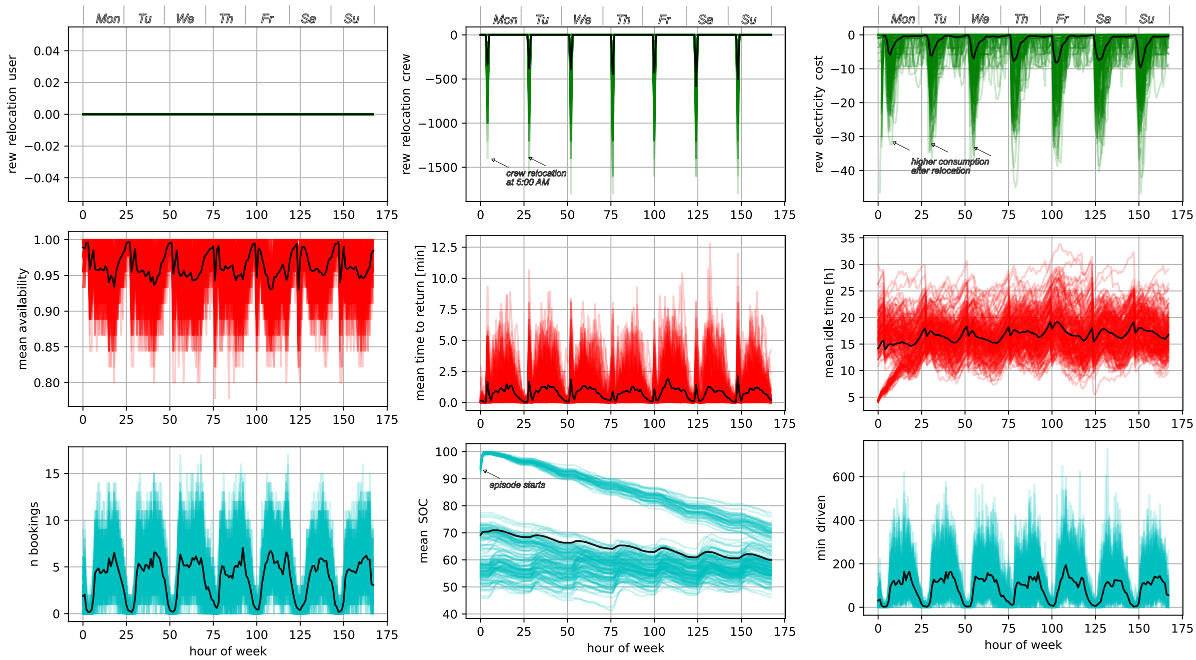

Figure 19: A visualization of the MC simulation results for the baseline 0% policy.

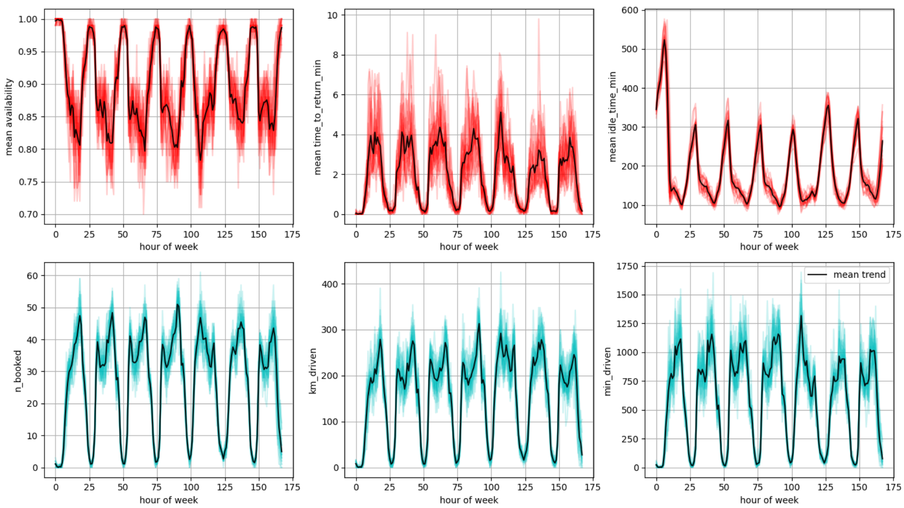

Figure 20: Simulated MC simulation of the same fleet with a higher utilization rate.

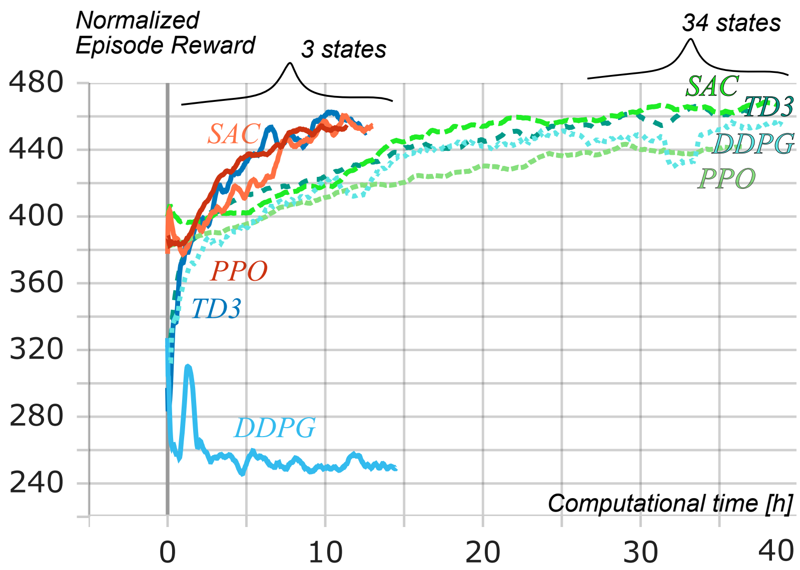

Figure 21: Episode roll-out reward for the four RL training algorithms. The plot illustrates the mean episode roll-out reward over the training iterations for SAC, TD3, PPO, and DDPG.

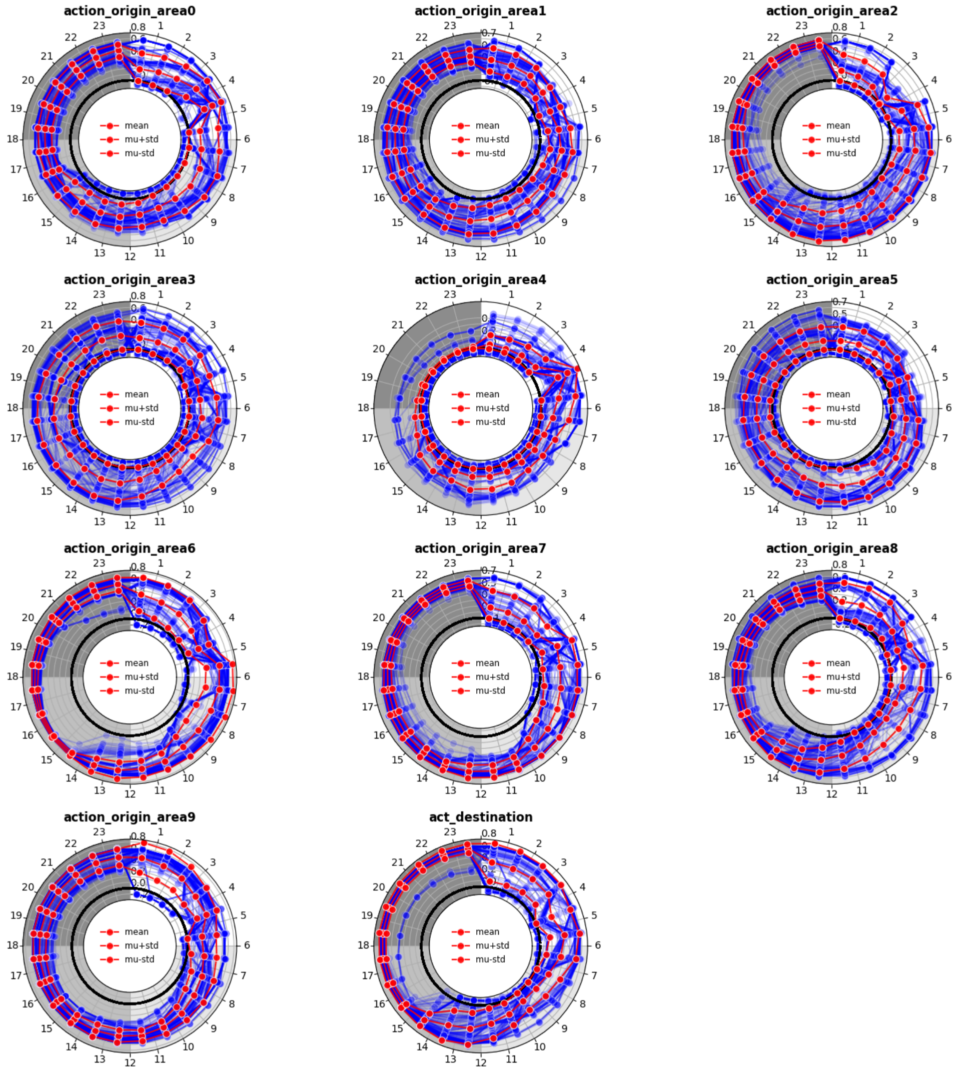

Figure 22: Daily action from the of the SAC policy. Each disk represents 24 hours while the blue lines correspond to the global destination incentive (a_d) and origin incentives for each zone $a_{o,k},\ k = 1, ..., 10$ on different days. The red curves display the average daily action and the mean action plus and minus one standard deviation.

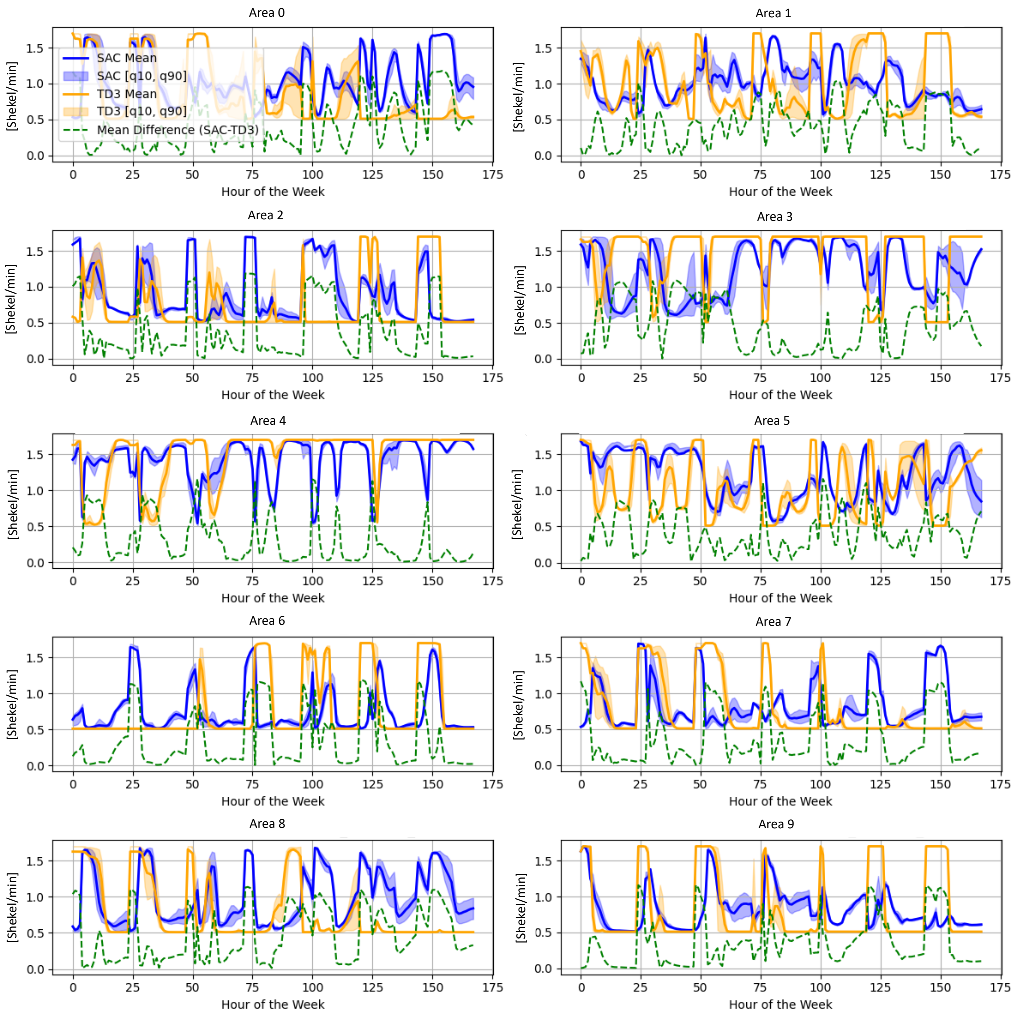

Figure 23: A comparison between the weekly distribution of discounted minute tariffs [Shekel/minute] for the SAC (in blue) and TD3 agent (yellow) for the 10 zones (the different panels). The absolute difference between the assigned minute tariff is presented by the dashed green line.

2.8 Demonstrative Digital Cross-Sector Platform

A digital platform that integrates electric vehicle (EV) information with grid information has the potential to enhance grid stability and optimize EV operations. Such a platform can either directly offer grid services or, at the very least, mitigate the impact of EV charging on the power grid. By bridging the gap between EV operations and grid requirements, this platform enables innovative use cases that leverage the flexibility of EV fleets. Gathering data and pertinent information essential for the implementation and to leverage EV flexibility for both the Distribution System Operator (DSO) and the Transmission System Operator (TSO).

In collaboration with our project partners, we have identified two primary use cases to study within the context of demonstrative digital cross-sector platforms:

1. **Ancillary Services Provisioning to a TSO with a Corporate Fleet**: Leveraging a corporate EV fleet to participate in ancillary service markets, explicitly focusing on the provision of tertiary control power (TRL) to Swissgrid. This use case addresses how EVs can serve as flexible energy resources for grid stabilization.
2. **Grid-Aware Relocation Optimization with a Free-Floating Shared Fleet**: This use case optimizes the relocation of vehicles in a free-floating shared EV fleet based on grid constraints and demand patterns. It seeks to balance fleet operational efficiency with minimizing grid stress during peak charging times.

In both use cases, understanding the information requirements is key to enabling the desired services. Our analysis focused on identifying the critical data points needed from the EV fleet to offer these services effectively.

2.8.1 Ancillary Services Provisioning to Swissgrid with a Corporate Fleet

Hive Power, leveraging its FLEXO platform developed to harness the flexibility of EV fleets, identified a practical and immediate use case for Switzerland: offering Tertiary Frequency Regulation (TRL) services to Swissgrid using a corporate fleet.

TRL, also referred to as Manual Frequency Restoration Reserve (mFRR), is used to restore system balance and release resources deployed for Frequency Containment Reserve (FCR) and automatic Frequency Restoration Reserve (aFRR). Unlike FCR and aFRR, mFRR is manually activated within 15 to 30 minutes and is typically employed for prolonged imbalances or to provide reserve capacity.

In Switzerland, mFRR is procured through both daily and weekly tenders, with minimum bid sizes starting at 5 MW and a maximum of 100 MW per bid. Energy delivery is offered in 15-, 30-, or 60-minute blocks, and remuneration follows a market-based pay-as-bid mechanism. Eligible technologies include generation, demand, and storage, offering flexible options for providers. Key features of the mFRR market in Switzerland are summarized in Table 8.

In order to bid EV flexibility on the Tertiary Frequency Regulation market in Switzerland, several critical requirements must be met. Participation necessitates being a registered Balancing Service Provider (BSP), and the minimum bid size is 5 MW, which typically requires pooling flexibility from multiple sources due to the limited capacity of individual EV fleets. Recognizing these prerequisites, Hive Power has collaborated with car fleet owners, a BSP, and an aggregator that provides pooling capabilities and acts as a technical service provider to the BSP[2]. This collaborative approach is indispensable for offering TRL services, as it combines the necessary regulatory qualifications and aggregated flexibility to meet market thresholds. Through this partnership, Hive Power is currently focusing on the daily market and specific 4-hour blocks. TRL services represent one of the "lowest hanging fruits" in terms of feasibility and market readiness for EV participation in Switzerland, as the requirements for real-time dispatch are less stringent.

[2] Due to privacy considerations, the names of the collaborating companies cannot be disclosed at this time.

Feature	Value
Service type (mandatory/optional)	Mandatory
Remuneration (market-based, pay-as-bid or pay-as-cleared)	Market-based, pay-as-bid
Technologies allowed for service provision (generation/demand/storage)	Generation, demand, and storage
Product granularity (delivery periods)	15-, 30-, or 60-minute periods
Minimum quantity per facility (MW)	5 MW
Maximum quantity per facility (MW)	100 MW per bid
Mode of activation (manual/automatic)	Manual
Price interval (€/MWh)	€0.01 increments
Gate opening time (GOT)	Weekly: Wednesday 13:00 (D-7) Daily: Earlier than 14:30 (D-1)
Gate closure time (GCT)	Weekly: Tuesday 13:00 (D-1) Daily: 14:30 (D-1)
Results publication	Results are published anonymously on Swissgrid's website
Schedule horizon	Weekly and Daily, divided into 4-hour blocks (00:00–04:00, 04:00–08:00, etc.)

Table 8: Key Features of the Tertiary Frequency Regulation Market in Switzerland

Swissgrid, as the transmission system operator, has published detailed guidelines for the prequalification of EV fleets to participate in ancillary service markets [46]. These guidelines outline the technical and operational requirements that EV aggregators must meet to qualify for market participation. Key aspects include:

- Ensuring reliable monitoring of individual EV states, such as the State of Charge and power availability.
- Providing aggregated fleet-level capabilities to meet the technical specifications of TRL services.
- Establishing bidirectional charging (Vehicle-to-Grid, V2G) where applicable.

To align with these requirements, Swissgrid's prequalification guidelines have been analyzed to understand how an aggregated EV fleet can be structured, monitored, and controlled to meet these standards. To enable the provisioning of ancillary services, it is essential to collect reliable and timely data from the electric vehicles (EVs) in the fleet. Two main approaches have been identified for obtaining the required data—State of Charge (SoC), charging and discharging capabilities, and vehicle identification. Each approach has distinct advantages and limitations, as outlined below.

1. **Via the Charging Station (OCPP + ISO 15118):** Modern charging stations that support the Open Charge Point Protocol (OCPP) and ISO 15118 facilitate data exchange between EVs and the grid infrastructure.
   - **Advantages:**
     - OCPP is widely adopted, enabling standardized communication with backend systems.
     - ISO 15118 enhances OCPP by allowing secure and advanced data exchange, including vehicle identification (EVCCID) and SoC.
     - Suitable for corporate fleets that utilize dedicated charging infrastructure with uniform capabilities.
     - Supports bidirectional charging (V2G) for additional flexibility in grid services.

- **Limitations:**
  - ISO 15118 implementation is not yet widespread, especially among AC charging stations.
  - Compatibility must be ensured between specific EV models and charging stations to enable successful data collection.
  - Requires investment in compliant charging infrastructure, which may not be feasible for fleets relying on public charging.

2. **Directly from the Vehicle** Alternatively, data can be retrieved directly from the EV using manufacturer-specific APIs.

   - **Advantages:**
     - Provides direct access to vehicle-specific data such as SoC, location, and charging status, independent of the charging station.
     - Can bypass limitations associated with older charging stations or those without ISO 15118 support.
   - **Limitations:**
     - Manufacturer APIs are diverse, requiring integration with multiple systems, each with its own protocols, data formats, and authentication methods.
     - Access to APIs may involve licensing fees or contractual agreements with manufacturers.
     - Real-time data retrieval may depend on network connectivity, introducing potential latency or availability issues.

Control capabilities are also a critical aspect of managing EV fleets for ancillary services. Using OCPP + ISO 15118, fleet operators can remotely initiate or stop charging sessions, adjust power levels, and implement bidirectional energy flows (V2G), provided the vehicle is connected to a compliant charging station. Alternatively, manufacturer APIs allow direct control of individual vehicles, such as setting charge limits, controlling charging schedules, or sending commands even when vehicles are not actively charging. While OCPP + ISO 15118 ensures standardized control for fleet-wide operations, APIs provide greater flexibility but may require more complex integrations to achieve uniform control across diverse vehicle models.

After evaluating both approaches, Hive Power opted to connect with the vehicle rather than the charging station, for several compelling reasons. Collecting data directly from the vehicles provides richer and more immediate information, such as the vehicle identification number and SoC, independent of the charging infrastructure. This method overcomes the limitations associated with charging stations that may not support protocols like OCPP or ISO 15118, or that lack internet connectivity. Additionally, the corporate fleet considered in this use case consists of a limited number of car models, making integrating with manufacturer APIs more manageable.

In pursuit of efficient data collection and control across various car manufacturers and models, platforms such as enode.com and smartcar.com have emerged as effective solutions. These platforms provide a unified interface to the vehicles, handling the complexities of integrating and managing the diverse APIs from different manufacturers internally. Hive Power opted to integrate with Enode's platform to manage data collection and control across diverse EV models. Key advantages of this solution include:

- **Unified Integration:** Enode's platform provides a single interface for seamless connection and management of various EV models, simplifying the integration process and reducing the complexity of dealing with manufacturer-specific APIs.

- **Enhanced Data Access:** The platform enables direct access to vehicle data, supporting functionalities such as geofencing based on vehicle location to determine when an EV is at a designated charging station. Data collection is performed using webhooks, which deliver updates immediately as they occur, significantly reducing network traffic compared to traditional API polling.

- **Direct Control Capabilities:** The API allows sending commands directly to vehicles, including initiating or stopping charging sessions and adjusting charging parameters, providing precise control over fleet operations.
- **Infrastructure Independence:** By interfacing directly with vehicles rather than relying on charging stations, this solution overcomes the challenges posed by non-compliant or offline charging infrastructure, ensuring an effective and scalable approach to data collection and control.

Despite choosing to communicate directly with the vehicles rather than relying on charging stations, Hive Power ensured that Swissgrid's requirements for TRL provisioning were met. While the alternative solution—monitoring and controlling the charging station, identifying the car, and communicating the SoC using the ISO 15118 standard—is, in principle, conceptually cleaner and better aligned with industry standards, it is currently impractical for large fleets due to limited deployment and compatibility. Furthermore, the direct vehicle communication approach offers an additional advantage for future developments: it enables access to the State of Charge (SoC) of vehicles even when they are not connected to a charging station, providing more comprehensive insights for forecasting flexibility. This pragmatic choice allows Hive Power to deliver a functional and scalable solution today while laying the groundwork for enhanced forecasting capabilities in the future.

System Architecture and Operation In this section, we focus on the details of the system architecture and its operation, highlighting how the implemented solution enables the provision of tertiary control power (TRL) services to Swissgrid. The block diagram (Figure 24) illustrates a schematic and simplified version of the system's architecture, its components, and their interactions.

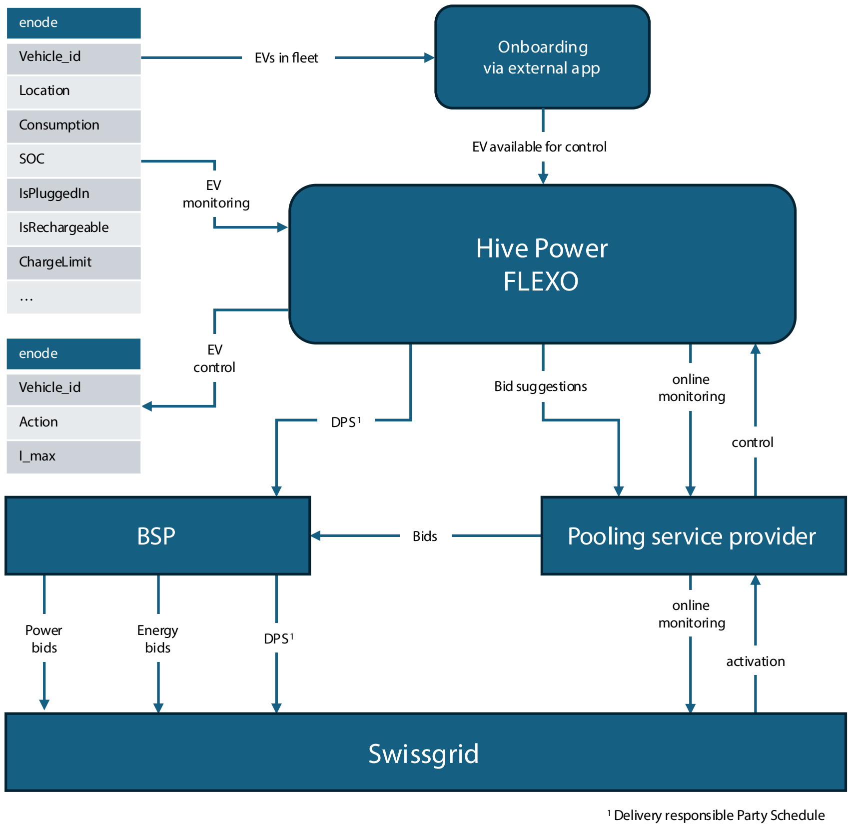

Figure 24: Hive Power's TRL platform architecture (simplified version)

Components of the System

1. **Enode Platform**
   - Acts as the intermediary between vehicles and Hive Power's FLEXO platform.
   - Collects essential vehicle data: identification, location, energy consumption, state of charge (SoC), connection status, and charge limits.
   - Enables direct control over vehicles, such as setting charge limits and, in the future, adjusting charging currents.
   - Data collection via webhook; control commands sent via API.
2. **External Onboarding App**
   - Handles vehicle onboarding to the flexibility service.
   - Ensures only vehicles explicitly enabled for flexibility services are available for control.
3. **Balancing Service Provider (BSP)**
   - Trades flexibility on the daily TRL market for agreed-upon time slots.
   - Transmits Delivery Responsible Party Schedules (DPS) to Swissgrid.
4. **Pooling Service Provider**
   - Aggregates flexibility from the EV fleet to meet Swissgrid's minimum requirement of 5 MW for TRL services.
   - Submits aggregated bids to the BSP.
   - Forwards monitoring data to Swissgrid.
   - Receives activation signals from Swissgrid and communicates them to Hive Power for execution.
5. **Swissgrid (Transmission System Operator)**
   - Purchases flexibility services.
   - Oversees activation and monitoring of services.
6. **Hive Power FLEXO Platform**
   - Aggregates data from Enode and manages fleet operations.
   - Monitors EV availability and forecasts available flexibility.
   - Prepares power and energy bid suggestions for the pooling service provider.
   - Provides online monitoring of total fleet power.
   - Executes control actions to meet service commitments.
   - Ensures compliance with Swissgrid's ancillary service requirements.

Operation of the System:

1. **Vehicle Onboarding**
   - Vehicles are onboarded to the platform through the external app.
   - Necessary identification and operational data are collected during onboarding.
2. **Real-Time Monitoring**
   - Enode continuously monitors the fleet in real-time.

- Updates on vehicle status—including SoC, location, and availability—are transmitted to the Hive Power FLEXO platform.

3. **Bid Preparation and Submission**
   - Hive Power prepares power and energy bid suggestions based on the collected data.
   - The pooling service provider aggregates these suggestions and submits bids to Swissgrid through the BSP.

4. **Activation and Control Execution**
   - When flexibility is called upon, Swissgrid sends activation signals to the pooling service provider.
   - The pooling service provider communicates these signals to Hive Power.
   - Hive Power uses Enode to execute commands directly on the vehicles, ensuring the requested reserve power is delivered.

5. **Performance Monitoring and Compliance**
   - The system monitors fleet performance during activation.
   - Compliance with Swissgrid's requirements is ensured.
   - Data is recorded for reporting and optimization purposes.

This architecture effectively combines Enode's integration capabilities with Hive Power's expertise in grid service provisioning. Its support for real-time control and monitoring enables the provision of ancillary services, even with a diverse fleet of vehicles and non-compliant charging infrastructure.

Current Operational Implementation The system described above is now fully operational and has been successfully integrated into the ancillary services market in Switzerland. Hive Power's platform is currently live and providing TRL services to Swissgrid using a corporate EV fleet.

At present, the focus is on offering negative TRL services in 4-hour blocks. Negative TRL, which requires increased consumption when activated, is advantageous for several reasons. Firstly, it generally offers higher remuneration compared to positive TRL. Secondly, it is less risky in terms of guaranteeing a minimum State of Charge (SoC) in the vehicles. Positive TRL, which involves stopping the charging process when activated, introduces challenges in ensuring sufficient SoC levels across the fleet. However, positive TRL services are foreseen to be activated in the future as the operational framework matures and reliability of SoC management improves.

Operationally, vehicles are allowed to charge up to a certain SoC threshold, typically 80%, and then charging is stopped before the period in which negative TRL could be activated. This strategy ensures that sufficient charging capacity is available if flexibility is activated by Swissgrid. When an activation signal is received, the charging sessions for the EVs that were held at 80% SoC are resumed, thereby increasing the total consumption to provide the required reserve power.

The number of activated charging sessions is dynamically adjusted based on the amount of requested power from the aggregrator. Hive Power's FLEXO platform ensures adherence to the power profile requested by the aggregator. Data is continuously collected from the vehicles, and monitoring information is sent to the aggregator in real-time.

Geofencing is used to determine whether the cars are connected to the intended charging station. Due to limitations in data frequency from the EVs, FLEXO incorporates predictive algorithms dedicated to nowcasting the fleet's total power consumption and SoC levels. Extensive testing was conducted to validate this approach, and Swissgrid has approved Hive Power's method for monitoring and control. SoC is closely monitored and communicated as per Swissgrid's requirements, ensuring compliance and reliability in service provision.

This implementation has already seen real-world activation of flexibility services, demonstrating the viability and effectiveness of the system in participating in the ancillary services market. As the system evolves, Hive Power plans to expand its capabilities to include positive TRL services, further enhancing its contribution to grid stability and flexibility.

User Interface The user interface plays a crucial role in monitoring and controlling the EV fleet for TRL service provision. In this report, two versions of the user interface are presented: the current operational interface based on Grafana, and a design prototype illustrating a possible future version.

The user interface currently in production is based on Grafana[3]. Although it does not yet support direct input or control actions from operators, it offers comprehensive real-time insights into the system. Designed to be highly informative, the dashboard enables operators to effectively monitor the fleet and system performance. The interface, represented as a mockup in Figure 25, is shown here instead of the actual deployment, as displaying the real interface is not possible due to privacy concerns.

The key features of the operational interface include:

- **EV Availability Monitoring:** Displays the number of EVs available on the Enode platform and those available for control, providing real-time visibility into the fleet's operational readiness.
- **Fleet Composition Statistics:** Presents statistics on brands and models within the fleet, aiding in understanding the diversity and capabilities of the vehicles.
- **Use Case Allocation[4]:** Displays the allocation of EVs to various use cases or specific 4-hour time slots within the TRL market, allowing operators to effectively manage how vehicles are utilised across services and time periods.
- **Geolocation Data:** Provides the last known location of each EV for geofencing purposes, ensuring that vehicles are within designated operational areas.
- **Real-Time TRL Power Availability:** Displays the real-time available TRL power, allowing operators to assess the current flexibility that can be offered to Swissgrid.
- **Total Fleet Power and Flexibility:** Shows the total fleet power with upward and downward flexibility, including activation signals, to monitor the fleet's capacity to respond to grid demands.
- **Flexibility Forecasts:** Presents probabilistic forecasts of available power for the upcoming days, aiding in bid preparation and operational planning.
- **Bid Proposals:** Provides suggested bid proposals for the upcoming days based on forecasted flexibility, streamlining the bidding process.
- **Performance Metrics:** Displays monthly Key Performance Indicators (KPIs) on fleet performance and system health, allowing for ongoing assessment and optimisation.

Building upon the insights gained from the current operational interface, Hive Power's interaction designer has developed a visually appealing prototype that envisions enhanced functionalities and an improved user experience for future iterations. Although no final decision has been made regarding the specific features of the new platform, it is clear that the next iteration will provide operators with greater control capabilities and more detailed monitoring functionalities. The prototype explores possibilities such as detailed real-time data visualization, flexible vehicle assignment to different services, comprehensive fleet analytics, and advanced notification systems.

While the current interface focuses primarily on monitoring, future developments will integrate control functionalities, enabling operators to interact directly with the system. This will include capabilities such

[3] `https://grafana.com/`

[4] While this report focuses on TRL provision, the system is designed with the flexibility to allocate vehicles to other services, such as optimising charging profiles against the spot market, or to different time slots within the TRL market.

as assigning vehicles to specific services, initiating control actions, and customizing operational parameters. These functionalities will be refined based on feedback received from our partners who are actively using the operational interface. Their insights are crucial for tailoring the platform to meet real-world operational needs effectively.

The overarching goal is to merge the comprehensive data visualization of the Grafana dashboard with the enhanced control features envisioned in the design prototype, creating an intuitive and user-friendly experience for operators. A depiction of this prototype is shown in Figure 26.

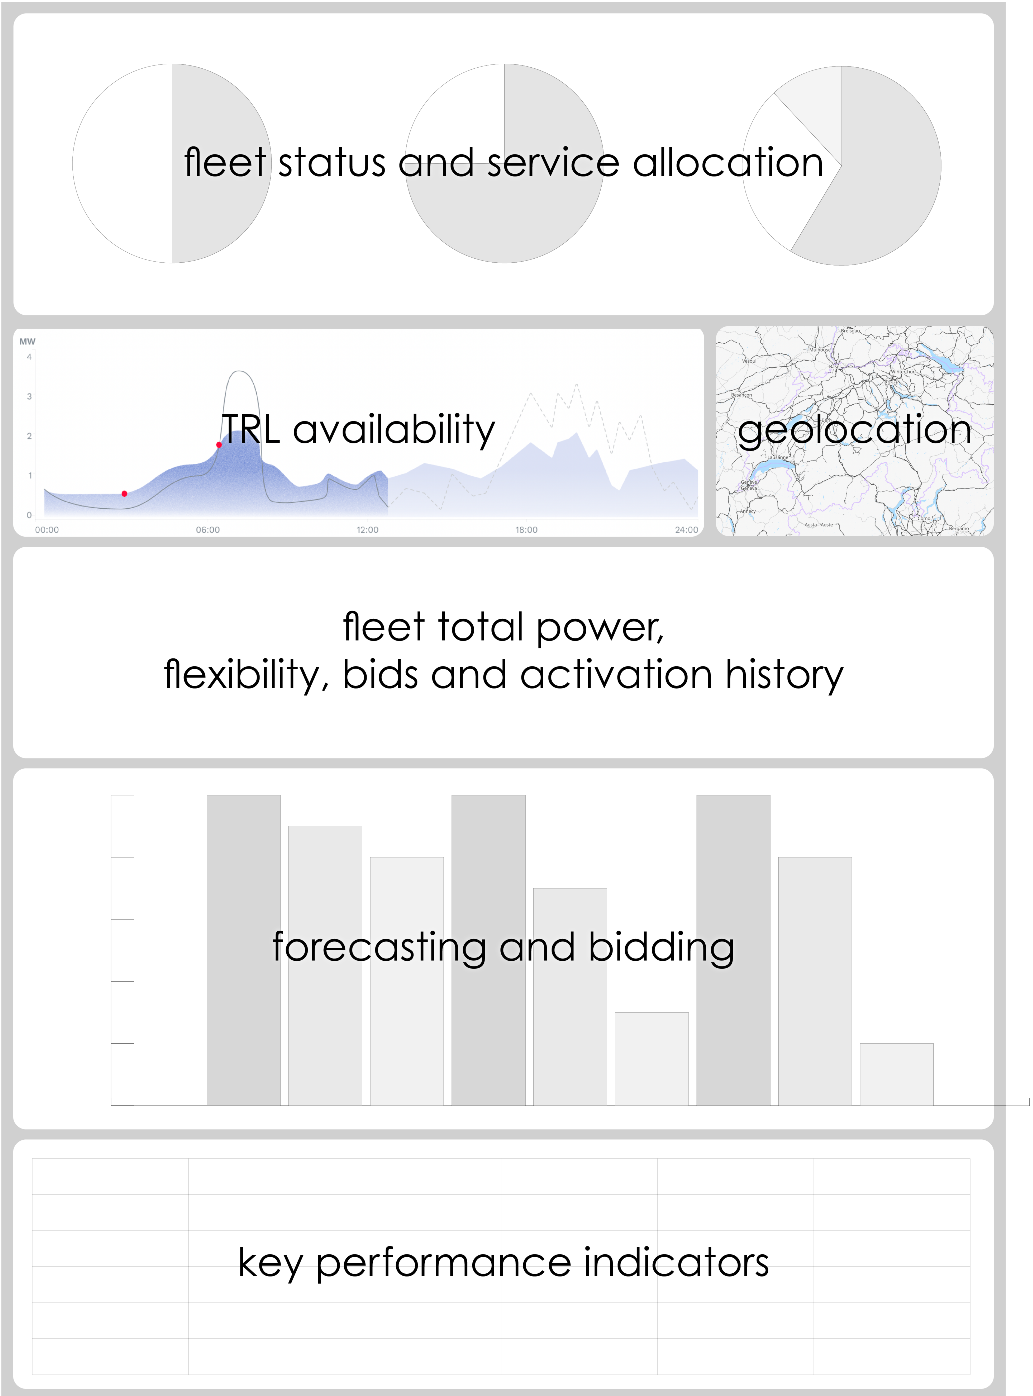

Figure 25: Mockup of the operational user interface, currently providing real-time monitoring of the fleet's status.

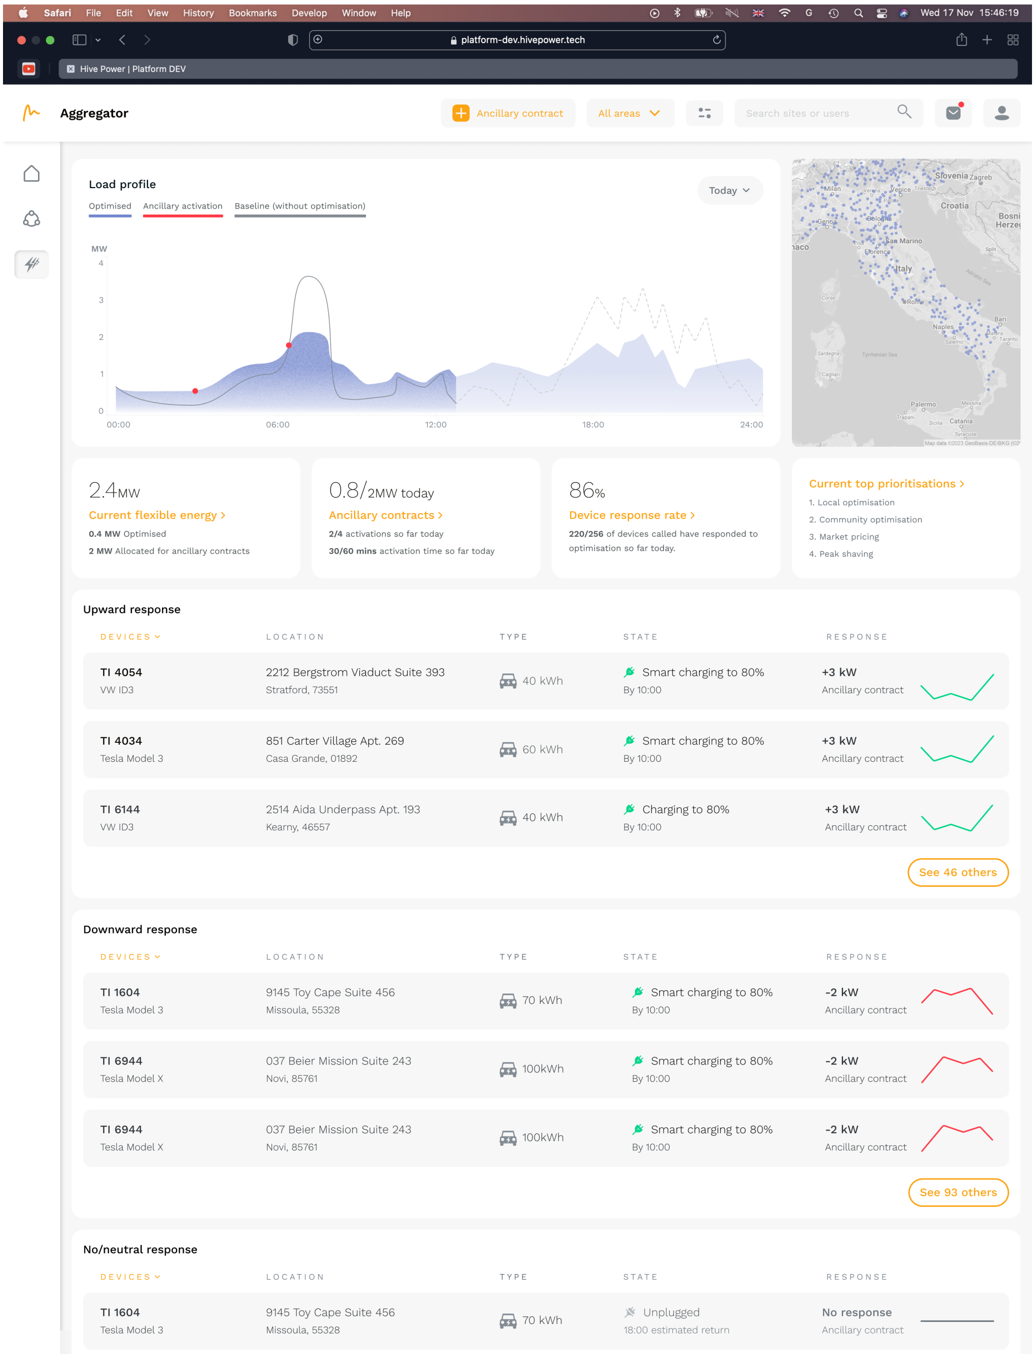

Figure 26: Design prototype of the user interface illustrating detailed monitoring and control capabilities.

2.8.2 Grid-Aware Relocation Optimization with a Free-Floating Shared Fleet

SUPSI in collaboration with Hive Power, AutoTel, and other project partners, developed a demonstrative platform to facilitate grid-informed relocation of EVs for the free-floating system in Tel Aviv. This platform provides visualization and decision support tools to support the fleet operators, and can be further extended to leverage flexibility and provide ancillary servies for the DSO and Israeli TSO.

System Architecture, Operations, and data collection This demonstrative platform, put emphasis on grid-informed mobility and relocation aspects, has been designed to prioritize the capture various features from charging stations and EVs. Data acquisition is largely contingent on the business model in question, hence, for a free-floating model, the following operational features are considered:

- Vehicle ID and characteristics
- Geographical positioning of EVs
- Connection status to a charger
- State of Charge
- Charging station availability
- Idle and charging duration indicators
- Availability of distributed generators and other power generation/ energy flexibility sources

Generally speaking, collecting the data directly from the vehicles is perceived to be more advantageous than sourcing data from charging stations, as it offers richer information, such as the car's identification number and its SOC. Such details can substantially enhance the precision in forecasting the vehicle's charging and mobility needs. One of the paramount challenges encountered so far is the real-time procurement of consumption data. This challenge extends to both the EVs and the overarching grid system. The availability of power data and prerequisites for participation in the ancillary service market remain under active investigation.

Implementation and design Building on insights from business model analyses, we propose a cross-sector platform designed to aggregate EV data and provide critical information about fleet status and relocation needs across various zones within a control area. The platform is implemented via the *streamlit-folium* python package, which builds upon the folium API for geospatial visualization. This platform addresses multiple requirements, enabling the identification of flexibility potential in specific zones and supporting grid-informed relocation strategies. Key functionalities of the platform include:

1. **Interactive mapping**: Provides detailed insights into EV-specific metrics such as battery status and SoC, along with real-time updates on charging station statuses.
2. **Zone-Specific Fleet Tracking and Visualisation :** Displays fleet status, relocation demands, and global fleet metrics for each zone in real-time.
3. **Relocation Alerts and Recommendations:** Generates warnings and control actions for fleet re-balancing and relocation using reinforcement learning policies, supporting operators in optimising fleet management.
4. **Inspection of Zone-Specific Fleet and Relocation Information:** Facilitates detailed analysis and visualisation of fleet distribution and movement across zones.
5. **Grid-Aware Insights and Forecasting:** Offers predictions of renewable energy production, aggregated consumption, and interactions between EV charging demands and the grid.

The primary objective of this platform is to enhance the efficiency of relocation efforts by enhancing grid-awareness, a critical aspect of free-floating car-sharing systems. By delivering timely and actionable insights, the platform helps optimise fleet operations, benefiting both fleet operators and the stability of the power grid. The interactive frontend of the platform is depicted in Figure 27, showcasing its application in the Tel Aviv use case. Figure 28 illustrates a sidebar editor used to adjust forecasting parameters and evaluate the impact of various policies through simulation-based analyses.

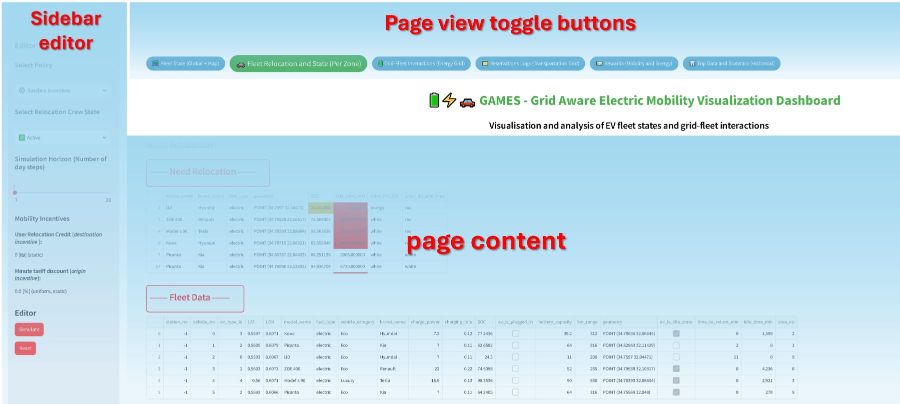

Figure 27: Example of the interactive platform and dashboard for the use case in Tel Aviv. The platform includes a sidebar editor, page toggle selector, and main page content.

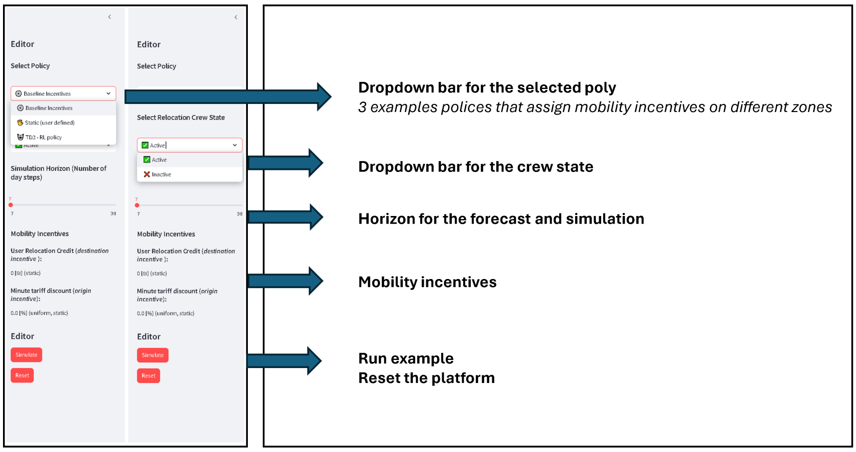

Figure 28: Example of editable sidebar for simulation-based and forecasting analysis for the use case in TelAviv.

Functionality and user interface Following the initial design provided by Hive Power, a cross-sector platform has been developed to support online monitoring and optimization of relocation tasks in the free-floating system in Tel Aviv.

1- Interactive mapping: This centralised dashboard displays the status of all vehicles in the system, showing SoC, availability, and current usage patterns across the service area. Global Fleet State, visualization of vehicles and charging stations are fundamental to support the ened user, the fleet operator, in relocation and rebalancing decision making. See Figure 5 for an example of a street map, fleet state visualisation.

2- Relocation Alerts and Warnings: Using dynamic data inputs, the platform alerts operators to areas and specific EVs that require intervention, helping them to manage inactive crew relocation and improve fleet distribution and management. Figure 29 presents a simple example of EV-specific features inspection, which can be carried out by the operators, and an example of severe (red) warning for relocation, due to a SOC lower than 20% and long idle time. Other minor warning messages are indicated by the orange car icons in the map.

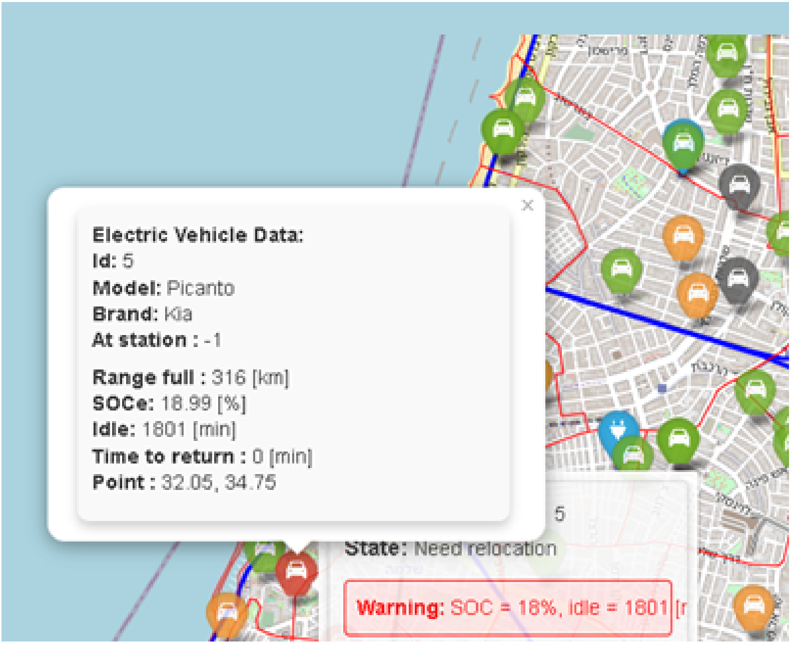

Figure 29: Example of severe warning (red) and moderate warnings (orange).

3 - Visualisation and inspection of zone-specific Fleet States and Relocation Information: This feature allows operators to track vehicle status in specific areas, offering insights into fleet utilisation patterns and real-time relocation needs. In particular, the combination of SOC and EV density in station-assigned zones allows identifying areas with potential for peak shaving and valley filling services. Figure 30 presents two examples for this.

4 - Grid-aware Global Consumption Analyses: Understanding the interplay between electricity demand, PV production, and EV fleet consumption is essential for developing resilient and efficient grid systems. This analysis considers how scaling up EV consumption (from a large fleet of more than 1000 EVs) impacts the demand and supply balance, particularly within urban environments. Figure 31 illustrates this with data from Tel Aviv, showcasing the relationship between foretasted PV production (blue on the top right panel), total city demand (red, right panels), and projected and aggregated EV consumption (in green). The EV consumption on individual zones (on the bottom left panel over the 10 station assigned zones of Tel Aviv) is also presented. Note that aggregated and uncontrolled EV demand, following a combined relocation effort, could have led to an unwarranted peak affecting the total consumption of the municipality. Such insights help identify potential stress periods on the grid and guide strategies for optimizing energy storage, distribution, and load balancing in real time.

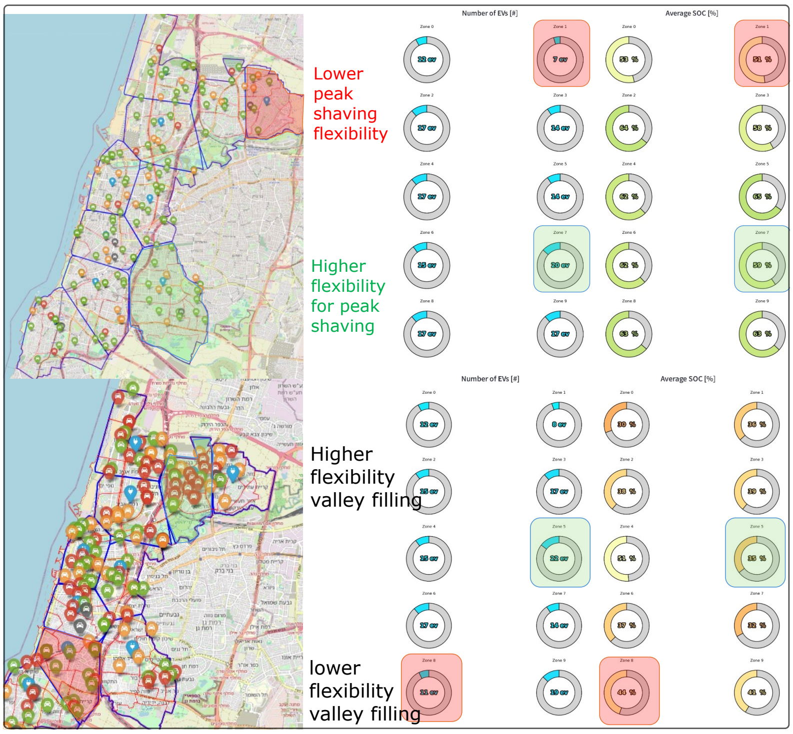

Figure 30: Aggregated average state of charge and vehicle density per station-assigned zones and for the use case in Tel Aviv. An example of energy flexibility awareness tools provided by the platform.

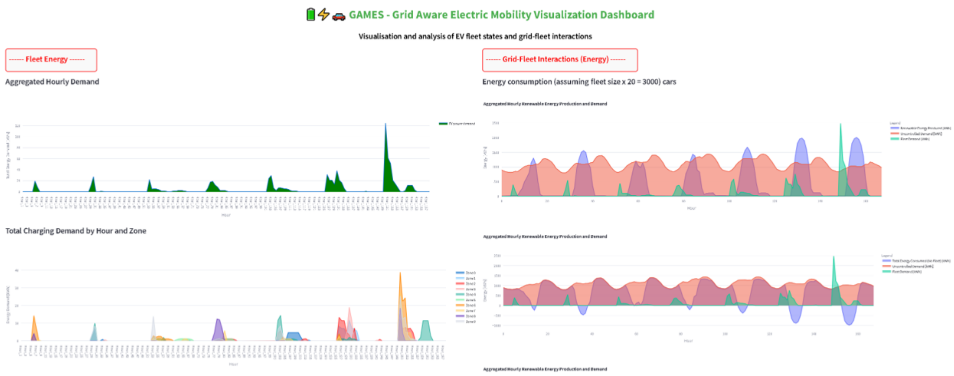

Figure 31: Example of global PV production, demand, and scaled-up EV fleet consumption in Tel Aviv.

3 Conclusions and outlook

GAMES aimed to assess revenue potential and opportunities for a grid-aware electrified carsharing system, focusing particularly on the use of forecasting tools and relocation strategies driven by energy and grid data. We tested various control policies, including user incentives, and explored fleet management strategies to evaluate the scalability of the proposed solutions.In conclusion, GAMES has made substantial progress and has achieved notable outcomes across several key areas.

A concise summary of the key activities and achievements is provided below:

- Development of **macroscopic traffic and simulation models** to evaluate various operational scenarios for EV fleets and interactions between the energy and mobility sectors.
- Development of **forecasting models and algorithms** to predict demand in both the energy grid and transportation network.
- **Optimization of relocation policies** π against baseline strategies, demonstrating the effectiveness of advanced user incentive strategies.
- **Evaluation and optimization of economic dispatch policies** π_{ED} based on optimal relocation policies π (find more details in section 4).
- Creation of **demonstrative cross-sector digital platforms** as a proof-of-concept.
- Conducting workshops, interviews, a focus group, and a survey to gather the perspective from market agents and citizens (find more details in section 4).
- Conducting various dissemination and communication activities (find details in section 5).

Simulation Environments and Macroscopic Traffic Models: Significant progress has been made in the development of advanced simulation environments and macroscopic traffic models. These tools are essential for analysing and understanding mobility and energy patterns across different use cases, carsharing systems, including both free-floating and station-based models.

Non-parametric Space-Time Probabilistic Demand Models and Mobility Forecasting: Parametric and non-parametric space-time probabilistic mobility and energy demand forecasters have been developed and rigorously tested. These data-driven, adaptable models provide critical insights into the expected state of transportation systems and energy needs, with applications extending beyond the specific use cases discussed in this report. By leveraging conditional probability methods, such as Kernel Density Estimators, alongside machine learning-based predictors, these models effectively capture the distribution of mobility demand. Particularly for free-floating carsharing systems, which are characterized by higher variability, this reveal complex but explainable usage patterns, later essential for system optimization and simulation-based analyses. These probabilistic models have been integrated into the mobility simulation environments, cross-sectoral platforms, and have been used to generate plausible scenario realizations, supporting optimization and decision-making processes. Their ability to model uncertainty and variability ensures robust analysis for a wide range of applications in mobility and energy systems.

Simulation and Optimization: The project has conducted extensive simulations to explore the impact of electric carsharing fleets and different operational policies on carsharing revenues. A Reinforcement Learning framework was developed to optimize vehicle relocation and user participation in the carsharing activities, hence testing advanced AI and ML-driven approaches to tackle this complex problem. This reinforcement learning framework was designed to tackle dimensionality and uncertainty challenges that impact the relocation optimization problem while offering an adaptive strategy. The framework optimizes actions applied to zones rather than individual cars, enabling efficient scaling for large fleets at minimal computational additional costs. User participation is facilitated via discounts and mobility credits, seeking to encourage departures from specific under-served areas and to promote the return of electric vehicles to charging stations when relocation is necessary. The proposed RL framework addresses challenges such as the efficacy of the mobility service, charging station utilization, electrical costs, and mobility incentives. User participation in relocation activities could be a viable solution to reduce relocation costs and increase the density of grid-connected EVs for energy-related services. The optimization results have provided valuable insights into fleet management strategies and helped identify optimal resource allocation to enhance overall system performance. Our findings indicate that well-designed adaptive incentive strategies can enhance user participation in relocation activities and improve fleet utilization. This can lead to significantly higher revenues for the fleet operator, but it can also enhance the satisfaction of the carsharing users, e.g., by assigning mobility credits and discounts. Benefits for the fleet operator include reduced relocation costs (by shifting part of relocation crew costs to users' credit assignment) and improved fleet utilization (by adaptively changing minute tariffs to promote departures). While charging costs and revenues were relatively minor in comparison, see results for the AutoTel use case, this is expected to change with the introduction of incentive schemes, rising electricity costs, and larger EV fleets. Extensive numerical tests have demonstrated the potential of these advanced approaches and the resulting RL policies can provide an optimized design of dynamic and space-heterogeneous incentives. In particular, the dynamic and zone-specific incentives provided by the RL agents demonstrated superior performance (compared to various baselines) in terms of revenue generation, achieving better mobility service, and substantial reduction in relocation costs.

Economic Dispatch: An innovative Economic Dispatch model aimed at optimizing the SOC of idle electric vehicles was developed. This model addresses a key challenge in electric mobility by efficiently managing the energy levels of EVs when not in active use, improving overall fleet management and reducing operational costs. The results from this work have shown that the aggregated demand from the fleet generally has a minimal impact on the overall grid, especially for smaller EV fleets. However, localized issues may arise if the density of EVs increases and there is poor coordination in charging behaviour. For instance, if all EVs in a given area charge simultaneously, this can lead to local congestion within the DSO network and result in unwarranted localized peaks in demand. Therefore, uncoordinated charging must be carefully avoided to prevent such disruptions. Note that results of this Economic Dispatch model are not described extensively in this report as they were mainly developed by E7, a GAMES Austrian partner, but are publicly available on the GAMES website.

Digital Cross-Sector Platforms: Two demonstrative cross-sector platforms have been developed to integrate mobility and energy data, incorporating forecasting and simulation tools within a common framework. These platforms align with industry requirements, promoting interoperability and grid awareness. By connecting transportation data with energy systems, they facilitate a more comprehensive understanding of local criticalities in both the grid and the transport network, energy flexibility potential, and opportunities for intelligent fleet management—including relocation and participation in ancillary services. The first platform focuses on Ancillary Services Provisioning with a Corporate Fleet. Leveraging a corporate electric vehicle (EV) fleet to provide Tertiary Frequency Regulation (TRL) services to Swissgrid, the transmission system operator in Switzerland, this platform collects real-time data directly from the vehicles and integrates with market mechanisms. Doing so enables the EV fleet to participate in ancillary service markets, enhancing grid stability and optimizing EV operations. The platform's architecture allows for direct communication with vehicles, overcoming limitations associated with charging infrastructure and providing richer data for forecasting and control. The second platform emphasizes Grid-Aware Relocation Optimization with a Free-Floating Shared Fleet. Developed in collaboration with project partners, this platform aggregates EV data to offer decision support tools for fleet operators in Tel Aviv. It enhances efficiency in relocation tasks by providing interactive mapping, relocation alerts, zone-specific fleet tracking, and grid-aware insights. Operators can monitor fleet status, receive warnings and recommendations for rebalancing, and analyze energy flexibility potential in specific zones. This supports more effective fleet management by considering both mobility demands and grid constraints. The ability to monitor and analyse energy demand from electric vehicle fleets in real time allows for better coordination of charging schedules and relocation strategies, reducing the risk of local grid congestion and optimizing energy distribution. Furthermore, the cross-sector approach ensures that these platforms can accommodate future developments, such as the expansion of electric vehicle fleets, fluctuating energy demands, and the integration of renewable energy sources. By providing a clearer view of both transport and grid conditions, these platforms lay the foundation for smarter, more efficient, and sustainable mobility solutions that can respond dynamically to local needs and challenges.

A concise summary of the results: In the context of fully electrified electric vehicle fleets, the results of GAMES suggest that there is a moderate to high potential for sustainable grid-aware management of small size electrified urban mobility services, especially if supported by smart operational policies and possibly higher density, design, and allocation of dedicated charging stations. The potential for sustainable grid-aware management, increases dramatically for large fleets, where it becomes evident that uncontrolled or poorly controlled management of the fleet can generate issues on both power and transportation systems. Operating such EV fleets is inherently more complex than traditional fleets, this is partially due to longer charging time, and particularly challenging for free-floating systems that require advanced relocation strategies. Station-based systems, with their longer idle times and more controllable bookings, offer greater opportunities for effective economic dispatch while optimize EV readiness to be booked. Effective management of idle "free-floating" EVs require smart relocation and re-balancing and suitable distributed charging points for distributed/centralized control. While the potential for fully electrified EV fleets is significant, addressing these complexities with innovative solutions is vital for their success in advancing sustainable transportation. These aspects may require further investigation. The proposed incentive optimization framework using reinforcement learning assigns incentive actions to zones rather than individual cars, enabling efficient scaling for large fleets at minimal computational additional costs. Discounts and credits were used to encourage departures from specific areas and the return of electric vehicles to charging stations and could be a viable solution to reduce relocation costs and increase densitiy of gird-connected EVs for energy-related services. The proposed RL framework addresses challenges such as the efficacy of the mobility service, charging station utilization, electrical costs, and mobility incentives. Extensive numerical tests have demonstrated the potential of these advanced approaches and the resulting RL policies can provide an optimized design of dynamic and space-heterogeneous incentives. In particular, the dynamic and zone-specific incentives provided by the RL agents, when compared to purely crew-led relocation strategies and combined relocation strategies with static incentives (user and crew relocation), demonstrated superior performance in terms of revenue generation, achieving better mobility service, and substantial reduction in relocation costs. The proposed framework is a significant advancement in optimizing FFEVCS systems. Despite its limitations and necessary modelling assumptions, it provides a foundation for further research and development in this area. By addressing the identified limitations and exploring future research directions, we can continue to advance urban mobility systems and grid-aware mobility solutions, contributing to the realization of more sustainable and efficient transportation solutions. In conclusion, the GAMES project has delivered valuable tools, models, and insights that enhance mobility services, support informed decision-making, and enable resource optimization in urban settings. The achieved results form a robust foundation for ongoing and future research endeavours, contributing to the advancement of sustainable and efficient transportation systems.

Future work: Switzerland's energy transition is characterised by increasing electrification of transport, growing PV penetration, and distribution-grid constraints. Building on GAMES, subsequent activities will target scalable, grid-aware coordination of EV charging and infrastructure deployment.

Within the Horizon Europe project InterPED[5], SUPSI will develop grid-aware control algorithms for public charging stations in Positive Energy Districts, building directly on GAMES' probabilistic space–time demand modelling and the reinforcement-learning framework. The work will formulate and implement controllers (e.g., MPC and RL) that schedule charging subject to transformer/feeder limits, voltage constraints, dynamic tariffs, and flexibility requests; specify and prototype interfaces with CPO/OEM/DSO/TSO systems; and validate the algorithms in pilots using quantitative indicators (grid loading, energy cost, and service quality). Methods and interfaces will be prepared for replication beyond the pilot sites.

Within the Horizon Europe project [6], SUPSI will develop a planning tool for public EV charging infrastructure, leveraging GAMES' city-scale mobility/traffic simulation stack and grid-aware fleet operations insights. The tool will integrate demand estimates (arrivals and dwell times) from mobility and traffic models with distribution-grid hosting capacity, reinforcement costs, and siting constraints (land use, accessibility, and equity). Multi-objective optimisation will support scenario analysis (e.g., baseline vs.

[5] https://interped.eu/ (project website); https://cordis.europa.eu/project/id/101138047 (CORDIS).

[6] https://ginnger-project.eu/ (project website); https://cordis.europa.eu/project/id/101123324 (CORDIS).

accelerated rollout, AC/DC mix, curbside vs. hubs) and generate GIS outputs and decision material for municipalities and DSOs.

Hive Power is already operating in production the FLEXO-based tertiary reserve aggregation stack delivered in GAMES with customers in Switzerland. Building on this live deployment, the roadmap includes scaling to larger and more heterogeneous fleets; enabling bidirectional TRL and additional ancillary services (e.g., aFRR, congestion management); strengthening integrations with CPMS/OEMs and charger-side control; enhancing forecasting and automated bidding; standardising onboarding, reporting, and settlement; and adding wholesale market arbitrage (day-ahead/intraday) with SoC- and grid-constrained scheduling. These developments are intended to support the secure integration of distributed flexibility and cost-effective operation of electrified mobility in Switzerland.

4 National and international cooperation

The project actively involved four research organizations; Scuola universitaria professionale della Svizzera italiana (SUPSI), Reichman University (RUNI), E7, and Salzburg Research Forschungsgesellschaft (SRFG). These partners meet once a month to update each other and organise the collaboration and exchanges of information/data. In addition, the partners organized national and international activities, the results of which contributed to the project as a whole, such as:

- On 23.11.2022 SUPSI organised at the SUPSI Mendrisio Campus a stakeholder workshop where the core swiss cooperation partners - Mobility, Hive power, and Sun2wheel - attended in person. During this meeting were discussed the concepts of carsharing, smart charging, and V2G in terms of the current status and potential, and the key challenges and use cases as perceived by the cooperation partners. Inputs from the stakeholderx were shared with all partners, and were especially of interest to the the Austrian partners, that lead the task of exploring new business strategies.
- Between Nov. 2022 and Feb.2023 was conducted by SRFG an idea-competition "Energy sharing with benefits" that was open to the general public. More information about the competition platform and the results can be found on the GAMES website. The idea competition was "advertised" through the networks of all GAMES partners, and attracted mostly participants from Austria and Switzerland.
- On 11.01.2023 was conducted an international workshop with participants from Switzerland, Austria, and Israel, during which were introduced conceptually some potential business models for controlled charging of EVs to provide flexibility and grid services. The introduction was followed by a group exercise during which the participants discussed sharing requirements, and barriers and opportunities for business models for novel bidirectional charging management of EVs. This international workshop helped not only to capture barriers and opportunities, but also notice some contextual differences in the different countries.
- On 23.02.2023 was organised an Austrian stakeholder workshop that focused on use cases and business opportunities for smart charging and V2G. To this event arrived many national stakeholders interested in a deep dive into this topic.
- During Q1/2023 RUNI organised 3 focus groups in which was investigated the concept of carsharing and EVs for residents of multifamily apartment buildings in urban environments. The advantages and disadvantages of these concepts were explored, systematically gathering the preferences of the participants.While the responders were from Israel, the challenges as well as the integration of shared vehicle services discussed are pertinent to densely populated urban areas worldwide. Results are disscussed in the journal publication [47]
- The RUNI team lead interviews with key stakeholders in Austria, Germany, Switzerland, and Israel, supporting the GAMES project effort to expose different perspectives in the context of sustainable mobility in general, and more specifically carsharing fleets and electric cars usage.

- As described in section 2.2, each research partner from CH, AT, and IL brought to the project an industrial stakeholder that provided invaluable data that was used for the three distinct case studies. This data was extensively and thoroughly analysed by SUPSI and used as a basis for much of the work described in this report.
- Moreover, GAMES research partners participate in the ERANet Knowledge Community activities, and are involved in two working groups; "Regulatory and Market Development" and "Consumer and Citizen Involvement". In the context of these working groups, GAMES researchers contributed to the "Joint Programming Platform Smart Energy Systems - Policy Brief ", and participated as well in the peer-to-peer ERANet Smart Energy Systems projects feedback session conducted on 27.06.2023.

5 Communications

In order to disseminate the activities of GAMES to a broader audience, several communication channels were established. These include a dedicated website, which can be accessed at (Games-innovation), as well as a presence on LinkedIn via the following link (GAMES-innovation-linkedin). A public repository for the project GAMES is published on GitHub (GAMES-GitHub-repository) and a Zenodo repository (GAMES-zenodo) was setup to archive all the deliverables, publications, code, and datasets of GAMES that are publicly available.

Additionally, GAMES gained some visibility through:

- 8 peer reviewed articles in published in journals and conferences (find more details in section 6)
- 5 deliverables in the form of reports uploaded to the GAMES website
- 35 media coverage articles (a sample communication is presented in Figure 32).
- 15 third-party events, presentations and panel debates
- 3 events organised by GAMES partners (find more details in section 4)

To date the GAMES research partners have presented results in the following conferences and panel discussions:

- Ignite the Spark (academia workshop), "Shared electric vehicles and the energy transition"(22 May 2023)
- International Society for Professional Innovation Management (ISPIM) Innovation Conference, Ljubljana, SI (4-7 June 2023)
- RD Management Conference 2023 "Responsible and Responsive Innovation for a Better Future", Seville, ES (19-21 June 2023)
- 8th International Conference on New Business Models (NBM2023), Maastricht, NL , "Multi-Stakeholder Collaboration in Grid-Aware Mobility: an Ecosystem Approach" (22-23 June 2023)
- Project presentation and knowledge exchange meeting with Austrian power grid and Svenska Kraftnät, Vienna, AT (22 August 2023)
- eMOKON B2B e-Mobility congress 2023, Teesdorf, AT (14 September 2023)
- EMC-Kongress at Johannes Kepler University Linz, Linz, AT (22 September 2023)
- Energy Reform Group Workshop organized by the Technical University of Munich (TUM), "Shared Mobility in Residential Buildings - Findings from Focus Groups in Israel" (10 September 2023)
- 12th DACH+ Conference on Energy Informatics, Vienna, AT (4-6 October 2023)
- Symposium Energieinnovation 2024, Graz, AT (14-16 February 2024)

- GAMES Project presentation for Verbund (largest Austrian energy supplier), Vienna, AT (19 March 2024)
- Panel discussion at Vehicle and Grid Forum 2024, Graz, AT (28 May 2024)
- Panel discussion at EV PV Power Day of Bundesverband e-Mobility Austria and Austrian Chamber of Commerce, Wien, AT (07 June 2024)
- CIRED 2024 "Increasing Distribution Network Hosting Capacity", Vienna, AT (19-20 June 2024)
- eMOKON B2B e-Mobility congress 2024, Teesdorf, AT (12 September 2024)

6 Publications and other communications

GAMES deliverables are freely available for download on the GAMES website (GAMES-Innovation-publication). These reports are:

- Policy Brief - Energy-mobility sector coupling through smart and bidirectional vehicle charging
- Shared mobility user characteristics and transport mode decision-making
- Industry Whitepaper - New Business Opportunities Leveraging the Flexibility Potential Of Electric Shared Vehicle Fleets
- Digital cross-sector platform proof-of-concept

In addition, various presentations and publications at international conferences and in academic journals include:

- Nespoli L., Wiedemann N., Suel E. et al. "*National-scale bi-directional EV fleet control for ancillary service provision*". Energy Inform 6 (Suppl 1), 40 (2023). Available at: link
- Thelen, M., Pressmair, G., Lassnig, M., Hornung- Prähauser, V. *Electric Vehicles as Flexibility Assets: Unlocking Ecosystem Collaborations*, New Business Models Conference Proceedings (2023), Available at: link
- Thelen, M., Hornung- Prähauser, V., Lassnig, M.,Pressmair, G. *Emergence of New Ecosystems for Innovative e-Mobility Services: Exploring Business Model Patterns for Vehicle-to-Grid Technology*, R&D Management Conference 2023, Seville, ES (2023), Available at: link
- Thelen, M., Hornung- Prähauser, V., Lassnig, M.,Pressmair, G. *How to Innovate Evehicle Charging in Support of the Energy Transition.*, ISPIM Innovation Management Conference 2023, Ljubljana, SI (2023), Available at: link
- Rocchetta, R., Nespoli, L., Medici, V., Shemesh, A., Parag, Y., and Maayan Tardif, J. (2024). *Optimization of Mobility Incentives in Electric Vehicle Car Sharing Systems: A Reinforcement Learning Framework*. Available at SSRN: link
- Pressmair, G., Papouschek, J., Mayr, M., Thelen, M.*Smart charging oder V2G? Das wirtschaftliche Flexibilitätspotential von e-Fahrzeugflotten*, Symposium Energieinnovation Conference, Graz, AT (2024), Available at: link
- Pressmair, G., Papouschek, J., Rocchetta, R. *A Business Model Or A Tariffing Issue? Peak Shaving In Local Distribution Grids Through Electric Vehicle Fleets*, CIRED 2024 Vienna, AT, (2024). Available at: link
- Shemesh, A., Parag, Y. *Sharing a car with your next-door neighbor: Motivations and barriers to adoption of low-carbon mobility in Israel*, Energy Research & Social Science, Volume 114, 2024, 103586, ISSN 2214-6296, Available at: link
- (pending) Papouschek, J.,Pressmair, G. *Electric vehicles bridging the storage gap in a fully renewable energy system? Results from Tel Aviv*, International Energy Economic Conference 2025, Vienna, AT (2025)

Wie E-Autos die Netze stabiler machen

Da immer mehr Haushalte Strom sowohl in das Netz einspeisen als auch aus diesem herausholen, wird das Netzmanagement komplizierter. In Salzburg werden nun Ideen gesucht, inwieweit Akkus von E-Autos als Zwischenspeicher helfen können.

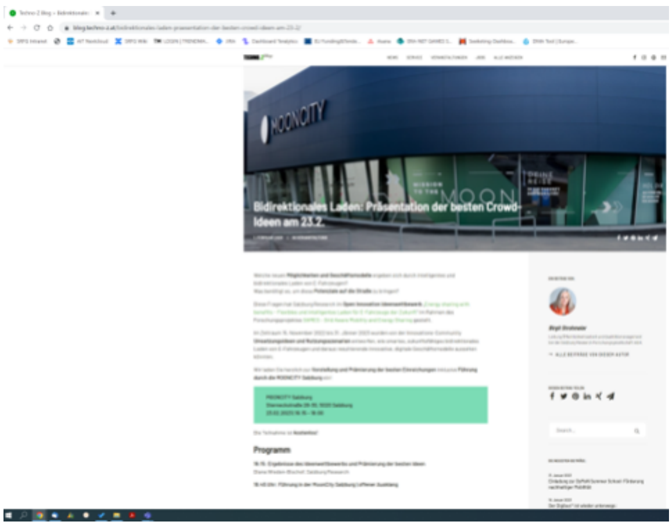

Figure 32: Press releases related to GAMES activities

7 Annex

7.1 Coordinates matching and shortest path

A coordinates matching approach is considered to ensure coherence between the geographical coordinates in the mobility data $\mathcal{D}_{tr}$ and the topological structure of $\mathcal{G}_{tr}$. This is necessary to ensure consistency of the mobility assessment and good results. The nearest graph's node to the starting and ending trip location x_s, x_e. Let x_e and x_s be the starting and ending points of a given trip, the closest node to the ending and starting location of the trip are given by $n_e \in \mathcal{N}_{tr}$ and $n_s \in \mathcal{N}_{tr}$, respectively. For all $(x_e, x_s) \in \mathcal{D}_{tr}$, the corresponding starting and ending nodes are identified numerically by minimizing a distance metric $d(n, x)$ between the graph nodes and trip sources and destinations, for instance, by minimizing the Euclidean distance between n and x. The shortest route between the starting and ending nodes (n_e, n_s) can be also analyzed using the macroscopic transportation model, and, in this work, are compared to the information available in the data set for preliminary screening of trips. This problem can be tackled by identifying a sequence of graph edges $p_i = (s_1, s_2, ...)$, i.e., the route with the shortest travelling time, for all trip samples i in the data set. Established graph algorithms, such as Dijkstra's method [48], are adopted to search $\mathcal{G}_{tr}$, which can efficiently find the shortest path between two nodes.

7.2 Allocation methods for charging stations

Approaches based on density-based clustering [13], zoning [14], and station-assigned zoning can be used to partition the operational area in zones. The operational area of the system is given by the set of coordinates $\mathcal{Z} := \left\{x \in \mathbb{R}^2 : g(x) = 1\right\}$, where $g(x) = 1$ indicates that the coordinate x belongs to the operational area of the FFEVCS system. Similarly to [1], station-assigned zones are defined here by solving a maximum demand coverage problem and dividing $\mathcal{Z}$ into $\mathcal{Z}_i, i = 1, ..., n_z$ non-overlapping areas such that $\bigcap_{i=1}^{n_z} \mathcal{Z}_i = \emptyset$ and the union is equal to the control area $\mathcal{Z} = \bigcup_{i=1}^{n_z} \mathcal{Z}_z$. Each zone is represented by a Voronoi cell built over a k-means cluster of the trip starting coordinate from $\mathcal{D}$. Algorithm 2 presents the allocation procedure for the charging stations. A station cs is assigned to one zone in correspondence with the cluster's centroid and a Voronoi quantization approach, *VQ*, applied for this task. A *cKDTree* model determine $d_{x \to \mathcal{C}}$ via efficient nearest-neighbour lookup, i.e., computation of the minimum distances between the coordinates in the set $\mathcal{C}$ and a query point x. After allocation and partitioning, each event in $\mathcal{D}$ is equipped with two additional indices, one for the origin area and one for the destination of the event.

7.2.1 Optimal coverage problem: an overview

The optimal coverage problem is a significant area within combinatorial optimization, focusing on selecting a subset of elements from a given set to maximize a specific coverage criterion while adhering to constraints. This problem is particularly relevant in the context of transportation systems, where the goal is often to ensure that various demand points are adequately served by resources, such as charging stations. Formally, let X_d represent the domain set of elements, such as demand points and their coordinates, and x_{st} denote a family of subsets of X_d, e.g., representing a vector of coordinates from the set X_d where to allocate stations. A function $f(X_{st})$ defines a coverage function, e.g., the number of demand points X_d covered the union of disks centred on the coordinates in X_{st}. The optimization problem can be mathematically defined as follows:

$$
\begin{aligned}
&\max_{x_{ev}} H(x_{st}) \\
\text{s.t. } & x_{st} \subseteq X_{ev} \\
& f_{eq}(x_{st}) = 0 \\
& fin(x_{st}) \leq 0
\end{aligned}
\tag{14}
$$

Here, $f_{eq}(x_{st}), fin(x_{st})$ represent a set of equality and inequality constraints imposed on the chosen subset x_{st}, such as cardinality constraints or capacity constraints. Solving optimal coverage problems involves finding an optimal subset X_{ev} that maximizes the specified coverage function while adhering to the given constraints. Recently, [49] proposed an interesting formulation for solving a mobility demand coverage problem dynamically. Specifically, the authors focuses on vehicle re-balancing/coverage and seek minimization of a coverage function $H : \mathcal{Z} \to \mathbb{R}$ specifically computed on areas defined by a new Voronoi-inspired partitioning scheme. The approach combines traditional Voronoi cells with a disk centred on an EV coordinate and introduce a treatable reformulation of the problem while accounting for the probability density $f_{dem}(x) : \mathcal{Z} \to [0, 1)$ of the mobility demand. By assuming that the distribution of departures can be used to approximate the demand Probability Density Function (PDF), the objective introduced by [49] reduces to:

$$H(f_{dem}(q), x_{ev}) = -\sum_{i=1}^{n_{EV}} \int_{q \in \mathcal{Z}_i \subset \mathcal{Z}} ||x_{ev,i} - q||^2 f_{dem}(q) dq, \tag{15}$$

where $\mathcal{Z}_i = S_i \bigcap V_i \mathbb{R}^2$ defines the integration domain centered around the i-th EV and it is defined by the intersection of $S_i = q \in \mathcal{Z} : ||x_{ev,i} - q|| \leq r$ (a disk centered in $x_{ev,i}$) and the Voronoi cell,

$$V_i(x_{ev,i}) = \{q \in \mathcal{Z} : ||x_{ev,i} - q|| \leq ||x_{ev,j} - q||, \ \forall j, j \neq i, \}.$$

The set $S = \bigcup_{i=1}^{n_{EV}} S_i$ is the union set (the area covered by the fleet), $\bigcup_{i=1}^{n_{EV}} V_i = \mathcal{Z}$, and $\bigcap_{i=1}^{n_{EV}} \mathcal{Z}_\rangle = \emptyset$. The radius r defines a radius of influence, e.g., $r = 500$ meters. The quantity $f(||x_{ev,i} - q||^2)$ defines a performance score dependent on the squared euclidean distance between a coordinate q and vehicle i to be integrated within $\mathcal{Z}_i$. By maximizing $H(f_{dem}(q), x_e v)$, a set of coordinates $x_e v$ is computed as follows (see [49] for an iterative solution approach):

$$\max_{x_e v} H(f_{dem}(q), x_{ev}) \quad s.t. \quad x_{ev,i} \in \mathcal{Z}_i, \ \forall i = 1, ..n_{EV} \tag{16}$$

The approach in [49] requires $f_{dem}(q)$ to be available to estimate H. Moreover, $f_{dem}(q)$ can be as we will see, time-variant such that an optimal coverage vector x_{ev} selected according to $f_{dem}(q)$ may result in poor coverage if compared to an optimal distribution selected according to a conditional distribution $f_{dem}(q|\tau)$. If an estimate for $f_{dem}(q|\tau)$ is available, such that $\tau \in [h, h + \Delta_t]$ belongs to the next time window Δ_t, a simple change to H which allows accounting for time-variability of the demand distribution is given by:

$$H(f_{dem}(q|h), x_e v) = \sum_{i=1}^{n_{EV}} \int_{q \in \mathcal{Z}_i} ||x_{ev,i} - q||^2 f_{dem}(q|h) dq. \tag{17}$$

Note that a maximization of the coverage $H(f_{dem}(q|\tau), x_e v)$, when $f_{dem}(q|\tau)$ is approximated by an empirical sample distribution $\hat{f}_{dem}$, has some analogy with an optimal transport problem, where the goal is to minimize a 'probabilistic distance' between the empirical distribution defined by $x_e v$ and the distribution of demand points:

$$\min_{x_e v} d(f_{x_{ev}}, \hat{f}_{dem}(q|\tau)) \quad s.t. \quad x_{ev,i} \in \mathcal{Z}_i, \ \forall i = 1, ..n_{EV} \tag{18}$$

There is an extensive academic literature on probabilistic distance metrics $d(\cdot, \cdot))$. These metrics encompass a diverse array of techniques, including but not limited to the Kullback-Leibler divergence [50], Wasserstein distance [51, 52], Bhattacharyya distance [53, 54], Shannon divergence [55], Hausdorff distance [56], Bray-Curtis distance [57, 58], Carathéodory distances [59], and Kolmogorov-Smirnov distance [60], to name some of the most applied distance metrics. This diverse set of metrics equips researchers and practitioners with valuable tools for both the assessment and fine-tuning of the distribution of EVs, thereby making substantial contributions to the ongoing progress in electric mobility solutions.

7.2.2 Optimized coverage of charging stations and EVs in Tel Aviv

An optimal spatial distribution of EVs across the Tel Aviv network has identified by applying the optimal coverage problem framework outlined in Section 7.2.1, program 14. Our primary objective was to achieve an optimal placement of coverage points, stations denoted as x_{st}, by minimizing a probabilistic distance metric between the distribution of demand points and the coordinate vector x_d. To accomplish this, we selected an radius of influence of each station $x_{st,i}$ of 500 meters and maximize the number of departure points covered by the union area. Figure 33 demonstrate the efficacy of the maximum coverage problem on the Tel Aviv network (the panel on the right-hand side), where red disks display the area covered by the designated stations (red crosses). The two panels on the left-hand side demonstrate an alternative allocation scheme where maximum coverage is achieved by allocating stations in the k-mean centroid of demand coordiantes.

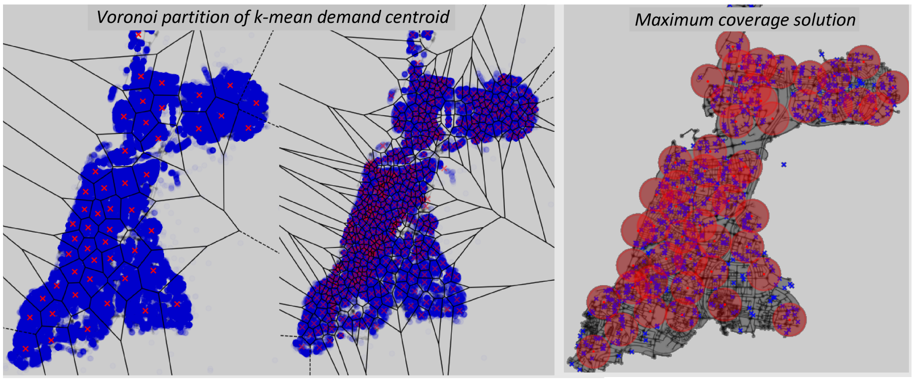

Figure 33: Left-hand side: Voronoi-inspired partitioning of k-mean centroid of departure coordinates with k =50 and k=500. Right-hand side: An example of optimal coverage for the Autotel case study in Tel Aviv. The red disks indicate the area covered by the designated location, e.g., a set of 50 designated spots where to relocate the EVs to guarantee high grid coverage.

7.3 Modelling the economic benefits

In the course of the GAMES-project, four different case studies are presented, analyzing the economic and ecological potential of stationary EV car-sharing fleets. The potential analysis is accomplished by E7 using mathematical optimization tools, i.e., employing the software *GAMS* (Generic Algebraic Modeling System). The optimization model is formulated as a linear optimization problem and is applied in different variations for the various case studies. Each version contains the following three core elements:

Objective Function The core formula of the optimization model, the objective function, is used to maximize or minimize the system's total cost to determine the maximum potential of the fleet:

$$\text{Cost} = \sum_{t,a} p_{\text{sup}}(t,a) \cdot E_{\text{sup}}(t,a) - \sum_{t,a} p_{\text{fi}}(t,a) \cdot E_{\text{fi}}(t,a) + \sum_{t,a,cs} \frac{E_{ev,d}(t,a,cs)}{\eta_{ev,d}(cs)} \cdot p_{\text{deg}} + \sum_{t,a,cs} \frac{E_{ev,c}(t,a,cs)}{\eta_{ev,c}(cs)} \cdot p_{\text{deg}}$$

The Cost function calculates the total economic cost of operating the system, considering multiple factors:

– The first term represents the cost of energy supplied by the grid, where $p_{\text{sup}}(t,a)$ is the price of the supplied energy and $E_{\text{sup}}(t,a)$ is the amount supplied at time t for actor a.

- The second term accounts for revenue generated by feeding energy back into the grid via vehicle-to-grid (V2G) technology, where $p_{\text{fi}}(t, a)$ is the price of fed-in energy and $E_{\text{fi}}(t, a)$ is the amount fed back.
- The third term captures the cost of battery degradation from energy discharged, with $E_{ev,d}(t, a, cs)$ representing the energy discharged, $\eta_{ev,d}(cs)$ the efficiency of discharging, and p_{deg} the battery degradation cost per unit of energy cycled.
- The fourth term represents the cost of battery degradation during charging, with $E_{ev,c}(t, a, cs)$ denoting the energy charged and $\eta_{ev,c}(cs)$ the efficiency of charging.

The variables and parameters in the equation depend on three indices: t (time step), a (actor, such as a charging station), and cs (charging session). This formulation ensures that all costs associated with grid interaction and battery usage are included, allowing for optimization of the economic performance of stationary EV car-sharing fleets. In addition to the core equations described above: 1. Energy Balance Equations 2. EV battery equations (SoC consistency equations) and other constraints are included based on the different use cases addressed. For example, in the second case study, the costs of a peak power tariff are included in the objective function. Constraints generally ensure realistic charging behaviour, for example, they force the EVs to charge and discharge energy within their battery capacities. Figure 34 shows a schematic illustration of the core elements of the optimisation model. It represents the maximum possible applications of the model. Depending on the case study and the scenario investigated, certain sub-elements of the optimisation model are applied, and the results are analysed.

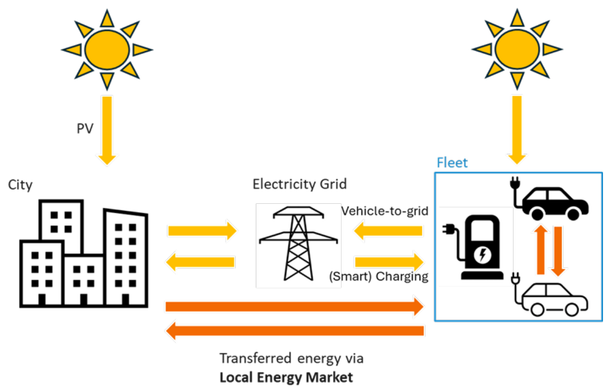

Figure 34: Overall model outline showing the maximum potential applications of the optimization model.

7.4 Effect of fleet size

This analysis examines nine different fleet size scenarios within the Tel Aviv carsharing system to evaluate their performance in terms of booking revenue. Three key metrics are assessed: the number of bookings, total kilometers traveled, and total minutes in use. It's worth noting that larger fleets may incur additional costs, such as maintenance, which are proportional to the number of electric vehicles (EVs) in the fleet. To account for this, we also investigated the corresponding performance metrics normalized per vehicle. The fleet sizes under consideration span from 5 vehicles to larger fleets comprising 500 cars. For each fleet size, we conducted 100 simulation episodes utilizing the gym-based simulation environment. Each episode consists of five steps, with each step representing 12 hours. At the outset of each episode, we initialized the distribution of EVs across 12 charging stations with predefined locations. Performance scores were collected at the conclusion of each episode step, providing a comprehensive evaluation of fleet performance across various sizes.

Figure 35 presents the mean values of these diverse performance metrics for the nine fleet sizes under consideration. It offers insights into the average performance, both in its raw unnormalized

	n_booked	km_driven	min_driven	km_driven_per_car	min_driven_per_car	n_booked_per_car
5	21.33200	123.95200	483.47200	24.79040	96.69440	4.26640
25	98.90800	575.44400	2275.21600	23.01776	91.00864	3.95632
50	179.84800	1036.89600	4088.32400	20.73792	81.76648	3.59696
75	243.45600	1410.97600	5581.40400	18.81301	74.41872	3.24608
100	292.14800	1694.74400	6711.74800	16.94744	67.11748	2.92148
150	365.15600	2115.71200	8377.08000	14.10475	55.84720	2.43437
200	415.02400	2411.35200	9570.10000	12.05676	47.85050	2.07512
250	447.49200	2594.50400	10265.25200	10.37802	41.06101	1.78997
500	512.86800	2987.87200	11779.56000	5.97574	23.55912	1.02574

Figure 35: The mean number of bookings, km and minutes driven (unormalized and normalized) for 9 different fleet sizes, from 5 to 500.

form and as normalized scores per car. This table enables us to discern the patterns in these metrics as fleet size varies. Figure 36 further enriches the analysis by examining the empirical cumulative distribution functions (ECDFs) for critical performance metrics. The top panel of the figure displays the ECDFs for the number of booking events, total distance traveled (in kilometers), and total duration (in minutes) across different fleet sizes. These results stem from 50 Monte Carlo simulations In the bottom panel, these results are normalized by fleet size, offering insights into the average performance per vehicle within the fleet. These visualizations provide a detailed perspective on how fleet size influences the distribution of performance metrics, aiding in making informed decisions about fleet management strategies.

7.5 Modeling the free-floating carsharing simulation environment

This section outlines the key components that make up the gymnasium simulation environment and the discrete-event simulation model. These elements form the foundation for accurately replicating the dynamics of an electrified carsharing system, including interactions with the energy grid, user demand, and fleet management strategies.

7.5.1 A binary transportation mode choice

The BMC model is defined as follows:

$$mc = \begin{cases} 1 & \text{if } \exists\, ev \in \mathcal{E} : c_{ev}^{\text{carsharing}} \\ 0 & \text{otherwise} \end{cases} \tag{19}$$

where the logical condition $c_{ev}^{\text{carsharing}} = c_{1,ev} \wedge c_{2,ev} \wedge op_{ev} \neq 0$ determines as subset of $ev \in \mathcal{E}$ that are suitable for the users. If at least one among the non-booked EVs ($op_{ev} \neq 0$) meets the following two conditions, then a carsharing transportation mode is accepted, mc set to 1 and the closest among the available cars selected for the trip. The conditions for carsharing are defined as follows:

$$\begin{aligned} c_{1,ev} &: d(x, x_{ev}) \leq d_{walk}(a_o), \\ c_{2,ev} &: SoC_{ev} \geq SoC_s^{min}, \end{aligned} \tag{20}$$

where $c_{1,ev}$ determines the set of $ev \in \mathcal{E}$ that are close enough to the renter, i.e., if an incentive-dependent maximum walking distance $d_{walk}(a_o)$ is lower than the Euclidean distance between the user's coordinates (origin of the trip) and the vehicle, $d(x, x_{ev})$. The second condition $c_{2,ev}$, checks weather a condition on the minimum SoC is met by a car. This can be seen as a proxy for the user's perceived reliability of the car, i.e., sufficient fuel to move from x_s to x_e.

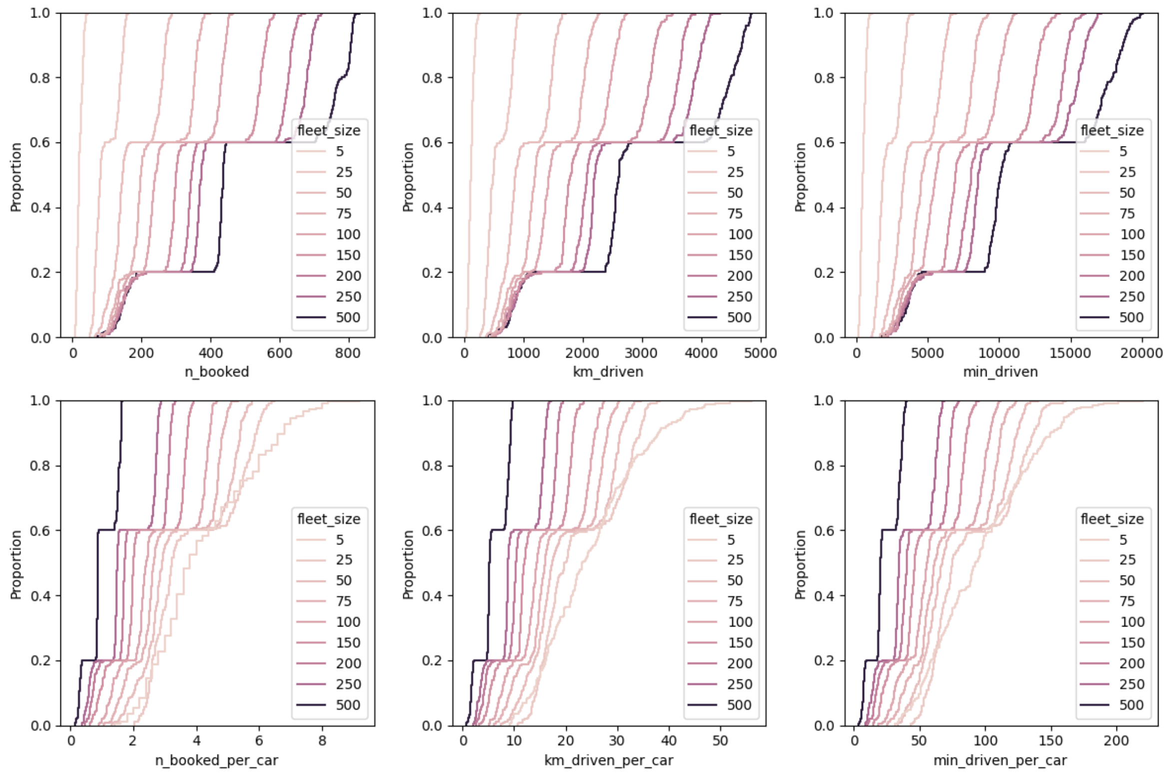

Figure 36: The effect of fleet size, 5 to 500, on the empirical CDF the number of booking events, total distance traveled [km], and total duration [min]. The top panel show the normalized results from 50 MC samples of a 5-steps episode trajectory. The bottom panel show the same results but normalize by the fleet size, e.g., average number of booking events for each EV.

7.5.2 A model for Hybrid relocation strategies

This study considers HR in conjunction with varying incentives to boost user participation in relocation activities and improve fleet utilization. Hybrid relocation strategies leverage both crew and user participation in relocation activities. Within the simulation environment, it is assumed that carsharing staff (the crew) perform daily relocation activities. Crew interventions are scheduled once a day at a fixed hour, h^{crew}, with a simple rules-based strategy determining which vehicles should be relocated based on their SoC and idle time. In contrast, user-led relocation activities rely on active user participation. In this framework, this is achieved by offering relocation credits (destination incentives). While user-led relocation is cheaper but less reliable, crew-led relocation is more dependable but costly. The combination of these two approaches defines a hybrid relocation strategy in the proposed environment.

Crew-led relocation: To model crew-led activities, a Boolean indicator function is introduced to identify vehicles for relocation by the crew according to the following rules:

$$c_{ev}^{crew} = \begin{cases} 1 & \text{if } h = h_{crew} \land c_{ev}^{crb} \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

where c_{ev}^{crew} is true if the car indexed by ev will be relocated by the crew and c_{ev}^{crb} defines the simple relocation rule:

$$c_{ev}^{crb} : (SoC_{ev} \leq SoC_{cb}^{min} \ \land \ \tau^{crew1} \leq \tau_{ev}^{idle}) \lor (\tau^{crew2} \leq \tau_{ev}^{idle}),$$

where SoC_{cb}^{min} defines a lower bound on the SoC for crew relocation activities, τ^{crew1} is a warning idle time, e.g., 24 hours, and τ^{crew2} is a maximum allowable idle time, e.g., more than four days inactive. Hence, the crew becomes active only if the simulated hour matches the scheduled hour for crew activities and if the SoC and idle time conditions are met. For example, if a car is parked for less than a day, it will not be relocated. If it exceeds this time, it will be relocated if the SoC is lower than SoC_{cb}^{min}. Finally, if the idle time is more that four days, the car will be relocated regardless of its state of charge. For simplicity, it is assumed that a sufficiently large crew can accommodate any number of relocations. The total number of relocations performed by the crew at each time step, denoted as $n_t^{crew} = \sum_{ev \in \mathcal{E}} c_{ev}^{crew}$, determines the relocation cost for the crew, see eq. (12). We assume no delays in the relocation costs, that is, r_t^c is observed in correspondence of the designated hour h_{crew}.

User-led relocation: Complementing the crew-led relocation activities, a user-led relocation strategy is also considered, which encourages carsharing customers to drop off and plug in EVs at designated charging stations. If a car is rented, see the BMC model in equation (19), the user may be offered a mobility credit to relocate the vehicle. This offer is extended if the state of charge of the rented EV is below a minimum value, $SoC_{ev} \leq SoC_{ub}^{min}$, and based on the current value of the destination incentive a_d. The acceptance (or rejection) of this offer is determined by a simple acceptance model as follows:

$$c_{ev}^{user} = \begin{cases} 1 & \text{if } d_{x_e \to x_C} \leq d_{walk}(a_d) \wedge a_d > 0 \wedge op_{ev} = 0 \\ 0 & \text{otherwise} \end{cases} \tag{22}$$

where $d_{x_e \to x_C} = \min_{cs \in \mathcal{C}} d(x_e, x_{cs})$ is the distance between the user's destination x_e and the closest station $x_{cs} \in \mathcal{C}$. An acceptable relocation distance $d_{walk}(a_d)$, which depends on the assigned destination incentive, determines whether the relocation offer is accepted or rejected by the user. If the indicator c^{user} equals one, the user accepts the offer and relocates and plugs-in the vehicle ev to the closest charging station at the end of the trip. Otherwise, the offer is rejected, and the EV is dropped off at the original destination (disconnected). At each simulation step, the total number of relocations performed by users is computed as $n_t^{user} = \sum_{ev \in \mathcal{E}} c_{ev}^{user}$, which is then used to calculate the user relocation cost in the reward function. Like for crew-led activities, we assume zero delays in the relocation costs r_t^u.

Jointly, equations (21) and (22) define various hybrid relocation strategies based on the applied incentive a_d. For instance, for $a_d \leq 0$, this model yields crew-led relocation only. Conversely, by applying positive credits, various HR strategies can be defined. It is important to note that the RL policies introduced in this work also represent hybrid relocation strategies, where RL agents dynamically and heterogeneously apply incentives to modulate user participation.

7.5.3 Apply ED policy and update SOC

The charging of EVs is assumed governed by a battery energy storage module. Battery control actions within this module are determined by an economic dispatch policy π_{ED}, which determines the charging strategy for the subset of EVs parked and connected to charging stations. For this, a versatile and general battery model is adopted as in [61]:

$$SoC_{ev}^{new}(t + \delta_t) = SoC_{ev}^{old}(t) + \frac{\pi_{ED} \eta_{ev} \delta_t}{E^{rd}} \text{ if } op_{ev} = 2, \tag{23}$$

where SoC_{ev}^{new} represents the updated state of charge of the EV (capped to 100), SoC_{ev}^{old} is its previous state of charge, the charging policy π_{ED} determines the positive charging power of the battery in kW, δ_t is the time difference in hours between consecutive events in the simulated dataset of demands $\mathcal{D}_{sim}$, η_{ev} is the charging efficiency, and E^{rd} is the energy storage capacity of the battery kWh installed on $ev \in \mathcal{E}$. Note that charging operations are executed only if the EV is parked and grid-connected, i.e., $op_{ev} = 2$, and if $SC < 100\%$. That is, charging is applied only to plugged-in EVs and if not fully charged. Based on (23), the energy demanded for charging the fleet is computed as follows:

$$E_{ev}^{+}(\pi_{ED}) = E^{rd}(SoC_{ev}^{new}(\pi_{ED}) - SoC_{ev}^{old}), \tag{24}$$

where $E_{ev}^{+}(\pi_{ED})$ represents a positive increment in the energy stored in an EV. A simple policy $\pi_{ED} = P_{ev}^{rd}$ is applied and sets the charging power to a fixed and vehicle-dependent rated power, i.e., the SoC of an EV is adjusted according to a "charge as soon as possible" principle, which dictates that the vehicle should be charged at the earliest opportunity. The SoC of an ev decreases if utilized by a user to carry out a trip during the reservation process. At the end of a trip event in $e \in \mathcal{D}_{\text{res}}$, which is characterized by duration τ and distance traveled l, a linear SoC loss is computed as follows [62]:

$$SoC_{ev}^{new}(t+\tau) = SoC_{ev}^{old}(t) - \frac{100l}{d_{ev}^{range}} \tag{25}$$

where the state-of-charge decreases proportionally to the distance traveled in kilometers l divided by the maximum driving range at full storage capacity, d_{ev}^{range}.

7.6 Algorithms

Algorithm 2 Charging Station Allocation for Optimal Demand Coverage

Function `AllocateCS`($\mathcal{D}, Conf$)**:**

- **Input:** $\mathcal{D}$ (trip data), $Conf$ (configuration)
- **Output:** $\mathcal{C}$ (charging stations), $\mathcal{D}$ (mobility data), $d_{x\to\mathcal{C}}$ (distance model)
- $n_z \leftarrow Conf$; `// Get number of stations/zones`
- $(x_s, x_e) \leftarrow \mathcal{D}$; `// Get trip coordinates`
- $(x_1, \ldots, x_{n_z}) \leftarrow \mathsf{KM}(x_s, n_z)$; `// Get` n_z`-means clusters`
- $x_\mathcal{C} \leftarrow (x_1, \ldots, x_{n_z}) \in \mathbb{R}^{2\times n_z}$; `// Cluster coordinates`
- $\mathcal{C} \leftarrow x_\mathcal{C}$; `// Assign stations coordinates`
- $s, e \leftarrow \mathsf{VQ}(x_e, x_s, x_\mathcal{C})$; `// Get zone indices for trips`
- $\mathcal{D} \leftarrow (s, e)$; `// Append zone indices to trips`
- $d_{x\to\mathcal{C}} \leftarrow \mathsf{cKDTree}(x, x_\mathcal{C})$; `// Save distance lookup model`
- **return** $\mathcal{C}, \mathcal{D}, d_{x\to\mathcal{C}}$; `// Return allocated stations, updated demand, and distance model`

Allocation method for CS The feature δ_t between events is exploited within the simulation loop to efficiently simulate reservations events, see Section 7.5 for details. Algorithm 3 illustrates the event-based simulation approach that determines reservations $\mathcal{D}_{\text{res}}$ and revenues together with relocation and charging costs at each step of an episode.

Discrete event simulation algorithm for FFEVCS

Simulation environment for training and testing reinforcement learning agents **INITIALIZE:** The environment loads a configuration file, a transportation graph $\mathcal{G}_{tr}$, an operating area $\mathcal{Z}$, and the dataset $\mathcal{D}_{tr}$. The simulated fleet $\mathcal{E}$ is initialized, and charging stations $\mathcal{C}$ are allocated to n_z non-overlapping zones (see the allocation procedure in the appendix, Algorithm 2). **RESET:** An episode starts by sampling the state s_0 using this method, which reset the state of $\mathcal{E}$ and $\mathcal{C}$ to their initial states, e.g., moving EVs to their starting locations. The SoC and τ^{idle} of the fleet are randomized according to a stationary uniform distribution f_{s_0}. **RUN EPISODE:** This step of Algorithm 4 simulates a episode for the fleet. First, the environment reset its initial state. Then mobility incentives a_t are applied for all $t = 0, ..., T$ in consecutive steps and step size defined by Δ_t, e.g., one hour, samples of demand $\mathcal{D}_{\text{sim}}$ are drawn from the probabilistic forecaster (see Algorithm 1) and carsharing reservations, relocation activities, and energy for charging the fleet estimated (see definition of a discrete-event model in the next section). A new state vector s_{t+1} and the total reward signal r_t are then obtained from the new system stat and the individual rewards. **DISCRETE-EVENT SIMULATOR** $\mathcal{M}$**:** The discrete-event simulator $\mathcal{M}$, as Algorithm 3 and Figure 16, handles system updates, transportation and relocation mode choices, and SoC updates based on the charging policy. Simulated reservations $\mathcal{D}_{\text{res}}$ depend on user willingness to walk, influenced

Algorithm 3 DISCRETE-EVENT SIMULATOR: Pseudo code

Input: $\mathcal{D}_{sim}$ (demand samples), a_t (actions), $\mathcal{E}$ (fleet of electric vehicles)
Output: $\mathcal{D}_{\text{res}}$ (reservations), r_t^{el} (electricity cost), r_t^{rel} (relocation cost)
Function $\mathcal{M}(\mathcal{D}_{sim}, a_t, \mathcal{E})$**:**

```
  $\mathcal{D}_{res}, r_t^{el}, r_t^{u}, r_t^{c} \leftarrow \emptyset, 0, 0, 0$ ;            // initialize variables
  $(a_o, a_d) \leftarrow a_t$ ;                                   // get origin-destination incentives
  $h \leftarrow t$ ;                                              // get current time step hour
  // Apply crew-based relocation and calculate cost as in eq. (21)
  foreach $ev \in \mathcal{E}$ do
  |   $c_{ev}^{crew} \leftarrow SoC_{ev}, \tau_{ev}^{idle}$ ;          // calculate relocation parameters
  $n_t^c \leftarrow \sum_{ev\in\mathcal{E}} cb_{ev,t}$ ;                     // total number of relocations
  $r_t^c \leftarrow -\rho^c n_t^c$ ;                                  // relocation cost (crew)
  foreach $e \in \mathcal{D}_{sim}$ do
  |   $(x_s, x_e, l, \tau, \delta_t) \leftarrow e$ ;                    // get event details
  |   // Update time stamps and energy demand
  |   if $\delta_t > 0$ then
  |   |   foreach $ev \in \mathcal{E}$ do
  |   |   |   if $op_{ev} = 0$ then
  |   |   |   |   $\tau_{ev}^{return} \leftarrow \max(0, \tau_{ev}^{return} - \delta_t)$ ;   // update time-to-return
  |   |   |   if $op_{ev} \neq 0$ then
  |   |   |   |   $\tau_{ev}^{idle} \leftarrow \tau_{ev}^{idle} + \delta_t$ ;          // update idle time
  |   |   |   if $\tau_{ev}^{return} = 0$ then
  |   |   |   |   $op_{ev} \leftarrow 1$ ;                        // update operational state
  |   |   |   if $op_{ev} = 2$ then
  |   |   |   |   $E_{ev}^{+} \leftarrow \pi_{ED}$ ;                    // apply charging policy
  |   |   $r_t^{el} \leftarrow r_t^{el} - \sum_{ev} \rho_t^{el} E_{ev}^{+}$ ;          // add electricity cost
  |   // Assign transportation mode as in eq. (19)
  |   $mc \leftarrow BMC(x_s, a_o)$ ;                   // apply origin incentives and mode choice
  |   if $mc = 1$ then
  |   |   $\mathcal{D}_{res} \leftarrow e$ ;                          // append event to reservation set
  |   |   $ev \leftarrow \arg\min_{ev\in\mathcal{E}} d(x_e, x_{ev})$ ;     // select vehicle based on distance
  |   |   $op_{ev} \leftarrow 0$ ;                                // update operational state
  |   |   $SoC_{ev} \leftarrow SoC_{ev} - l \cdot d_{ev}^{rate}$ ;          // update SoC based on distance
  |   |   // Offer relocation credit to user and add cost, see eq. (22)
  |   |   $c_{ev}^{user} \leftarrow x_e, a_d$ ;                  // calculate user relocation incentive
  |   |   if $c_{ev}^{user} = 1$ then
  |   |   |   $r_t^u \leftarrow r_t^u - \rho^{user}(a_d)$ ;              // add user relocation cost
  |   |   |   $cs \leftarrow \arg\min_{cs\in\mathcal{C}} d(x_e, x_{cs})$ ;     // select closest charging station
  |   |   |   $\mathcal{E} \leftarrow x_{ev} = x_{cs}$ ;                   // assign station to vehicle
  // Compute final relocation cost as in eq. (12)
  $r_t^{rel} \leftarrow r_t^c + r_t^u$ ;                              // total relocation cost
  return $\mathcal{D}_{res}, r_t^{el}, r_t^{rel}$ ;                          // return results
```

Algorithm 4 The Simulation Environment (Gym.Env)

Input: $Conf$
Output: env
Function: INITIALIZE$(Conf)$; `// Initialize the environment`
$\mathcal{G}_{tr}, \mathcal{D}_{tr}, \mathcal{E} \leftarrow Conf$; `// Load streets, trips, fleet`
$\mathcal{C} \leftarrow$ ALLOCATE-CS$(Conf, \mathcal{D}_{tr})$; `// Allocate charging stations`
$env \leftarrow \mathcal{G}_{tr}, \mathcal{D}_{tr}, \mathcal{E}, \mathcal{C}$; `// Initialize environment`
return env ; `// Return the initialized environment`

Input: $Conf$
Output: s_0
Function: RESET$(Conf)$; `// Reset the environment`
$t, t_{dw}, t_{hd}, \mathcal{E}, \mathcal{C} \leftarrow Conf$; `// Reset time indices and sets`
$SoC_{ev} \sim \mathcal{U}(70, 100)$; // Reset SoC for all $ev \in \mathcal{E}$
$\tau_{ev}^{idle} \sim \mathcal{U}(120, 600)$; // Reset idle duration for all $ev \in \mathcal{E}$
$\tau_{ev}^{return} = 0$; // Reset rent duration for all $ev \in \mathcal{E}$
$s_0 \leftarrow env$; // Get state variable at $t = 0$
return s_0 ; `// Return initial state`

Function: RUN EPISODE$(\pi_\theta, Conf)$; `// Run episodic loop`
Input: $\pi_\theta, Conf$
Output: R_{ep}
$T, \Delta_t \leftarrow Conf$; `// Get episode parameters`
reset$(Conf)$; `// Reset initial state`
$t = 0$
while $t < T$ **do**
  $a_t \leftarrow \pi_\theta(s_t)$; `// Apply incentives`
  $\mathcal{D}_{\text{sim}} \leftarrow$ FORECASTER.sample$(\mathcal{D}_{tr}, \Delta_t)$; `// Sample demand from FORECASTER`
  $\mathcal{D}_{\text{res}}, r_t^{el}, r_t^{rel} \leftarrow \mathcal{M}(\mathcal{D}_{\text{sim}}, a_t, \mathcal{E})$; `// Run DISCRETE-EVENT SIMULATOR`
  $r_t^{\text{min}}, r_t^{\text{res}} \leftarrow \mathcal{D}_{\text{res}}$; `// Compute rewards`
  $r_t = r_t^{\text{min}} + r_t^{\text{res}} + r_t^{el} + r_t^{rel}$; `// Total reward`
  $s_t, r_t \leftarrow \mathcal{E}$; `// Get new system state and reward`
  $R_{ep} \leftarrow r_t$; `// Append reward`
  $t = t + 1$; `// Next step`
return R_{ep} ; `// Return episodic rewards`

by incentive a_t. A binary transportation mode choice model iteratively matches user demands with the fleet, moving cars to the trip's ending location when a carsharing trip is accepted.

8 References

[1] S. Weikl and K. Bogenberger, "A practice-ready relocation model for free-floating carsharing systems with electric vehicles – mesoscopic approach and field trial results," *Transportation Research Part C: Emerging Technologies*, vol. 57, pp. 206–223, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X15002399

[2] D. Freund, A. F. Raab, T. Küster, S. Albayrak, and K. Strunz, "Agent-based integration of an electric car sharing fleet into a smart distribution feeder," in *2012 3rd IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*. IEEE, 2012, pp. 1–8.

[3] T. Yoon, C. R. Cherry, and L. R. Jones, "One-way and round-trip carsharing: A stated preference experiment in beijing," *Transportation Research Part D: Transport and Environment*, vol. 53, pp. 102–114, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920917303139

[4] F. Ciari, B. Bock, and M. Balmer, "Modeling station-based and free-floating carsharing demand: Test case study for berlin," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2416, pp. 37–47, 12 2014.

[5] D. J. Mühlematter, N. Wiedemann, Y. Xin, and M. Raubal, "Spatially-aware car-sharing demand prediction," *arXiv preprint arXiv:2303.14421*, 2023.

[6] L. Nespoli, N. Wiedemann, E. Suel, Y. Xin, M. Raubal, and V. Medici, "National-scale bidirectional ev fleet control for ancillary service provision," 2022.

[7] S. Kaushik, H. Kawsar, M. Rouhani, and M. Aiello, "Optimal joint operation of coupled transportation and power distribution urban networks," *Energy Informatics*, vol. 5, p. 35, 2022.

[8] F. Rosato, "Heuristic graph partitioning with preferred cluster sizes and application to the generation of realistic distribution grid topologies," in *2021 IEEE 15th International Conference on Compatibility, Power Electronics and Power Engineering (CPE-POWERENG)*, 2021, pp. 1–7.

[9] A. DImovski, E. Ragaini, D. I. Edeme, S. Corigliano, and M. Merlo, "Gis approach for modeling the electricity grid of small islands using open-source data: The case of isola del giglio," in *2022 IEEE 21st Mediterranean Electrotechnical Conference (MELECON)*, 2022, pp. 536–541.

[10] L. Li and X. M. Chen, "Vehicle headway modeling and its inferences in macroscopic/microscopic traffic flow theory: A survey," *Transportation Research Part C: Emerging Technologies*, vol. 76, pp. 170–188, 2017.

[11] S. Weikl and K. Bogenberger, "A practice-ready relocation model for free-floating carsharing systems with electric vehicles–mesoscopic approach and field trial results," *Transportation Research Part C: Emerging Technologies*, vol. 57, pp. 206–223, 2015.

[12] H. Beenish, T. Javid, M. Fahad, A. A. Siddiqui, G. Ahmed, and H. J. Syed, "A novel markov model-based traffic density estimation technique for intelligent transportation system," *Sensors*, vol. 23, no. 2, 2023. [Online]. Available: https://www.mdpi.com/1424-8220/23/2/768

[13] S. Atev, G. Miller, and N. P. Papanikolopoulos, "Clustering of vehicle trajectories," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 647–657, 2010.

[14] F. Heymann, C. Pereira, V. Miranda, and F. J. Soares, "Spatial load forecasting of electric vehicle charging using gis and diffusion theory," in *2017 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*. IEEE, 2017, pp. 1–6.

[15] S. Ke, L. Chen, J. Yang, G. Li, F. Wu, L. Ye, W. Wei, and Y. Wang, "Vehicle to everything in the power grid (v2eg): A review on the participation of electric vehicles in power grid economic dispatch," *Energy Conversion and Economics*, vol. 3, no. 5, pp. 259–286, 2022.

[16] H. Patil and V. N. Kalkhambkar, "Grid integration of electric vehicles for economic benefits: A review," *Journal of Modern Power Systems and Clean Energy*, vol. 9, no. 1, pp. 13–26, 2020.

[17] O. Frendo, J. Graf, N. Gaertner, and H. Stuckenschmidt, "Data-driven smart charging for heterogeneous electric vehicle fleets," *Energy and AI*, vol. 1, p. 100007, 2020.

[18] M. S. Mastoi, S. Zhuang, H. M. Munir, M. Haris, M. Hassan, M. Alqarni, and B. Alamri, "A study of charging-dispatch strategies and vehicle-to-grid technologies for electric vehicles in distribution networks," *Energy Reports*, vol. 9, pp. 1777–1806, 2023.

[19] Y. Yoshida and Y. Takano, "Linear control policies for online vehicle relocation in shared mobility systems," *Expert Systems with Applications*, vol. 210, p. 118417, 2022.

[20] R. Zhang, H. Kan, Z. Wang, and Z. Liu, "Relocation-related problems in vehicle sharing systems: A literature review," *Computers & Industrial Engineering*, vol. 183, p. 109504, 2023.

[21] C. V. Beojone and N. Geroliminis, "Relocation incentives for ride-sourcing drivers with path-oriented revenue forecasting based on a Markov Chain model," *Transportation Research Part C: Emerging Technologies*, vol. 157, p. 104375, 2023.

[22] G. Guo and T. Xu, "Vehicle rebalancing with charging scheduling in one-way car-sharing systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 5, pp. 4342–4351, 2020.

[23] A. Bogyrbayeva, S. Jang, A. Shah, Y. J. Jang, and C. Kwon, "A reinforcement learning approach for rebalancing electric vehicle sharing systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 8704–8714, 2022.

[24] M. Schiffer, G. Hiermann, F. Rüdel, and G. Walther, "A polynomial-time algorithm for user-based relocation in free-floating car sharing systems," *Transportation Research Part B: Methodological*, vol. 143, pp. 65–85, 2021.

[25] X. Liu, S. W. Kim, and C. Kwon, "An adaptive large neighborhood search method for rebalancing free-floating electric vehicle sharing systems," *Computers & Operations Research*, vol. 155, p. 106220, 2023.

[26] M. P. Fanti, A. M. Mangini, M. Roccotelli, and B. Silvestri, "Innovative approaches for electric vehicles relocation in sharing systems," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 1, pp. 21–36, 2022.

[27] E. Correia Sinézio Martins, J. Lépine, and J. Corbett, "Assessing the effectiveness of financial incentives on electric vehicle adoption in europe: Multi-period difference-in-difference approach," *Transportation Research Part A: Policy and Practice*, vol. 189, p. 104217, 2024.

[28] K. Sevdari, L. Calearo, P. B. Andersen, and M. Marinelli, "Ancillary services and electric vehicles: An overview from charging clusters and chargers technology perspectives," *Renewable and Sustainable Energy Reviews*, vol. 167, p. 112666, 2022.

[29] A. Angelopoulos, D. Gavalas, C. Konstantopoulos, D. Kypriadis, and G. Pantziou, "Incentivized vehicle relocation in vehicle sharing systems," *Transportation Research Part C: Emerging Technologies*, vol. 97, pp. 175–193, 2018.

[30] T. Chouaki, L. M. R. Madrigal, and S. Hörl, "Assessing the impact of monetary incentives for walking using agent-based mobility simulations and discrete mode choice models," in *The 103rd Transportation Research Board (TRB) Annual Meeting*, 2024, pp. 1–20.

[31] R. Curtale, F. Liao, and P. van der Waerden, "Understanding travel preferences for user-based relocation strategies of one-way electric car-sharing services," *Transportation Research Part C: Emerging Technologies*, vol. 127, p. 103135, 2021.

[32] D. Wang and F. Liao, "Incentivized user-based relocation strategies for moderating supply-demand dynamics in one-way car-sharing services," *Transportation Research Part E: Logistics and Transportation Review*, vol. 171, p. 103017, 2023.

[33] J. Alonso-Mora, S. Samaranayake, A. Wallar, E. Frazzoli, and D. Rus, "On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment," *Proceedings of the National Academy of Sciences*, vol. 114, no. 3, pp. 462–467, 2017. [Online]. Available: https://www.pnas.org/doi/abs/10.1073/pnas.1611675114

[34] X. Chang, J. Wu, G. H. de Almeida Correia, H. Sun, and Z. Feng, “A cooperative strategy for optimizing vehicle relocations and staff movements in cities where several carsharing companies operate simultaneously,” *Transportation Research Part E: Logistics and Transportation Review*, vol. 161, p. 102711, 2022.

[35] Y. Yoshida and Y. Takano, “Linear control policies for online vehicle relocation in shared mobility systems,” *Expert Systems with Applications*, vol. 210, p. 118417, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417422015214

[36] A. Bogyrbayeva, S. Jang, A. Shah, Y. J. Jang, and C. Kwon, “A reinforcement learning approach for rebalancing electric vehicle sharing systems,” *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[37] B. Kiumarsi, K. G. Vamvoudakis, H. Modares, and F. L. Lewis, “Optimal and autonomous control using reinforcement learning: A survey,” *IEEE Transactions on neural networks and learning systems*, vol. 29, no. 6, pp. 2042–2062, 2017.

[38] O. Nachum, S. S. Gu, H. Lee, and S. Levine, “Data-efficient hierarchical reinforcement learning,” *Advances in neural information processing systems*, vol. 31, 2018.

[39] N. Wiedemann, V. Wüest, A. Loquercio, M. Müller, D. Floreano, and D. Scaramuzza, “Training efficient controllers via analytic policy gradient,” 2022. [Online]. Available: https://arxiv.org/abs/2209.13052

[40] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” 2017.

[41] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” 2018.

[42] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” 2019.

[43] S. Fujimoto, H. Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods,” in *International conference on machine learning*. PMLR, 2018, pp. 1587–1596.

[44] M. Soppert, C. Steinhardt, C. Müller, J. Gönsch, and P. M. Bhogale, “Matching functions for free-floating shared mobility system optimization to capture maximum walking distances,” *European Journal of Operational Research*, vol. 305, no. 3, pp. 1194–1214, 2023.

[45] M. Ampudia-Renuncio, B. Guirao, and R. Molina-Sanchez, “The impact of free-floating carsharing on sustainable cities: analysis of first experiences in Madrid with the university campus,” *Sustainable Cities and Society*, vol. 43, pp. 462–475, 2018.

[46] Swissgrid Ltd, “Guidelines and Recommendations for the Prequalification of Electric Vehicles,” Swissgrid Ltd, 2023, accessed: 2024-11-25. [Online]. Available: https://www.swissgrid.ch/content/dam/swissgrid/customers/topics/ancillary-services/prequalification/2/Richtlinien-Praequalifikation-e-Autos-de.pdf

[47] A. Shemesh and Y. Parag, “Sharing a car with your next-door neighbor: Motivations and barriers to adoption of low-carbon mobility in israel,” *Energy Research Social Science*, vol. 114, p. 103586, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2214629624001774

[48] E. W. Dijkstra, “A note on two problems in connexion with graphs,” *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.

[49] P. Zhu, I. I. Sirmatel, G. F. Trecate, and N. Geroliminis, “Idle-vehicle rebalancing coverage control for ride-sourcing systems,” in *2022 European Control Conference (ECC)*, 2022, pp. 1970–1975.

[50] S. Kullback and R. A. Leibler, “On information and sufficiency,” *The annals of mathematical statistics*, vol. 22, no. 1, pp. 79–86, 1951.

[51] N. Fournier and A. Guillin, “On the rate of convergence in wasserstein distance of the empirical measure,” *Probability theory and related fields*, vol. 162, no. 3-4, pp. 707–738, 2015.

[52] A. Gray, A. Wimbush, M. de Angelis, P. O. Hristov, D. Calleja, E. Miralles-Dolz, and R. Rocchetta, “From inference to design: A comprehensive framework for uncertainty quantification in engineering with limited information,” *Mechanical Systems and Signal Processing*, vol. 165, p. 108210, 2022.

[53] A. Bhattacharyya, “On a measure of divergence between two multinomial populations,” *Sankhyā: the indian journal of statistics*, pp. 401–406, 1946.

[54] S. Bi, M. Broggi, and M. Beer, “The role of the bhattacharyya distance in stochastic model updating,” *Mechanical Systems and Signal Processing*, vol. 117, pp. 437–452, 2019.

[55] B. Fuglede and F. Topsoe, “Jensen-shannon divergence and hilbert space embedding,” in *International symposium onInformation theory, 2004. ISIT 2004. Proceedings.* IEEE, 2004, p. 31.

[56] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, “Comparing images using the hausdorff distance,” *IEEE Transactions on pattern analysis and machine intelligence*, vol. 15, no. 9, pp. 850–863, 1993.

[57] W. Zhao, L. Yang, C. Dang, R. Rocchetta, M. Valdebenito, and D. Moens, “Enriching stochastic model updating metrics: An efficient bayesian approach using bray-curtis distance and an adaptive binning algorithm,” *Mechanical Systems and Signal Processing*, vol. 171, p. 108889, 2022.

[58] K. R. Clarke, P. J. Somerfield, and M. G. Chapman, “On resemblance measures for ecological studies, including taxonomic dissimilarities and a zero-adjusted bray–curtis coefficient for denuded assemblages,” *Journal of experimental marine biology and ecology*, vol. 330, no. 1, pp. 55–80, 2006.

[59] J. Mitchell, “On carnot-carathéodory metrics,” *Journal of Differential Geometry*, vol. 21, no. 1, pp. 35–45, 1985.

[60] F. J. Massey Jr, “The kolmogorov-smirnov test for goodness of fit,” *Journal of the American statistical Association*, vol. 46, no. 253, pp. 68–78, 1951.

[61] Y. Yang, S. Bremner, C. Menictas, and M. Kay, “Modelling and optimal energy management for battery energy storage systems in renewable energy systems: A review,” *Renewable and Sustainable Energy Reviews*, vol. 167, p. 112671, 2022.

[62] F. Hipolito, C. Vandet, and J. Rich, “Charging, steady-state SoC and energy storage distributions for EV fleets,” *Applied Energy*, vol. 317, p. 119065, 2022.