

PAVING THE PATH TOWARD SMART SUSTAINABLE MOBILITY

Data-Driven Operations Management for Emerging Mobility Systems:
Sharing, Automation, and Electrification

Ramin Ahadi

2025

Paving the Path toward Smart Sustainable Mobility

Data-Driven Operations Management for Emerging Mobility Systems: Sharing,
Automation, and Electrification

Inauguraldissertation
zur
Erlangung des Doktorgrades
der
Wirtschafts- und Sozialwissenschaftlichen Fakultät
der
Universität zu Köln

2025

vorgelegt
von

Ramin Ahadi, M. Sc.

aus

Ghochan

Promotionsausschuss (Doctoral Committee)

Referent (Referee):

Prof. Dr. Wolfgang Ketter

Koreferent (Co-Referee)

Prof. Dr. Konstantina Valogianni

Vorsitz (Chair):

Prof. Dr. Markus Weinmann

Tag der Promotion (Date of Defense):

June 03, 2025

To my family

Acknowledgments

I began my academic journey with a bachelor's degree in mechanical engineering before switching to industrial engineering and operations management for my master's degree. While my master's program provided a solid foundation, it did not offer the depth in data-driven decision sciences needed to make a meaningful contribution to building a sustainable society - one of my lifelong goals. To fill this gap, I pursued a Ph.D. focused on integrating data science and machine learning to improve the efficiency of critical systems in urban mobility and energy.

For the past five years, I have dedicated myself to this challenging yet rewarding mission. This journey, like any ambitious endeavor, has presented numerous obstacles that I have embraced as opportunities for growth. Through these experiences, I have developed expertise in advanced machine learning, deep reinforcement learning, and mathematical programming, as well as essential professional skills such as project management and communication. Beyond the technical knowledge, this period was transformative, providing me with invaluable insights that continue to shape my personal and professional growth.

None of this would have been possible without the support and guidance of many individuals, including my advisors, colleagues, and co-authors, many of whom I am now fortunate to call friends. I am deeply grateful for their invaluable contributions throughout my Ph.D. journey.

First and foremost, I extend my gratitude to **Prof. Wolfgang Ketter**, as my advisor, mentor, and supporter. Thank you, Wolf, for supporting my work and providing the guidance, resources, and freedom essential to the development of my research. Your optimism, coupled with your high standards, has been a constant source of inspiration, motivating me to strive for excellence. I also appreciate the opportunities you have created through your extensive network of academics and practitioners, including attending top conferences, teaching advanced machine learning courses, reviewing for leading journals, and contributing to scientific grant proposals.

I am also deeply grateful to **Prof. Konstantina Valogianni** for her precious support and guidance in shaping some of my research projects. Dear Konstantina, your insightful discussions on technical aspects and strategic advice on positioning my work have been immensely beneficial. I am especially grateful for the opportunity to work with you at IE Business School, where I completed the final chapter of my thesis. This experience provided valuable connections with supportive scholars and deepened my understanding, significantly enriching my research journey.

I would like to express my gratitude to **Prof. Alok Gupta** for his continuous guidance and support. Dear Alok, I deeply appreciate your thoughtful feedback and critiques, which have greatly improved the quality of my work. Our discussions have provided invaluable insights into positioning my research within the field of Information Systems and have equipped me with essential skills for developing impactful research. Although my planned research visit to the Carlson School of Management at the University of Minnesota was unfortunately hindered by visa issues, your mentorship was truly invaluable. I look forward to future opportunities for face-to-face collaboration after my Ph.D. is completed.

I would like to thank **Prof. Dr. Markus Weinmann** who kindly agreed to chair my PhD defense. Dear Markus, it is an absolute honor to have a highly accomplished IS scholar like you chair my PhD committee. I am very grateful for the time and effort you dedicate toward this.

Throughout my Ph.D., **Prof. John Colins**, **Dr. Nicolò Daina**, and **Dr. Karsten Schroer** have been invaluable mentors and co-authors. Dear John, your guidance at the beginning of my Ph.D. was instrumental in teaching me to think creatively while maintaining a deep focus on details. Your technical expertise in simulation optimization and the life lessons you shared with me have had a lasting impact on my journey. Dear Nicolò, your deep domain knowledge in transportation systems and electric mobility played a pivotal role in shaping my first research project, which was successfully published in a high-impact journal. Dear Karsten, as a trusted colleague, co-author and friend, I have learned so much from you. Thank you for your constant support both during your time at the chair and after graduation.

I would like to thank my doctoral colleagues at the University of Cologne, who have been the most reliable and pleasant companions during the past five years. **Nastaran, Philipp, Karsten, Fred**, and **Muhammad**: I fondly remember my early days when you, as senior colleagues, helped me shape my research ideas, manage teaching and administrative tasks, and navigate the challenges of academic life. The great times we shared, and the fact that we continue to stay in touch and support each other even after your graduation, mean a lot to me. **Anna, Janik, Philipp Kai, Lennard**, and **Julius**: It has been a privilege to work alongside you at the chair. From co-authoring papers to teaching together, providing feedback, and growing together, these shared experiences have been incredibly rewarding. I'm truly grateful for the companionship, encouragement, and moments of fun we've shared. You have always helped me recharge and stay motivated to pursue my goals, and I'm lucky to have had you as colleagues and friends.

During my Ph.D. journey, I also benefited from several other formative experiences that have undoubtedly contributed to my development as an independent researcher on the one hand, and as a more rounded, open-minded, and reflective person.

First and foremost, I would like to highlight the unforgettable and deeply enriching three months of research at IE University in Madrid. I am deeply grateful to **Prof. Konstantina Valogianni** for hosting me and ensuring that my stay was both enjoyable and highly productive.

I would also like to extend my sincere thanks to **IE Business School** and **C-SEB - Center for Social and Economic Behavior at the University of Cologne** for their scholarships and supports which made this invaluable experience possible.

Another transformative experience was attending the AIS Doctoral Consortium of 2022 in Timisoara, Romania, to which I was honored to be invited. Interacting with leading scholars in the field of information systems and working with some of the brightest students in the discipline was inspiring. I am deeply grateful to the **Association for Information Systems (AIS)** for sponsoring my stay at the consortium and for selecting me as a member from a competitive pool of applicants.

I was privileged to attend many academic conferences in some of the most exciting places in the world (Timisoara, Kristiansand, Florence, Paphos, Munich, etc.). I would like to thank the **Faculty of Management, Economics and Social Sciences at the University of Cologne** for funding my position and travel.

Finally, I am deeply grateful to my family - my father, mother, and brother - for their unwavering support and encouragement throughout this journey. From the beginning, they believed in me, offered their guidance, shared my excitement, and provided reassurance during challenging times. A special thanks to my life partner, **Sepideh**, who has been by my side every step of the way. Her unwavering support and love have been my greatest strength during these years. I am truly grateful to all of you!

Cologne, March 04, 2025

Ramin Ahadi

Summary

Despite significant advances in mobility systems, they continue to perform inefficiently, contributing to a range of global challenges. To address these issues, many researchers predict that the future of transportation will be characterized by shared, autonomous, and electric vehicles (SAEVs). However, this transformation poses several managerial and societal challenges that must be addressed to fully realize its potential benefits. This thesis highlights the potential of research at the intersection of information systems and operations management to improve the performance of next-generation mobility systems (Chapter 1). By integrating data-driven decision models four research projects underscore this potential.

Chapter 2 studies the cooperative charging management of SAEVs to maximize profits and service quality while accommodating uncertain demand, limited charging infrastructure, and time-varying electricity prices. A distributed approach using cooperative multi-agent reinforcement learning is proposed that outperforms centralized static charging strategies and provides insights to improve the performance of SAEVs.

Chapter 3 explores the impact of emerging technologies on the physical world and their interaction with user adoption behavior. Specifically, it evaluates a hybrid system that combines autonomous and human-operated ride-hailing services. The study first identifies mobility user preferences toward autonomous services, and using an agent-based model predict and analyze the future of such hybrid service platforms, and evaluate potential changes. The agent-based model enables analysis of the end-to-end impact of key factors, such as trust in the technology.

Chapter 4 proposes a method for leveraging data-driven digital twin frameworks to design large-scale charging hubs. Such problem classes are difficult to solve with traditional mathematical programming optimization. This study shows how high-fidelity, data-driven simulation environments coupled with reinforcement learning can achieve arbitrary scalability and high modeling flexibility. The benchmark experiments show that the proposed model designs result in superior cost and service-level performance under real-world operating conditions.

Chapter 5 supports the operational management of EVCHs through dynamic pricing. Drawing on cutting-edge deep reinforcement learning algorithms, a model-free solution is provided to find optimal pricing policies. The proposed pricing policy is a time-dependent function of the service rate, called dynamic capacity-based pricing. Benchmark analysis shows that the proposed model not only ensures high profits for EVCHs, but also successfully reshapes aggregated demand as desired, even in environments with high variability in supply and demand.

Table of Contents

Acknowledgments	ix
Summary	xiii
1 Introduction – Integrating Operations Management and Information Systems to Harness the Full Potential of Technological Advancements in Transportation Systems	1
1.1 Motivation and Background	1
1.2 Research Objective and Questions	3
1.3 Problem Statement and Research Gaps	6
1.3.1 Current State of the Literature	6
1.3.2 Need for Interdisciplinary Approaches	12
1.4 Summary of Research Articles	14
1.4.1 Research Article 1	16
1.4.2 Research Article 2	16
1.4.3 Research Article 3	17
1.4.4 Research Article 4	18
1.5 Conclusion	19
1.5.1 Theoretical Implications	21
1.5.2 Practical Implications	21
1.5.3 Limitations and Future Research	22
1.6 Declaration of Contributions	23
2 Cooperative Learning for Smart Charging of Shared Autonomous Vehicle Fleets	25
2.1 Introduction	25
2.2 Related Work	28
2.2.1 Decision Making for Shared Electric Vehicles	28
2.2.2 Reinforcement Learning Applications to Shared Mobility Systems	29
2.3 Problem Description	31

2.4	Model	32
2.4.1	Multi-agent Reinforcement Learning Model	34
2.4.2	Dealing with Temporal Abstraction	37
2.4.3	Overview: Solving an MDP with Deep Q-learning	37
2.4.4	Extension to a Hierarchical Approach	38
2.4.5	Benchmark and Upper Bound Models	40
2.5	Experiments	42
2.5.1	Mobility Simulation	42
2.5.2	Experimental Setup	43
2.5.3	Experimental Results	44
2.6	Discussion and Future Work	51
2.7	Appendix	53
2.7.1	Agent-based Simulation	53
2.7.2	Configuration Data	56
2.7.3	Results and Sensitivity Analysis of Strategic Factors	58
3	Adoption of Autonomous Vehicles in Ride-Hailing Services: The Role of User Preferences	67
3.1	Introduction	67
3.2	Literature Review	70
3.2.1	Next-Generation Urban Mobility Systems	71
3.2.2	User Behavior and Human-AI Interactions	72
3.3	Research Method	73
3.3.1	Shared Mobility Demand	74
3.3.2	Agent-Based Simulation for Hybrid Autonomous Ride-Hailing	76
3.3.3	Key Performance Indicators	80
3.3.4	Technical Implementation and Simulation Validation	81
3.4	Simulation Experiment Setup	82
3.5	Results	84
3.5.1	Latent Class Model of DCE	84
3.5.2	Simulation Analysis Results	86
3.6	Discussion	91
3.7	Appendix	95
3.7.1	Design of the Discrete Choice Experiment	95
3.7.2	Descriptive Statistics Discrete Choice Experiment	99
4	Data-driven Planning of Large-Scale Electric Vehicle Charging Hubs using Deep Reinforcement Learning	103
4.1	Introduction	103

4.2	Background	106
4.2.1	Electric Vehicle Charging Hubs (EVCHs)	106
4.2.2	Reinforcement Learning and its Application to Planning Problems	108
4.2.3	Digital Twins (DTs) and their use in Operations Management	109
4.2.4	Reinforcement Learning and its Application to Planning Problems	111
4.2.5	Digital Twins (DTs) and their use in Operations Management	113
4.3	Model	114
4.3.1	Defining the EVCH Planning Scope	114
4.3.2	Parameterizing the Digital Twin Simulator	118
4.3.3	Setting up the SAC RL Framework	126
4.4	Performance Evaluation and Benchmark	128
4.4.1	Evaluating Convergence and Optimality Properties	129
4.4.2	Evaluating Scalability to Real-Sized Problems	130
4.5	Scenario Analyses	132
4.5.1	Impact of Variations in User Preferences	132
4.5.2	Impact of Operational Policy Choices	134
4.6	Discussion	135
4.7	Appendix	141
4.7.1	Operational Modeling and Algorithms	141
4.7.2	Preference Modeling	142
4.7.3	Solution Frameworks and Models	143
4.7.4	Experimental Setup	148
4.7.5	Supplemental Results	151
5	A Pricing Decision Framework for Electric Vehicle Charging Hubs	155
5.1	Introduction	155
5.2	Literature Review	159
5.2.1	Information Systems and Sustainable Electric Mobility Systems	159
5.2.2	Management of Electric Vehicle Charging Hubs	161
5.3	Model	162
5.3.1	EV User Problem	163
5.3.2	Pricing Problem with Perfect Information	164
5.3.3	Proposed Method: Learning Near-Optimal Dynamic Capacity-Based Pricing Policies	166
5.3.4	Simulation Environment	172
5.3.5	Input and Data Preparation	172
5.3.6	Agent-based Modeling of EVCHs	173
5.4	Simulation Experiments	176
5.4.1	Hyperparameters of SAC Model	177

5.4.2	Upper-Bound Benchmarking Experiments	177
5.4.3	Benchmark Pricing Policies	178
5.5	Discussion and Conclusion	181
5.6	Appendix	184
5.7	Analytical Solutions for EV User Decision-Making Problem	184
5.7.1	Solving the optimization problem using the Lagrangian method	184
5.7.2	The optimal requested energy can not be greater than the raw demand	185
5.8	Mathematical Programming Notations	185
5.9	Additional Information for Agent-Based Modeling of EVCHs	186
5.9.1	Estimating charging demand	186
5.9.2	EVCH physical equipments	187
5.9.3	Electricity Tariff for EVCH	188
	References	189
	Declaration (Eidesstattliche Erklärung)	207
	About the Author	209
	Research Portfolio	211
	Teaching Portfolio	213
	CV	217

Chapter 1

Introduction – Integrating Operations Management and Information Systems to Harness the Full Potential of Technological Advancements in Transportation Systems¹

1.1 Motivation and Background

Technological advances continue to reshape modern society, driving significant change across multiple sectors. Today, more than half of the world's population lives in urban areas, and this number is expected to grow significantly in the coming decades (Ritchie et al. 2024). This growing urbanization emphasizes the critical need to integrate emerging digital and technological innovations to ensure sustainable growth and efficient resource management. One of the most impactful developments in this regard is the emergence of *smart cities*, where data-driven technologies improve urban life, drive economic progress, and promote sustainability (Batty et al. 2012, Albino et al. 2015). Among the most pressing challenges, urban mobility systems are under increasing strain due to rapid urban expansion, rising population density, and growing environmental concerns.

¹In accordance with the rules of the Faculty of Management, Economics, and Social Science of the University of Cologne regarding self-quotation in the introduction of a cumulative thesis, parts of Chapter 1 have been adopted from the research articles.

Despite growing awareness of these challenges, global transportation networks remain disconnected, inefficient, and environmentally harmful (Bruun and Givoni 2015). Conventional travel options continue to drive climate change, produce harmful pollutants that threaten public health (World Health Organization 2018), and contribute to systemic inefficiencies (Cramton et al. 2018, Cheng et al. 2020). The reliance on privately owned, petrol-powered vehicles escalate road congestion, rise CO₂ emissions, and increase accident rates. In fact, the transportation sector accounts for about 24% of global carbon emissions, with cities bearing the greatest economic and social burden (Agency 2022). With estimates suggesting that 68% of the world’s population will live in urban areas by 2050 (Nations 2018), the need for more sustainable and resilient transport systems has never been more urgent.

Emerging technologies provide powerful tools to address these challenges and transform urban mobility. Smart cities use real-time data and digital platforms to coordinate infrastructure and planning, enabling AI-driven traffic control that optimizes flow by analyzing sensor data from vehicles and road networks (Zheng et al. 2020, Wang et al. 2023b). Coordinated parking solutions further reduce search times and emissions, while digital twin simulations allow planners to test and refine mobility scenarios with greater precision (Qi and Tao 2018). At the same time, innovations in electric vehicles (EVs), autonomous systems, the Internet of Things (IoT), and artificial intelligence (AI) are creating cleaner, more adaptable, and more sustainable transportation options (Ketter et al. 2023). Among these advances, the Connected, Autonomous, Shared, and Electric (CASE) paradigm has gained prominence for its potential to alleviate congestion, reduce environmental impact, and provide fair and convenient mobility services (Sperling 2018b, Zhang et al. 2022).

Each component brings benefits and has synergies with the others. EVs help mitigate the impact on the climate change, especially as electricity generation becomes more reliant on renewable sources (Nanaki and Koroneos 2016). At the same time, shared mobility reduces the total number of vehicles on the road compared to private ownership (Shaheen et al. 2016). Autonomous vehicles (AVs) further drive the adoption of shared electric mobility by increasing efficiency, reducing costs, and removing barriers such as range anxiety (Pevac et al. 2020). By integrating these benefits, shared autonomous electric vehicles (SAEVs) offer on-demand mobility services that improve safety, sustainability, and operational efficiency (Qi et al. 2022). Already, major urban centers - including San Francisco, Shanghai, Phoenix, Singapore, Tokyo, and Moscow - are demonstrating the real-world feasibility of autonomous taxis through pilot projects (Dong et al. 2022).

While the widespread adoption of SAEVs holds great promise for achieving a sustainable and efficient mobility system in the future, they also present new challenges. Technological complexities, such as ensuring the reliability and cybersecurity of SAEVs, are compounded by environmental considerations such as battery production and grid stability. Regulatory constraints and broader governance and societal issues, including trust in self-driving technology

and equitable mobility access, introduce additional uncertainty (Mahdavian et al. 2021, Sheldon and Dua 2024). Moreover, inconsistent implementation of advanced SAEV solutions may be counterproductive. As Sperling (2018a) warns, AVs risk aggravating traffic congestion and urban sprawl if not carefully integrated. As a result, system-level approaches that consider managerial, technical, environmental, and social factors are critical to effectively harnessing these innovations (Ketter et al. 2023).

1.2 Research Objective and Questions

Building on the emphasis within Information Systems (IS) and Operations Management (OM) on the social and managerial dimensions of next-generation mobility, this work focuses on the associated challenges. On the managerial side, key concerns include optimizing fleet operations, efficiently allocating scarce resources such as charging stations, and implementing dynamic pricing strategies to improve overall performance. Social considerations are also critical and include user acceptance, equitable access to services, and efforts to bridge the digital divide. Addressing these multifaceted issues requires a comprehensive research agenda that integrates IS, data analytics, modeling, and advanced decision-making techniques – such as agent-based modeling, digital twins, and machine learning – to improve urban mobility (Rahman et al. 2021).

This research addresses critical challenges affecting the adoption and integration of shared, autonomous, and electric mobility which require IS-enabled solutions due to the complex interactions among various entities. The research objectives consist of two main areas: (a) the operations management of shared electric autonomous fleets and the influence of user behavior on their performance, and (b) the development of charging infrastructure through optimized planning and operational decision-making.

In the context of SAEV operations management, the first research project addresses the challenge of optimizing SAEV charging management. This issue is especially critical compared to other operational decisions, such as vehicle repositioning, due to the limited availability of charging infrastructure and the lengthy charging process. Therefore, the first research objective (RO1) is to keep SAEVs adequately charged for mobility services while leveraging their connectivity and cooperation to cope with charging infrastructure constraints. The adoption of SAEVs also heavily depends on social factors such as mobility users' reactions and public acceptance of AVs. Therefore, the second research objective (RO2) is to analyze user preferences toward AVs and assess their evolving impact on SAEV performance throughout the adoption period.

The second part of this thesis focuses on the development of charging infrastructure, a key requirement for the widespread adoption of electric mobility. This research argues that centrally operated, large-scale charging hubs—especially in workplaces and commercial areas—are crucial in urban settings where home charging is not widely available. In this context, I examine the operations management of large-scale charging facilities through optimizing two strategic and

tactical decision-making problems. The primary goal is to facilitate charging infrastructure development by improving the economic performance while considering system-level factors such as power grid integration and sustainability. Accordingly, the third research objective (RO3) focuses on investment planning decisions for large-scale charging hubs from a managerial perspective, incorporating user preferences, asset modeling, and system operations. Finally, recognizing the pricing module as a crucial link between supply and demand, the fourth research objective (RO4) is to develop a decision framework that enables charging operators to optimize both economic and system-level performance through advanced dynamic pricing models, enhancing profitability and sustainability. RO3 and RO4 are connected to both managerial and social challenges. Improving the economic performance of charging stations could lead to a well-distributed and widely accessible charging network that plays a crucial role in EV adoption by alleviating drivers' range anxiety, one of the primary barriers to widespread adoption. Additionally, RO4 is closely tied to social challenges, as pricing serves as a key mechanism to influence EV users' charging behavior which is an essential factor for the sustainable development of charging infrastructure, preventing strain on the power grid, and promoting the use of renewable energy sources.

Figure 1.1 summarizes the key obstacles facing next-generation mobility solutions, along with this dissertation's research framework, objectives, and guiding questions aimed at addressing key management and societal challenges. Note that all these problems need IS-enabled approaches to deal with the complex relationship between different entities. I explain the research questions in the following.

A key priority for operators of shared autonomous fleets is to ensure that vehicles consistently make effective operational decisions to maintain an adequate supply of mobility services. The first study addresses these operational complexities in SAEVs, with a particular focus on charging management. The goal is to establish a scalable, decentralized decision-making framework that enables SAEVs to learn and adopt optimal operational policies toward a common objective (e.g., fleet economic performance).

Research Question 1 *How and to what extent can shared autonomous electric fleet operators improve their performance through cooperative, decentralized charging coordination?*

While operations management is critical, socio-behavioral factors are equally influential in shaping the future of next-generation mobility. Variations in user acceptance of SAEVs are particularly significant, prompting the second study to identify distinct user segments and explore how changing attitudes affect ride-hailing platforms. To explore these dynamics, I propose two sub-questions.

Research Question 2.a *How do users differ in their preference for autonomous versus human-driven ride-hailing services, based on their trip and user characteristics?*




Figure 1.1: Thesis Framework: Challenges in Next-generation Mobility Systems and Research Objectives

Research Question 2.b *How do anticipated shifts in user preferences influence ride-hailing platforms' outcomes as they move from traditional human-driven to autonomous services?*

In the field of electric mobility, accessible charging infrastructure is critical for widespread EV adoption. Currently, home charging is the dominant approach in many markets (Lee et al. 2020, Hoover et al. 2021). However, as more consumers without access to home charging switch to EVs, there is a need to expand charging opportunities at workplaces, popular destinations (e.g., supermarkets), and fleet depots (Babar and Burtch 2024, Lee et al. 2020). We call these high-density charging solutions EV Charging Hubs (EVCHs). In addition to encouraging broader EV adoption, EVCHs can facilitate daytime charging that takes advantage of peak solar power output - an opportunity not available with overnight charging (Lee et al. 2018). To support the development of EVCHs, the third study in this thesis focuses on optimizing investment decisions for the expansion of large-scale charging hubs.

Research Question 3 *How can digital twins², in combination with machine learning methods, be harnessed to enhance multi-stage investment decisions for building large-scale EV charging hubs?*

While large-scale charging facilities are critical to advancing EV adoption, they can be economically not viable and stress the power grid if not properly managed (Engel et al. 2018,

²A digital twin is a real-time virtual replica of a physical system that leverages sensor data, simulation, and analytics to optimize performance and decision-making.

Valogianni et al. 2020b). To encourage infrastructure investment while mitigating adverse electricity consumption patterns, the fourth study in this thesis proposes a decision framework to help large-scale stations optimize both economic and system-level performance. By incorporating dynamic pricing strategies, EVCHs can improve profitability and make overall energy demand more sustainable.

Research Question 4 *How can a machine learning-based decision framework be designed to improve the economic performance of EV charging hubs while also shaping the aggregated load profile?*

1.3 Problem Statement and Research Gaps

This section provides a comprehensive review of the literature on next-generation mobility systems, highlights key research gaps linked to the central questions of this thesis, and explains why an interdisciplinary and multi-method approach is essential for effectively addressing these gaps.

1.3.1 Current State of the Literature

This section provides an integrated review of key literature on the design, operation, and deployment of next-generation mobility systems, with a particular focus on SAEVs and large-scale EVCHs. The review covers five main research streams: (1) smart and sustainable mobility systems, (2) operations management of shared autonomous electric fleets, (3) user behavior in hybrid autonomous ride-hailing systems, (4) EVCH planning and operations, and (5) pricing management of EVCHs through IS-enabled decision frameworks.

Smart and Sustainable Mobility Systems

Traditional mobility systems have led to negative impacts such as environmental damage, public health hazards, and economic inefficiencies due to congestion and excessive resource use (Pal et al. 2023). As a result, improving the efficiency and sustainability of transportation has become a critical priority. A mobility system is considered *smart and sustainable* when it uses digital technologies and data-driven strategies to balance user needs, organizational goals, and environmental considerations. While the word smart refers to real-time information flows and the use of new innovations such as CASE technologies for automation and coordination, sustainability encompasses environmental responsibility as well as socio-economic factors such as safety, resource efficiency, and viable business models (Ketter et al. 2023). Related to this thesis, researchers widely recognize the potential of the CASE paradigm to combine smartness and sustainability to address current transportation challenges, with each dimension offering unique benefits. However, the most significant benefits are realized when these dimensions are integrated into a cohesive, data-driven framework (Sperling 2018a).

The CASE paradigm has emerged as a key approach to addressing the environmental and operational challenges of urban transportation. Each of its four components offers distinct advantages while complementing the others. *Connected* (C) vehicles use communication technologies to facilitate data exchange and improve real-time information about traffic conditions and infrastructure status (Batty et al. 2012). *Autonomous* (A) capabilities enable vehicles to operate with minimal or no human intervention, incorporating various levels of automation to improve safety and efficiency (Yurtsever et al. 2020). *Shared* (S) mobility optimizes resource utilization by allowing multiple users to access a shared fleet of vehicles, thereby reducing congestion and private vehicle ownership (Shaheen et al. 2016). Finally, *Electric* (E) propulsion, including battery electric and fuel cell technologies, minimizes pollution and noise, especially when integrated with renewable energy sources (Tran et al. 2012).

A prominent example of CASE mobility is *shared autonomous electric vehicle fleets*, often referred to as SAEVs, which have attracted considerable interest in engineering, management, and sociology (Narayanan et al. 2020, Dlugosch et al. 2020). SAEVs hold great promise for reducing greenhouse gas emissions, reducing traffic congestion, and improving resource utilization. Several studies suggest that a single SAEV, if efficiently operated and routed, can replace multiple conventional vehicles (Chen et al. 2016, Fagnant and Kockelman 2014). However, challenges remain in scaling these systems, particularly in terms of infrastructure requirements, user acceptance, and cost-effective operation (Mahdavian et al. 2021).

Building on the CASE paradigm, *smart and sustainable* mobility systems use digital technologies, data analytics, and innovative service models to address broader transportation challenges, including congestion, emissions, and equity of access. These systems often integrate intelligent transportation technologies - such as sensors and vehicle-to-infrastructure communications - to optimize routing and traffic management in real time, reducing travel time and fuel consumption (Golbabaei et al. 2021). In addition, shared mobility services such as bike-sharing and car-sharing can increase resource efficiency by promoting multimodal travel and reducing private vehicle ownership (Eren and Uz 2020).

Sustainability is at the core of this approach, with a growing emphasis on low-emission vehicles (e.g., electric or hybrid) and infrastructure that supports non-motorized transportation (e.g., bicycle lanes, walkable urban design) (Benevolo et al. 2016). These measures aim to reduce both greenhouse gas emissions and local air pollutants, in line with broader climate goals. In addition, the integration of big data analytics and machine learning enables proactive policy interventions, such as dynamic tolling or congestion pricing, to nudge users toward more sustainable modes of transportation (Shaheen et al. 2020).

Despite significant progress, several obstacles remain. Ensuring equitable access to advanced mobility solutions, protecting privacy, and achieving interoperability across technologies and jurisdictions are among the major concerns. To address these issues, ongoing interdisciplinary research focuses on harmonizing governance structures, standards, and technology deployment

strategies. As part of these challenges, the following subsections review the operations management of SAEVs, the social factors and user behavior toward the acceptance of SAEVs, and optimizing the investment and management of non-home charging facilities to ensure a rapid adoption of EVs while aligning it with system preferences such as the capacity of the power grid and resources.

Operations Management of Shared Autonomous Electric Fleets

Although SAEVs have the potential to greatly benefit urban mobility, they also present a number of decision complexities at the strategic, tactical, and operational levels. At the strategic level, scholars have examined the appropriate fleet size for SAEVs (Lokhandwala and Cai 2018, Levin et al. 2019) and the associated charging infrastructure needs (Lokhandwala and Cai 2020). Meanwhile, at the operational level, decisions about dynamic routing and vehicle assignment are critical to aligning service availability with fluctuating demand (Dong et al. 2022, Ho et al. 2018). I focus on the operational aspects due to the focus of RQ1.

Previous literature on repositioning and charging in ride-hailing typically adopts either a centralized, fleet-centric model or an uncoordinated approach focused on driver incentives (Kullman et al. 2021a). Centralized frameworks often face scalability barriers, while decentralized methods lack the cooperative mechanisms needed to optimize overall system performance. While a non-cooperative decentralized model (a common model in the literature (e.g., Liang et al. 2020)) may be appropriate for e-taxi services (where individual drivers pursue their own profit), it fails to reflect the priorities of a single fleet operator focused on maximizing profitability and service quality. Another common oversight is the assumption of unlimited or uniform charging infrastructure, which overlooks differences in station capacity and performance. Recent studies (Guillet et al. 2022, Pantelidis et al. 2022, Froger et al. 2022) highlight how station capacity constraints significantly affect scheduling decisions. Moreover, while several papers (e.g., He et al. 2021, Iacobucci et al. 2019) discuss EV charging as a means to meet demand or reduce energy costs, few consider both objectives simultaneously.

Compounding these challenges, simple rule-based charging strategies—such as always sending low-battery vehicles to the nearest station—prove inadequate for real-world scenarios, underscoring the need for large-scale, data-driven methods to decide which vehicles should charge, at which station, and when (Kullman et al. 2021a, Liang et al. 2020). In this context, multi-agent reinforcement learning (MARL) has emerged as a promising solution that allows for decentralized coordination while still allowing vehicles to cooperate and anticipate future demand (Shou and Di 2020). Building on these insights, I propose a decentralized cooperative framework that identifies vehicles requiring charging and assigns them to limited, heterogeneous charging stations under time-varying energy prices, thereby jointly optimizing service quality and operating costs.

User Behavior in Hybrid Autonomous Ride-Hailing Systems

As fully autonomous fleets remain in development, recent research has explored *hybrid* mobility systems where human-driven and AVs coexist. During the transition, large ride-hailing services are expected to integrate both, raising challenges related to driver compensation, fleet optimization, and user trust (Ao et al. 2024, Adam et al. 2022, Siddiq and Taylor 2022).

The success of AVs depends heavily on user acceptance, which remains uncertain (Ketter et al. 2023). While AVs promise operational benefits (e.g., Yao et al. 2020, Chen et al. 2024b), widespread skepticism may hinder their effectiveness. Studies highlight trust as a key barrier, with users often preferring human interaction (Glikson and Woolley 2020, Gnewuch et al. 2023). Research suggests that most people remain hesitant about fully driverless mobility, with only a small fraction expressing confidence in such systems (Shariff et al. 2017). Technology adoption is dynamic, evolving with user familiarity, societal influence, and advances in capabilities (Komiak and Benbasat 2006, Venkatesh and Davis 2000). However, much research takes a static perspective and overlooks the gradual nature of this change (e.g., Fagnant and Kockelman 2018, Dong et al. 2022). A dynamic approach considers how hybrid fleets affect user preferences and operational outcomes over time.

User perception is critical: while automation increases efficiency, trust and perceived usefulness determine adoption (Dietvorst et al. 2014, Jabbari et al. 2022). Adoption varies across demographics and is influenced by price sensitivity and openness to new technologies (D’Acunto et al. 2019, Curtale et al. 2022). To model these behaviors, researchers use *agent-based modeling*, which integrates user decisions with supply-side policies (Jing et al. 2020). Agent-based modeling helps ride-hailing platforms anticipate responses to pricing, incentives, and operational changes, facilitating a smoother transition to autonomous mobility. Future research should refine adaptive strategies that foster user trust while ensuring profitability and operational efficiency in hybrid autonomous-human fleets.

Electric Vehicle Charging Hubs: Planning and Operations

The future of EV charging networks, particularly in urban areas where home charging is not widely accessible, depends on high-density charging facilities, known as EV charging hubs (EVCHs). EVCHs pose distinct operational and planning challenges due to their centralized management and integration with the buildings to which they are attached. Unlike other charging use cases, EVCHs represent large, concentrated loads that often require network expansion and load shaping (Lee et al. 2019). They also offer opportunities to integrate on-site energy generation (e.g., PV, storage) to reduce peak loads and operating costs (Nunes et al. 2016, Ferguson et al. 2018). In addition, user behavior at EVCHs varies significantly by facility type (e.g., workplace, shopping mall) which must be taken into account for planning investment decisions. Unlike dispersed charging stations, all chargers in an EVCH are centrally located, allowing end-to-end control of the parking and charging process through intelligent parking systems (Babic

et al. 2022a). This enables optimized vehicle allocation and efficient scheduling of charging equipment (Ferguson et al. 2018).

Regarding EVCHs, most of the existing research concentrates on the operational decision-making problems such as load management. Research on EVCH operations builds on extensive work on smart charging and scheduling (see Mukherjee and Gupta (2015) for a review). However, EVCHs introduce additional complexities, particularly the need to integrate charging management with building energy loads and site-level constraints. For example, Huang and Zhou (2015) proposes a mixed-integer optimization framework for workplace charging, while Wu et al. (2017) develops a two-stage energy management system for office buildings with EV charging. Nunes et al. (2016) explore the coordination of solar-powered parking lot charging, and Ferguson et al. (2018) propose an integrated load management approach that considers building base loads and PV generation. Practical implementations of site-level load management have been demonstrated by Jun and Meintz (2018). In addition, Lee et al. (2019) examine optimization-driven approaches to EVCH operations, highlighting the complexity introduced by parallel-use charging docks that require simultaneous allocation and charging decisions.

In the operations management of EVCHs, the design problem (e.g., planning and strategic decisions) has received less attention. It is a multi-stage stochastic decision problem requiring simultaneous infrastructure and operational decisions while accounting for interdependence with heterogeneous user preferences. Traditional methods such as standard stochastic programming and approximate dynamic programming approaches struggle with such complexity (Powell 2014, Hannah 2015). Simulation-based methods have been explored to address this challenge. For example, Kazemi et al. (2016) use a genetic algorithm on a simplified simulation model to determine the optimal EV parking lot size, and Babic et al. (2022a) use a greedy search approach to configure the investment decisions for a large charging station. With a simplification assumption, Li et al. (2020) develop a deterministic optimization framework for sizing and operating a EV parking facility which could perform poorly in realistic cases where many components (e.g., EV drivers, on-site electricity generation, grid costs) follow stochastic patterns. The literature shows that the few existing studies on EVCH design incorporate significant simplifications, such as assuming a single-period design and neglecting building loads or the parallel use of charging docks.

Addressing RQ3 fills the above mentioned gap in EVCH operations management research and expands on the design challenge by optimizing investment planning decisions while accounting for the interdependencies between asset modeling, demand preferences, and operational policies. Unlike prior work, it incorporates detailed preference modeling using real-world parking and charging data, allowing for preference-aware infrastructure sizing. Additionally, existing building load profiles are integrated into investment and operational decision-making, offering a more holistic approach. The proposed model accounts for parallel charging infrastructure, enhancing asset efficiency despite increased operational complexity. Furthermore, by aligning infrastruc-

ture provisioning with user preferences, this approach has significant social and sustainability implications, ensuring efficient and user-centric charging solutions.

Pricing Management of EVCHs through IS-enabled decision frameworks

Revenue management and dynamic pricing strategies can align economic incentives with sustainability objectives, particularly in EV charging hubs (EVCHs). Such approaches adjust prices based on real-time or forecasted demand, charging capacity, and electricity market signals. Traditional pricing models in the energy sector have explored auction mechanisms, time-of-use rates, and subscription-based models (Hou et al. 2019, Valogianni et al. 2020a). More recent research has introduced *capacity-based* and *deadline-differentiated* pricing, which tailors fees to service availability or waiting times (Moradipari and Alizadeh 2019, Lin et al. 2023). Studies indicate that dynamic pricing can not only increase revenue but also incentivize off-peak charging, mitigating local grid stress (Lee and Choi 2021). Cui et al. (2021) propose an optimal pricing strategy to balance demand across multiple EVCHs, while Luo et al. (2017) introduce a stochastic dynamic pricing model that accounts for uncertainty in charging demand and renewable energy availability. Other studies have explored vehicle-to-grid (V2G) pricing strategies to utilize EV batteries as storage systems (Mao et al. 2017) and competitive pricing schemes among multiple EVCHs (Lu et al. 2018).

Several studies have linked pricing to service quality. For example, Lin et al. (2023) incorporate waiting time into dynamic pricing for fast-charging services to optimize queuing efficiency at public stations. Valogianni et al. (2015) propose a capacity-based pricing model aimed at reducing peak demand from EV loads. Other research has explored menu-based pricing strategies that take advantage of flexible charging demand, such as Moradipari and Alizadeh (2019), which designs an optimal pricing mechanism that prioritizes users with urgent needs. Similarly, Zeng et al. (2021) differentiates pricing for flexible and inflexible charging demands, assuming that users choose between regular and priority services rather than changing their energy demands. Lu et al. (2023) introduce a time-differentiated pricing model, offering discounts to users willing to park longer to reveal their actual departure time.

Research study 4 enhances dynamic pricing for capacity-based EV charging by integrating user behavior, acknowledging that charging at EVCHs is typically a secondary activity (e.g., during work or shopping) rather than the main reason for a trip (Daina et al. 2017, Lee et al. 2019). Unlike previous studies, it assumes that users modify their energy intake in response to price signals but do not significantly change departure times. IS-enabled decision systems can optimize operations by leveraging data-driven approaches and incorporating novel decision-making algorithms such as deep reinforcement learning algorithms for sustainable energy and mobility management (Watson et al. 2010, Seidel et al. 2013b, Ketter et al. 2023, 2018b).

Summary and Research Gaps

From this review, several gaps become apparent:

- **Cooperative Decision-Making Under Scale:** Centralized models for routing or charging of SAEVs often fail to scale in realistic settings, while purely distributed approaches neglect system-level benefits. There is a need for cooperative decentralized approaches such as MARL-based frameworks that balance distributed execution with common objectives.
- **Predicting the Impacts of Diverse User Preferences in Hybrid Autonomous Mobility:** Although there has been considerable research on SAEVs, few studies have examined hybrid human-operated and autonomous fleets or the phase-in of full autonomy. A deeper understanding of user adoption, trust, behavioral dynamics, and evolving system impacts is essential.
- **Preference-Aware Planning of EVCHs:** Existing research on EVCH design often reduces complexity by ignoring operational details, neglecting building energy interactions, or simplifying stochastic user preferences. More robust, data-driven approaches –possibly via digital twins– are needed to reflect real-world usage patterns and operational details.
- **Operations Management of EVCHs through Dynamic Pricing:** While existing research has independently focused on optimizing EV charging station pricing for either profit maximization or alleviating stress on the power grid, a scalable decision support system is needed to simultaneously achieve high economic performance and system-level objectives, such as peak shaving and load reshaping.

1.3.2 Need for Interdisciplinary Approaches

Addressing the challenges of next-generation urban mobility requires an integrated framework that draws from operations management (OM), information systems (IS), data-driven methods, and behavioral science. This dissertation adopts a multidisciplinary perspective by incorporating agent-based modeling, advanced machine learning and optimization (e.g., reinforcement learning and mathematical programming), digital twin simulation, and user behavior analysis to create robust decision support tools for operators, policy makers, and end users. Although OM provides sophisticated optimization strategies for transportation, it often makes simplifying assumptions about the interplay between mobility components, user demand, and decision making. In contrast, IS research focuses on information exchange, interactions between system components, technology adoption, and user behavior, providing essential socio-technical perspectives. The combination of these methods provides a powerful tool to address the complex socio-technical challenges of new mobility systems (Ketter et al. 2023).

IS research typically revolves around using information technology (IT) to create value for individuals and organizations. While IT may constitute the central focus, it is more impactful

when designed for people/users and capable of delivering tangible economic or societal benefits. Many successful IT applications in business, such as AI-driven analytics or digital platforms, inherently involve operational tasks. This synergy has led to increasing interdependence between IS and OM, with researchers contributing by either applying IS solutions to OM problems or employing OM techniques to address IS questions (Kumar et al. 2018).

This thesis mainly focuses on incorporating IS-enabled solutions to solve operational decision-making problems of next-generation mobility systems by offering beyond traditional optimization techniques. Traditionally, IS have played a pivotal role in OM problems by providing the data and insights necessary for improved operational decision making. This includes sharing information across supply chain departments to optimize processes and decision making (e.g., Cachon and Fisher 2000, Demirezen et al. 2016), the technological transition in healthcare with the adoption of electronic health records and the foundation of health information exchanges (Yaraghi et al. 2015, Bhargava and Mishra 2014), and the advancement of multi-channel retailing and recommendation systems (Kumar et al. 2019, Zhang et al. 2020a). In addition to supply chain management, IS also plays a critical role in service operations, ranging from traditional web-based interfaces for customer care to more advanced solutions involving artificial intelligence and knowledge-based systems (Bensoussan et al. 2009, Setia and Patel 2013). Beyond these traditional use cases, Kumar et al. (2018) highlights other cross-disciplinary applications, including smart city management, healthcare operations, blockchain, Industry 4.0, and digital platforms.

Next-generation mobility systems share characteristics with smart city initiatives, given transportation serves as the backbone of modern societies, and with digital platforms, given mobility-on-demand services use a digital layer to match supply and demand. In this data-rich context, a multidisciplinary approach is critical to effectively monitor performance, streamline processes, and enable real-time operational decisions. According to Yoo et al. (2010b), the digitization of mobility results in a layered modular architecture in which a digital coordination layer is superimposed on the physical infrastructure. This layered design spans multiple stakeholders—regulators, fleet operators, infrastructure providers, and users—each with different responsibilities, but all relying on an integrated socio-technical platform for cohesive coordination.

Research on IS-enabled cases in smart city management has been growing recently (Ismailova et al. 2019). Cranefield and Pries-Heje (2023) discuss IS as key enablers in smart city management, providing real-time data integration, advanced analytics, and collaborative platforms that improve urban infrastructure and services. For example, Kahlen et al. (2018a) proposes a virtual power plant system composed of EVs to illustrate how digital tools and OM methodologies can be combined to balance energy needs with profit opportunities. Valogianni et al. (2020b) propose a dynamic pricing model that avoids additional peak loads from EVs. There is also an intersection between digital platforms and intelligent mobility systems. Emerging issues such as the on-demand and shared economy using digital tools pose various operational chal-

lenges, as shown by recent studies on ride-hailing (Banerjee et al. 2022, Bai et al. 2019, Guda and Subramanian 2019, Taylor 2018, Siddiq and Taylor 2022, Benjaafar et al. 2024).

Researchers in IS are uniquely qualified to tackle complex socio-technical mobility problems for two main reasons. First, IS has a robust set of methods in its toolbox, ranging from algorithmic and mechanism design (Bichler et al. 2010, Kahlen et al. 2024) to machine learning and optimization techniques (Meyer et al. 2014, Abbasi et al. 2015, Barfar and Padmanabhan 2017, Guo et al. 2019, Dlugosch et al. 2020, Zhang et al. 2024). Second, IS approaches combine these computational competencies with deep-rooted experience in conducting large-scale behavioral studies (Babar and Burtch 2020, Burtch et al. 2018, Oestreicher-Singer and Zalmanson 2013, Osterwalder et al. 2005), enabled by a distinct socio-technical perspective (Sarker et al. 2019a). This combination places IS professionals at the forefront of initiatives aimed at addressing major social and environmental challenges (Ketter et al. 2023), as evidenced by the broad involvement of IS research in sustainable energy (Dedrick 2010, Melville 2010a, Watson et al. 2010, Seidel et al. 2013b, Ketter et al. 2018a) and emissions management (Corbett 2013).

The integration of information systems (IS) and operations management (OM) offers a robust framework for addressing mobility challenges, which extend beyond environmental concerns to social and economic factors such as traffic disruptions and safety risks. Mobility requires real-time coordination, where failures can disrupt services, delay transport, and pose safety hazards. Additionally, diverse stakeholders with independent decision-making add complexity to operations. IS-enabled tools help navigate these challenges. Agent-based modeling captures human factors in policy decisions, improving estimates of shared autonomous electric vehicle (SAEV) adoption (Haki et al. 2020). Machine learning, particularly reinforcement learning, has proven effective in vehicle routing (Dong et al. 2022), SAEV charging (Ahadi et al. 2023), and dynamic EV pricing (Lee and Choi 2021). Digital twins further enhance decision-making by simulating EV charging hubs (EVCHs), integrating sensor data with behavioral modeling to optimize investments and operations (Cimino et al. 2019, Jones et al. 2020, van der Valk et al. 2020).

Following Ketter et al. (2023) and Dhanorkar and Burtch (2022), this study underscores the need for an IS-OM synergy to tackle mobility challenges, from equitable access to efficient resource allocation. By bridging human-technology interactions with advanced optimization, IS research drives sustainable and socially inclusive urban transportation solutions.

1.4 Summary of Research Articles

Table 1.1 details the set of articles that support the research contributions and goals of this dissertation. This table lists peer-reviewed conference papers, peer-reviewed journal articles, and papers under review for journal publication. In the following chapters, only the content of the journal articles is presented.

RQ	Chapter	Bibliographic Data	Publication Status
1	Chapter 2	Ahadi, R. , Ketter, W., Collins, J., & Daina, N. (2023). <i>Cooperative Learning for Smart Charging of Shared Autonomous Vehicle Fleets</i> . <i>Transportation Science</i> , 57(3), 613–630.	Published journal article
		Ahadi, R. , Ketter, W., Collins, J., & Daina, N. (2021). Siting and Sizing of Charging Infrastructure for Shared Autonomous Electric Fleets. In <i>Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems</i> .	Published conference article
2	Chapter 3	Ahadi, R. , Taudien, A., Ketter, W., & Gupta, A. (2024). Adoption of Autonomous Vehicles in Ride-Hailing Services: The Role of User Preferences. (1 st -round)	Under review
		Ahadi, R. , Taudian, A., & Ketter, W. (2023). Human versus automated agents: how user preferences affect future mobility systems. In <i>Proceedings of the European Conference on Information Systems</i> .	Published conference article
3	Chapter 4	Schroer, K., Ahadi, R. , Lee, T. Y., & Ketter, W. (2024). “Data-driven Planning of Large-Scale Electric Vehicle Charging Hubs using Deep Reinforcement Learning.” (2 nd -round)	Under review
		Schroer, K., Ahadi, R. , Lee, Y.T., Ketter, W. (2021). “Preference-aware Planning and Operations of Electric Vehicle Charging Clusters: A Data-Driven Prescriptive Framework.” In <i>Proceedings of the SIG GREEN Workshop 2021</i> .	Published conference article
		Schroer, K., Ahadi, R. , Lee, Y.T., Ketter, W. (2021). “Preference-Aware Planning and Operations of Electric Vehicle Charging Clusters: A Prescriptive Framework.” In <i>Workshop on Information Systems and Technology (WITS) 2021</i> .	Published conference article
4	Chapter 5	Ahadi, R. , Valogianni, K., Schroer, K., & Ketter, W. (2025). “A Pricing Decision Framework for Electric Vehicle Charging Hubs.” (1 st -round)	Under review
		Ahadi, R. , Schroer, K., & Ketter, W. (2024). “Managing Electric Vehicle Charging Hubs Through Dynamic Capacity-Based Pricing.” In <i>Proceedings of the European Conference on Information Systems</i> .	Published conference article

Table 1.1: Research Articles

1.4.1 Research Article 1

Research Objective

Chapter 2 focuses on how operators of shared autonomous electric fleets address day-to-day operational challenges, particularly in handling charging tasks. The primary goal is to develop scalable, cooperative strategies that allow the fleet to optimize charging schedules while dealing with factors such as limited charging capacity, long charging times, demand uncertainty, and fluctuating electricity prices.

Methodology

Drawing on agent-based modeling, a comprehensive simulation environment of SAEVs is modeled and a cooperative multi-agent reinforcement learning framework is developed to facilitate decentralized decision making. This approach allows SAEVs to communicate in real time to share charging station resources and predict their energy needs, effectively synchronizing across the fleet. A hierarchical decision framework improves the charging policy by organizing decisions into multiple levels. At the top level, the system determines *when* vehicles should charge, while the lower level decides *where* they should be routed for charging. These levels interact dynamically to coordinate their respective decisions. In addition, reward-shaping strategies are introduced to reinforce collaborative behaviors and promote overall resource efficiency and service reliability. The study examines how these strategies vary based on fleet size, vehicle distribution, and different demand conditions.

Contribution

This research introduces a decision framework that achieves higher profitability and operational reliability than established approaches. Unlike traditional centralized models, it is significantly more scalable by coordinating decentralized agents through information sharing. These agents are trained in a simulation environment calibrated with real-world data, providing practical guidance for SAEV adoption by fleet operators and policy makers. The results underscore the effectiveness of decentralization and cooperation in addressing operational bottlenecks in electrified urban mobility. In particular, the study shows that under this framework, SAEVs can achieve profitability levels comparable to those of gasoline-powered shared vehicles, highlighting their potential to support sustainable mobility in the future.

1.4.2 Research Article 2

Research Objective

Chapter 3 explores the user-oriented and societal challenges associated with the rise of AVs in ride-hailing platforms. Specifically, it examines how individual differences influence a rider's

preferences for autonomous versus human-operated services, and assesses the implications of these preferences for ride-hailing platforms as they transition to fully autonomous operations.

Methodology

This study presents a multi-method framework that integrates a discrete choice experiment (DCE) with agent-based modeling, bridging behavioral research and design science. The stated preference (SP) approach of the DCE identifies user preferences for AVs in shared mobility and clusters users accordingly. Using these insights, along with empirical data and simulation techniques, a behaviorally based agent-based modeling is developed to assess how different user behaviors influence hybrid autonomous ride-hailing services. Constructing a realistic simulation for such a complex socio-technical system requires extensive domain expertise and empirical validation.

Contribution

This study integrates behavioral and design science to advance research on hybrid autonomous ride-hailing. It identifies four user segments based on preferences for autonomous versus human-driven services, and shows that only one segment inherently favors AVs. However, price and wait time significantly influence all segments, providing fleet operators with strategic levers to drive AV adoption. A feature-rich, behaviorally-informed agent-based model is developed to simulate hybrid fleets, incorporating both demand- and supply-side dynamics. User behavior, population composition, and AV trust levels are derived from a discrete choice experiment, while historical data calibrates travel demand. The supply-side model integrates autonomous and human-driven fleets, validated by operations management rules and driver behavior modeling. The results show that even a modest share of AVs can reduce CO₂ emissions and increase revenues, although adoption remains constrained by trust issues. Increasing trust amplifies these benefits, underscoring the need for targeted strategies to build user confidence. Fleet management measures, such as AV ride discounts and optimized fleet sizing, further enhance system performance.

1.4.3 Research Article 3

Research Objective

Chapter 4 focuses on how to effectively plan large-scale EVCHs to support the growing population of EVs. The goal is to determine the optimal size, resource allocation, and operational strategies for EVCHs that both control investment costs and maintain high levels of user satisfaction.

Methodology

The multi-stage and stochastic nature of EVCH planning poses significant challenges to conventional optimization methods. Ensuring tractability often requires simplifying the problem by reducing planning horizons, using broader time intervals, or making deterministic assumptions. This work presents a novel approach that leverages the wealth of granular operational and preference data now available through pervasive IoT sensor technologies. Specifically, parking and energy datasets are integrated with high-resolution asset models and real-world operational policies in a detailed simulation environment to form a nearly exact digital twin of the planned EVCH. An actor-critical reinforcement learning framework then interacts iteratively with this environment to learn an optimal configuration policy over multiple simulated epochs.

Contribution

Methodologically, it establishes a framework that combines reinforcement learning with large-scale data-driven simulations (digital twins) to facilitate ex-ante risk management and decision support in the design of service systems such as EVCHs. By circumventing common limitations such as simplifying assumptions, our approach enables more realistic, data-driven modeling of stochastic dynamics and operational details, provides computational scalability over traditional optimization, and maintains a flexible structure for testing diverse operational policies. Through extensive simulation experiments, the study demonstrates that the proposed method achieves near-optimal scheduling solutions for EVCHs and outperforms alternative techniques in both speed and scalability. In addition, it exploits the adaptability of the digital twin to evaluate different preference profiles and operating conditions, providing practitioners with multiple domain-specific insights.

1.4.4 Research Article 4

Research Objective

Chapter 5 develops a dynamic pricing mechanism to improve the financial performance of EV charging hubs while influencing overall energy consumption patterns to promote grid stability and sustainable energy use.

Methodology

This study develops a decision support system that incorporates a deep reinforcement learning approach to identify near-optimal dynamic pricing policies for EVCH operators. Traditional optimization becomes intractable in such environments due to the curse of dimensionality and is unable to incorporate uncertainties in demand and supply. Dynamic pricing in EVCHs requires addressing three critical factors: (a) sequential decision making, which requires pricing

decisions at each time step that affect subsequent decisions; (b) uncertainty about demand preferences and supply constraints, with unknown probability distributions; and (c) the large scale of EVCHs, which involve multiple interdependent components and processes. To address these complexities, a deep reinforcement learning algorithm is introduced. In particular, a soft actor-critical (SAC) model is used because it does not rely on explicit transition probabilities and can handle continuous action spaces (i.e., price signals). By merging reinforcement learning with deep learning, the algorithm can handle both large state and action spaces. Training takes place within a detailed digital twin of an EVCH, which incorporates rich operational and behavioral data to capture realistic system dynamics. This simulation, calibrated with unique real-world data, enables the pricing agent to iteratively refine its decisions and maximize revenue.

Contribution

This study presents a machine learning-based decision support system that combines deep reinforcement learning with capacity-based pricing to help large EVCHs manage demand, maximize profitability, and ensure stable grid usage. Specifically, an advanced reinforcement learning algorithm identifies near-optimal pricing signals under realistic, stochastic, and large-scale conditions, eliminating the need for extensive modeling assumptions. A detailed agent-based simulation, calibrated with real-world data on arrivals, departures, and energy consumption, mimics daily operations to train the learning algorithm and reveal user responses to pricing incentives. Empirical results show that the proposed approach closely approximates optimal strategies and significantly outperforms static pricing, providing crucial benefits such as peak shaving and improved revenue. In practical terms, this decision framework can promote financially viable EVCH operations, thereby accelerating the widespread adoption of EVs and reducing energy infrastructure costs by mitigating peak loads.

1.5 Conclusion

This dissertation examines the operational, behavioral, and infrastructural challenges of shared autonomous electric mobility and sustainable energy systems. Utilizing IS-enabled, data-driven approaches-including multi-agent reinforcement learning, agent-based modeling, digital twins, and dynamic decision making-it demonstrates how advanced tools combined with well-designed information pipelines can effectively manage the increasing socio-technical complexity of next-generation transportation systems. As it is illustrated in Figure 2.1, this thesis focus mostly on the managerial and social challenges of shared autonomous electric mobility and the four provided research projects divide into main two areas: (a) the operations management of shared autonomous electric fleets and the reliance of the user acceptance and trust, and (b) facilitating the development of EVCHs by optimizing investment planning and pricing management. I should point out that although the development of EVCHs is not directly related to the oper-

ations management of SAEVs, both fields contribute to overcoming the barriers in front of the smart sustainable mobility systems. In the following, I summarize the research studies and their relationships together.

A primary focus is on optimizing the day-to-day operations of shared autonomous fleets, particularly the charging processes of SAEVs. Through cooperative, decentralized coordination, this research shows that fleets can achieve higher profits and more reliable service levels, especially under constraints such as limited charging stations, uncertain travel demand, and fluctuating energy costs. These findings underscore the value of real-time, adaptive decision-making enabled by IS for managing electrified urban mobility at scale.

Yet the successful deployment of SAEVs is not only a matter of operational efficiency; social and behavioral factors play an equally important role. By identifying different user segments and analyzing trust in autonomous services, the dissertation highlights the conditions under which users are likely to adopt or reject new mobility technologies. This behavioral perspective provides ride-hailing platforms and public agencies with strategies to build trust, personalize services, and more effectively promote the adoption of AVs.

In the design and operation of large-scale EV charging hubs (EVCHs), the dissertation demonstrates the utility of combining digital twins with deep reinforcement learning. This integrated approach is shown to handle the complexities of multi-stage, stochastic planning, providing a richer alternative to traditional optimization models that must often simplify real-world conditions. By accurately modeling user behavior, building loads, and intermittent renewable energy, digital twins shed light on investment decisions, site configuration, and long-term infrastructure needs.

Finally, dynamic pricing emerges as a powerful mechanism for synchronizing economic performance with sustainable energy goals. The proposed capacity-based pricing models, driven by deep reinforcement learning, learn to manage peak demand, smooth energy consumption, and ultimately increase the profitability of charging facilities. These results support the broader potential of price signals to guide consumer behavior and orchestrate the interplay between supply and demand in energy-constrained environments.

Taken together, the contributions of this dissertation form a cohesive framework for building and managing smart mobility systems that are both user-centric and environmentally aware. By weaving together methods from operations management, information systems, machine learning, and behavioral science, the research underscores the critical need for interdisciplinary solutions that integrate technical optimization with socio-technical insights. This work lays the groundwork for further scientific and practical innovations that can make next-generation mobility equitable, efficient, and sustainable.

1.5.1 Theoretical Implications

Theoretically, this dissertation extends the scientific discourse at the intersection of information systems, operations management, and sustainable urban mobility. By employing cooperative multi-agent reinforcement learning strategies, it goes beyond conventional single-agent or centralized optimizations and introduces a data-centric, decentralized framework for large-scale decision making. This contribution adds a new dimension to existing theories of dynamic multi-agent systems, especially in contexts where real-time coordination and adaptability are essential. In addition, by introducing a new hierarchical design for making operational decisions at different levels, this work contributes to the solution of large-scale sequential stochastic decision problems.

This dissertation advances the understanding of technology acceptance and behavioral operations by examining user behavior and trust in the adoption of autonomous mobility. By analyzing how factors such as perceived safety and willingness to pay influence ride-hailing adoption, it highlights the impact of different user segments on system performance. This perspective extends traditional operational models by incorporating socio-technical factors into strategic and tactical decision making. In addition, the integration of discrete choice experiments with agent-based modeling provides a methodological contribution that enables the prediction of evolving user behavior in autonomous mobility services.

In the field of infrastructure design, the application of digital twin and reinforcement learning to the planning and operation of EVCHs provides a novel methodological perspective. The integration of real-world data and high-fidelity simulations overcomes the limitations of static or deterministic optimization techniques and enhances the robustness of multi-stage infrastructure planning. This approach contributes to the theoretical literature on stochastic, data-driven planning and reveals new ways to deal with uncertainty and resource allocation in complex mobility and energy systems.

Finally, the introduction of dynamic pricing frameworks-particularly those focused on capacity-based pricing-illustrates the interplay between market-based incentives and environmental goals, and provides new theoretical insights into how revenue management can coexist with sustainability goals. These findings underscore the feasibility of coordinating economic, social, and environmental interests through advanced pricing strategies that balance user demand with power system constraints. Taken together, the contributions of this dissertation lay the groundwork for further research on the socio-technical, operational, and environmental facets of next-generation mobility systems.

1.5.2 Practical Implications

This dissertation offers valuable insights for fleet operators, policymakers, and urban planners seeking to modernize transportation systems in a sustainable and user-centered manner. It provides operational guidance for managing shared autonomous electric vehicle (SAEV) fleets,

emphasizing decentralized, real-time coordination to improve economics and efficiency under uncertain travel demand and charging constraints. These findings inform strategic decisions about vehicle deployment, charging infrastructure, and scheduling. In addition, the research demonstrates that with proper management, SAEVs can effectively address the technological challenges EV while maintaining service availability comparable to conventional vehicles.

Second, the user-focused analyses - especially those that reveal heterogeneous consumer preferences - provide guidance for both public agencies and mobility platforms. Understanding how factors such as cost, wait time, and trust in autonomous technologies influence passenger choices enables more targeted marketing and service design. Such insights help mitigate potential resistance to AVs, paving the way for broader adoption and a smoother transition to mixed human and self-driving fleets. Key findings show that confidence in AVs should increase as they become more prevalent in the market to ensure smooth adoption. In addition, even a small share of AVs can significantly reduce carbon emissions and improve the performance of ride-hailing platforms.

The introduction of digital twin-driven approaches to the design and operation of large-scale EVCHs has significant implications for city planners and energy providers. By combining high-fidelity simulations with deep reinforcement learning, this dissertation demonstrates how to optimize charging infrastructure under realistic conditions, taking into account fluctuating demand, renewable integration, and building energy loads. These data-driven methods support policy decisions on zoning, grid expansion, and public-private partnerships, reducing investment risk and improving long-term viability. The results suggest that standard chargers can adequately meet EVCH demand, minimizing the need for fast charging, while smart operations management can reduce costs and effectively integrate renewable energy sources.

Finally, the development and application of dynamic pricing techniques demonstrates how EVCH operators can address both financial and system-level objectives. By shifting charging demand away from peak periods and adjusting variable power supply, operators can reduce grid stress while increasing revenues. These pricing models are particularly valuable in urban contexts, where growing EV adoption and limited grid capacity require innovative strategies to balance supply and demand.

Overall, the dissertation provides actionable insights that bridge technological solutions with social and economic considerations. It underscores the value of interdisciplinary methods for designing more resilient, profitable, and inclusive urban mobility systems-ultimately contributing to a more sustainable and interconnected transportation ecosystem.

1.5.3 Limitations and Future Research

While this dissertation provides substantive advances in shared autonomous electric mobility, limitations provide avenues for future work. First, although the simulation-based approaches employed here allow for rich exploration of operational and behavioral complexities, real-world

validation through pilot studies or field experiments remains limited. Future research could integrate real-world metrics, user feedback, and partial implementations to refine the models and validate their applicability.

Second, while the focus on social and managerial aspects is consistent with information systems and operations management, less attention is paid to the technical and policy dimensions—for example, the technical requirements of vehicle-to-infrastructure communications, cybersecurity challenges, or the nuances of multi-stakeholder regulatory frameworks. Collaboration with engineering and policy researchers would deepen the understanding of how different institutional and technological forces shape SAEVs deployment at scale.

Third, the work emphasizes single-mode vehicle services (ride-hailing fleets or EV charging hubs) without considering the interplay of multiple modes (e.g., public transit, micromobility) in a more holistic ecosystem. Future research could explore how the integration of multimodal transportation options could further optimize resource allocation, user accessibility, and environmental impacts.

Fourth, the socio-technical transitions examined here are typically viewed at the city level, yet mobility behavior and policy contexts can vary dramatically by region. Cross-comparative studies or international collaborations would help assess the extent to which these findings generalize to different cultural, economic, or regulatory environments.

Finally, while the research highlights trust and behavioral factors as critical to user adoption of autonomous services, more questions remain. For example, how might shifts in public perceptions of safety or privacy affect SAEV adoption, or how will competitive dynamics among multiple operators shape both pricing and sustainability outcomes? Addressing these and related questions may require deeper forays into behavioral economics, sociology, and market design. By tackling these challenges, subsequent research can both extend and refine the lessons learned from this dissertation and move us closer to resilient, equitable, and highly efficient urban mobility systems.

1.6 Declaration of Contributions

I (the author) would like to express my gratitude to the co-authors and collaborators who have contributed, to varying extents, to the individual research projects (Chapters) presented in this cumulative doctoral dissertation (the Thesis). Below, I outline my own contributions to these projects and acknowledge the valuable roles my co-authors played in their development and success.

Chapter 1: The author is the primary contributor to this chapter, having performed the majority of the work. This includes conceptualizing the interdisciplinary IS-OM research framework for next-generation mobility systems, defining the research objectives, and structuring the overall narrative. The author also conducted a comprehensive review of the relevant literature,

synthesized key theoretical and methodological findings, and positioned the research within the broader academic discourse. In addition, the final manuscript was developed to provide a cohesive foundation for the dissertation, ensuring clarity in the articulation of the research vision and objectives.

Chapter 2: This chapter was written solely by the author of the dissertation and encompasses the full scope of research, analysis, and writing. The work has greatly benefited from the constructive feedback and guidance provided by **Prof. Dr. Wolfgang Ketter**, **Dr. John Collins**, and **Dr. Nicolò Daina** throughout the writing and revision process. Their insights helped to refine the research approach, strengthen the theoretical foundations, and improve the clarity and impact of the final manuscript.

Chapter 3: This chapter was co-authored with **Anna Taudien**. For the modeling **Anna Taudien** conducted the online survey experiments and extracted the user classes and I developed the agent-based modeling, designed the simulation experiments and conducted the simulation analysis. For the rest, both authors contribute equally to all aspects of the research, including conceptualization, positioning, and writing the final paper. The work has also benefited from valuable feedback and guidance provided by **Prof. Dr. Wolfgang Ketter** and **Prof. Dr. Alok Gupta** throughout the writing and revision process, which helped to refine the research approach and improve the clarity of the final manuscript.

Chapter 4: This chapter was co-authored with **Dr. Karsten Schroer**, with both authors contributing equally to the conceptualization, algorithmic and simulation modeling, analysis, positioning, and writing of the final manuscript. The research greatly benefited from the insightful feedback and guidance of **Prof. Dr. Wolfgang Ketter** and **Prof. Dr. Thomas Y. Lee**, who served as academic advisors on this project and provided valuable direction throughout its development.

Chapter 5: The author is the primary contributor to this chapter, overseeing all aspects of its development. The work benefited greatly from the valuable feedback and guidance of **Prof. Dr. Konstantina Valogianni**, **Dr. Karsten Schroer**, and **Prof. Dr. Wolfgang Ketter**, whose insights were instrumental in refining the manuscript and strengthening its positioning within the research field.

Chapter 2

Cooperative Learning for Smart Charging of Shared Autonomous Vehicle Fleets¹

2.1 Introduction

Transportation systems are evolving at a fast pace through electrification, automation, and business model innovations brought by the sharing economy (Mahmassani 2016, Sperling 2018b). Electric vehicles (EVs) promise better air quality in urban areas and mitigate the climate impact of transport systems as electricity generation is increasingly decarbonized. Simultaneously, shared mobility provides on-demand transport with fewer vehicles on the road than the traditional personal transportation paradigm based on private car ownership. Autonomous vehicles (AVs) are prominent in stimulating shared electric mobility adoption by boosting efficiency, reducing operational costs, and mitigating electrification barriers such as range anxiety. Combining the advantages, shared autonomous electric vehicles (SAEVs) allow mobility-on-demand (MoD) fleets to streamline their services by improving safety, sustainability, and efficiency (Qi et al. 2022). SAEVs could also increase the long-term profitability of shared fleets due to eliminating driver payments, which encourage ride-hailing platforms (e.g., Lyft and Uber) to employ AVs in their fleets (Siddiq and Taylor 2022). Real-world use cases can be found in large cities (e.g., San Francisco, Shanghai, Phoenix, Singapore, Tokyo, Moscow) deployed as autonomous taxis (Dong et al. 2022).

¹This Chapter has been published in its entirety in the following peer-reviewed academic journal: Ahadi, R., Ketter, W., Collins, J., & Daina, N. (2023). Cooperative learning for smart charging of shared autonomous vehicle fleets. *Transportation Science*, 57(3), 613-630. Earlier versions of this Chapter have also appeared in a (non-copyrighted) peer-reviewed academic conference: Ahadi, R., Ketter, W., Collins, J., & Daina, N. (2021, May). Siting and Sizing of Charging Infrastructure for Shared Autonomous Electric Fleets. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (pp. 88-96).

Despite the benefits, SAEVs bring multiple managerial challenges for fleet operators (Loeb et al. 2018). This study sets out to support operational decision-making solutions for a fleet operator, assuming higher-level decisions (e.g., fleet size and charging infrastructure) to be exogenously determined. The objective is to maximize service quality (the rate of served requests) and profitability (profits from serving trips minus energy cost). We design a realistic comprehensive package of operations but zoom in on smart charging, which substantially challenges the fleet managers of shared electric vehicles due to the technological challenges of EVs (e.g., long charging time and infrastructure scarcity) to which their non-EV counterparts are immune (Pantelidis et al. 2022).

Recent works (e.g., Kullman et al. 2021a) address the dynamic and stochastic charging management of driver-based ride-hailing fleets; they mostly overlook the role of autonomous and connected mobility by considering either a centralized control that is not scalable or a decentralized model that does not exploit cooperation among vehicles. A decentralized approach without cooperation suits fleets of e-taxis or (to a lesser extent) e-ride-hailing services, where drivers act to maximize their profit: these are indeed the applications found in recent literature. However, cooperation is crucial for a fleet that owns vehicles, and where what matters is the overall fleet operator profit (and the quality of service provided to the clients to retain them). Another common unrealistic core assumption of the existing research is an unlimited (i.e., charging stations (CSs) are available anytime) and/or homogeneous (i.e., the same capacity and power rate) charging network. At the same time, Guillet et al. (2022), Pantelidis et al. (2022), and Froger et al. (2022) demonstrate the significant impact of capacitated CSs on the charging scheduling of electric fleets. Moreover, most researchers (e.g., He et al. 2021, Iacobucci et al. 2019) manage the charging behavior of shared EVs to meet mobility demands or minimize energy costs but fail to treat both goals simultaneously. To cover these gaps, we propose a decentralized, cooperative model that optimally identifies vehicles that need charging and assigns them to CSs in a limited and heterogeneous charging network with time-varying energy prices. Overall, we show that for addressing the charging management of SAEVs with capacitated CSs, it is not only vital to jointly model the charging scheduling and CSs allocation problems but also beneficial to learn optimal policies hierarchically (i.e., distinguishing decisions to different levels). For brevity, we term this problem as *Cooperative Charging of Shared Autonomous Electric Vehicles* (SAEV-CC).

We define our main contributions as follows:

1. We relax common critical and unrealistic assumptions of most existing research by addressing SAEV-CC with capacitated charging infrastructure (i.e., a few CSs with heterogeneous capacity and power rates) and time-varying energy prices, which leads to a joint charging scheduling and resource allocation problem. CSs are privately managed, but because of uncertain charging interruptions (flexibility of SAEVs) and new charging demand arrivals, vehicles' charging/waiting process takes stochastic periods. We demonstrate that correct-

ing the unlimited charging network assumption alters optimal charging policies due to interactions between proper charging time and location.

2. We propose a fully-cooperative and decentralized charging management model for SAEVs. Although all vehicles belong to a single fleet (i.e., a central fleet management problem), even cutting-edge central dynamic control approaches are computationally intractable for nontrivial cases. Therefore, we decentrally model SAEV-CC as a Markov decision process (MDP) and solve it using deep Q-network (DQN) algorithms. In contrast with recent decentralized dynamic control in e-taxis and e-ride-hailing applications (e.g., Kullman et al. 2021a, Liang et al. 2020), we address two major challenges by training agents to cooperate and mitigating the non-stationarity issue in multi-agent control systems (i.e., agents changing their policies constantly). Instead of off-the-shelf independent learning, our multi-agent reinforcement learning (MARL) agents forgo selfishness with the aid of reward-shaping techniques (Oroojlooyjadid et al. 2022). To mitigate an unstable environment, we apply mean-field approximation by which agents make adjusted decisions with respect to a proxy of the system state.
3. Our model builds another methodological advantage by using hierarchical learning to distinguish decision levels while preserving interactions. We clarify that hierarchical models have superior efficiency for multi-level decisions as they scale down action spaces, ensure an unbiased exploration of decision levels, and identify each level reward function.

In addition to our major contributions, we also deduce some managerial implications.

4. As our proposed model includes decomposition assumptions, we provide an upper bound (a non-EV scenario) and two benchmark models (central reinforcement learning and online reoptimization). Our decentralized model performs and behaves close-to-the central dynamic model while dominating the reoptimization model. Compared to the upper bound, there is only a slight reduction in fleet performance, validating our model performance and encouraging operators willing to electrify their fleets.
5. To evaluate policies, we simulate the mobility environment using an agent-based model (ABM) calibrated by actual data from ShareNow (a leading carsharing platform) in Berlin, Germany.
6. To provide further managerial insights, we conduct several sensitivity analyses on strategic and tactical factors and show how the fleet size, charging infrastructure, and energy tariffs affect fleet performance and charging policies.

The rest is organized as follows. Section 2.2 shortly reviews the related literature. We describe SAEV-CC in Section 2.3. In Section 2.4, we propose our model to address the problem, then apply it to real-world numerical examples and discuss results in section 2.5. Finally, Section 2.6 concludes the paper and presents limitations and future works.

2.2 Related Work

This work contributes to two streams of literature: (1) operations management of shared electric mobility platforms and (2) distributed autonomous agent controllers applied to dynamic decision-making of shared vehicles.

2.2.1 Decision Making for Shared Electric Vehicles

Many researchers expect that technological trends in mobility (sharing, automation, and electrification) could dominate the future. On the one hand, Fagnant and Kockelman (2014), and Chen et al. (2016) reveal that a shared autonomous vehicle (SAV) can replace up to 10 conventional vehicles, while it is limited to 3.7-6.8 for an SAEV, relying on charging rate and battery capacity. Their findings have motivated us to focus on the charging management of SAEVs to mitigate the technological challenges of electrification. On the other hand, these promising characteristics bring multiple operations management challenges at strategic, tactical, and operational levels. Strategically, Levin et al. (2019), and Lokhandwala and Cai (2018) study the fleet sizing of SAEVs and carsharing systems, and (Lokhandwala and Cai 2020), and Ahadi et al. (2021) study the impact of autonomous mobility on charging infrastructure development. Most recently, Chen and Liu (2022) investigate the charging facility development and smart charging management of autonomous carsharing platforms as a joint problem, where one of their key findings approves that the deployment of both standard, and fast CSs can improve the fleet performance.

Operationally, there is growing interest surrounding the dynamic vehicle routing problem of SAEVs (Dong et al. 2022). Our problem resembles ride-hailing platforms that are recently reviewed by Ho et al. (2018). Precisely, Hyland and Mahmassani (2018) explore matching vehicles and requests; and Rossi et al. (2018), and Schroer et al. (2022) study the repositioning of empty vehicles. Using EVs in MoD services presents recharging challenges, overlapping with smart charging problems. From the power grid perspective, there are two modes of load control. The first is top-down coordination, where a central operator directly controls charging loads, like the model studied by Jian et al. (2017). The second is bottom-up coordination, where market mechanisms incentivize users to charge in a desirable pattern employing various methods, such as a mean-field game theoretic approach proposed by Zhu et al. (2016), a novel dynamic model introduced by Valogianni et al. (2020b), and a two-stage mechanism design and charge scheduling model proposed by Wu et al. (2022). In our work, we consider a combination. The grid operator provides electricity with time-of-use (ToU) prices (indirect control), while the fleet operator can directly control its EV charging loads. In another segment of smart charging literature concerning demand side preference analysis, Daina et al. (2017) examine individual responses to price signals while considering coupled travel and charging requirements. Integrating routing and charging decisions, Sweda et al. (2017) find optimal CSs for an EV along a given path, and

Florio et al. (2021), and Kullman et al. (2021b) address the uncertainty of CS availability and queue status. We follow this research problem to find optimal policies for charging and routing shared EVs to CSs by solving SAEV-CC.

With the growth of shared mobility, the average vehicle utilization and need for electrification mushroomed. Therefore, recent works have examined recharging for shared mobility services. Abdelwahed et al. (2020) optimize the charging of electric buses assuming perfect route schedules information. For rental networks like free-floating carsharing with uncertain demand, He et al. (2021) and Roni et al. (2019) address a joint charging infrastructure planning and operating problem. Kahlen et al. (2018b) study electric fleets as virtual power plants to offset the inflexibility of renewable energy sources. However, less focus has been given to the smart charging of SAEVs and ride-hailing systems. Iacobucci et al. (2019) optimize the (dis)charging of SAEVs as well as rebalancing and assignment decisions, but they fail to deal with uncertain demands. Zhang and Chen (2020) virtually generate a CS wherever a demand occurs and dynamically schedule the charging events using heuristic approaches. Regarding the synergy of SAEVs and microgrids, Qi et al. (2022) propose a time-space-energy network to explore how connected SAEVs can enhance the self-sufficiency and resilience of future microgrids while modeling a very abstract mobility system. In another work, Ma and Xie (2021) employ rule-based strategies to determine vehicles needing for charging and separately develop an online reoptimization model to assign vehicles to CSs. We consider their approach as a benchmark for our proposed model.

2.2.2 Reinforcement Learning Applications to Shared Mobility Systems

Reinforcement learning (RL) is a competing method recently applied to capture the dynamics and uncertainties of shared fleet management problems. Qin et al. (2020) develop a MARL model to assign vehicles to requests dynamically. The authors consider an event-based problem and formulate it as a semi-MDP (Sutton et al. 1999), where decisions take stochastic time steps to terminate. With a similar approach, Shou and Di (2020) study the repositioning of a multi-driver MoD system. Since a direct application of MARL might lead to optimal policies from the drivers' perspective, they propose a reward design scheme (i.e., reward-shaping Devlin and Kudenko (2011)) to achieve the desired equilibrium from a fleet perspective. They deploy a mean-field RL approach to overcome the non-stationarity issue. The mean-field RL is introduced by Yang et al. (2018) and applied in a ride-sharing order dispatching problem by Li et al. (2019), in which agents are aware of an approximation of neighbors' actions and states. In the present study, we adapt these augmentations for SAEV-CC to stabilize our learning algorithms.

A smaller subset of studies regards mobility service fleets' dynamic and stochastic charging decisions. Dong et al. (2022) explore a dynamic vehicle allocation problem for SAEVs with respect to vehicles' energy levels and travelers' requirements. However, similar to many other related works, they use a rule-based (sending vehicles with low energy levels to the closest CS) charging strategy to recharge vehicles. Al-Kanj et al. (2020a) use approximate dynamic

programming to determine when a vehicle needs to charge and when it should relocate. To match vehicles and trips while considering charging decisions, Shi et al. (2020) decentrally learn state values using RL and solve a mathematical model to make operational decisions centrally. Similarly, Liang et al. (2020) use independent deep RL agents for shared EVs to make rebalancing and charging decisions. Although the mentioned works find the optimal policies determining when vehicles should charge, they disregard optimizing the CS allocation problem as they assume the nearest CS is always optimal. In the domain of e-taxis, Wang et al. (2020) also address the joint charging and repositioning recommendation problem using hierarchical MARL, where the manager and worker models act hierarchically. In their model, each taxi driver aims to maximize his cumulative rewards, distinct from our problem, as taxi drivers do not cooperate for a joint goal. Moreover, their model only recommends whether drivers charge their batteries or not while not recommending the proper CS (i.e., a homogeneous charging facility assumption). In a recent work sharing many similarities to ours, Kullman et al. (2021a) address operational decision-making (assignment, rebalancing, and charging) for dynamic ride-hailing with EVs. Their approach merges charging and rebalancing decisions, meaning that if a vehicle is not assigned to a trip, it must relocate to a station. To do so, they develop a central RL agent to find optimal policies and extend it to MARL for the sake of scalability.

We remark on a few works (e.g., Pantelidis et al. 2022, Froger et al. 2022) considering capacitated CSs in other applications. The former solves the repositioning of shared cars for a limited charging network (still, CSs have the same power) while disregarding the optimal charging scheduling of vehicles (vehicles charge when their energy level is below 20%). The latter addresses how to route EVs considering routes charging capacity and vehicle states to mitigate charging facilities scarcity. Although they include a multitude of charging technologies and modeling advanced (non-linear) charging constraints, they neglect to optimize the charging schedules and restrict their mobility environment to a small number of predefined customers with known requirements.

Overall, previous research in the charging management of shared vehicles develops either a central control approach that is not applicable for realistically-sized problems or a multi-agent control approach while disregarding agents' cooperation (e.g., independent ride-hailing or taxi drivers). Therefore, ours mark the first application of the dynamic charging management of non-trivial SAEV fleets, bringing together cooperative multi-agent decision-making and considering the charging decisions' knock-on effects in anticipation of future mobility demands. Another void in the literature is separating the problems of vehicles' charging scheduling and allocating to capacitated CSs. This separation is primarily considered in the literature in favor of an unrealistic assumption of unlimited charging infrastructure. However, we signify a joint approach that optimally schedules charging times and allocates vehicles to proper CSs. Another novelty that singles out our methodology is the employment of hierarchical decision-making by which agents make decisions at different levels while preserving their interactions. This extension con-




Figure 2.1: Operations Problem of a Shared Autonomous Electric Fleet

siderably builds up the scalability, breaking down an enormous action space into sub-action spaces and stabilizing learned policies, which is a major challenge of learning algorithms. Finally, although our focus is narrowed on the charging management of SAEVs, our approach is easily extensible for various stochastic and dynamic operational decision-making problems dealing with distinguished decision levels.

2.3 Problem Description

We develop an operational decision support system for a fleet of SAEVs focusing on the smart charging management over an operating period \mathcal{T} . Although our focus is narrowed on developing a dynamic charging management model, first, we describe the overall fleet operations and then zoom in on the charging section. A big picture of the problem is summarized in Figure 2.1. The fleet owns $\mathcal{J} = \{1, 2, 3, \dots, J\}$ homogeneous EVs (the same size, energy consumption, and battery capacity), fully automated (able to make/execute decisions), and connected (communicating with the fleet operator and other vehicles). Vehicles provide MoD services for customers ($i \in \mathcal{I}$) requesting a point-to-point trip within a geographical region. Based on spatio-temporal arrival rates, requests arise randomly in multiple zones ($z \in \Gamma$) across a given region, specified by an origin (u_i), a destination (d_i), a price (p_i), and an uncertain patience time (b_i).

Once a trip is requested, the fleet operator searches for an appropriate vehicle, during which its traveler might cancel if the waiting time takes longer than the patience threshold (b_i). Vehicles are eligible to serve a request only if they are in the vicinity and have enough energy. Therefore, to achieve higher performance, the operator must continuously check and rebalance (charge and reposition) vehicles in anticipation of future demands. Ideally, the fleet operator should make all decisions (assignment and rebalancing) jointly, considering decisions interactions and vehicles cooperation. However, since the arrival rates of requests and their requirements (destination and energy consumption) are dynamic and uncertain, a holistic decision-making problem becomes computationally intractable for nontrivial cases, especially if precise micro

decisions (e.g., charging scheduling and CSs allocation) are included. Therefore, to analyze the details of operational decisions, we separate the assignment from rebalancing problems. Also, to refine the scalability, we assume that vehicles cooperate to make distributed rebalancing decisions. This is only a decomposition approach, and vehicles seek a joint goal, maximizing fleet profitability (trip revenues minus operational costs) and service quality (acceptance rate).

Every idle vehicle should choose between charging, repositioning, and serving at every decision time, which happens after completing a request or remaining idle longer than a decision waiting threshold (δ^{idle}). If the vehicle decides to serve, it will be assigned to a request when a match happens; otherwise, it will relocate to a parking spot after a period. If the decision is to charge or reposition, the vehicle relocates to a CS or a target zone. Note that vehicles can stay idle in large parking spots within the service area and charge their batteries in restricted charging facilities $\mathcal{C} = \{1, 2, 3, \dots, C\}$ with identified location, capacity (number of docks), and power rate.

Figure 2.2 depicts the operations of an individual vehicle and its connection to the fleet operator (the fleet and vehicles have separate agents). The fleet agent constantly tracks the system state (CSs, open requests, and vehicles) and is responsible for assigning vehicles to requests. It periodically pools open requests and matches them with available vehicles using an online optimization algorithm (details in Online Appendix A.3). Vehicle agents receive the assigned requests requirements and start serving them accordingly. They also make rebalancing decisions (charging, repositioning, and parking in order) while having complete access to the fleet state and their individual information. From now on, we use the agent term to represent a vehicle agent.

According to our focus, we provide precise information on the charging process. To make charging decisions, agents cooperate to maximize fleet performance in anticipation of future demands. The charging decision entails time and location; i.e., vehicles must decide when and in which CS to charge. We assume the fleet operator privately manages CSs. They are limited and distributed within the service region, which could vary in queue state, distance from vehicles, charging rate, and location. Each CS manages its queue using a lowest-energy-highest-priority strategy, meaning that low-energy vehicles get connected earlier than others even though they arrive later. The CS operator charges connected vehicles with full available power. Finally, we assume that waiting and charging vehicles in CSs can interrupt the process and serve urgent mobility demands if they have enough energy (details in Online Appendix A).

2.4 Model

SAEV-CC is stochastic and sequential. The fleet state (including supply and demand) changes continuously, and the operator (or vehicles) needs to make charging decisions accordingly while accounting for subsequent states and decisions in addition to anticipating uncertain mobility




Figure 2.2: Vehicle Operations and Interaction with Fleet Operator

demands. As a tangible example, if a vehicle decides to charge, the fleet might miss some immediate trip requests; if not, it might face a lack of supply for expected future demands. With perfect information relaxation, the operator can schedule charging decisions using deterministic methods for trivial instances. In contrast, planning is peculiar in stochastic and real-size cases and needs approximate solutions. Since the environment’s dynamics are unknown in SAEV-CC, we use RL, a model-free approach, to learn optimal policies through interactions with the environment.

Adopting a single decision-maker is infeasible for nontrivial fleets as the joint action space becomes intractable. Therefore, we design a distributed strategy by which each vehicle (agent) makes individual charging decisions (when and where to charge) while simultaneously learning optimal policies. In our model, a decision epoch begins when a vehicle completes a request or remains idle longer than the waiting threshold ($\delta^{idle} = 15$ minutes) from its last update. It leads to a temporally abstracted and discrete event problem where decisions take stochastic time steps to terminate. We formulate the problem as a semi-MDP to cope with this challenge.

Although vehicles make decisions individually, they jointly seek the fleet goal, leading to a fully cooperative MARL problem. We use reward-shaping (i.e., penalizing responsible vehicles for unserved trips) to nudge vehicles to cooperate. Also, to tackle the non-stationarity issue of MARL (agents change their policies continuously), we employ a mean-field approximation to stabilize the agents by improving the state space definition to encompass an approximation of the fleet state. A charging decision includes a lower sub-action (determining the CS). Therefore, we apply a hierarchical RL algorithm to make sub-decisions using distinguished but connected

controllers. It helps define independent reshaped reward functions and scale down the action space. Finally, we use deep learning to boost our model’s generalizability.

2.4.1 Multi-agent Reinforcement Learning Model

To decentrally address SAEV-CC, we develop a MARL model, capturing the agents’ cooperation. We model it as a partially observable MDP (Littman 1994), represented by a tuple $(\mathcal{S}, \mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_J, \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_J, P, R_1, R_2, \dots, R_J, J, \gamma)$. \mathcal{S} is the environment state space (fleet state in SAEV-CC). J is the number of vehicle agents. Each agent has an action space \mathcal{A}_j (whether and where to charge) and a local observation space $\in \mathcal{O}_j$ (time, mode, location, state of charge (SoC)) as a part of the system state $s \in \mathcal{S}$, yielding joint observation ($\mathcal{O} = \mathcal{O}_1 \times \mathcal{O}_2 \times \dots \times \mathcal{O}_J$) and action ($\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_J$) spaces. $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ represents the state transition probability, $R_j : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function of the agent j (accumulated profits), and γ is a discount factor. The agent j , given an observation $o_j \in \mathcal{O}_j$ takes an action $a_j \in \mathcal{A}_j$ following a policy $\pi_j : \mathcal{O}_j \times \mathcal{A}_j$. The system state transits to s' based on a state transition probability $P(s'|s, a)$, and the agent j receives a reward $r_j(s, a, s')$. Since vehicles are homogeneous we consider the same policy for all agents (i.e., $\pi_j = \pi, \forall j \in \mathcal{J}$). The identification of a policy requires defining action-state values ($Q^\pi(s, a)$) representing the expected cumulative reward by taking action a in state s , and following policy π , which can be represented recursively using the Bellman equation $Q^\pi(s, a) = \mathbb{E}_{s' \sim P(\cdot|s, a)}[r(s, a, s') + \gamma \max_{a' \in \mathcal{A}(s')} Q^\pi(s', a')]$. Finding optimal values ($Q(s, a) = \max_{\pi} Q^\pi(s, a), \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$) identifies the optimal actions (Sutton and Barto 2018b).

Applying Mean-field Approximation

Since other agents affect the environment, the Q-value function ($Q_j = Q(s, a)$) relies on the system (fleet) state s and the joint action a (charging decision of all SAEVs). This dependency makes the learning process computationally intractable due to huge-sized state and action spaces in addition to our problem’s asynchronous nature (i.e., varied and stochastic number of agents take actions between two consecutive epochs). To overcome this, we simplify Q-values dependencies on the joint state s and action a . Instead of an independent setting ($Q_j = Q(o_j, a_j)$), we apply mean-field approximation to combat the non-stationarity issue (Yang et al. 2018, Wang et al. 2020). The idea is to simplify the interaction between the agent j and its neighbors by aggregating them to a virtual mean agent. Thus, we redefine Q-values as $Q_j(o, a) = Q_j(o_j, o_{-j}, a_j, a_{-j})$, where o_{-j} and a_{-j} denote the joint observation and action of agents except the agent j , respectively. Since the agent j majorly interacts with its neighbors, we denote them as N_j and approximate the Q-values as $Q_j \approx \frac{1}{|N_j|} \sum_{n \in N_j} Q_j(o_j, o_n, a_j, a_n)$. Further, we use a Taylor ex-

pansion to simplify the Q-values to:

$$Q_j \approx \frac{1}{|N_j|} \sum_{n \in N_j} Q_j(o_j, o_n, a_j, a_n) \approx Q_j(o_j, \bar{o}_j, a_j, \bar{a}_j) \approx Q_j(o_j, \bar{o}_j, a_j), \quad (2.1)$$

where \bar{o}_j, \bar{a}_j are respectively the proxies of neighbors' states and actions. In our problem, the mean observation \bar{o}_j is a vector containing the number of open requests within the vehicle j vicinity (local demand), the number of neighboring available vehicles (local supply), and CSs status. The neighbors' mean action (\bar{a}_j) is unknown and needs to be predicted in advance, which is very uncertain. Hence, the agents do not regard it while making charging decisions. However, the neighbors' mean action (\bar{a}_j) affects the next state and consequently influence the estimated Q-values. To simplify writing, we use the traditional form of Q-values $Q(s, a)$ and define our MDP's components in the following.

States

We define the modified vehicle's state as $s = (s_t, s_v, s_l, s_c)$. The state of time (time of the day in minutes) is denoted by s_t . The vehicle state $s_v = (m, SoC, l)$ is a vector of: mode ($m \in \{\text{idle, serve, reposition, queue, charge, en-route}\}$) describing the current vehicle job, energy status ($SoC \in [0, 1]$) indicating the battery percentage, and location ($l \in \Gamma$) representing the zone where the vehicle is located. Note that job types idle, reposition, queue, and charge are preemptable, meaning that they can be interrupted to serve urgent requests, whereas serving and en-routing jobs are non-preemptable. The vehicle's local state $s_l = (n_s, n_d)$ entails the number of neighboring available vehicles/supply (n_s) and the number of open requests/demands (n_d) within the vehicle's vicinity. Finally, the CSs status $s_c = (q_1, q_2, \dots, q_C)$ indicates the number of vehicles in each CS. All vehicles start an episode with an initial state s_v^0 at time 0: a random location, an idle mode, and a random SoC . They finally terminate the episode when the operations time is over.

Actions

The action $a \in \mathcal{A}(s)$ indicates a charging decision. At any epoch, each agent takes only a single action. The action space ($\mathcal{A}(s) = \{c_1, c_2, \dots, c_C, \emptyset\}$) entails charging at different CSs (e.g., $a = c_c \in \mathcal{C}$) or doing nothing ($a = \emptyset$). Doing nothing means remaining available for the upcoming demands. Some actions are not allowed. A vehicle is allowed to choose among eligible CSs, which have at least one free parking spot (not necessarily a free charger) and are energy-feasible, meaning that the energy consumption driving to the CS must be less than the vehicle's battery state minus a safety energy amount for unexpected events.

Reward Function

To eliminate the credit assignment challenge of multi-agent models (i.e., distinguishing each agent’s share out of a joint reward), we consider individual rewards for agents. The reward $r(s, a)$ corresponds to the accumulated revenues (positive rewards) and costs (negative rewards) incurred during a subsequent epoch. Positive rewards are the revenues of serving trips (r_s) excluding their driving costs. Negative rewards account for charging costs (r_c): the cost of charged energy with a time-varying electricity tariff, driving costs (r_d): the cost of relocating to the assigned CS (proportional to distance), and waiting costs (r_w): the parking costs is the assigned CS, which is a function of queue time. The cost calculation details are in Online Appendix A.4.

Considering individual vehicle rewards might not lead to a desirable equilibrium from the fleet’s perspective due to vehicles’ selfishness (Shou and Di 2020). To tackle this challenge, operators can modify/shape the reward function (Oroojlooyjadid et al. 2022). Therefore, we include a penalty factor in the reward function to bypass selfishness. Since the objective is to maximize profits and acceptance rate, we assign a penalty r_p to agents once a trip remains unsatisfied. The penalty r_p also represents the acceptance rate in our reward function.

Penalizing all agents might mislead them as they do not receive a proper critic according their individual action (i.e., they do not realize if their action caused the missed trip or not). Thus, we define and only penalize the responsible agents for a missed trip, which are those within the coverage area (Δ) of the missed trip that: a) had high *SoC* (set to 50%) and decided to charge (to reduce unnecessary charging when the supply is low), or b) had low *SoC* (less than covering the trip) and did not decide to charge (to incentivize vehicles with a lack of energy to charge instead of seeking immediate profits). Assigning excessively penalties might yield sub-optimal policies since it is possible to gain more profits from charging immediately (even with high *SoC*) and serve more future demands. Therefore, we conduct simply a search to find a balance between the profit and penalty components. We formulate the reward function as:

$$r(s, a) = w_s r_s - w_c r_c - w_d r_d - w_w r_w - w_p r_p. \quad (2.2)$$

Each reward component has a corresponding adjustable weight (w) that is quantified based on the agent’s goal. Therefore, we set the weights of profit components to one, but empirically tune the penalty weight (w_p) to relate profits and service quality and achieve better performance.

Transition Function

A transition from one state to another is a function of the selected action a and its termination duration τ . Since executing the action, the agent might have completed several jobs (e.g., charging, and serving trips). The selected action and the behavior of other agents until the termination affect the system state. So, when an epoch starts (an agent makes a charging

decision), we update the subsequent state s' by observing changes in the vehicle status (s'_v), time (s'_t), nearby supplies (s'_s) and demands (s'_l), and CSs status (s'_c). Note that the system dynamics are unknown, and we use a simulation to track them.

2.4.2 Dealing with Temporal Abstraction

As epochs start at different times in our problem (an event-based RL model), we face the challenge of temporal abstraction; i.e., actions take different and stochastic time steps to terminate, and the reward is distributed over this period. The easiest way to extend the RL framework with temporal abstractions is to retain a semi-MDP (Sutton et al. 1999). The major difference is an accumulated reward function that is discounted over the action duration time:

$$\hat{r} = \sum_{h \in \mathcal{H}} \gamma^{\tau_{e_h}} \left(\frac{r_h}{\tau_{o_h}} + \gamma \frac{r_h}{\tau_{o_h}} + \dots + \gamma^{\tau_{o_h} - 1} \frac{r_h}{\tau_{o_h}} \right) = \sum_{h \in \mathcal{H}} \gamma^{\tau_{e_h}} \frac{r_h (\gamma^{\tau_{o_h}} - 1)}{\tau_{o_h} (\gamma - 1)}. \quad (2.3)$$

Agents receive multiple rewards for several tasks (e.g., charging and serving trips). After taking action, each task (h) takes a period to start (τ_{e_h}) and a period to complete (τ_{o_h}). We split and discount each task's reward over the time steps. Also, the action duration (τ) must be considered while updating state values, which leads to $Q(s, a) = \mathbb{E}_{s' \sim P(\cdot | s, a)} [\hat{r}(s, a, s') + \gamma^\tau \max_{a' \in \mathcal{A}(s')} Q^\pi(s', a')]$.

2.4.3 Overview: Solving an MDP with Deep Q-learning

To identify optimal policies, we use a value-based method (Q-learning), learning the optimal policy π through updating state-action value $Q(s, a)$. Therefore, each agent maps its state s to values, takes an action a , moves to a new state s' , receives a reward \hat{r} , and updates Q-values by:

$$Q(s, a) \leftarrow Q(s, a) + \alpha [\hat{r}(s, a, s') + \gamma^\tau \max_{a'} Q(s', a') - Q(s, a)], \quad (2.4)$$

where α is a learning rate and $0 < \alpha \leq 1$. We use a semi-MDP, where the next state-action value is discounted based on the action duration time (τ). To eliminate the limitations of tabular methods, we employ a deep Q-network (DQN) to estimate $Q(s, a)$. We refer the readers to Mnih et al. (2015) for more details. Generally, DQN uses a neural network with weights θ as a Q-network (i.e., $Q(s, a; \theta) \approx Q(s, a)$) for approximation. To train, agents store transitions (state, action, reward, duration) in a single memory (experience replay). The Q-network parameters get updated using a random batch of transitions from the experience replay by minimizing the loss value:

$$\mathcal{L}(\theta) = \mathbb{E}_{s, a, s'} [(\hat{r}(s, a, s') + \gamma^\tau \max_{a'} Q_\theta(s', a') - Q_\theta(s, a))^2]. \quad (2.5)$$




Figure 2.3: Hierarchical Reinforcement Learning for SAEV-CC Problem

To compute the target $(r(s, a, s') + \gamma \max_{a'} Q_{\bar{\theta}}(s', a'))$, we use a target network with parameters $\bar{\theta}$ that revamps the stability of DQN.

2.4.4 Extension to a Hierarchical Approach

SAEV-CC contains two decision levels (charging scheduling and allocating to CSs). After deciding to charge, vehicles need to take a charging destination. Instead of including all CSs in the same action space (see section 2.4.1), we use a hierarchical learning model, distinguishing reward functions for each decision level, which stabilizes the learning process and reduces the action space size significantly. Precisely, a high-level controller learns a policy over intrinsic sub-goals (whether to charge), and a low-level controller learns a policy over atomic actions (where to charge).

We develop a similar model to Kulkarni et al. (2016) (following the same terminology) and present high-level and low-level decisions as sub-goals and actions, respectively. The meta-controller takes a sub-goal, and the sub-controller takes the associated action (see Figure 2.3).

The sub-controller state is $s_S = (s_t, s_v, s_l, s_c)$ (the same as MARL agent state), while a redundant component s_d is included in the meta-controller state $s_M = (s_t, s_v, s_l, s_c, s_d)$, that is an estimated charging destination. We consider s_d to link controllers since the sub-controller policy changes over time, and without knowing the current sub-policy, agents might take non-optimal actions. The sub-goal ($a_{meta} \in \{c, \emptyset\}$) is among charging (c) and doing nothing (\emptyset); and the atomic action ($a_{sub} \in \{c_1, c_2, \dots, c_C\} \cup \{l\}$) is among CSs locations and the current vehicle location.

The extrinsic (sub-goal) reward r_{ex} is the same as in Equation (2.2). However, we adopt the penalty term for the intrinsic (atomic action) reward r_{in} slightly. Here, the responsible vehicles are charging and waiting vehicles that: a) are within the missed trip coverage area (Δ) but are not energy-feasible to serve the trip even after a long time (set to 30 minutes) since they decided to charge, or b) were positioned in the missed trip coverage area (Δ) when decided to charge, have enough energy to serve the trip but are located in a remote CS. A summary of the

Algorithm 1: Learning how to Make Charging Decisions

```

1 1: Initialize  $D_M, D_S, \theta_M, \theta_S, \epsilon_M$  and  $\epsilon_S$ ,
2 2: for  $i \in \text{episodes}$  do
3   3: Initialize the environment and get the start state  $s_S$  and  $s_M$ ,
4   4: Estimate the Q-value of current state using meta-controller for each sub-goal,
5   5: Take a sub-goal  $a_{meta}$  using an  $\epsilon - greedy$  approach,
6   6: if  $a_{meta} = \text{charging}$  then
7     7: Estimate the Q-value of the current state using the sub-controller for each
       action,
8     8: Determine a destination using an  $\epsilon - greedy$  approach,
9     9: else the destination is the current location,
10  10: Execute the action  $a_{sub}$ , receive rewards ( $r_{ex}, r_{in}$ ) and obtain next states  $s'_S$  and
       $s'_M$ ,
11  11: if  $a_{meta} = \text{charging}$  then
12    12: Store the transition  $(s_S, a_{sub}, r_{in}, s'_S)$  and update  $\theta^S$  using mini-batches from
         $D_S$ ,
13  13: Store the transition  $(s_M, a_{meta}, r_{ex}, s'_M)$  and update  $\theta^M$  using mini-batches from
         $D_M$ .

```

learning process for the hierarchical-MARL (HMARL) model is represented in algorithm 1. For controllers, we use separate neural networks (with parameters θ_M, θ_S) to make charging and destination decisions, replay memories (D_M, D_S) to store observations, and exploration rates (ϵ_M, ϵ_S) to explore the environment by making random decisions. For each episode, we initialize the environment. Agents receive states s_M and s_S , estimate action-state values, and determine a sub-goal using an $\epsilon - greedy$ approach (i.e., taking a random action with a probability of ϵ). The sub-controller determines a destination to fulfill the charging sub-goal by applying an $\epsilon - greedy$ policy over eligible stations (energy feasible with at least one free parking spot). Whereas the destination is the current vehicle location for doing nothing sub-goals. The agent executes the action and receives extrinsic r_{ex} and intrinsic r_{in} rewards. In the case of charging,

the agent stores the transition and updates the parameters of the meta-controller (θ_M) and sub-controller (θ_S); otherwise, it only saves the transition on the meta-memory (D_M) and updates the associated network.

2.4.5 Benchmark and Upper Bound Models

Assessing the learned policies is essential since distributed approaches only guarantee sub-optimal solutions. Due to a lack of optimal analytical solutions, we benchmark our proposed model against two central decision-making models: a) a single RL agent and b) a reoptimization strategy. Further, we consider an upper-bound scenario of non-EVs without technical charging restrictions.

Single-agent Reinforcement Learning Benchmark Policy

Our multi-agent decision-making approach contains some decomposition assumptions: a) a decentralized cooperative MARL to enhance scalability, b) a mean-field approximation of the fleet state to boost stability and break the tendency to overestimate, and c) a hierarchical decision-making approach to de-escalate the action space. To relax them, we design a single-agent dynamic charging (SADC) benchmark policy that makes operational decisions (charging, relocating, and serving) centrally on behalf of the fleet operator. To test it for the default size of the problem, we must make decisions for vehicles in multiple batches (e.g., 20 vehicles) and limit the number of destination options. We assume that the number of idle vehicles rarely exceeds the batch size.

We formulate the problem as an MDP. SADC is a discrete-time model that centrally makes decisions at equal time steps. The state contains the day time, a vector of vehicle states (mode, location, *SoC*, destination, expected task duration), a vector of associated vehicles in the decision batch, a vector of open trips in all zones, and a vector of CSs status (number of charging and waiting vehicles). For each vehicle in the decision batch, SADC takes action among serving, relocating, charging in the closest CS, and charging in the closest fast CS. If there are more idle vehicles than the batch size, SADC iteratively checks all of them. Here, we consider the reward function exactly aligned with the fleet objective, which is maximizing profitability and service quality ($r_{SADC} = r_{profit} - r_{missed}$), where r_{profit} is the accumulated fleet profit (served trips profits minus all operational costs), and r_{missed} is the penalty for missed trips during the epoch. We use DQN algorithm to train SADC.

Reoptimization Benchmark Policy

As another benchmark policy, we develop a consolidated rule-based, and reoptimization model based on Ma and Xie (2021). Regarding highly uncertain mobility demands, to avoid arbitrary assumptions of future revenue estimations, we employ a rule-based strategy to determine the

charging demands. We modify their model by using an hourly charging threshold (see Online Appendix B.3 for the temporal *SoC* threshold) to charge vehicles mostly when demands are low. Vehicles with *SoC* below the threshold need to charge and will be allocated to CSs using a mathematical programming model. From now on, we use REOPT to refer to the reoptimization benchmark model. The objective is to allocate the possible number of vehicles (charging demands) by considering a high penalty ($\rho^{Charging}$) for unserved demands while minimizing the estimation of aggregated total charging time. Although this is not completely aligned with the proposed model's objective function, similar to models in the literature (Ma and Xie 2021, Kullman et al. 2021a), we also consider a cost minimization objective function due to a lack of revenue estimation as a function of charging decisions.

$$\min_y \sum_{c \in \mathcal{C}, j \in \mathcal{J}_c} (t_{j,c}^{Driving} + t_{j,c}^{Charging} + t_{j,c}^{Queue}) + \sum_{j \in \mathcal{J}_c} (1 - \sum_{c \in \mathcal{C}} y_{j,c}) \rho^{Charging}. \quad (2.6)$$

where \mathcal{C} and \mathcal{J}_c respectively denote the set of CSs and charging vehicles. The decision variable $y_{j,c}$ is a binary assignment variable, which is set to 1 if vehicle j allocates to CS c . Based on assignment decisions, the model computes the estimated driving time $t_{j,c}^{Driving}$, charging time $t_{j,c}^{Charging}$, and waiting time $t_{c,j}^{Queue}$ for vehicle j and CS c . The parameters are as follows. $\rho^{Charging}$ is a penalty rate if a charging demand remains unserved. β^{Energy} and β^{Time} are respectively the energy consumption and duration of time ratios which are direct proportional to driving distances. $D_{j,c}$ is the distance between vehicle j and CS c . SoC_j is the vehicle j energy state. N_c , $\kappa_c^{Charging}$, $\kappa_c^{Parking}$, and κ_c^{Power} are respectively the current number of vehicles, chargers, parking spots, and the charging rate of CS c . Regarding queue management in CSs, $Z_{j,j'}$, a binary parameter, indicates the waiting priority of vehicle j over vehicle j' and only sets to one if the energy level of vehicle j is lower than vehicle j' . Finally, $N_{j,c}$ indicates the number of charging vehicles in CS c plus the number of waiting vehicles that have lower *SoC* than vehicle j . We define the constraints as follows.

$$y_{j,c} D_{j,c} \beta^{Energy} \leq SoC_j \kappa^{Battery}, \quad \forall j \in \mathcal{J}_c, c \in \mathcal{C}, \quad (2.7)$$

$$N_c + \sum_{j \in \mathcal{J}_c} y_{j,c} \leq \kappa_c^{Parking}, \quad \forall c \in \mathcal{C}, \quad (2.8)$$

$$\sum_{c \in \mathcal{C}} y_{j,c} \leq 1, \quad \forall j \in \mathcal{J}_c, \quad (2.9)$$

$$t_{j,c}^{Driving} = \beta^{Time} D_{j,c} y_{j,c}, \quad \forall j \in \mathcal{J}_c, c \in \mathcal{C}, \quad (2.10)$$

$$t_{j,c}^{Charging} = \frac{100\% - SoC_j}{\kappa_c^{Power}} \kappa^{Battery} y_{j,c}, \quad \forall j \in \mathcal{J}_C, c \in \mathcal{C}, \quad (2.11)$$

$$t_{j,c}^{Queue} = y_{j,c} \max\left\{ \frac{N_{j,c} + \sum_{j' \in \mathcal{J}_C} y_{j',c} (1 - Z_{j,j'}) - \kappa_c^{Charging}}{\kappa_c^{Charging}} E / \kappa_c^{Power}, 0 \right\}, \quad \forall j \in \mathcal{J}_C, c \in \mathcal{C} \quad (2.12)$$

Constraint (2.7) guarantees the energy feasibility for allocating vehicle j to CS c . The model does not assign vehicles to CSs more than their parking capacities employing Constraint (2.8). Constraint (2.9) allocates each vehicle to at most one CS. Constraint (2.10) computes the driving time as a function of $D_{j,c}$. Constraint (2.11) estimates the charging time of vehicle j at CS c and constraint (2.12) estimates the waiting time for vehicle j in CS c (E is the average charging demand). This applies a lowest-energy-highest-priority queue management. Vehicle j only needs to wait for the charging vehicles and waiting vehicles with higher priority. Also, $t_{j,c}^{Queue}$ only gets a non-zero value if the vehicle j is allocated to CS c ($y_{j,c} = 1$) and the number of more prioritized vehicles exceeds the capacity of CS c . We linearize Constraint (2.12) using a big M trick (See Online Appendix A.5).

Upper Bound

As an upper bound for the fleet performance, we consider a fleet of autonomous non-EVs, assuming that they can refuel all over the service area (every zone has at least one gas station) in a very marginal time. As we cannot access optimal solutions, this upper bound helps assess our proposed algorithms more accurately. Moreover, it deduces some managerial insights regarding the fleet performance of EVs and non-EVs.

2.5 Experiments

In this section, we apply the proposed models to a set of experimental instances, for training and evaluating which we simulate the transportation environment.

2.5.1 Mobility Simulation

We develop an ABM to simulate the mobility environment, wherein vehicles, CSs, and the fleet operator are interactive agents. We assume the system configuration (fleet size, charging infrastructure, pricing scheme, and electricity tariff) to be exogenously determined. Figure 2.4 illustrates the simulation flowchart. At operational decision times (e.g., 15 minutes), all idle vehicles run the operations module to make charging and repositioning decisions. Also, at assignment time steps (e.g., two minutes), open requests are matched with available vehicles using a reoptimization algorithm. Matched vehicles serve their assigned trip requests and follows

the operations module to determine their next operational actions. For details, please see Online Appendix A.




Figure 2.4: Agent-based Simulation Flowchart of SAEVs' Operations

2.5.2 Experimental Setup

First, we adopt a discretization approach in the underlying spatio-temporal network. It divides the service region into same-sized zones using hexagonal grids, leading to 89 hexagons (represented by Γ) with an edge length of 1.22 km. These zones have distinguished time-varying arrival rates, origin-destination patterns, and specific charging capacities. Regarding time, each episode takes five days (business days), assuming that vehicles get restarted during the weekend when the demand is low. Also, each time step in our simulation equals a minute in the real world.

A historical trip dataset of a leading free-floating carsharing company (ShareNow) from Berlin, Germany, is used to generate trips. The data was collected from November 2019 until February 2020, comprising 684,229 trips by 897 vehicles. Each observation consists of trip start and end date, time, location (latitude and longitude), fuel level, and vehicle ID. We generate trips for each zone ($z \in \Gamma$) and time bucket ($t \in \mathcal{T}$) using a Poisson distribution with arrival rate ($\lambda_{t,z}$). After generating a request with origin u_i at time t_i , we specify its destination (d_i) employing a multi-nominal distribution with $d_i \sim Mult_{|\Gamma|}(1, \{\Omega_1, \dots, \Omega_{|\Gamma|}\} | t_i \in \mathcal{T}, u_i \in \Gamma)$, considering the possible destinations as outcome categories (Ω_z). A spatial description of trip origins is illustrated on the left side of Figure 2.5 as a heat map (the darker color, the higher demands); trips mostly start from central areas. Also, we visualize an average of rentals for each

hour of the day on the right side of Figure 2.5. From early working hours (7 a.m), the number of rentals piles up and reaches a peak in the evening and drops off over night hours.




Figure 2.5: Spatial Demand Pattern and Distribution of CSs (left), and Hourly Demand Pattern (right)

To set up the charging infrastructure, we consider similar features as the results of Ahadi et al. (2021). The locations and capacities of CSs are presented in the left plot of Figure 2.5, where numbers denote CSs id and capacity, respectively. There are two types of chargers: standard (11 kW) and fast (50 kW), the current charging technologies of public CSs in Berlin (Open charge map 2020). To characterize the fleet, we use fewer SAEVs (e.g., 150, 200, and 300) than conventional shared cars due to their high utilization. Vehicles are homogeneous similar to Tesla Model 3 (Fuel economy guide 2020), which are initialized across the service area with an arbitrary *SoC* between 70% and 80%. We assume that each vehicle has an approximate average speed of 20 km/hr and travel costs of \$0.53/km (Bösch et al. 2018). A summary of parameters is provided in Table 2.1.

Parameter	Value
Episode length	5 days
Number of zones/hexagons	89
Number of CSs (fast CSs)	16(5)
Charging rate (fast charging rate)	11(55) kW
SAEV battery capacity (energy consumption)	50 kWh (0.15 kWh/km)
Vehicle speed	20 km/hour
Trip request coverage area	10 km
Travel cost	\$0.53/km

Table 2.1: Instance Parameters

2.5.3 Experimental Results

We train our agents with the above-mentioned initialization. After comparing the policies learned by four proposed algorithms, we assess our superior policy against the benchmark models (fleet performance and charging behavior) and the upper bound (only fleet performance).

We measure the performance by two key indicators: (1) total fleet profit, trip revenues minus operational costs (noted, we do not include the missed trips penalties as they are only virtual costs), and (2) service quality, the ratio of accepted trips. For analyzing charging behaviors, we consider three indicators: (1) a mean hourly utilization of all CSs to track the charging schedules, (2) a mean utilization of each individual CS over the whole operations period to observe charging allocations, and (3) an hourly *SoC* of vehicles to track the energy level.

Evaluating the Multi-Agent Charging Algorithms




Figure 2.6: Fleet Profitability (left) and Service Quality (right) for the Proposed MARL Algorithms

We compare four modified MARL models (Hyperparameters of DQN are given in Online Appendix B.1). These settings are as follows: (1) full-MARL, a single controller for each vehicle considering full options of CSs, (2) full-HMARL, hierarchical controllers with full CSs, (3) limited-MARL, a single controller considering limited options of CSs, and (4) limited-HMARL, hierarchical controllers with limited CSs. By limited CSs, we mean a list entailing the closest, the closest free, and the closest fast CS to the vehicle. We plot the fleet profitability and service quality in Figure 2.6. All agents except full-MARL converge relatively to the same numbers for both indicators. The poor performance of full-MARL is likely an artifact of considering all CSs in a single action space, which makes a biased exploration and disables identified reward-shaping for scheduling and allocating decisions. Among three other algorithms, limited-HMARL converges faster due to a hierarchical setting (distinguished rewards) and less need for exploration. We use limited-MARL for further analysis, and we call it multi-agent dynamic charging (MADC).

To ascertain that deep learning does not compromise the results, we retrain the agents with a tabular solution, Q-learning (see Online Appendix B.5 for details). A comparison in Figure 2.7 ensures that both algorithms converge to the same objective values, while it takes longer for Q-learning due to the need for experiencing all probable states and actions.




Figure 2.7: Evaluating the Superior Policies against the Tabular Approach, Benchmark Models, and Upper Bound

Assessing against the Upper Bound and Benchmark Models

We assess our superior model with an upper bound (details in Online Appendix B.4) and two benchmark models. The results of the upper bound scenario reveals that using non-EVs raises profitability and service quality by approximately 6% and 2% (to \$703,591 and 96%), respectively (see Figure 2.7). In other words, even with non-EVs, there are still a few unserved trips caused by a lack of supply or very low waiting thresholds of travelers. Also, the rise in profitability is less than the increase in service quality, caused by a lower energy efficiency of non-EVs and higher petrol prices.




Figure 2.8: Cost Components for the Proposed and Benchmark Models

A comparison with SADC explains that using a central dynamic operational decision-making could augment the fleet performance slightly. Although SADC achieves less than 1% higher profitability and service quality (to \$686,832 and %91, see Figure 2.7), it decreases the scalability due to a giant joint action range. Precisely, SADC convergence time is approximately 12 times (see Figure 7) of MADC and will increase for larger fleet sizes. Also, with the same computational resources, the size of the joint action space is restricted roughly to a batch size of 20 vehicles,

even if we limit the charging destination options to only two alternatives. This means that for large fleet sizes (e.g., 500), we cannot assume that a joint action for only 20 agents can cover the whole idle vehicles, yielding poor performances. However, compared to REOPT, MADC increases service quality and profits roughly by 5% and 7% (84% (\$628,879) for REOPT and 90% (\$680,635) for MADC). This indicates that learned dynamic policies reduce energy costs in addition to serving more requests.

We also illustrate the proposed and benchmark models' waiting, driving, and charging costs in Figure 2.8. As expected, SADC and MADC have very similar cost values; the only difference is that the waiting costs for SADC are slightly higher. However, comparing the static and dynamic models, the driving costs are lower using REOPT than MADC. This means that the number of charging decisions is lower, and REOPT allocates vehicles to the closer CSs. Interestingly, the charging costs of MADC are lower, highlighting that the dynamic model takes the advantage of time-varying electricity prices, even though the number of served trips, and consequently, consumed energy are higher for MADC (higher charging costs are expected with flat electricity tariffs).




Figure 2.9: Individual CS Utilization for the Proposed and Benchmark Models

Regarding charging behaviors, Figure 2.9 illustrates that vehicles take similar allocating (destination) decisions following either of these models. There is a gap between fast CSs and standard CSs utilization, meaning all strategies allocate more vehicles to fast CSs (the gap is smaller for REOPT). However, as Figure 2.10 displays, the average hourly utilization of CSs follows unlike patterns for dynamic (MADC and SADC) and static (REOPT) models. With MADC, CSs have higher utilization during the night and early morning (off-peak hours); after a reduction by noon, it continues to rise until midnight (vehicles have low *SoC* and need to recharge in anticipation of future demands). SADC acts similar while the utilization during the night is a bit higher. Using REOPT, CSs are not that utilized during off-peak hours, and it levels off after 8 a.m until midnight.

Figure 2.11 compares the use of fast and standard CSs at different times. Almost in all hours, fast chargers are more utilized than standard ones following both learned and reoptimization




Figure 2.10: Average Hourly Utilization of CSs for the Proposed and Benchmark Models

policies, while the learned policies take more advantage of fast chargers. Comparing the learned policies, SADC prefers fast charging even during the night more than MADC. Regarding standard CSs, all models use standard chargers considerably during off-peak hours. However, when the mobility demand is high, the learned policies rarely allocate vehicles to standard CSs, unlike REOPT.




Figure 2.11: Hourly Utilization of CSs (distinguished by charging rate) for the Proposed and Benchmark Models

We also show how the smart charging models alter vehicles' average *SoC*. With MADC, vehicles have a lower battery level when they start the day; it reaches a peak of approximately 85% by early morning and drops gradually to around 30% until midnight. SADC acts very similar while having a higher peak (roughly 90%), caused by more utilization of fast charging during the off-peak hours. Employing REOPT, the peak is limited to almost 70%, and after a drop, the average *SoC* fluctuates around 50%. These results demonstrate that with dynamic charging, vehicles charge less during the day when demands are high and mainly recharge their batteries during off-peak hours.




Figure 2.12: Average SoC of Vehicles for the Proposed and Benchmark Models

Sensitivity Analysis

We analyze the proposed model’s responsiveness to strategic decisions. Here, we only present a summary and refer to Online Appendix C.2 for details and visualizations. We exclude SADC due to a very similar behavior to MADC, and compare the results with REOPT for only charging infrastructure scenarios. Numerical results are provided in Table 2.2, summarizing the objective function indicators and charging behavior (SoC of vehicles, fast charging occupancy (FCO), and standard charging occupancy (SCO)), distinguished for various time-buckets regarding the mobility demands (low/medium/high).

Scenarios	Profit USD	Service Quality	SoC(low/medium/high)	FCO(low/medium/high)	SCO(low/medium/high)
Full Charging Capacity	680,635	0.90	0.63/0.61/0.45	0.38/0.51/0.71	0.33/0.03/0.01
No Fast Charging	552,180	0.76	0.42/0.47/0.32	0.74/0.48/0.63	0.74/0.48/0.63
Half Charging Capacity	641,752	0.85	0.49/0.51/0.35	0.74/0.72/0.92	0.70/0.22/0.15
150 Vehicles	522,878	0.71	0.57/0.46/0.35	0.44/0.51/0.71	0.10/0.02/0.01
200 Vehicles	680,635	0.90	0.63/0.61/0.45	0.38/0.51/0.71	0.33/0.03/0.01
250 Vehicles	772,029	0.98	0.74/0.68/0.62	0.38/0.61/0.63	0.22/0.13/0.07
300 Vehicles	788,104	0.99	0.73/0.69/0.65	0.27/0.41/0.50	0.28/0.32/0.35
50 kWh Battery	680,635	0.90	0.63/0.61/0.45	0.38/0.51/0.71	0.33/0.03/0.01
75 kWh Battery	703,841	0.92	0.56/0.63/0.47	0.64/0.47/0.40	0.25/0.06/0.01
100 kWh Battery	723,648	0.94	0.55/0.60/0.49	0.72/0.41/0.35	0.29/0.07/0.01
Electricity Tariff 1	772,029	0.98	0.63/0.61/0.45	0.38/0.61/0.63	0.22/0.13/0.07
Electricity Tariff 2	759,043	0.98	0.68/0.62/0.48	0.22/0.53/0.57	0.27/0.20/0.11
Electricity Tariff 3	760,591	0.98	0.70/0.62/0.58	0.30/0.57/0.58	0.23/0.17/0.08

Table 2.2: Numerical Results for Different Strategic Scenarios using MADC

First, we examine the charging infrastructure impact. Without fast charging, the service quality (profits) drops significantly to 76% (\$552,180) and 66% (\$476,274) using MADC and REOPT, respectively. The utilization gap between fast and standard CSs decreases, and the average charging occupancy increases due to long charging sessions. Vehicles’ *SoC* distribution

alters too; the peak drops and occurs later. However, a reduction in CSs capacity (half-sized) marginally affects the fleet performance and charging behavior. The service quality (profits) drops to 85% (\$641,752) and 78% (\$594,141) for MADC and REOPT, respectively.

Concerning fleet characteristics, larger fleet sizes of 150, 200, 250, and 300 SAEVs boost the service quality non-linearly to 71%, 90%, 98%, and more than 99%. This rise flattens the vehicles' *SoC* as they have lower utilization and higher flexibility to charge. Also, with more vehicles, fast CSs gain less attraction and the hourly CSs utilization variance reduces as well. Regarding battery capacity, different sizes of 50, 75, and 100 (kWh) lead to service quality of 90%, 92.5%, and 94%, while the charging behaviors are similar.

The final analysis is on electricity tariffs. We consider a flat tariff and two ToU tariffs (see Online Appendix B.2). Different tariffs do not change the fleet performance significantly but change charging strategies slightly. With a flat tariff (Tariff 1), CSs utilization only follows the mobility demand (high utilization in off-peak hours and low utilization in peak hours). Using a ToU tariff with relatively high prices in peak hours (Tariff 2) does not change the utilization significantly, while super-peak prices (Tariff 3) shift some charging decisions to low-price hours.




Figure 2.13: Cost Components for Considered Sensitivity Analysis Scenarios using MADC

Figure 2.13 shows the cost components of sensitivity analysis scenarios. Regarding charging infrastructure features, without fast charging, charging costs reduce due to the lower served demands (lower energy consumption), and waiting costs increase due to the lower charging power. These changes are similar for the half CS capacity scenario, while the growth in waiting costs is more visible. Driving cost decreases in the no fast CS scenario as vehicles do not drive longer for fast charging, but it increases in the half CS capacity scenario since vehicles select remote areas or fast CSs to avoid long queues. Regarding fleet configurations, increasing the

number of vehicles raises charging costs (more served demands) and decreases driving costs since vehicles charge less frequently. Higher battery capacities lead to lower driving costs and higher waiting costs as charging events occur less frequently but are more time-consuming. Finally, different electricity tariffs do not change the costs components significantly, except for marginal increases in charging costs for Tariffs 2 and 3 due to relatively higher electricity prices. We consider 250 SAEVs for this comparison (see Online Appendix C.5).

2.6 Discussion and Future Work

We explore the dynamic smart charging management of SAEVs. Two major sub-problems are determining proper charging time and selecting the best charging destination, which are coupled and should be solved simultaneously. We formulate this stochastic dynamic decision-making problem as a semi-MDP to cope with the temporal abstraction of actions and adopt a deep RL approach to extract optimal charging policies. For scalability, we propose a decentralized cooperative model in which vehicle agents make individual charging decisions while maximizing fleet service quality and profitability. Further, to break the tendency of selfish policies, similar to Shou and Di (2020), Oroojlooyjadid et al. (2022) we use reward-shaping techniques to penalize the responsible agents for unserved demands. To robust the model against the changes in the mobility environment, we adjust the state space by including an approximation of neighboring agents' observations. The other methodological contribution is identifying the decision levels (high level: charging and low level: allocating to CSs) using hierarchical learning. It helps define assorted reward functions for each level and augments the scalability by reducing the action space size.

To evaluate the proposed models, we simulate the mobility platform using an advanced ABM where the fleet operator, vehicles, and CSs are represented as interactive agents. We calibrate the trip characteristics using historical rental data from a worldwide leading carsharing company (ShareNow in Berlin, Germany). The key findings are as follows. Hierarchical learning revamps the performance of the MARL model, specifically when the action space is enormous. Also, the results of a tabular algorithm (Q-learning) show that using deep learning does not diverge from the optimal policies. Moreover, we assess our proposed method (MADC) with a centralized RL agent and a reoptimization strategy. Similar fleet performance and charging policies using SADC and MADC ensure that our decentralized and hierarchical decision-making assumptions of MADC do not yield sub-optimal policies. Note that SADC suffers from the curse of dimensionality and is not applicable for realistic problem sizes. The second benchmark results declare that learned dynamic policies have superior performance. Due to the lack of global optimum, we compare the learned policies with an upper-bound scenario (a fleet of non-EVs). It declares that using our proposed model, fleet operators can tackle the technological challenges of EVs and achieve comparable performance with non-EV fleets. Noted, our results

do not guarantee that using petrol cars is more profitable in real-world cases as we disregard the investment and maintenance costs.

To provide more managerial implications, we expound that fleet operators can adopt our model under existing restricted CSs to build a profitable advantage by optimizing dynamic charging policies even without the need for fleet expansion or massive charging facilities. Furthermore, we explore the impacts of exogenous features. Although MADC accounts for charging rate, fast charging technology dramatically impacts the fleet performance and charging behavior. Fast CSs are preferred destinations if standard chargers might not meet future demands. Reducing the capacities of CSs does not change the charging patterns significantly. However, it causes a marginal drop in the fleet's service quality, which is in line with the results of Abouee-Mehrizi et al. (2021).

Regarding fleet characteristics, increasing the number of vehicles has a non-linear positive effect on fleet performance. It changes the optimal charging policies: the more vehicles, the less need for fast charging. Adding vehicles could not be economically beneficial, but might reduce the need for fast CSs. Enlarging battery capacities also positively impacts fleet performance but is not as influential as the fleet size. There is still a high demand for fast charging, even with big batteries. Finally, we demonstrate that the electricity tariff could be more important for grid operators to design appropriate tariffs and encourage desired load profiles. From the fleet's perspective, since serving the trips is crucial, the marginal price differences in ToU tariffs would not be encouraging to shape charging loads considerably and learned charging policies follow mostly the mobility demand patterns. Finally, we show that limiting charging destinations to more probable options could converge faster to similar policies. From a fleet operator's perspective, sending vehicles to remote CSs without an additional value is nonsense. However, it could vary if we regard other perspectives such as a distribution system operator that aims to shift charging loads to low-demand areas.

Our work is not free of limitations. Regarding the proposed model generalizability, we employ function approximation and test different fleet characteristics and charging infrastructure scenarios. However, it still needs to be examined for other shared fleets and geographical areas. The problem also could cover more aspects. We focus on the smart charging, which is not optimally integrated with other operational decisions. Moreover, we assume a privately-managed charging infrastructure, which is not always the case, as fleets might use public/shared CSs. These limitations present opportunities for future works. To progress generalizability, optimal policies can be smoothed by transfer learning approaches. Further, to approach more realistic assumptions, one could consider shared CSs and analyze its effects on the optimal charging policies. Finally, to integrate all operational decisions, agents could use a joint decision system where they start from high decision levels (e.g., charging, relocating) and hierarchically make the lower level decisions.

2.7 Appendix

2.7.1 Agent-based Simulation

In this appendix, we explain the details of our multi-agent simulation for shared autonomous electric vehicles (SAEVs). Each module of the simulation is described in the following.

Trip Module

Figure 2.14 shows the process of serving a trip. When a vehicle is assigned to a trip, it starts relocating to the trip's origin. Meanwhile, the traveler might cancel her request due to long waiting time or urgent events. Thus, we consider the probability of trip cancellation, which is a function of the driving distance between the trip's origin and the assigned vehicle. Once a traveler cancels her request, the assigned vehicle runs the parking module (explained later on). Otherwise, the vehicle picks up the traveler, drives to the destination, and drops it off (we assume no cancellation during serving time). After serving the trip, the vehicle updates its state and runs the operations module to take the best operational action.




Figure 2.14: Trip Module of SAEVs: the Process of Serving a Trip after Matching to a Vehicle

Operations Module

The operations module comprises three sub-modules: charging, repositioning, and parking; executed sequentially (see Figure 2.15). First, the vehicle checks whether it should charge or not. If the decision is to charge, the vehicle must also choose a charging destination, relocate to the selected CS, wait in a queue (if there is any), and start charging. Instead of having a first-in-first-serve strategy, the CS prioritizes vehicles based on their *SoC* (the lowest *SoC*, the highest priority). We also assume that charging and waiting vehicles with enough *SoC* can be assigned to requests. We avoid interrupting charging/waiting processes if idle vehicles are available to serve all open requests, which is considered in our matching strategy (see 2.7.1). If the vehicle does not interrupt charging, it finishes the task and follows the repositioning module. Also, when the vehicle decides not to charge, it checks the repositioning conditions.

The repositioning module distributes available vehicles according to supply and demand. We use a simple decentralized heuristic approach to reposition vehicles. Each vehicle first compares

model assigning vehicles to trips. We define the sets, parameters, and decisions to do so. The set of available vehicles and the set of open requests are indicated by \mathcal{J}_S and \mathcal{I}_O , respectively. The term $\beta^{Energy} D_{i,j}$ represents energy that vehicle j needs to serve request i , $C_{i,j}$ is the cost of assigning vehicle j to trip i (a function of distance $D_{i,j}$ and the vehicle's mode), $D_{i,j}$ is the distance between vehicle j and the origin of trip i , P_i is the price of serving request i , SoC_j is the state of charge of vehicle j , δ is safety energy after serving requests. Finally, the decision variable ($x_{i,j}$) is a binary variable set to one if vehicle j is assigned to request i .

$$Max_x \sum_{j \in \mathcal{J}_S} \sum_{i \in \mathcal{I}_O} x_{i,j} (P_i - C_{i,j}). \quad (2.13)$$

The objective is to match vehicles with trips while maximizing total profits. Each trip's price varies according to its length and duration in addition to a base fair. Assigning a vehicle to a trip entails different costs, depending on the vehicle mode (e.g., charging, parked) and the distance to the trip. A higher cost is designated to charging vehicles to allocate them with trips only when there is a lack of supply and to reduce charging interruptions. The assignment costs are set to be less than trip profits to only distinguish vehicles according to their distances with trips and avoid affecting the acceptance rate (i.e., not rejecting trips if there is enough supply).

$$\sum_{j \in \mathcal{J}_S} x_{i,j} \leq 1, \quad \forall i \in \mathcal{I}_O, \quad (2.14)$$

$$\sum_{i \in \mathcal{I}_O} x_{i,j} \leq 1, \quad \forall j \in \mathcal{J}_S, \quad (2.15)$$

$$x_{i,j} (\beta^{Energy} D_{i,j} - SoC_j \kappa_j^{Battery} - \delta) \leq 0, \quad \forall j \in \mathcal{J}_S, i \in \mathcal{I}_O, \quad (2.16)$$

$$x_{i,j} D_{i,j} \leq \Delta, \quad \forall j \in \mathcal{J}_S, i \in \mathcal{I}_O, \quad (2.17)$$

$$x_{i,j} \in \{0, 1\}, \quad \forall j \in \mathcal{J}_S, i \in \mathcal{I}_O. \quad (2.18)$$

As we match vehicles and open requests simultaneously, the model must ensure to assign each vehicle to at most one trip and each trip to at most one vehicle; guaranteed by Constraints (2.14) and (2.15), respectively. Constraint (2.16) ensures that a vehicle can be assigned to a trip only if it has enough energy to pick up the traveler, take her to the destination, and reach a CS. As a simplification, we consider the same safety threshold throughout the business area, assuming that CSs are evenly distributed. Moreover, to avoid prolonged waiting time after assignments, we add Constraint (2.17) to only allow matching when driving distances to trips'

origin is not too large. Although we aim to serve all trips in the earliest time window, sometimes demands exceed supplies, and unassigned trips shift to the next time window.

Revenue and Costs

To calculate trip revenues, we use a base fare (2 USD) plus a variable fare, which is a linear function of trip distance and duration. Other factors, such as the time of day and traffic congestion, are excluded for simplicity. Each trip also has a travel cost that must be taken into account. We show the profit calculation for a trip in Equation (2.19), where ω is the revenue per distance (1.3 USD/km), ϕ is the revenue per duration (0.35 USD/min), and ψ is travel cost per distance (0.53 USD/km). Note that the minimum charge for each trip is 5 USD (Taxi costs in Berlin 2022).

$$\text{TripProfit} := \min(\text{BaseFare} + \omega * \text{Distance} + \phi * \text{Duration}, \text{MinimumFare}) - \psi * \text{Distance}. \quad (2.19)$$

Regarding other costs, we calculate the charging cost based on an hourly energy fee:

$$\text{ChargingCost} = \sum_{t \in \mathcal{T}_{\text{charging}}} \text{EnergyCharged}_t * \text{EnergyPrice}_t, \quad (2.20)$$

Where $\mathcal{T}_{\text{charging}}$ is the set of periods (minutes) during a charging session. It means that charging costs are independent of charging rates, which are aligned with current electricity tariffs for large-scale EV fleets (Lee et al. 2019). Moreover, the driving cost is a linear function of the distance ($\text{DrivingCost} = \psi * \text{Distance}$), and the waiting cost is also based on a fixed ratio per time ($\text{WaitingCost} = \Omega * \text{WaitingTime}$), where Ω is the parking cost per hour (3 USD/hr).

Linearizing Queue Management of the Reoptimization Benchmark Policy

To linearize the waiting constraint of the reoptimization benchmark, we use a big M trick. Therefore, we can replace the nonlinear constraint as follows.

$$t_{j,c}^{\text{Queue}} \geq \max\left\{\frac{N_{j,c} + \sum_{j' \in \mathcal{J}_C} y_{j',c} (1 - z_{j,j'}) - \kappa_c^{\text{Charging}}}{\kappa_c^{\text{Charging}}} E / \kappa_c^{\text{Power}}, 0\right\}, \forall j \in \mathcal{J}_C, c \in \mathcal{C} \quad (2.21)$$

$$t_{j,c}^{\text{Queue}} \leq M^{\text{Queue}} y_{j,c}, \forall j \in \mathcal{J}_C, c \in \mathcal{C} \quad (2.22)$$

This pair of constraints guarantees that the waiting time of vehicle j in charging station c only could get a non-zero value if vehicle j is allocated to CS c (i.e., $y_{j,c} = 1$).

2.7.2 Configuration Data

In this subsection, we quantify the mobility environment parameters in addition to the charging agents' hyperparameters.

Hyperparameters

Here, we specify the hyperparameters of smart charging agents and additional details regarding the training and implementation of MARL and HMARL agents. Hyperparameters, chosen based on different experiments, are given in Table 2.3. The HMARL agent has two separate networks (high-level and low-level), and the MARL agent has one network with the same structure and hyperparameters as the high-level network of the HMARL agent.

Parameter	High-level	Low-level
Optimizer (learning rate)	Adam(0.001)	Adam(0.001)
Loss function	MSE	MSE
Discount factor γ	0.99	0.99
Memory capacity	1,000,000	1,000,000
Steps prior to learning	1000	500
Training frequency	50	20
Batch size	32	32
Initial ϵ	0.6	0.6
Final ϵ	0.01	0.05
Target network update frequency	1000	1000
DQN activation functions	ReLU	ReLU
Number of hidden layers (nodes)	2 (256, 512)	2 (256, 512)

Table 2.3: Agents’ Hyperparameters

Energy Price

We use an actual EV-focused time-of-use (ToU) electricity tariff (Tariff 2 in Table 2.4), that is used by Southern California Edison (SCE) and in recent related works (e.g., Lee et al. 2019). In addition, we test two other tariffs in this paper. The first is a flat tariff (Tariff 1), where electricity prices are the same for all hours of the day. The second one (Tariff 3) is similar to Tariff 2 with a larger price variation, where the price in super on-Peak hours is costly.

	Tariff 1	Tariff 2	Tariff 3
Super Off-Peak (9pm-6am)	0.23 USD/kWh	0.08 USD/kWh	0.08 USD/kWh
Off-Peak (6am-4pm)	0.23 USD/kWh	0.08 USD/kWh	0.23 USD/kWh
On-Peak (4pm to 6pm)	0.23 USD/kWh	0.23 USD/kWh	0.23 USD/kWh
Super on-Peak (6pm to 9pm)	0.23 USD/kWh	0.23 USD/kWh	0.50 USD/kWh

Table 2.4: Time-of-use Tariff for Large-scale EV Charging Customers

	1	2	3	4	5	6	7	8	9	10	11	12
SoC threshold % (morning)	0.45	0.60	0.65	0.62	0.58	0.55	0.52	0.50	0.40	0.40	0.40	0.40
SoC threshold % (afternoon)	0.38	0.35	0.32	0.25	0.25	0.20	0.20	0.25	0.27	0.35	0.35	0.40

Table 2.5: Hourly Charging Threshold for the Reoptimization Benchmark Policy

SoC Thresholds of the Reoptimization Benchmark Model

To make charging decisions using REOPT, we determine vehicles needing to charge using an hourly *SoC* threshold. Whenever a vehicle’s energy level falls below this threshold, it is marked as a charging demand. We consider an hourly threshold with the opposite pattern of mobility demands; i.e., when the demand is low (e.g., during the night), the threshold is high, meaning that more vehicles will charge. The values are provided in Table 2.5.

Details of the Upper Bound Scenario

For the upper bound scenario, we use the Audi 3 with a fuel efficiency of 7.29 liter/km Fuel economy guide (2020), and consider the average price of gasoline in Germany in 2021 (\$1.8) Global petrol prices (2021) to compute energy costs (i.e., the driving cost of petrol vehicles equals \$0.13/km).

Modification for Tabular Solutions

We compare our proposed model (MADC) with a tabular version (Q-learning) to ensure that deep learning does not compromise the learned policies. Therefore, we discretize the state space; hexagons and hours represent locations and time, respectively, *SoC* is divided into ten levels (10%, 20%, ..., 100%), and local supply and demand are categorized into three levels (low, medium, and high). Low/medium/high supply means less than 3/between 3 and 10/more than 10 available SAEVs within the coverage area of the vehicle (10 km radius). Low/medium/high demand means that there are less than 5/between 5 and 10/more than 10 open requests within the coverage area of the vehicle (10 km radius). We also identify the CSs with at least one free spot using binary variables and restrict the action space to only limited charging destinations.

2.7.3 Results and Sensitivity Analysis of Strategic Factors

This subsection provides detailed results of the scenario analyses using the proposed models.

The Impact of Fast Charging

We resolve SAEV-CC assuming there is "no" fast charging technology. The results show that without fast charging, the fleet performance decreases by 14% and 24% using MADC and REOPT, respectively. We display vehicles’ *SoC* in Figure 2.16. The disparity between MADC and REOPT is wider when there is "no" fast charger, which can be seen from the gap between the vehicles’ *SoC*. Comparing the results with the fast charging scenario, vehicle batteries have lower energy since they need more time to charge, during which they might interrupt the charging process to meet urgent mobility demands.

Regarding the charging strategy, Figure 2.17 shows that there is no longer a significant gap between the fast CSs and standard CSs occupancy (CSs locations cause the difference). Another change occurs in the hourly utilization of CSs, which is visualized in Figure 2.18. As can be seen, although both strategies follow the same pattern, the average utilization is higher compared to the fast charging scenario due to longer charging sessions.




Figure 2.16: No Fast Charging Scenario: Average SoC of Vehicles for the Proposed and Benchmark Models




Figure 2.17: No Fast Charging Scenario: Individual CS Utilization for the Proposed and Benchmark Models




Figure 2.18: No Fast Charging Scenario: Hourly CSs Utilization for the Proposed and Benchmark Models

The Impact of Charging Infrastructure Capacity

Another influential charging-related factor is the size of each CS. The more charging capacity, the higher fleet performance since vehicles have more resources to charge their batteries. Results show that with half of the charging capacity, service quality reduces by 5% and 12% for MADC and REOPT, respectively, while vehicle *SoC* patterns do not change significantly (see Figure 2.19). The shape of SoC curves is the same. However, the peaks drop and occur later compared to the full charging capacity scenario. We also plot the utilization of CSs in Figure 2.20, where still fast chargers are more attractive for both charging strategies. Figure 2.21 depicts a very similar occupancy pattern for both strategies while the values are almost doubled.




Figure 2.19: Half Charging Capacity Scenario: Average SoC of Vehicles for the Proposed and Benchmark Models




Figure 2.20: Half Charging Capacity Scenario: Individual CS Utilization for the Proposed and Benchmark Models

The Impact of Fleet Size

One critical factor in shared autonomous fleet performance is the number of SAEVs. It gains more importance when the fleet is electric, and vehicles need more time to refuel their batteries. We test our proposed charging strategy (from now on, we will only analyze results using MADC and exclude REOPT) for four different fleet sizes (150, 200, 250, and 300 SAEVs). By increasing the number of SAEVs, the service quality increases, but not linearly (71%, 90%, 98%, and more




Figure 2.21: Half Charging Capacity Scenario: Hourly CSs Utilization for the Proposed and Benchmark Models

than 99%, respectively, for the mentioned fleet sizes). Finding the optimal fleet size is beyond the scope of this paper; we merely aim to show its impacts on performance and charging policies. Therefore, if the goal is maximizing fleet performance, it makes no sense to facilitate more than 250 SAEVs to cover demands for this case. In more detail, vehicles' *SoC* is illustrated in Figure 2.22 and what is clear is that by increasing the number of vehicles, the *SoC* distribution gets flatter (i.e., there is no such a gap between the vehicles' *SoC* in the early morning and evening when there are 300 SAEVs), which is caused by the low utilization of vehicles.




Figure 2.22: Different Fleet Sizes Scenario: Average SoC of Vehicles for Four using the Proposed Model

More importantly, different fleet sizes alter the learned charging strategies. Figure 2.23 demonstrates that the smaller fleet sizes, the more fast charging demands. Indeed, when the number of vehicles is low, MADC takes the fast CSs as the best destination to ensure enough *SoC* for providing high service quality. On the other hand, the utilization distribution gets flat when there are many SAEVs (i.e., there is no need for urgent charging sessions). Regarding the hourly utilization, a rise in the number of vehicles creates a peak during midday hours

(Figure 2.24). This needs to be interpreted cautiously as many factors might be influential, like electricity price, mobility demand, and charging capacity.




Figure 2.23: Different Fleet Sizes Scenario: Individual CS Utilization using the Proposed Model




Figure 2.24: Different Fleet Sizes Scenario: Hourly Utilization of CSs using the Proposed Model

The Impact of Battery Capacity

Another impactful fleet configuration is the vehicles' battery capacity. We test MADC for three different battery sizes (50 kWh, 75 kWh, and 100 kWh), yielding service levels of 90%, 92.5%, and 94%, respectively. It shows that a bigger battery size could increase fleet performance and is worth considering (further analysis and the cross effects with fleet size could also be interesting,

which is beyond the scope of this study). Figure 2.25 shows that the *SoC* of vehicles with smaller batteries reaches a higher peak occurring earlier than the case of larger batteries (due to the shorter charging time). Figure 2.26 displays that fast CSs are still more popular destinations for different battery sizes. Thus, larger batteries do not remove the need for fast chargers in our case. Hourly patterns are also similar. However, increasing the battery capacity decreases the reduction rate (from the peak during night hours until the valley during day hours) as it takes longer to charge the batteries (see Figure 2.27).




Figure 2.25: Different Battery Capacity Scenario: Average SoC of Vehicles using the Proposed Model




Figure 2.26: Different Battery Capacity Scenario: Individual CS Utilization using the Proposed Model




Figure 2.27: Different Battery Capacity Scenario: Hourly Utilization of CSs using the Proposed Model

The Impact of Electricity Tariff

We check three different tariffs to understand how electricity prices affect charging behaviors (See Table 2.4). We consider 250 vehicles for this comparison to avoid the bias of supply scarcity (i.e., with lower supply than demand, the charging behavior only follows the mobility demand and disregards the electricity price). Figure 2.28 illustrates that vehicles increase their *SoC* during the night using the flat tariff (Tariff 1), and that it will drop relatively until the end of the peak hours. On the other hand, vehicles' *SoC* for two other tariffs (peak and off-peak prices) have periodic levels. Vehicles recharge mainly during the night (due to the low electricity price and mobility demand) and exactly before high price hours. This is more visible for Tariff 3, which has higher prices during super-peak hours (18-21).




Figure 2.28: Different Electricity Tariff Scenario: Average SoC of Vehicles using the Proposed Model




Figure 2.29: Different Electricity Tariff Scenario: Hourly Utilization of CSs using the Proposed Model

Symbol	Description	Unit
Sets & Spaces		
\mathcal{T}	Set of operations time with index t	set
Γ	Set of hexagonal zones within the service region with index z	set
\mathcal{H}	Set of tasks while executing an operational decision with index h	set
\mathcal{C}	Set of charging stations with index c	set
\mathcal{I}	Set of all trip requests over the operations horizon with index i	set
\mathcal{I}_O	Set of open trip requests at the current time with index i	set
\mathcal{J}	Set of all vehicles (vehicle agents) with index j	set
\mathcal{J}_S	Set of available vehicles for serving trip requests with index j	set
\mathcal{J}_C	Set of vehicles with the need for charge with index j	set
\mathcal{N}_j	Set of neighboring vehicle agents to vehicle j	set
\mathcal{S}	The state space of the whole system with index s	space
\mathcal{O}_j	The observation space of agent j with index o_j	space
\mathcal{A}	The joint action space of all agents with index a	space
\mathcal{A}_j	The action space of agent j with index a_j	space
Parameters		
α	Learning rate of Q-learning	ratio
γ	Reward discount rate	ratio
ϵ	Probability of taking a random action using an ϵ_{greedy} policy	ratio
θ	Parameter of action-value function approximations (neural network)	float
δ	Minimum safety energy threshold for reaching the closest CS	kWh
δ^{idle}	The maximum waiting time for vehicles before checking their status	minutes
Δ	The maximum coverage area of trip request or SAEVs	km
E	Average charging demand of vehicles	kWh
$\rho^{Charging}$	Penalty for an unassigned charging vehicle to CSs	USD
β^{Energy}	Energy consumption per driving distance	kWh/km
β^{Time}	Driving time per driving distance	minute/km
SoC_j	State of charge of vehicle j	%
N_c	Number of present vehicles in CS c	unit
$N_{j,c}$	Number of vehicles in CS c with lower SoC than vehicle j	unit
$\kappa_c^{Power}/\kappa_c^{Parking}$	Maximum charging power/number of parking spots of CS c	kW/unit
$\kappa_c^{Charging}$	Maximum number of charging docks in CS c	unit
$\kappa_j^{Battery}$	Maximum battery capacity of vehicle j	kWh
$D_j, i/Dj, c$	Driving distance between vehicle j and request i / CS c	km
$w_s/w_c/w_d/w_w/w_p$	Weight for reward of serving/charging/driving/waiting/penalty of a vehicle	float
$u_i/d_i/p_i/b_i$	Information of trip request i (origin/destination/price/patience-time)	mixed
$Z_{j,j'}$	A binary indicator if vehicle j has lower SoC of vehicle j'	boolean
Variables		
$SoC/m/l$	Energy/mode/location of the corresponding vehicle	%/integer/integer
s	General state of the system regarding the corresponding vehicle	mixed
$s_t/s_v/s_l/s_c$	State of time/vehicle/local supply & demand/CSs	mixed
n_s/n_d	Number of local supply/demand around the corresponding vehicle	integer
q_c	Number of waiting vehicles in CS c	integer
c_c	Decision variable of charging the corresponding vehicle in CS c	boolean
$r_s/r_c/r_d/r_w/r_p$	Reward of serving/charging/driving/waiting/penalty of a vehicle	USD
r_{profit}/r_{missed}	Reward of profits and missed trips penalty of the fleet	USD
τ	Termination duration of an operational action	period
τ_{eh}	Period between triggering and starting task h	period
τ_{oh}	Period between starting and finishing task h	period
a	General operational action (charging and allocating to CS) of vehicles	integer
a_{meta}	Action of the meta controller (charging or not)	integer
a_{sub}	Action of the sub controller (which CS)	integer
$t^{Driving}$	Driving time between CS c and vehicle j	period
$t_{j,c}^{Charging}$	Charging time of vehicle j at CS c	period
$t_{j,c}^{Queue}$	Waiting time of vehicle j at CS c	period
$y_{j,c}$	Indicator whether vehicle j is assigned to CS c	boolean
$x_{i,j}$	Indicator whether vehicle j is assigned to trip request i	boolean

Table 2.6: A Summary of Notations

Chapter 3

Adoption of Autonomous Vehicles in Ride-Hailing Services: The Role of User Preferences¹

3.1 Introduction

Socio-economic and technological advancements in transportation systems play a pivotal role in shaping the development of smart cities. With the rise of the sharing economy, ride-hailing companies such as Uber and Lyft are transforming mobility-on-demand services worldwide by connecting travelers with drivers through digital platforms (Wei et al. 2022). Despite growing demand, these platforms face challenges of profitability and supply uncertainty, largely due to their reliance on human drivers (Heineke et al. 2021). To address these issues, companies are investing substantially in autonomous vehicles (AVs) (Siddiq and Taylor 2022, Chen et al. 2024b). AVs promise lower labor costs and higher fleet utilization through full-time, driverless operations (Al-Kanj et al. 2020b), while their electric nature offers potential reductions in carbon emissions (Zhuge and Wang 2021, Zhang and Chen 2020).

Whether investments in AVs will deliver on their promises remains unclear, particularly from a user adoption perspective (Ketter et al. 2023). While research highlights potential operational improvements for platform providers (e.g., Yao et al. 2020, Chen et al. 2024b), these gains ultimately depend on users' willingness to embrace the new technology. If a significant proportion of users are hesitant or unwilling to use autonomous services, the anticipated benefits for ride-hailing platforms could be compromised. Prior studies have shown that users often struggle to trust new autonomous technologies, especially when the technology is perceived as imperfect

¹This Chapter is currently under review at a leading peer-reviewed academic journal. Earlier versions of this Chapter have also appeared in a (non-copyrighted) peer-reviewed academic conference: Ahadi, R., Taudien, A., & Ketter, W. (2023). Human versus automated agents: How user preferences affect future mobility systems. ECIS 2023 Research Papers, 382.

(Glikson and Woolley 2020). In the context of autonomous services, users generally prefer that at least some aspects of the service involve human interaction (Gnewuch et al. 2023). Initial studies on user attitudes toward AVs suggest that the majority remain skeptical about relying on driverless mobility, with only a small minority showing trust in these systems (Shariff et al. (2017)).

Given this backdrop, understanding variations in user preferences for AVs across different groups is essential for addressing adoption barriers. Identifying these differences can help ride-hailing platforms target user groups that are hesitant to AV adoption. Moreover, analyzing the conditions under which certain users may hesitate or refuse to switch to AVs can provide strategic insights for platforms as they transition to autonomous fleets. We, therefore, aim to explore heterogeneity in user preferences for human-driven versus autonomous ride-hailing services and evaluate the resulting impacts on platform performance, including profitability, service quality, and environmental outcomes. We pose the following research question (RQ):

RQ1: *How do users differ in their preferences for autonomous versus human-driven ride-hailing services considering trip and user characteristics?*

In addition to understanding current user preferences, it is essential to recognize that both technology adoption and user attitudes are dynamic and evolve over time. The transition from human-driven to fully autonomous fleets is likely to occur gradually. Previous research has often adopted a static perspective, overlooking the potential implications of a phased transition from human to autonomous services (e.g., Fagnant and Kockelman 2018, Dong et al. 2022). In contrast, our study adopts a dynamic approach, acknowledging that the coexistence of human and autonomous services during this transition will influence both user behavior and platform outcomes. As AVs become more prevalent, user preferences are likely to evolve, shaped by factors such as increased familiarity with the technology (Komiak and Benbasat 2006), societal influences (Venkatesh and Davis 2000), and ongoing technological advancements. Understanding how these preferences evolve and their implications for the adoption of autonomous services is crucial for predicting the long-term success of AV adoption in ride-hailing fleets. This requires a thorough examination of how the interplay of evolving user preferences and technology adoption affects the profitability and operational efficiency of ride-hailing platforms. This leads to our second research question:

RQ2: *How does the likely evolution of user preferences affect ride-hailing platforms' outcomes during the transition from purely human to autonomous services?*

To answer these questions, we employ a novel multi-method approach that combines a discrete choice experiment (DCE) and agent-based modeling, connecting behavioral and design science research. The DCE captures current heterogeneity in user preferences, while the ABM simulates dynamic interactions between users, services, and their characteristics. Given the limited actual deployment of AVs, particularly outside pilot zones (e.g., San Francisco, Phoenix), field data remains scarce. A stated preference (SP) approach, such as DCE, allows us to over-

come this limitation by presenting participants with hypothetical scenarios. In these scenarios, participants choose between alternatives (human-driven vehicles (HVs) or AVs), with each option defined by specific attributes. This method allows us to estimate utility functions and predict user trade-offs between attributes, offering key insights into future decisions (Hensher et al. 2005). DCEs are widely used in transportation (e.g., Li and Kamargianni (2020)) and behavioral research (Schlereth and Skiera 2017, wiakowski et al. 2016), making them a robust tool for this study.

To predict the performance of autonomous ride-hailing platforms and assess the impact of user characteristics, it is essential to control for factors such as user population composition, trust in technology, and service features (e.g., the proportion of AVs and operational policies). These factors interactively influence the system’s dynamics; for instance, users’ decisions depend on vehicle availability and can alter the spatio-temporal distribution of vehicles within the service area. To address this complexity, we use an ABM, which allows for modeling the relationships between different entities in the mobility environment while controlling for various counterfactual variables. ABM enables the computational simulation of socio-technical systems by designing interactive agents that represent entities of the environment (e.g., ride-hailing users), with the aim of predicting and evaluating complex phenomena (Miller and Page 2009). In mobility-on-demand platforms, ABM offers a cheaper, faster, and safer (i.e., risk-free) alternative to methods like field experiments. It also provides the flexibility to isolate and explore the impact of multiple factors (e.g., trust and penetration of user clusters). Particularly, given the limited use cases of AVs, a method capable of predicting the comprehensive impacts of anticipated system changes is crucial.

By combining behavioral and design science, we make primary contributions to the literature. First, we analyze user preferences for autonomous versus human-driven ride-hailing services. We find four distinct groups of users, of which only one prefers AVs over human-driven ride-hailing services. However, factors such as price and waiting time influence user preferences in all classes, offering platform providers an important tool to encourage users to use their autonomous services.

We develop a feature-rich, and behaviorally informed ABM of hybrid autonomous ride-hailing services. Designing an accurate simulation tool for such a complex socio-technical environment requires in-depth domain knowledge and empirical analysis. Specifically, we must intricately design both the supply and demand sides and their interactions. On the demand side, our DCE analysis informs user behaviors, population composition, and trust in autonomous mobility, while historical travel data calibrates demand generation in our ABM. On the supply side, we implement validated operations management of the ride-hailing fleet and model the behavior of human drivers to produce reliable results. Thus, we build upon prior studies that simulate and/or optimize the performance of autonomous ride-hailing platforms (e.g., Fagnant et al. 2015, Dlugosch et al. 2022). These works mostly focus on the fleet management and often overlook the

gradual adoption of AVs in shared mobility which yields hybrid autonomous services. This work is one of the first studies of predicting end-to-end impacts of heterogeneous user characteristics and their likely evolving behaviors on hybrid ride-hailing systems. We also demonstrate the usefulness of ABM to explore the impacts of specific factors on the adoption of AVs while controlling for counterfactual variables which is often very difficult to conduct analytically or experimentally (Zhang et al. 2020a).

Furthermore, we utilize our ABM to systematically evaluate and predict the longitudinal dynamics of ride-hailing platforms, focusing on how user interactions influence system performance over time. Predicting the end-to-end impacts of AV adoption within shared mobility platforms, as well as evolving user characteristics, necessitates a design science approach, as autonomous services are not yet widely implemented (Hevner et al. 2004). Even when such a system exists, assessing the impact of demand characteristics by controlling user preferences (e.g., trust in AVs) would be prohibitively expensive, if not impossible, through field experiments. Our findings show that, given user characteristics from our DCE analysis, even a small proportion of AVs in a ride-hailing fleet can significantly reduce CO₂ emissions and increase provider revenues. Despite these benefits, user acceptance rates remain relatively low, even with a high AV penetration rate that should ensure high service availability. This indicates that the deployment of AVs alone is insufficient to fully optimize service quality. Our further analysis demonstrates that increasing trust in AVs can amplify these positive outcomes, emphasizing the importance for platforms to foster trust in their AV services. Notably, the impact of trust varies across different user groups, suggesting that service providers should target trust-sensitive users and focus on improving safety measures and educating users to boost confidence in AVs. Additionally, our supply-side analysis highlights that effective fleet management—such as offering AV discounts and optimizing fleet size—can further enhance system performance.

The rest of paper is structured as follows. We position our paper within the existing literature of information systems (IS) and shared mobility. Next, we provide an in-depth explanation of our research approach, and finally present results and discuss our managerial insights.

3.2 Literature Review

While mobility provides benefits for our society, it also creates environmental, social, and economic challenges (Whittle et al. (2019)). These challenges, inherent in socio-technical systems as complex as transport, are often classified as wicked problems (Ketter et al. (2016)). The mobility transformation - driven by the sharing economy, electrification, and automation - has the potential to address these issues (Dlugosch et al. (2022)). As part of the sharing economy, ride-hailing platforms leverage digitally-enabled business models (de Reuver et al. 2018), creating a digital layer atop traditional physical mobility systems and interweaving transportation with IS (Yoo et al. 2010a). While digital platforms are grounded in technical infrastructure, IS

literature emphasizes their socio-technical dimensions and the importance of considering their functionalities within a broader social context (Bonina et al. 2021). In the following, we review related literature on AVs in shared mobility and the interaction of users with autonomous technologies.

3.2.1 Next-Generation Urban Mobility Systems

The future of urban mobility is forecasted to be shared, electric, and autonomous (Sperling 2018b). Shared mobility systems, in particular, have drawn research attention (e.g., Rhee et al. 2022, Kahlen et al. 2024) due to their transformative impacts on transportation and the potential enhancements derived from the integration with theoretical studies. Shared mobility platforms are categorized into two types: (i) business-to-customer vehicle sharing platforms that privately own and manage their resources, and (ii) consumer-to-consumer ride-sharing platforms that coordinate private drivers with travel requests. Both face operations management challenges, from strategic issues like fleet sizing to operational decisions like rebalancing policies (Gansterer et al. 2022). The use of autonomous vehicles (AVs) can enhance the performance of both ride-hailing and carsharing systems by lowering costs associated with driver payments (in ride-hailing fleets) and service worker expenses (in carsharing fleets) (Dong et al. 2022). As a result, large mobility service providers (e.g., Lyft (2019), Uber (2019)) are working toward the adoption of AVs, converging on the concept of shared autonomous electric vehicles (SAEVs) (Siddiq and Taylor 2022).

Recent studies increasingly explore various aspects of SAEV integration, including operational management, environmental impact, and societal factors such as user adoption. For details, we refer to a comprehensive review provided by Narayanan et al. (2020). Many works such as Fagnant and Kockelman (2014), Chen et al. (2016), Loeb and Kockelman (2019) evaluate the use of SAEVs under different circumstances and show that relying on technological capacities, SAEVs can replace multiple HVs if operated correctly. Dhanorkar and Burtch (2022) examine both the advantages and drawbacks of SAEVs, emphasizing how improper use can exacerbate traffic congestion. Simoni et al. (2019) suggest a pricing strategy to mitigate unwilling driving mileage issues, and Wei et al. (2022) explore the integration of SAEVs with public transportation to reduce unnecessary trips. Methodologically, ABM is frequently employed to capture complex relationships in mobility systems (Jing et al. 2020). In the design phase, Lokhandwala and Cai (2018), Zhang and Chen (2020) study the optimal fleet sizing and charging infrastructure development of SAEVs, respectively. From an operational perspective, Ahadi et al. (2023) examine the dynamic allocation and charging strategies for SAEVs, Chen and Liu (2022) propose integrated approaches for planning and operations, while Siddiq and Taylor (2022) investigate the profitability of competitive ride-hailing platforms adopting AVs and analyze the societal and economic impacts of shared AVs on drivers and the community.

As highlighted in the studies mentioned above, a diverse research community, encompassing fields such as operations management and economics, acknowledges the pressing need for academic research on SAEVs. However, except some recent works (Babar and Burch 2020, Dlugosch et al. 2022, Zhang et al. 2020b, Rhee et al. 2022, Kahlen et al. 2024) the IS community has not reacted enough to these socio-technical mobility systems. Therefore, this work responds to the call by Ketter et al. (2023) to explore the potential of SAEVs using data-driven and IS-enabled techniques. Moreover, a domain void in transportation literature is studying an anticipated hybrid human-driven and autonomous shared mobility system (Ao et al. 2024). This gradual transition phase has not been explored enough since most researchers assume a fully automated fleet. IS community is well-suited to address this gap by exploring the interactions of users and hybrid autonomous mobility platforms and the adoption of AI-enabled new technologies.

3.2.2 User Behavior and Human-AI Interactions

As AVs are introduced into ride-hailing platforms, they create what Rai et al. (2019) term next-generation digital platforms, where interactions occur between humans and AI. While fleet operators provide the technical infrastructure, user acceptance ultimately determines the success of these developments. Factors influencing user acceptance in IS literature include the perceived usefulness (Davis 1989), ease of use (Venkatesh et al. 2012) and trust (Glikson and Woolley 2020). Research highlights that while users may benefit from autonomous systems, adoption is often accompanied by challenges, such as psychological barriers and a need for control (Rijsdijk and Hultink 2003). Studies in transportation on the adoption of AVs align with these findings. AVs are expected to be a more radical and less familiar prospect than shared and electric vehicles for the majority of people (Whittle et al. 2019). This is reflected in a lack of user trust due to safety (Jabbari et al. 2022) and security concerns (Wali et al. 2021). These findings are in line with literature stressing that trust in technology automation may be hard to establish (Lee and See (2004)) or can easily be lost after seeing systems fail (Dietvorst et al. 2014).

User behavior towards autonomous systems is heterogeneous, with some groups more willing to adopt technology than others (D'Acunto et al. 2019). Early research on user perceptions of SAEVs reveal that willingness to use these services varies across demographic and psychological factors, such as price sensitivity (Jabbari et al. 2022) and performance expectations (Curtale et al. 2022). Most studies, however, assume an abrupt transition to autonomous systems, overlooking the likely interim phase where human-driven and autonomous systems coexist. In contrast, Adam et al. (2022) emphasize the importance of studying user behavior in such hybrid scenarios, finding that user intentions can shift across different stages of interaction with autonomous systems. In the following we address the gap of studying user behavior in a hybrid autonomous transportation system in a dynamic environment.

3.3 Research Method

The primary objective of this paper is to examine the preferences of mobility users within hybrid autonomous ride-hailing systems and to develop an IS-enabled simulation that aids shared mobility platforms in predicting the impact of important factors such as AV adoption and user characteristics. Our research approach, illustrated in Figure 3.1, consists of two main components: a) analyzing and modeling user behavior within shared mobility systems, and b) modeling hybrid autonomous ride-hailing systems to assess system performance. Given the limited implementation of AVs and the scarcity of empirical data on user preferences, we conduct an online DCE to identify and analyze ride-hailing users' preferences for autonomous versus human-driven mobility services. For other user characteristics such as arrival time and origin-destination distributions, we perform empirical analysis on observational taxi trip data. Using our ABM, we simulate a hybrid autonomous ride-hailing system, incorporating insights from our DCE and empirical analysis to calibrate mobility user agents realistically. By integrating real-world data, survey experiments, and rigorous software testing and validation, we ensure that the simulation artifact accurately reflects real-world settings. Finally, we design simulation experiments to examine the interactions between heterogeneous user behavior and the gradual adoption of AVs in ride-hailing services.




Figure 3.1: Research Process

3.3.1 Shared Mobility Demand

Survey Overview

To analyze user preferences towards autonomous mobility, we conducted an online choice experiment as part of a larger survey. The survey was administered via Prolific, an online platform that provides access to a diverse participant pool with quality filters (Peer et al. 2017). Participants were required to have prior experience using ride-hailing services to ensure they could meaningfully engage with the choice scenarios. The survey began with an explanation of ride-hailing services, followed by questions about participants’ most recent ride-hailing trip, including its purpose and timing. This information was used to customize the scenarios and make them relatable, minimizing response bias (Papu Carrone et al. 2020). Participants were then introduced to the experimental task, imagining they were using a ride-hailing app to select from two available vehicle options (i.e., an AV or a HV). An example scenario was provided, along with a detailed explanation of the vehicle types and attributes which are included in the Appendix. An attention check was included to ensure participants understood the task before proceeding. The survey concluded with questions on participants’ environmental concerns (Haboucha et al. 2017), trust in AVs (Yagoda and Gillan 2012), general interest in technology (Haboucha et al. 2017), familiarity with AVs, ride-hailing frequency, and demographic information.

Design of the Discrete Choice Experiment

The DCE is an econometric method used to elicit consumer preferences and predict choice behavior (Ben-Akiva and Lerman 1985, Liu et al. 2018). Typically, participants are presented with a set of product choices (in our case ride-hailing options) and shall decide which option they prefer. The options differ in their attribute levels, requiring subjects to make trade-offs between the options (Liu et al. 2018). Based on related literature, we included attributes identified as most important for consumers’ travel and ride-hailing choices: price (Whittle et al. 2019, Hong et al. 2020), trip length, waiting time (Hong et al. 2020), and the vehicle’s propulsion (electric or petrol powered). A pilot study with 29 participants confirmed the relevance of the selected attributes and their levels, ensuring realistic and comprehensible scenarios. Table 1 summarizes the attributes and levels used in the DCE.

Attribute	Levels	
	Autonomous Vehicle	Human-driven Vehicle
Duration of the trip (min)	10, 20, 30	
Price	-20%, -10%, 1	1, +10%, +20%
Waiting time (min)	5, 7, 10	
Power	Electric	Electric, Petrol

Table 3.1: Overview of Attributes and Levels used in DCE

To reduce cognitive load, we applied a D-optimal design, which minimizes the D-error and improves the efficiency of parameter estimates compared to classic orthogonal designs (Caussade et al. 2005). The final design included 54 scenarios split into nine blocks, with each participant randomly assigned to one block. Each block contained seven scenarios, including one consistency check to ensure reliability. A separated dual response (SDR) approach was implemented to address potential biases. After completing the initial scenarios, participants were shown the same scenarios with the option to choose "none of the above." This approach reduces extreme response behavior and improves the ability to capture participants' willingness to pay (Schlereth and Skiera 2017).

Discrete Choice Model

Data from the DCE were analyzed using random utility theory (Liu et al. 2018, McFadden 1972). Based on the assumption that individuals maximize their utility when making choices, we estimate an individual's utility for a ride-hailing service as a function of the attribute levels of that service (Liu et al. 2018, Leitham et al. 2000). The utility (U) that an individual i assigns to choice scenario j takes the following form:

$$U_{ij} = V_{ij} + \epsilon_{ij}. \quad (3.1)$$

with

$$V_{ij} = \beta_0 ALTERNATIVE_{ij} + \beta_1 PRICE_{ij} + \beta_2 WAIT_{ij} + \beta_3 DURATION_{ij} + \beta_4 ALTERNATIVE_{ij} \times TRUST_i. \quad (3.2)$$

We estimate V_{ij} as the utility's deterministic element using a latent class approach. This allows us to address potential heterogeneity in subjects' choices regarding autonomous compared to HVs for ride-hailing. The idea behind latent class models (LCM) is that individual behavior depends on observable attributes as well as latent heterogeneity varying with factors unobserved by the researcher (Greene and Hensher 2003). LCMs incorporate heterogeneity through distinct preference classes and generate probability estimates for included attribute parameters that vary by class (McFadden 1972). Within a class, the correlation across an individual's choices are all attributed to the class membership. After conditioning upon the class, all responses are considered independent across choice pairs (McFadden 1972). Essentially, the conditional logit model is a regression model for choice data which determines the probability that an individual i selects alternative j from choice set k (McFadden 1972). Further, as subjects evaluated multiple scenarios, our data have a panel structure.

Empirical Mobility-on-Demand Data Analysis

To accurately analyze the impact of user behavior on hybrid autonomous ride-hailing systems, it is essential first to estimate shared mobility demand distributions realistically. We empirically analyze historical trip data that includes comprehensive information of mobility requests, such as origin and destination locations, start and end times, price, vehicle ID, and energy consumption. In line with related works (e.g., Shortle et al. 2018, Demircan et al. 2022), we assume that demand for ride-hailing services within the service area follows exponential spatio-temporal distributions. To determine the distribution parameters, we divide the geographical service area into uniform hexagonal-shaped zones, referred to as hexagons, each with a radius of approximately 1.3 km, and segment the day into hourly time intervals. For each time bucket $t \in \mathcal{T}$ (the set of time intervals) and hexagon $h \in \mathcal{H}$ (the set of zones), we model the inter-arrival time between successive ride requests using an exponential distribution with the rate parameter $\lambda_{t,h}$. We use these parameters to generate trip requests in our ABM. Each trip request i includes an origin o_i , an arrival time t_i , and a destination d_i . To determine the destination d_i , we use origin-destination patterns by drawing from a multinomial distribution, $d_i \sim \text{Mult}_{|\mathcal{H}|}(1, \{p_1, \dots, p_{|\mathcal{H}|}\} \mid t \in \Gamma, o \in \mathcal{H})$, where the possible destinations (p_h) are treated as outcome categories.

3.3.2 Agent-Based Simulation for Hybrid Autonomous Ride-Hailing

To examine the temporal dynamics of a hybrid autonomous ride-hailing system, we use an ABM, where all entities (e.g., mobility users, vehicles, and the ride-hailing operator) are represented as interactive agents. This approach is particularly suitable for our study because the limited adoption of autonomous ride-hailing fleets precludes field experiment analysis. Additionally, ABM allows us to predict the impact of specific factors while controlling others, enabling shared mobility platforms to proactively prepare for and integrate autonomous mobility into their systems. Agent-based systems are also cost-efficient, low-risk, and ideal for studying non-existent environments, such as shared autonomous fleets (Miller and Page 2009). ABM has been used in the social-science community to model complex dynamic systems (Zhang et al. 2020a). In the IS literature, Haki et al. (2020) presents an ABM to explore the evolution of information systems architecture, focusing on how different agents interact and influence the development and adaptation of system structures over time. Bapna et al. (2008) use ABM to simulate interactions between auctioneers and bidders within auction mechanisms; Ren and Kraut (2014) model online communities to study how design factors influence their success; Ketter et al. (2016) simulate the uncertain environment of trading agents, demonstrating the value of competitive benchmarking in addressing complex socio-technical challenges; and (Zhang et al. 2020a) highlights the advantages of ABM for analyzing the performance of recommender systems over time under diverse conditions.

Inspired by Ketter et al. (2023), Dlugosch et al. (2022), digitalized transportation system problems such as evaluating autonomous ride-hailing fleets are considered wicked problems where

we need to consider the interaction between mobility users, vehicles, infrastructure resources and road networks. For such complex problems, ABMs that employ self-operated agents with an bottom-up approach are appealing to quantitatively study the system performance under different scenarios of platform design, fleet configurations, and user characteristics (Yao et al. 2020, Fagnant and Kockelman 2018). Below, we outline the different components of our simulation framework, which integrates digital and physical layers to model the interactions required for on-demand mobility services.




Figure 3.2: A Framework for Hybrid Autonomous and Human-driven Ride-hailing Fleets

Drawing inspiration from the architecture of MatSim (W. Axhausen et al. 2016), a renowned agent-based urban transportation simulation framework, we structure our model around three main components: the user’s perspective (demand side), the operator’s perspective (supply side), and the underlying mobility infrastructure. We focus on a hybrid ride-hailing fleet of AVs and HVs, operated by a single provider delivering on-demand mobility services to different types of users. We model a complete package of both demand and supply sides while zooming in on the user characteristics to answer our research questions regarding the influence of user behaviors on the adoption of AVs in shared mobility platforms.

Figure 3.2 presents an abstract of our proposed ABM, highlighting its primary components and their interactions. Our model’s inputs include service region, demand characteristics, and fleet configuration, to initiate the simulation environment. A service region defines the geographical boundaries where the fleet provides services. The business area is divided into uniform zones with identical demand distributions and infrastructure capacity. The fleet configuration includes details such as the total number of vehicles, the proportion of autonomous and electric vehicles, driver availability, infrastructure resources (e.g., charging stations), and operational strategies like trip matching and rebalancing policies. The third input component is demand characteristics, encompassing arrival patterns, origin-destination probabilities, and users’ mode

choice preferences, which are derived from our empirical and DCE analysis. Our ABM supports multiple agents: a fleet operator, vehicles (AVs and HVs), charging stations (CSs), and users, with individual attributes and decision-making/executing models. The arrows in the simulation framework box in Figure 3.2 indicate the flow of information between agents and the sequence of execution. Central to the simulation is the fleet agent, which systematically gathers all relevant information at each decision time step to oversee fleet management. This involves matching available vehicles with ride-hailing requests and making decisions related to fleet rebalancing. It identifies AVs that require charging, assigns them to CSs, and communicates this information to charging agents. The fleet agent matches vehicles with users by presenting each passenger with two vehicle options. The user agent selects one based on its mode choice preference and receives service from the chosen vehicle. This process repeats throughout the simulation’s operational period, recording key outputs such as aggregated data on served and canceled requests, vehicle utilization, and charging infrastructure status. Analyzing the logged data enables the quantification of key metrics, including fleet revenue, service quality (acceptance rate), and carbon emissions.

Supply Side

We assume a hybrid mobility supply of AVs and HVs, with AV penetration increasing gradually over time until it fully replaces HVs. In order to narrow our focus on examining the interaction of user behavior and AV penetration in ride-hailing systems, we fix the fleet size (total number of vehicles) and charging resources over time. While all AVs are electric (reflecting the synergy between electrification and automation), only a portion of HVs is electric, with this share increasing alongside technological advancements. All electric vehicles (EVs) are permitted to charge in the fleet’s privately-managed charging facilities, which are restricted in terms of parking capacity and charging power to consider the challenges of mobility electrification. AVs are always available for mobility services, with the exception of recharging or relocating periods. In contrast, as highlighted by Zwick et al. (2022), human drivers are only available partially. To account for driver availability, our model adjusts the actual number of HVs based on a distribution of driver shift hours, which is derived from historical trip data.

To match vehicles with trip requests, the fleet agent pools all open requests and available vehicles between two consecutive decision time slices (e.g., one minute), following the first-come-first-serve priority. For each request, the fleet suggests at most two vehicles (one AV and one HV) within the trip coverage boundary, maximum driving distance (e.g., 5 km). Depending on the system status, trips might receive two, one, or no options. If they receive at least one ride-hailing option, based on their preferences, they can choose one or reject the offer. However, due to the lack of supply, some trips might not be served in the current time window and will shift to the next one. We assume a patience time threshold for passengers, after which they will cancel their requests.

For service prices, we use a base fare plus a variable fare, which is a linear function of trip distance and duration. Other factors, such as the time of day, traffic congestion, and dynamic pricing, are excluded for simplicity. We show the price calculation for a trip in Equation (3.3), with ω as the revenue per distance, and ϕ as the revenue per duration. We consider a minimum price for all trips and assume that AV prices are relatively lower (reduced by α) than that of HVs due to the elimination of driver payments.

$$\textit{TripPrice} := \min(\textit{BaseFare} + \omega * \textit{Distance} + \phi * \textit{Duration}, \textit{MinimumFare}). \quad (3.3)$$

The fleet operator also centrally makes rebalancing decisions for all AVs, but HVs individually decide when and where to charge or relocate. Regarding charging decisions, all AVs with an energy level below a certain threshold (see Table 3.2) need to charge. The threshold has an opposite pattern compared to the average temporal demand of all zones, derived from Ahadi et al. (2023). All CSs use the same technology and have the same charging rate. Thus, each vehicle that needs to charge is assigned to the closest CS with at least one free charger (inspired by queuing theory extensions). The fleet operator also relocates the AVs in order to achieve a better spatial supply and demand balance. It first compares the supply and demand for each idle AV for its current zone and neighbors. If the supply exceeds the anticipated demand, the vehicle relocates to a target zone (the closest low-supply hexagon).

	1	2	3	4	5	6	7	8	9	10	11	12
Charging threshold % (morning)	0.45	0.60	0.65	0.62	0.58	0.55	0.52	0.50	0.40	0.40	0.40	0.40
Charging threshold % (afternoon)	0.38	0.35	0.32	0.25	0.25	0.20	0.20	0.25	0.27	0.35	0.35	0.40

Table 3.2: Hourly Charging Threshold for Determining Charging Vehicles

Demand Side

The demand factory in our ABM generates trips using estimated spatio-temporal exponential distributions derived from observational trip datasets. At each time step, ride-hailing requests are generated in parallel for each hexagon, with destinations assigned based on multi-nomial origin-destination patterns. The simulation then determines the user class according to their mode choice behavior. User classes, identified from our survey experiment analysis, are used to cluster users, with their penetration rates factored in. Each user class has distinct parameters for the mode-choice utility function, leading to varied mode-choice behavior. For instance, given the same AV and HV options, a user from class A might choose the AV, while a user from another class might choose the HV or reject both options. We assume that the penetration rates of user classes are consistent across all hexagons and time buckets. Users receive at most two ride-hailing options and can choose of them or reject both. Sometimes, in rush hours or when the local demand is higher than local supply, users might not receive any ride-hailing offer

from the fleet. We assume individual waiting time threshold (e.g., 10 minutes) of users after which they cancel their request (Yao et al. 2020).

As explained in Section 3.1.3, we assume that users are aware of their utility for each option and choose the vehicle with the highest utility. To calculate the utility we consider the length of the trip, alternative-specific factors including the price and waiting-time as well as a constant and the trust on AVs. Therefore in the simulation, whenever a user receives mobility service options, first, we calculate the utility for each option based on the aforementioned variables, and second, we use a Logit model (Hensher and Greene 2003) to calculate the probability of choosing one alternative over the other, according to the utilities:

$$P_{ij} = \frac{\exp(U_{ij})}{\sum_i \exp(U_{ij})} \quad (3.4)$$

The probability P_{ij} shows the likelihood of choosing alternative i for user j in a stochastic way. For example, if the probability of choosing AV is 75% for one trip, it does not mean that AV is certainly chosen (which may be the case in a deterministic process). In a stochastic process, there is 75% likelihood of having an AV trip and 25% likelihood of other choices.

3.3.3 Key Performance Indicators

To compare results across different user behavior scenarios, we quantify system performance in a way to cover our multiple objectives. For economic metrics, we measure the fleet's total revenue, which includes:

$$FleetRevenue = \sum_{r \in \mathcal{R}_A} (p_r(1 - \alpha)) + \sum_{r \in \mathcal{R}_H} p_r \zeta, \quad (3.5)$$

Revenues consist of serving requests using AVs and HVs, where \mathcal{R}_A and \mathcal{R}_H represent the sets of requests served by AVs and HVs, respectively. If an AV serves a trip, the total price ($p_r(1 - \alpha)$) is directly assigned to the fleet. It is important to note that the price of ride-hailing autonomous services is assumed to be marginally lower than that of human-based services due to the absence of driver payment. In order to account for this difference, we include a reduction of $\alpha\%$ of the price that users must pay for services from HVs. When a human driver serves a trip, only a portion of the trip price ($p_r \zeta$) is assigned to the fleet, and the driver receives her share directly from the user's payment. In addition to fleet revenue, we measure service quality through two metrics: the fraction of fulfillment requests (acceptance rate) ($ServiceQuality = \frac{|ServedRequests|}{|Requests|}$) and average *WaitingTime* for all served requests. The *WaitingTime* for each request is the summation of the assignment and pick-up duration. From an environmental perspective, we track overall *CO₂Emissions*, distinguishing between green transportation modes (EV) and non-electric vehicles. In other words, *CO₂Emissions* per driving distance differs for electric and petrol cars.

3.3.4 Technical Implementation and Simulation Validation

For the technical implementation of the simulation we use SimPy, a Python package based on object-oriented programming to develop discrete-time and discrete event stochastic simulation processes. The digital representation of the environment is defined on this simulation platform, allowing us to define multiple interactive agents. Agents execute parallel processes whenever an event occurs, triggering the subsequent process. For instance, when a user selects an AV from the available ride-hailing options, the fleet operator agent initiates a request serving event, passing it to the chosen AV. This action triggers the AV's serving process, which includes sub-processes such as relocating to the user's origin, picking up the user, and driving to the destination. The vehicle's status updates with each new task, preventing it from being assigned other tasks. For example, a vehicle that is charging cannot serve a mobility request unless an exception interrupts the charging process and triggers a different task. We also consider a sequential job queue in certain scenarios. For example, when matching vehicles and requests, we include not only idle vehicles but also AVs currently serving other requests that will soon be free and are near the matching user's origin. If a serving AV matches with another user, it adds this new task to its job queue, beginning the new task after completing the current one (e.g., dropping off the current user).

The simulation platform must closely mirror reality, ensuring that only feasible actions and decisions are allowed. To achieve a realistic simulation, we use real-world taxi trip data from the City of Chicago to estimate user arrival times at various locations and their likely destinations. For user preferences regarding autonomous versus HVs, we incorporate findings from our DCE. Additionally, to model the adoption rate of AVs, we rely on results from existing literature and industry reports (Talebian and Mishra 2018). To ensure the accuracy of the software platform, we implement validation techniques recommended by Sargent (2010) to guarantee a bug-free simulation. First, we trace all agents to avoid infeasible circumstances. For example, if a vehicle has a negative energy level, its location is outside the service area, or it's executing in parallel processes (e.g., serving two users or charging and relocating simultaneously) the simulation will abort with an error. Additionally, we conduct extreme condition tests. For instance, if the fleet lacks access to CSs, its performance rapidly declines as many vehicles run out of energy and cannot serve trip requests. To validate the simulation software further, we perform sensitivity analyses; for example, increasing the fleet size exponentially improves the acceptance rate and reduces average waiting time, but these benefits plateau beyond a certain point. Finally, we compare the simulated ride-hailing demand with observational trip data to ensure they exhibit similar patterns. Demand distributions are very similar. The only difference is that the generated trips by simulation have less spatio-temporal variance compared to the empirical data.

3.4 Simulation Experiment Setup

To discretize the spatio-temporal network, we divide the service region (city of Chicago) into 117 hexagons with an edge length of approx. 1.3 km. These zones are characterized by distinct hourly arrival rates, origin-destination patterns, and limited charging facilities. Each simulation time step corresponds to one minute in the real world, and each simulation experiment represents one business day.

We use observational taxi trip data from the City of Chicago to feed the simulation and generate realistic mobility requests. The dataset includes 6.38M trip observations from January 2022 to December 2022. Each observation specifies a trip start date/time and end date/time, origin location (latitude and longitude), destination location, distance, fuel consumption, payment details, and vehicle ID. This allows us to find the hourly arrival rates of each hexagon ($\lambda_{h,t}$) by fitting a Poisson distribution for trips originating from hexagon h at time t . The hourly-origin-destination probabilities are calculated as an average fraction of trips from hexagon i at time h that ends at hexagon j . Figure 3.3 visualizes the frequencies of demand origins on the left side and an hourly demand distribution on the right side. For the origin hexagons, a heat map compares the arrival rate of all hexagons within the service area. A darker color indicates higher arrival rates. As it shows, trips mostly start from central areas, whereas suburbs have considerably lower arrival rates with some exceptions such as trips started from the airport. Trip destinations follow the same pattern such that we do not visualize them here. Concerning the temporal distribution, we compute an average of rentals for each hour among all hexagons. The boxplot on the right side of Figure 3.3 shows that demand is deficient from midnight to early morning. Still, from the start of working hours in the morning, the average and variance of rentals pile up and peak in the evening and drop off during night hours. Variance in the boxplot reflects fluctuations across different days.




Figure 3.3: Demand Distribution for Mobility-on-Demand Services in the city of Chicago

Another outcome of the empirical analysis of observational trip data is the availability of drivers at different times of the day. To identify these patterns, we track drivers using their vehicle identifiers in the historical data to generate hourly availability data across the entire time horizon. We consider a driver active or available for a time slot if they serve a trip within a time window that overlaps with that slot. Conversely, drivers are marked as idle during slots when they do not serve any trips. It is assumed that drivers go offline if they remain idle for a while. Using the estimated availability data, we can describe the actual fleet size of HVs in different time slots. Figure 3.4 shows the estimated hourly patterns of the proportion of active drivers compared to the total number of vehicles. As illustrated, the actual fleet size closely mirrors the hourly trip patterns, with higher numbers in the afternoon and lower numbers during the night and early morning hours. For each hour, we fit a normal distribution to the observed fleet size and use this distribution in the simulation to approximate driver behavior accurately.




Figure 3.4: Hourly Driver Availability Distribution Driven from Observational Trip Data

To initiate the charging infrastructure, we consider 16 CSs distributed in different locations of the service region based on the trip demands. We sort the hexagons based on the demand and locate these 16 CSs in those with the highest mobility demand. All chargers are homogeneous with 20 number of fast charging (55 kW) connectors. To characterize the fleet, we consider a fixed number of 500 hybrid AVs and HVs. It should be noted that the determination of the fleet size and the number of CSs is conducted manually in order to ensure an acceptable level of mobility service quality. It is therefore evident that the aforementioned parameters do not affect the validity of the main findings if they are selected in a reasonable manner. As described before, over time, the penetration of AVs will increase but the overall fleet size remains the same (see section 5.2 for details). AVs are homogeneous and similar to Tesla Model 3 (Fuel economy guide 2020), which are allocated at the beginning of the simulation time across the service area with

an energy level between 50% and 70% of their battery capacity (50 *kwh*). However, HVs are heterogeneous in terms of vehicle brand, size, and fuel type. Regarding trip price components, we consider a base fare of 3.25 USD, a revenue per distance of 2.25 USD/mile, and a revenue per duration of 0.33 USD/min. Note that the minimum charge for each trip is 5 USD (Taxi costs in Berlin 2022). Concerning the differences between autonomous and human services, we consider a 10% price reduction (α) for AVs and a driver share of 75% (ζ) from the trip price. We assume that each vehicle has an approximate average speed of 20 km/hr for moving in the road network in the simulation. Finally, we quantify the CO₂ emissions of EVs and non-EVs to 75gr/km and 241gr/km, respectively (Transportation & Environment 2022). A summary of parameters is provided in Table 2.1.

Parameter	Value	Parameter	Value
Simulation length	1 day	Number of zones/hexagons	117
Number of charging stations	16	Total fleet size	500
Charging power rate	55 kW	EV battery capacity (energy consumption)	50 kWh (0.20 kWh/km)
EV (non-EV) CO ₂ emission	75 (241) gr/km	Vehicle speed	20 km/hour
Driving cost	0.43 USD/km	AV price reduction	10%
Driver revenue share (1- ζ)	75%		

Table 3.3: A Summary of the Key Parameters of Agent-Based Simulation

3.5 Results

We first analyze the outcomes of our DCE to estimate the utility function of ride-hailing users, allowing us to understand how users choose among the available services. Additionally, we investigate the heterogeneity in user preferences for AVs and HVs, identifying different user groups with distinct characteristics. Finally, we incorporate the results of our DCE analysis into our ABM to calibrate user behavior and investigate the impact of heterogeneous preferences on the environmental and economic outcomes of ride-hailing providers as the penetration of autonomous ride-hailing services grows over time.

3.5.1 Latent Class Model of DCE

In total, we recruited 595 subjects for the DCE of which 59 were excluded for selecting the dominated alternative in the choice set (i.e., an option inferior in all attributes) (Liu et al. 2018). This left a final sample of 536 participants: 52% identified as female, 47% as male, and 1% as diverse. The average age was 29 years ($SD = 8$). We provide the details in the online appendix.

Table 3.4 shows the results of our latent class model with four user classes. The reference category in our model is the opt-out option, representing participants choosing another mode of transportation or not taking the trip. The coefficients reflect how a given attribute affects the utility of a transportation mode, holding all other attributes constant. We find that subjects

	Class 1	Class 2	Class 3	Class 4
	AV Skeptical:	AV Enthusiasts	Trip Conscious	AV Skeptical:
	Trust Insensitive			Trust Sensitive
Price	-0.30*** (0.05)	-0.31*** (0.07)	-1.25*** (0.29)	-0.20*** (0.04)
Waiting time	-0.09* (0.05)	-0.52*** (0.10)	-0.62*** (0.04)	-0.20** (0.04)
Power	-0.82*** (0.28)	-0.53 (0.56)	-0.05 (0.35)	-0.09 (0.17)
Trip duration	0.16*** (0.04)	0.03 (0.06)	1.18 (0.85)	0.09*** (0.03)
Constant AV	-1.91** (0.84)	8.57*** (1.89)	1.44 (8.43)	-3.44*** (0.73)
Constant HV	2.10*** (0.57)	7.13*** (1.12)	8.22 (8.21)	5.26*** (0.55)
Constant AV \times Trust	0.68*** (0.12)	0.13 (0.22)	1.28*** (0.30)	1.65*** (0.14)
Constant	-0.48** (0.20)	-0.66*** (0.19)	-0.46*** (0.17)	0.00 (0.00)
Share of subjects	0.22	0.19	0.23	0.36

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; Power is a binary variable with 0 = electric, 1 = combustion engine. Standard errors are in parentheses.

Table 3.4: Latent Class Conditional Logit Model

vary in their preferences regarding AVs and HVs, as well as the importance they place on different attributes, since the coefficients of the mode-specific constants and attributes differ between classes. In all classes, the coefficients for price and waiting time are significant and negative. This indicates that subjects experience a lower utility from an option with increasing price and waiting time, which is to be expected. However, the magnitude of change in utility differs between the classes, indicating that the user classes value them differently. Below, we describe the classes based on their most prominent characteristics.

Class 1 (AV Skeptical: Trust Insensitive) is the only class with a statistically significant coefficient for power ($\beta = -1.13$, $p < 0.05$). Subjects in this class prefer EVs over conventional vehicles, suggesting they are more conscious of the environmental aspects of a ride-hailing option. They also derive higher utility from trips with an increasing trip duration ($\beta = 0.16$, $p < 0.05$), indicating a preference for ride-hailing for longer trips, while for shorter trips, they might prefer a different mode of transportation. Importantly, they derive a negative utility from AVs in general ($\beta = -1.91$, $p < 0.05$). Increasing trust in AVs also increases their utility from AVs but with a comparably small magnitude ($\beta = 0.68$, $p < 0.01$). Subjects in **Class 2 (AV Enthusiasts)** are the only ones in our sample who prefer autonomous over HVs. This class is the smallest of all classes, comprising 19% of our sample. **Class 3 (Trip Conscious)**

exhibits the highest price and waiting time sensitivity of all classes, with the largest absolute coefficient for price ($\beta = -1.25$, $p < 0.01$) and waiting time ($\beta = -0.62$, $p < 0.01$). This indicates that subjects in this class are highly sensitive to cost and waiting time, and these factors greatly influence their utility. Subjects in **Class 4 (AV Skeptical: Trust Sensitive)**, which is the largest of all classes (36%), experience the most negative utility from AVs in general of all classes ($\beta = -3.44$, $p < 0.01$). As such, they would only choose an AV over a HV if the price and waiting time are very high for HVs, or if their trust in AVs is established ($\beta = 1.65$, $p < 0.01$). As shown, the preferences of users regarding autonomous mobility are highly diverse, influenced by factors such as service characteristics (e.g., price) and trust in technology. This heterogeneity, combined with the complexity of the system, makes it challenging to conduct a comprehensive analysis. To address this, an agent-based simulation approach is essential, allowing for flexible exploration of the temporal effects of various factors while maintaining control over others.

3.5.2 Simulation Analysis Results

We evaluate the performance of hybrid ride-hailing systems over 10 phases of increasing AV penetration, following the S-curve pattern identified by Talebian and Mishra (2018). This pattern is defined by slow growth during the initial and final phases, with rapid, exponential growth in the middle stages. To isolate the influence of individual factors, we maintain all other variables constant within a baseline scenario. Separate scenarios are then constructed to assess the impact of specific factors, such as trust in AVs. Table 3.5 outlines the parameters for the baseline and other scenarios. In the baseline scenario, user behavior and fleet configuration remain unchanged throughout all phases. Additionally, trust in AVs and user class penetration are determined based on the DCE results specific to each user class. In Scenario A, trust in AVs increases for all user groups over time in alignment with the growth of AV adoption. In Scenarios B-E, we examine the impact of user group penetration within the population by analyzing how fleet performance changes when the entire population is assumed to belong to a single user class.

Figure 3.5 represents the results for the baseline scenario. Revenue and service quality (measured as acceptance rate) improve substantially during the early phases when only a small percentage of the fleet is replaced with AVs. Several factors contribute to this growth. First, AVs are operational full-time, except during charging periods, which greatly increases supply availability even with limited AV penetration. Second, ride-hailing with AVs is more cost-effective than with HVs, potentially attracting more users to accept ride-hailing mobility services. Third, some users inherently prefer autonomous services, and a mixed fleet helps improve overall acceptance rates. Lastly, AVs generate higher revenue for ride-hailing platforms since all trip earnings are retained without needing to pay drivers. After phase two, corresponding to the onset of exponential growth in AV penetration, revenue and service quality spike initially but plateau or decline beyond phase four. From phases five through ten, daily revenue shows a

Phase	1	2	3	4	5	6	7	8	9	10
	Baseline Scenario									
AV Adoption	0.02	0.05	0.15	0.42	0.70	0.85	0.91	0.95	0.98	1.00
AV Trust	[4.45, 4.76, 4.81, 4.00]	"	"	"	"	"	"	"	"	"
User Share (%)	[0.25, 0.23, 0.19, 0.33]	"	"	"	"	"	"	"	"	"
	Scenario A									
AV Trust	+0.06	+0.15	+0.44	+1.26	+2.09	+2.55	+2.7	+2.82	+2.91	+3.0
	Scenario B									
User Share (%)	[1.0, 0.0, 0.0, 0.0]	"	"	"	"	"	"	"	"	"
	Scenario C									
User Share (%)	[0.0, 1.0, 0.0, 0.0]	"	"	"	"	"	"	"	"	"
	Scenario D									
User Share (%)	[0.0, 0.0, 1.0, 0.0]	"	"	"	"	"	"	"	"	"
	Scenario E									
User Share (%)	[0.0, 0.0, 0.0, 1.0]	"	"	"	"	"	"	"	"	"

Table 3.5: Phase-Dependent Parameters




Figure 3.5: Hybrid Autonomous Ride-hailing Fleets Performance for Baseline Scenario

slight upward trend, but of particular interest is the decline in acceptance rate observed after phase four. This indicates that despite the anticipated increase in the acceptance of AVs due to their high availability and lower costs, we observe a contrasting trend as some users remain skeptical of AVs (see results in Table 3.4). Revenue continues to grow even when the acceptance rate decreases over the phases, as the majority of trips are now served by AVs, which are more profitable for ride-hailing fleets. The figure also shows an exponential reduction in CO₂ emissions as the fleet transitions from conventional to fully automated vehicles. This reduction underscores

the environmental significance of adopting AVs, even in the initial phases where user acceptance rates are lower. Regarding the average waiting times, there is a consistent decline throughout the phases, inversely mirroring AV penetration rates. Notably, the reduction in waiting times becomes more pronounced after phase three and continues at a slower rate through phase ten.





Figure 3.6: The Impact of Increasing Trust Along with AV Adoption Rate on Fleet Performance (Scenario A)

In Scenario A, we examine whether the fleet performance changes if users' trust in AVs increases over time, considering the same AV adoption rate as in the previous analyses. In each phase, user trust in AVs increases proportionally to AVs penetration, reaching its maximum when HVs are completely replaced (phase ten). The results for Scenario A are illustrated in Figure 3.6. Unlike the baseline scenario, service quality (acceptance rate) continues to rise beyond phase four as growing trust reduces user skepticism. However, even with maximum trust, the fleet cannot reach very high acceptance rates, which shows that other factors such as price, waiting time, and user types play an important role in the decision-making process of travelers. This figure also shows that the increase in revenue is higher if trust in AVs is growing compared to the baseline scenario, yielding almost 18% higher revenue.



1

Figure 3.7: The Impact of User Penetration on Fleet Performance (Scenario B-E)

To further analyze the impact of user characteristics on the performance of hybrid autonomous ride-hailing fleets, we simulate the system for each user class separately. As scenarios are shown in Table 3.5, in each scenario we generate all travelers from one user class. Each scenario associates with a specific user class (e.g., scenario B examines the user class 1). Figure 3.7 shows the acceptance rate for these scenarios under two conditions: steady trust in AVs and increasing trust over time. Here we only focus on requests acceptance rate since it explicitly shows the behavior of each user class. In the steady trust condition, differences in acceptance rates are significant during early phases when AV penetration is low, with acceptance rates increasing for all scenarios until phase four. Afterwards, acceptance rates for classes 2 (AV Enthusiasts) and 3 (Trip Conscious) continue to rise, while more users from classes 1 (AV Skeptical: Trust Insensitive) and 4 (AV Skeptical: Trust Sensitive) refuse ride-hailing offers after high adoption of AVs. This behavior aligns with the findings from the DCE analysis, where classes 1 and 4 exhibited AV skepticism. However, users from Class 2, identified as AV enthusiasts, may still reject autonomous ride-hailing offers due to other factors, such as price or trip characteristics. This results show the importance of using simulation experiments to analyze the user behavior impacts. Increasing trust in AVs over time alters the results. Still, the performance of the fleet is maximized when the majority of the population is from user Class 4, but interestingly, it surges along with the AV adoption. The reason could be the comparably high coefficient of trust in AVs in their utility function, which becomes more important when trust is increasing. For other scenarios, the fleet performance is almost merging to the same value, particularly when the penetration of AVs is high. It means that the positive effect of increasing trust in AVs covers other factors and more users would accept autonomous ride-hailing services.




Figure 3.8: The Impact of Trust Level on AV Skeptical: Trust Sensitive Users

Given that previous results show user Class 4 (AV Skeptical: Trust Sensitive) is more responsive to an increase in trust to AVs compared to other groups, we conduct a more detailed analysis of this interaction. Figure 3.8 illustrates the outcomes for different baseline trust levels under steady and increasing trust scenarios. For steady trust, acceptance rates vary significantly across trust levels. Low trust leads to a steep decline in acceptance rates after phase three, reaching

only 25% in a fully autonomous fleet. Medium trust results in approximately 50% acceptance, while high and very high trust levels enable the majority of ride-hailing offers to be accepted in later phases. Conversely, when trust in AVs increases over time, acceptance rates rise across all baseline trust levels. Except for the low-trust scenario, all other cases converge to nearly identical values, underscoring the critical role of trust-building initiatives for AV adoption.

To investigate other factors influencing fleet performance, we conducted an analysis on the supply-side determinants. For all supply-side analysis, we use the the baseline scenario configuration (see Table 3.5). Specifically, we focus on waiting time and price as these are two factors that can be adjusted by the ride-hailing platform operator. Note that we cannot directly control waiting time in our simulation experiments, so we need to analyze its effects through other variables. Since a higher number of available vehicles decreases average waiting time, we indirectly examine the impact of waiting time by running simulation experiments with different fleet sizes. Figure 3.9 shows the results, indicating that larger fleets improve acceptance rates, particularly in early phases with low AV penetration. Revenues do not increase as much as the acceptance rate in early phases, because most requests are served by HVs that impose high driver payment costs. In later phases, a higher number of vehicles also increases revenue and service quality, but this improvement is generally limited. In other words, after a certain level of availability, fleet performance does not change significantly. Therefore, ride-hailing operators cannot simply maximize performance by expanding fleet size alone.




Figure 3.9: The Impact of Fleet Size on the Fleet Performance

Regarding pricing, Figure 3.10 shows that reducing AV prices has minimal impact in early phases due to the limited number of AVs but increases acceptance rates and lowers revenues in later phases. Interestingly, revenues are highest during the middle phases when AVs and HVs are equally distributed and AV rides are discounted by (10%).

Finally, we examine the impact of communicating with users. In the primary configuration of the hybrid autonomous ride-hailing system, it is assumed that the platform will offer both autonomous and human-driven services when available. The results in Figure 3.11 compares scenarios where users are offered either the closest vehicle (AV or HV) or two options (one AV




Figure 3.10: The Impact of Price Reduction of Autonomous Services

and one HV). In the initial and concluding stages of the analysis, no discernible differences are observed, as the majority of the fleet comprises either HVs or AVs. However, in the middle phases, when both vehicle types are widely available, offering two options increases acceptance rates by accommodating diverse user preferences. This is because users who have different preferences for autonomous and human-based services would accept an offer with a higher probability. However, the revenue impact is smaller, as HV trips generate less profit for the platform.




Figure 3.11: The Impact of Interactive Communication with Users

3.6 Discussion

We investigate how the adoption of AVs into ride-hailing fleets affects their economic, environmental, and service outcomes under heterogeneous user preferences using a multi-method approach. First, we conduct a DCE to assess current user preferences for AVs and HVs under varying attribute levels, including price, waiting time, and trip length. We also examine the influence of user trust in AVs on mode choices. Second, we develop a multi-agent simulation of a hybrid ride-hailing fleet, where AVs and HVs serve a heterogeneous user population through

a digital platform. The simulation parameters are calibrated using real-world taxi trip data from Chicago and the results of DCE. This approach allows us to control for counterfactual variables while analyzing the impact of user behavior on the performance of hybrid fleets during the transition to fully autonomous systems.

The DCE results reveal that ride-hailing users have heterogeneous preferences for AVs versus HVs. We identify four distinct user groups of which the majority (>80%) generally prefers HVs over AVs when attribute levels are not considered. However, all user groups exhibit sensitivity to price and waiting time, offering providers opportunities to adjust these factors to make autonomous services more appealing. Additionally, most users (>80%) do not consider the power source of a vehicle (electric or gasoline) as a key determinant in their mode choice. Hence, while being electric is one advantage AVs can bring for the environment (Ketter et al. 2023), it does not seem to be an influencing factor for most of users' mode choices in a hybrid ride-hailing system. Still, even considering the existing user groups, introducing even a small proportion of AVs (15%) into a fleet can reduce its CO₂ emissions by nearly 40% compared to a non-automated fleet. In terms of economic outcomes, we find that ride-hailing providers' revenues increase logarithmically with AV penetration, assuming a 10% price reduction for AV rides compared to HVs. This price reduction aligns with forecasts suggesting that AVs can be offered at lower costs due to reduced operating expenses (Heineke et al. 2022). However, under the existing user preferences, ride-hailing providers need to make a trade-off between revenues and acceptance rate of users in later stages of AV adoption (AV penetration $\geq 40\%$). While revenues increase, service quality (acceptance rate) decreases as there remains a proportion of users who refrain from using AVs in general. Hence, even though providers' revenues may increase due to reduced operating costs of AVs (Al-Kanj et al. 2020b), they may lose a share of customers if little or no HVs are offered. Yet, it should be noted that these findings are based on current user preferences in a state in which subjects only had some experience with AVs. Of our subject pool, 70% stated that they were slightly familiar with AVs. As AV penetration increases, user familiarity and trust in AVs are expected to grow (Whittle et al. 2019), potentially reducing skepticism. However, similar to other automation technologies, the development of trust may be slow (Lee and See 2004, Ge et al. 2021) and may be undermined by visible system failures (Dietvorst et al. 2014). We, thus, analyze a further scenario in which we consider incrementally increasing user trust in AVs. We find that if user trust is increasing along with the AVs penetration, not only revenues but also the acceptance rate show a positive trend in later stages of AV adoption. As such, even in these stages, the provider could achieve increasing acceptance rates making trust an important factor for customer retention if ride-hailing fleets transition towards being purely autonomous.

Our results have implications for academic research, shared mobility providers, and city planners. One major challenge in encouraging AV adoption is building user trust, consistent with prior research (e.g., Shariff et al. (2017)). If ride-hailing providers proceed with the implemen-

tation of AVs into their fleet too fast such that user trust is still behind, this could lead to users refraining from using their services. Thus, providers should conduct market research to monitor trust levels during the transition to AVs. Fleet performance also varies significantly based on the user population's composition. User groups highly differ in their perception of AVs and willingness to use them. Therefore, platform providers should carefully think about customizing their communication strategy to certain customer groups. Potentially, for those users with low trust in AVs, stressing their safety benefits may at least mitigate some fears of users and lead to a higher acceptance rate. Alternatively, to address environmentally conscious users, platforms could display the CO₂ emissions of each vehicle type to nudge them towards taking an AV or an electric human-driven car and reduce their carbon footprint. From an operations management perspective, we show that a larger fleet at early phases increases the fleet performance due to the higher availability of vehicles that leads shorter waiting time, while not affecting the service quality at later phases. Additionally, price reductions for AV rides can boost acceptance rates, but operators must balance these discounts with revenue optimization. Dynamic pricing schemes that account for vehicle availability and demand fluctuations could enhance fleet profitability. However, indiscriminate price reductions may attract users away from public transportation, with potentially adverse effects on urban mobility. City planners should therefore regulate AV pricing to align with sustainability goals. Finally, platform providers should not only invest in the development of AVs but also in the application technology which connects customers to vehicles. We have demonstrated that through a simple change of operation mode, namely offering two vehicle options, the fleet provider's profitability can improve even when a share of customers is technologically skeptical. This change in the platform's customer communication design outperforms the status quo design of current ride-hailing platforms regarding its profitability.

To our knowledge, this study is the first to incorporate user behavior toward AVs into the analysis of their phased introduction in ride-hailing fleets. By doing so, we extend research on user behavior and interaction with AI occurring on the next generation of digital platforms (Rai et al. 2019). Further, we establish a simulation environment which serves as a testbed for a range of variations in future ride-hailing markets. Thus, we also contribute to the literature on mobility transition and green IS. However, our work does not come without limitations. The current simulation does not account for varying types of charging infrastructure or fluctuating fleet sizes. In reality, the speed and availability of charging infrastructure can significantly impact fleet performance, as EV range remains a key limitation (Al-Kanj et al. 2020b). Furthermore, while AVs reduce operating costs, providers must initially invest in their acquisition (Heineke et al. 2022), which we did not include in our analysis. Future research should explore the implications of these initial costs on the profitability of an autonomous ride-hailing fleet. We also assume that drivers behave in a uniform and predictable manner by accepting all customers' requests. In practice, drivers may reject requests based on factors such as the distance to the

pick-up point or the expected fare. We encourage future research to shed light on the influence of driver behavior on ride-hailing platforms transitioning towards autonomous services.

3.7 Appendix

3.7.1 Design of the Discrete Choice Experiment

Subjects were instructed that they should imagine they plan to go on a trip similar to their most recent ride-hailing trip. They were told that they were using their ride-hailing app to find the nearest car available and two options were offered to them which are the only options available for the presented trips. Subjects' task was to choose one of the options offered.

Explanation of Vehicle and Trip Attributes

Each scenario was described by the duration of the trip which is the same for both of the options. Additionally, the proposed vehicle options possess attributes which may differ between them (price, waiting time, power).

Vehicle Options

Human-driven vehicle: The car will be driven by a human driver registered at the ride-hailing provider.

Autonomous vehicle: Autonomous vehicles can drive by themselves. No human driver will be present in the car, and you will not be able to drive the car yourself.

Scenario-related descriptions

Trip duration: This attribute describes the time (in minutes) the ride will take from your origin to your destination.

Option-related descriptions:

Price: Shows the costs for the proposed vehicle option in US\$. Please imagine you would need to spend your own money for the trip.

Waiting time: Shows the time (in minutes) that it will take until your chosen vehicle option will arrive to pick you up at your origin.

Power: Shows how the vehicle offered to you for the trip will be powered. It can take one of two options:

- a) Electric: The vehicle is solely powered by electricity, using a battery
- b) Gasoline: The vehicle is solely powered by gasoline and has an internal combustion engine

Choice Set Design

Depicted in Figure 1 is an example of a choice set. Subjects were asked to choose one of the alternatives. We combined the choice sets using d-efficiency as a criterion and show all combinations in Table 1.

Block	Choice Set	Preis	Waiting time	Power	Trip duration	Mode
1	1	8.8	5	0	10	AV
1	1	12.1	5	0	10	HV

Block	Choice Set	Preis	Waiting time	Power	Trip duration	Mode
1	2	12.0	5	0	20	AV
1	2	15.0	5	0	20	HV
1	3	9.9	7	0	10	AV
1	3	12.1	7	1	10	HV
1	4	20.0	5	0	30	AV
1	4	25.0	5	0	30	HV
1	5	22.5	7	0	30	AV
1	5	27.5	5	0	30	HV
1	6	13.5	7	0	20	AV
1	6	15.0	7	1	20	HV
2	7	13.5	5	0	20	AV
2	7	16.5	7	1	20	HV
2	8	13.5	10	0	20	AV
2	8	16.5	5	1	20	HV
2	9	9.9	5	0	10	AV
2	9	12.1	7	1	10	HV
2	10	9.9	10	0	10	AV
2	10	12.1	5	1	10	HV
2	11	22.5	5	0	30	AV
2	11	30.0	7	0	30	HV
2	12	25.0	5	0	30	AV
2	12	25.0	10	0	30	HV
3	13	15.00	5	0	20	AV
3	13	16.50	7	0	20	HV
3	14	13.50	7	0	20	AV
3	14	18.00	5	0	20	HV
3	15	22.50	5	0	30	AV
3	15	27.50	5	1	30	HV
3	16	11.00	5	0	10	AV
3	16	12.10	7	0	10	HV
3	17	22.50	7	0	30	AV
3	17	25.00	7	1	30	HV
3	18	9.90	7	0	10	AV
3	18	13.20	5	0	10	HV
4	19	15.00	7	0	20	AV
4	19	15.00	10	1	20	HV
4	20	12.00	5	0	20	AV

Block	Choice Set	Preis	Waiting time	Power	Trip duration	Mode
4	20	18.00	10	1	20	HV
4	21	25.00	10	0	30	AV
4	21	25.00	5	1	30	HV
4	22	25.00	7	0	30	AV
4	22	30.00	5	0	30	HV
4	23	11.00	7	0	10	AV
4	23	12.10	10	1	10	HV
4	24	8.80	5	0	10	AV
4	24	13.20	10	1	10	HV
5	25	11.00	10	0	10	AV
5	25	12.10	5	0	10	HV
5	26	15.00	10	0	20	AV
5	26	18.00	7	1	20	HV
5	27	15.00	10	0	20	AV
5	27	15.00	5	0	20	HV
5	28	25.00	7	0	30	AV
5	28	27.50	10	1	30	HV
5	29	20.00	10	0	30	AV
5	29	30.00	5	1	30	HV
5	30	11.00	10	0	10	AV
5	30	13.20	7	1	10	HV
6	31	8.80	7	0	10	AV
6	31	12.10	5	1	10	HV
6	32	8.80	10	0	10	AV
6	32	12.10	7	0	10	HV
6	33	12.00	7	0	20	AV
6	33	16.50	5	1	20	HV
6	34	12.00	10	0	20	AV
6	34	15.00	7	0	20	HV
6	35	20.00	7	0	30	AV
6	35	30.00	10	0	30	HV
6	36	20.00	5	0	30	AV
6	36	27.50	10	1	30	HV
7	37	8.80	7	0	10	AV
7	37	13.20	7	0	10	HV
7	38	8.80	10	0	10	AV
7	38	12.10	10	1	10	HV

Block	Choice Set	Preis	Waiting time	Power	Trip duration	Mode
7	39	12.00	7	0	20	AV
7	39	18.00	7	0	20	HV
7	40	12.00	10	0	20	AV
7	40	16.50	10	1	20	HV
7	41	20.00	7	0	30	AV
7	41	25.00	7	1	30	HV
7	42	20.00	10	0	30	AV
7	42	27.50	7	0	30	HV
8	43	11.00	7	0	10	AV
8	43	12.10	10	0	10	HV
8	44	9.90	5	0	10	AV
8	44	12.10	10	0	10	HV
8	45	15.00	7	0	20	AV
8	45	16.50	10	0	20	HV
8	46	13.50	5	0	20	AV
8	46	15.00	10	0	20	HV
8	47	25.00	5	0	30	AV
8	47	30.00	7	1	30	HV
8	48	25.00	10	0	30	AV
8	48	27.50	7	0	30	HV
9	49	11.00	5	0	10	AV
9	49	13.20	5	1	10	HV
9	50	9.90	10	0	10	AV
9	50	13.20	10	0	10	HV
9	51	15.00	5	0	20	AV
9	51	18.00	5	1	20	HV
9	52	13.50	10	0	20	AV
9	52	18.00	10	0	20	HV
9	53	22.50	10	0	30	AV
9	53	30.00	10	1	30	HV
9	54	22.50	10	0	30	AV
9	54	25.00	10	0	30	HV

Table 3.6: Experimental Design Table

Imagine you want to use a ride-hailing service for this trip with the following trip duration:

Duration of the trip	30 min
----------------------	--------

You are using your ride-hailing app to find available vehicles and the following are offered to you. Among the following travel options, which one do you prefer?

	Autonomous Vehicle	Human-driven Vehicle
Price	\$ 22.50	\$ 30.00
Waiting time	10 min	10 min
Power	Electric	Gasoline

	Autonomous Vehicle	Human-driven Vehicle
Your choice:	<input type="radio"/>	<input type="radio"/>

Figure 3.12: Exemplary Choice Set

3.7.2 Descriptive Statistics Discrete Choice Experiment

Variables	
Age	
Mean	29.40
Standard Deviation	8.69
Trust	
Mean	4.38
Standard Deviation	1.29
Technology Interest	
Mean	4.10
Standard Deviation	0.96
Environmental Friendliness	
Mean	4.14
Standard Deviation	0.71

Table 3.7: Demographics and Questionnaire Variables (Continuous)

Variables	Shares
Gender	
Female	53.28%
Male	45.55%
Diverse	0.17%
Preferred not to report	0.67%
Ride-hailing frequency (past 12 months)	
0 times	3.03%
1-3 times	22.35%
4-6 times	22.02%
7-9 times	13.45%
10 times or more	39.16%
Car owner	
Yes	58.99%
No	41.00%
Education	
Less than high school	0.17%
High school graduate	14.62%
College	17.98%
Bachelor's degree	45.55%
Master's degree	18.32%
Professional degree	1.85%
Doctorate	1.18%
Preferred not to report	0.34%
Income	
Less than \$10,000	32.10%
\$10,000 - \$19,999	22.86%
\$20,000 - \$29,999	13.78%
\$30,000 - \$39,999	8.57%
\$40,000 - \$49,999	5.04%
\$50,000 - \$59,999	3.36%
\$60,000 - \$69,999	1.85%
\$70,000 - \$79,999	1.85%
\$80,000 - \$89,999	0.84%
\$90,000 - \$99,999	0.67%
More than \$100,000	2.18%
Prefer not to report	6.89%

Variables	Shares
------------------	---------------

Table 3.8: Demographics and Questionnaire Variables (Categorical)

Chapter 4

Data-driven Planning of Large-Scale Electric Vehicle Charging Hubs using Deep Reinforcement Learning¹

4.1 Introduction

With the proliferation of electric vehicle (EVs) arises the need for charging infrastructure that enables users to make the switch to EVs with minimal impact on lifestyle and behavior. Policy-makers have traditionally assumed that users would primarily charge their EVs overnight and at home. Indeed, home charging is currently the preeminent charging use case in many markets (Lee et al. 2020, Hoover et al. 2021). As more and more consumers without access to residential charging adopt EVs, charging opportunities at the workplace, at popular destinations such as supermarkets, and at fleet depots are needed (Jun and Meintz 2018, Lee et al. 2019, 2020). We refer to the systems that afford such high-density EV charging use cases as EV Charging Hubs (EVCHs). Apart from enabling widespread EV adoption, EVCHs can also play an important systems integration role by enabling daytime charging that takes advantage of high solar energy production, which is unavailable when charging overnight (Lee et al. 2018)². EVCHs constitute a novel (and under-researched) operational system class with cross-system interfaces (e.g., with attached buildings or the electricity grid) and a large number of strategic and operational de-

¹This Chapter is currently under review (second round) at a leading peer-reviewed academic journal.

Parts of this Chapter have appeared in the following (non-copyrighted) peer-reviewed academic conferences and workshops: Schroer, K., Ahadi, R., Lee, T. Y., & Ketter, W. (2021). Preference-aware Planning and Operations of Electric Vehicle Charging Clusters : A Data-Driven Prescriptive Framework. In Proceedings of the SIG GREENWorkshop (pp. 1–10).

Schroer, K., Ahadi, R., Lee, Y.T., & Ketter, W. (2021). Preference-Aware Planning and Operations of Electric Vehicle Charging Clusters: A Prescriptive Framework. In Workshop on Information Systems and Technology (WITS) 2021 (Austin, TX).

²This is particularly relevant for energy systems with high solar energy share such as California or Germany

cision variables (size and configuration of charging stations, on-site storage, charging decisions, etc.) that result in a highly complex planning challenge (Ferguson et al. 2018).

Operations managers traditionally approach the strategic planning of operational systems like EVCHs via mathematical programming methods (e.g., He et al. 2017, for on-demand vehicle sharing service region design) or queuing models (e.g., Wang et al. 2016, for last-mile delivery networks). However, the multi-stage and stochastic nature of the EVCH planning challenge makes it notoriously challenging for optimization-based methods (Powell 2014, Hannah 2015). Computational tractability in the traditional frameworks is only achieved at the expense of detail and scope of the planning problem. For example, shorter planning horizons, coarser temporal discretization, simplified operational detail or deterministic parameter assumptions are adopted to significantly reduce problem sizes.

This work develops a novel method that makes use of the fine-grained operational and preference data that has become abundant in this age of pervasive IoT³ sensor technology. As such it responds to calls from the Operations Management (OM) community to incorporate such data into OM frameworks (Qi and Shen 2018, Cohen 2018, Choi et al. 2022) and data-driven decision support systems (Ketter et al. 2023). Specifically, we leverage fine-grained sensor data from parking lots and energy consumption and production data and combine it with high-resolution asset models and real-world operational policies into a detailed simulated environment. This environment is a close-to-exact digital representation – i.e., a Digital Twin (DT) (Choi et al. 2022, Grieves and Vickers 2017) – of the EVCH that is to be planned. We then develop an actor-critic reinforcement learning (RL) framework that interacts with this environment to learn an optimal planning configuration policy by iterating over many simulated epochs.

Our work offers a number of contributions. Methodologically, we propose a framework for the effective use of RL in combination with large-scale data-driven simulation frameworks (i.e., DTs) for ex-ante de-risking and decision support in the design phase of service systems such as EVCHs. Our method circumvents the need for simplification and problem size reduction, among other theoretical benefits. These include (1) more realistic, data-driven modeling of stochasticity and operational detail of the EVCH, (2) computational scalability compared to mathematical optimization, and (3) flexible model setup that allows for easy evaluation of many different operational policies. In extensive simulation experiments, we show that our method achieves near-optimal EVCH planning results. We also show that it outperforms alternative candidate solution approaches such as DQN and DDPG in terms of solution speed and scalability. Finally, we make use of the flexible nature of the DT to evaluate different preference and operational regimes, thus deriving numerous novel domain-specific insights for practitioners.

The remainder of this work is structured as follows. In Section 4.2, we review the relevant literature to our work. We then set up and parameterize our model for data-driven planning of Electric Vehicle Charging Hubs using actor-critic RL (Section 4.3). In Section 4.4 we evaluate

³Internet of Things

our model in terms of its ability to converge to the global optimum solution and its scalability and performance characteristics versus other candidate solutions, such as Deep Q-Learning. We then use the evaluated model to run comprehensive scenario analyses (Section 4.5) to (1) demonstrate flexibility benefits and (2) to obtain interesting and actionable policy insights. We end with a discussion of implications and contributions to theory and practice (Section 4.6).

Operations managers traditionally approach the strategic planning of operational systems like EVCHs via mathematical programming methods (e.g., He et al. 2017, for on-demand vehicle sharing service region design) or queuing models (e.g., Wang et al. 2016, for last-mile delivery networks). However, the multi-stage and stochastic nature of the EVCH planning challenge makes it notoriously challenging for optimization-based methods (Powell 2014, Hannah 2015). Computational tractability in the traditional frameworks is only achieved at the expense of detail and scope of the planning problem. For example, shorter planning horizons, coarser temporal discretization, simplified operational detail or deterministic parameter assumptions are adopted to significantly reduce problem sizes.

This work develops a novel method that makes use of the fine-grained operational and preference data that has become abundant in this age of pervasive IoT⁴ sensor technology. As such it responds to calls from the Operations Management (OM) community to incorporate such data into OM frameworks (Qi and Shen 2018, Cohen 2018, Choi et al. 2022) and data-driven decision support systems (Ketter et al. 2023). Specifically, we leverage fine-grained sensor data from parking lots and energy consumption and production data and combine it with high-resolution asset models and real-world operational policies into a detailed simulated environment. This environment is a close-to-exact digital representation – i.e., a Digital Twin (DT) (Choi et al. 2022, Grieves and Vickers 2017) – of the EVCH that is to be planned. We then develop an actor-critic reinforcement learning (RL) framework that interacts with this environment to learn an optimal planning configuration policies by iterating over many many simulated epochs.

Our work offers a number of contributions. Methodologically, we propose a framework for the effective use of RL in combination with large-scale data-driven simulation frameworks (i.e., DTs) for ex-ante de-risking and decision support in the design phase of service systems such as EVCHs. Our method circumvents the need for simplification and problem size reduction, among other theoretical benefits. These include (1) more realistic, data-driven modeling of stochasticity and operational detail of the EVCH, (2) computational scalability compared to mathematical optimization, and (3) flexible model setup that allows for easy evaluation of many different operational policies. In extensive simulation experiments, we show that our method achieves near-optimal EVCH planning results. We also show that it outperforms alternative candidate solution approaches such as Deep Q-Learning in terms of solution speed and scalability. Finally, we make use of the flexible nature of the DT to evaluate different preference and operational regimes, thus deriving numerous novel domain-specific insights for practitioners.

⁴Internet of Things

The remainder of this work is structured as follows. In Section 4.2, we review the relevant literature to our work. We then set up and parameterize our model for data-driven planning of Electric Vehicle Charging Hubs using actor-critic RL (Section 4.3). In Section 4.4 we evaluate our model in terms of its ability to converge to the global optimum solution and its scalability and performance characteristics versus other candidate solutions, specifically Deep Q-Learning. We then use the evaluated model to run comprehensive scenario analyses (Section 4.5) to (1) demonstrate flexibility benefits and (2) to obtain interesting and actionable policy insights. We end with a discussion of implications and contributions to theory and practice (Section 4.6).

4.2 Background

Our work draws from three main bodies of literature, which we briefly review here. First, we discuss the problem class of EVCH planning and review traditional OM planning approaches. Second, we discuss reinforcement learning (RL) methods and their potential benefits for complex, multi-stage, stochastic planning problems like EVCH planning. Third, we review the extant work on DTs and their use for OM decision support.

4.2.1 Electric Vehicle Charging Hubs (EVCHs)

EV charging operations environments and use cases vary from fully distributed on-street charging, highway charging, and private home charging to charging in large-scale high-density parking lots. In this work, we focus on the latter use case, which we refer to as an EV charging hub (EVCH). EVCHs exhibit several unique features that distinguish them from other charging use cases. First, EVCHs typically represent large locally concentrated loads that may require significant local electricity grid extension making load shaping necessary (Lee et al. 2019). Second, integration with behind-the-meter loads (buildings) and generation units (PV, storage) may be desirable (Nunes et al. 2016) to reduce induced peak loads, drive sustainability and reduce costs (Ferguson et al. 2018). Third, EVCHs typically experience different user behavior compared to other charging use cases such as home charging, and this user behavior can vary substantially depending on the use case of the attached facility (workplace, mall, etc.). Fourth, siting of individual charging stations is of no concern in an EVCH context as all chargers will be located in the same space with users being largely indifferent between them. Finally, EVCHs allow for end-to-end control of the full vehicle-level parking and charging journey through what is sometimes referred to as smart EV-capable parking lots (Babic et al. 2022a). This enables the assignment of vehicles to chargers and central control over the charging process. It, thus, offers new scope for optimization, e.g., by leveraging parallel or sequential use of charging equipment in an optimal manner (Ferguson et al. 2018).

We briefly review state-of-the-art OM approaches in the realms of (1) operating and (2) planning EVCHs. In terms of EVCH operations, we acknowledge the extensive work on electric

vehicle charging scheduling and smart charging (see e.g., Mukherjee and Gupta (2015) for a recent review) on which most operations-focused EVCH research is based. A notable differentiator from the traditional smart charging literature is the inclusion of building/cluster-level constraints and optimization opportunities. Early examples include Huang and Zhou (2015) who develop a mixed-integer optimization framework for workplace charging strategies taking into account different eligibility levels and Wu et al. (2017) who propose a two-stage energy management framework for office buildings with workplace EV charging. Nunes et al. (2016) investigate how charging processes can best be coordinated to use parking lots for EV solar-charging. Ferguson et al. (2018) propose an integrated load management approach to optimize EV charging processes for minimum cost taking into account the building base load and PV generation. A similar approach to site-level load management was implemented in practice by Jun and Meintz (2018). Finally, Lee et al. (2019) explore several optimization-driven approaches to operational issues in charging hubs. Note that the inclusion of parallel-use charging docks that allow for simultaneous charging significantly complicates that EVCH management problem. In addition to the usual charging decisions, an assignment decision of vehicles to charging stations is required. Our notion of EVCHs considers this complication.

The design/planning of EVCH systems has received less attention. EVCH design is a multi-stage stochastic decision problem that requires large decisions (e.g., the number of charging docks to be installed at each stage in the planning horizon) and small decisions (e.g., charging individual vehicles) to be taken simultaneously. Such problems are notoriously difficult and cannot be solved efficiently with standard stochastic programming or even approximate dynamic programming (Powell 2014, Hannah 2015). Some research resolves the ensuing complexity using simulation-based approaches. For example, in Kazemi et al. (2016) the authors use a genetic search algorithm on top of a simplified simulation model to derive the optimal size of an EV parking lot. Babic et al. (2022a) also use a greedy search over a simulation of a parking lot to derive optimal infrastructure decisions. Naturally, optimality cannot be guaranteed with simple search approaches. Li et al. (2020) propose a mathematical deterministic programming framework for the joint optimization of the size and operations of a parking lot capable of 100 electric vehicles. Neither of these simulation- or optimization-based studies use high-granularity demand and/or operational data. In addition, extant EVCH design work exhibits relies on significant simplifications. For example, the studies cited here focus on a single planning period only, which reduces the problem to a single-stage planning challenge. In addition, the EVCH system scope tends to be considerably simplified (e.g., no consideration of attached building loads, single-use charging docks only, etc.).

Our work addresses several important gaps in the charging hub literature. We are the first to use detailed preference modeling in an extensive and novel set of real-world parking and charging data to ensure preference-aware sizing. In doing so, we explore the sensitivity of planning decisions to changes in user preferences, a point that has been completely neglected

by existing work. In addition we consider existing building load profiles in the operations and investment decision, taking a more comprehensive view compared to previous research. Our model also allows for parallel use of charging infrastructure which can significantly boost asset efficiency at the expense of higher operational complexity. Finally, our work has important social and sustainability implications insofar as it proposes a model for efficient provisioning of charging infrastructures that is aligned with customer preferences.

4.2.2 Reinforcement Learning and its Application to Planning Problems

Reinforcement learning (RL) represents a distinct class of machine learning that seeks to find an optimal policy which governs the behavior of an agent in an (emulated) environment such that a given objective is maximized (Sutton 2019). RL relies on the Markov property, meaning that future states in a stochastic process only depend on the current state (Sutton 2019). Popular examples of RL include an agent playing the game of Go (Silver et al. 2016) or Atari games (Mnih 2013). A policy, the goal of RL, can be understood as a function that takes an observed environment state as input and returns an action given the observed state. That policy is learned iteratively by interacting with the emulator through actions and observing the effect of these actions.

RL has received significant interest as a possible approach for dynamic optimization problems (Fu et al. 2015). Indeed, the method boasts several potential advantages over traditional mathematical programming approaches. First, RL does not require a model but instead relies on a simulated environment to interact with and learn from. This can be a major advantage, particularly in complex multi-stage settings (like EVCH planning) where developing a mathematical model that accurately reflects the behavior of physical assets, operational policies and individual preferences is impossible or extremely hard. This also means that RL requires fewer assumptions. In mathematical programming simplifying assumptions (e.g., coarser discretization, simplified operations, removing stochasticity of inputs, etc.) are often needed to achieve tractability resulting in a problem formulation that may not reflect reality accurately enough. Second, RL is flexible and readily adapts to changes in environmental conditions. Third, RL generally deals better with the curse of dimensionality and can scale to real-sized problems well beyond the tractability limits of optimization frameworks (van Hezewijk et al. 2022). RL also has disadvantages, most notably a lack of optimality guarantee. The real-world applicability of the RL solution will also have strong dependence on the quality of the emulator that is used for training.

RL has successfully been applied to a range of operational and strategic planning problems (Gijsbrechts et al. 2022). For example, van Hezewijk et al. (2022) examine the applicability of Proximal Policy Optimisation (PPO), a deep RL algorithm, to the stochastic capacitated lot sizing problem and show that the algorithm converges close to the global optimum and readily scales to problem sizes that are out of scope for traditional dynamic programming. Ahadi

et al. (2023) study the charging management of shared autonomous electric vehicles using a cooperative multi-agent reinforcement algorithm to simultaneously learn optimal scheduling and resource allocation policies. In a similar vein, Xie et al. (2023) consider a hybrid ride-hailing fleet of autonomous vehicles and conventional drivers and optimise the relocation policies using a two-sided deep RL design where the fleet operator makes central relocation decisions for autonomous vehicles and individual driver agents learn their non-cooperative relocation strategies.

There are two distinct routes for estimating the optimal policy in RL: (1) value-based approaches and (2) policy-based approaches. Value-based approaches estimate the total value associated with an action assuming the agent follows a given policy forever (e.g., the greedy policy of always selecting the action with the highest value). A value-based algorithm that has seen significant adoption is Deep Q-Learning (Gao et al. 2020). Deep Q-Learning uses Deep Q Networks (DQN) to estimate state-action values in a discrete action space. For a given state the DQN returns a Q-value for every possible action and the agent will pick a random (in the exploration phase) or the highest-valued action (in the exploitation phase). It is easy to see that DQN (and other value-based methods) can run into issues of scalability, especially if the action space is very large or even continuous, corresponding to a potentially infinite number of permutations of actions that each need to be evaluated for a given state (Dulac-Arnold et al. 2015). Policy-based methods, such as policy gradient-based algorithms can circumvent this issue and can work well in continuous action spaces (Sutton 2019). These methods estimate the policy function directly (typically using gradient descent-based optimization) without the need to evaluate each possible state-action pair. However, they tend to be inefficient and are susceptible to local optima as well as high variance (Sutton 2019). Actor-critic RL approaches combine value-based and policy-based approaches bringing together the benefits of both. Actor-critic algorithms consist of two main parts. First, the actor that takes decisions based on a learned policy function (policy-based). Second, the critic that determines the quality of the action using a value function (value-based). This actor-critic setup allows the actor to improve its policies more efficiently compared to pure policy-based methods. In this work, we use SAC (soft actor-critic) (Haarnoja et al. 2018), an actor-critic framework, which we adapt to work with large discrete action spaces (Dulac-Arnold et al. 2015).

4.2.3 Digital Twins (DTs) and their use in Operations Management

As mentioned, RL frameworks require a simulated environment (or simulator) to interact with and learn from. The closer this simulation comes to reality, the more likely the RL-derived solution is generalize to real-world conditions. Hence, researchers have called for the use of real-world system data to achieve more accurate representations of real-world conditions (Panzer and Bender 2022). Digital Twins (DTs), a form of data-driven simulation, provide an attractive solution. DTs have been hailed as a disruptive trend in OM of the IoT and Industry 4.0 era (Choi et al. 2022). At its most basic level, a DT is a digital representation of a specific physical asset

or system. DTs can be used as a decision support tool along the entire life cycle of that asset: from design, operations, and maintenance to disposal (Schleich et al. 2017). For the purpose of this research, we highlight several characteristics that we consider key and distinguish DTs from, e.g., traditional simulation frameworks used in OM (Boschert and Rosen 2016) and RL. The reader is referred to Jones et al. (2020) or Cimino et al. (2019) for comprehensive reviews.

First and foremost, DTs represent a close to the real-world representation of the physical system at a granular level using high-fidelity interconnected physical models of system components (Glaessgen and Stargel 2012). Note that this does not necessarily require identical accuracy but can also mean partial accuracy if this is sufficient for the DT to fulfill its intended use (van der Valk et al. 2020). Second, a DT is data-driven, meaning that it primarily relies on real-world operational data input acquired directly from sensors of the device and the intended application environment (van der Valk et al. 2020). It is often not the case that raw data on all required parameters is available. In such cases, synthetic data from statistical models or simulation frameworks can be used to supplement the data requirements (Sierla et al. 2018). Third, there is eventual synchronization of the DT and the physical asset. This can be achieved via one-directional (physical world to DT) or bi-directional data flows (Tao et al. 2018). DTs have been successfully applied in the use phase of operational systems (Jones et al. 2020). Examples include asset status and health monitoring (Glaessgen and Stargel 2012), asset optimization, maintenance planning (Cimino et al. 2019), and staff training (Choi et al. 2022).

As argued previously, a core feature of DTs is their reliance on real-world sensor data for asset and process representation in the virtual world. Use of real-world environmental context data and benchmark infrastructure sensor data is also useful for design/configuration applications of assets and service systems (Attaran and Celik 2023). Indeed, high-granularity data on many aspects of the intended application environment as well as on typical machine/process behavior is likely to already be available. This means that pre-use-phase DTs can be fed with live as well as realistic historic data (Boschert and Rosen 2016).

In this work, we leverage the DT concept to create a high-resolution, data-driven simulation environment that resembles real-world conditions as closely as possible. Wherever available, we leverage historic sensor data to model system dynamics. For example, we use real-world charging data, parking data, building data and photovoltaic production data to achieve granular patterns of expected power production and consumption in the EVCH. For system components where real-world sensor data is not available, we rely on close-to-exact simulations of their physical behavior in line with asset specification sheets and/or research findings. For example, to model the physical properties of an on-site battery storage asset, we draw on research by Ghiassi-Farrokhfal et al. (2016) to parameterize battery charge and discharge efficiency as well as minimum and maximum (dis-)charge rates. Hence, we follow Sierla et al. (2018) to supplement our DT simulation with data from simulation frameworks to fulfill our data requirements.

We show that DTs, in combination with scalable RL algorithms, can circumvent many of the model simplifications and problem reduction measures unavoidable in traditional mathematical programming or brute-force search strategies.

Our work addresses several important gaps in the charging hub literature. We are the first to use detailed preference modeling on an extensive and novel set of real-world parking and charging data to ensure preference-aware sizing. In doing so we explore the sensitivity of planning decisions to changes in user preferences, a point that has been completely neglected by existing work. In addition we consider existing building load profiles in the operations and investment decision, taking a more comprehensive view compared to previous research. Our model also allows for parallel use of charging infrastructure which can significantly boost asset efficiency at the expense of higher operational complexity. Finally, our work has important societal and sustainability implications insofar as it puts forward a model for efficient provisioning of charging infrastructures that is aligned with customer preferences.

4.2.4 Reinforcement Learning and its Application to Planning Problems

Reinforcement learning (RL) represents a distinct class of machine learning that seeks to find an optimal policy which governs the behavior of an agent in an (emulated) environment such that a given objective is maximized (Sutton 2019). RL relies on the Markov property, meaning that future states in a stochastic process only depend on the current state (Sutton 2019). Popular examples of RL include an agent playing the game of Go (Silver et al. 2016) or Atari games (Mnih 2013). A policy, the goal of RL, can be understood as a function that takes an observed environment state as input and returns an action given the observed state. That policy is learned iteratively by interacting with the emulator through actions and observing the effect of these actions.

RL has received significant interest as a possible approach for dynamic optimization problems (Fu et al. 2015). Indeed, the method boasts several potential advantages over traditional mathematical programming approaches. First, RL does not require a model but instead relies on a simulated environment to interact with and learn from. This can be a major advantage, particularly in complex multi-stage settings (like EVCH planning) where developing a mathematical model that accurately reflects the behavior of physical assets, operational policies and individual preferences is impossible or extremely hard. This also means that RL requires fewer assumptions. In mathematical programming simplifying assumptions (e.g., coarser discretization, simplified operations, removing stochasticity of inputs, etc.) are often needed to achieve tractability resulting in a problem formulation that may not reflect reality accurately enough. Second, RL is flexible and readily adapts to changes in environmental conditions. Third, RL generally deals better with the curse of dimensionality and can scale to real-sized problems well beyond the tractability limits of optimization frameworks (van Hezewijk et al. 2022). RL also has disadvantages, most notably a lack of optimality guarantee. The real-world applicability of

the RL solution will also have strong dependence on the quality of the emulator that is used for training.

RL has successfully been applied to a range of operational and strategic planning problems (Gijsbrechts et al. 2022). For example, van Hezewijk et al. (2022) examine the applicability of Proximal Policy Optimisation (PPO), a deep RL algorithm, to the stochastic capacitated lot sizing problem and show that the algorithm converges close to the global optimum and readily scales to problem sizes that are out of scope for traditional dynamic programming. Ahadi et al. (2023) study the charging management of shared autonomous electric vehicles using a cooperative multi-agent reinforcement learning to simultaneously learn optimal scheduling and resource allocation policies. In a similar vein, Xie et al. (2023) consider a hybrid ride-hailing fleet of autonomous vehicles and conventional drivers and optimise the relocation policies using a two-sided deep RL design where the fleet operator makes central relocation decisions for autonomous vehicles and individual driver agents learn their non-cooperative relocation strategies.

There are two distinct routes for estimating the optimal policy in RL: (1) value-based approaches and (2) policy-based approaches. Value-based approaches estimate the total value associated with an action assuming the agent follows a given policy for ever (e.g., the greedy policy of always selecting the action with the highest value). A value-based algorithm that has seen significant adoption is Deep Q-Learning (Gao et al. 2020). Deep Q-Learning uses Deep Q Networks (DQN) to estimate state-action values in discrete action space. For a given state the DQN returns a Q-value for every possible action and the agent will pick a random (in the exploration phase) or the highest-valued action (in the exploitation phase). It is easy to see that DQN (and other value-based methods) can run into issues of scalability, especially if the action space is very large or even continuous, corresponding to a potentially infinite number of permutations of actions that each need to be evaluated for a given state (Dulac-Arnold et al. 2015). Policy-based methods, such as policy gradient-based algorithms can circumvent this issue and can work well in continuous action spaces (Sutton 2019). These methods estimate the policy function directly (typically using gradient descent-based optimization) without the need to evaluate each possible state-action pair. However, they tend to be inefficient and are susceptible to local optima as well as high variance (Sutton 2019). Actor-critic RL approaches combine value-based and policy-based approaches bringing together the benefits of both. Actor-critic algorithms consist of two main parts. First, the actor that takes decisions based on a learned policy function (policy-based). Second, the critic that determines the quality of the action using a value function (value-based). This actor-critic setup allows the actor to improve its policies more efficiently compared to pure policy-based methods. In this work, we use SAC (soft actor-critic) (Haarnoja et al. 2018), an actor-critic framework, which we adapt to work with large discrete action spaces (Dulac-Arnold et al. 2015).

4.2.5 Digital Twins (DTs) and their use in Operations Management

As mentioned, RL frameworks require a simulated environment (or simulator) to interact with and learn from. The closer this simulation comes to reality, the better the RL-derived solution is likely to generalize to real-world conditions. Hence, researchers have called for the use of real-world system data to achieve more accurate representations of real-world conditions (Panzer and Bender 2022). Digital Twins (DTs), a form of data-driven simulation, provide an attractive solution. DTs have been hailed as a disruptive trend in OM of the IoT and Industry 4.0 era (Choi et al. 2022). At its most basic level, a DT is a digital representation of a specific physical asset or system. DTs can be used as a decision support tool along the entire life cycle of that asset: from design, operations, and maintenance to disposal (Schleich et al. 2017). For the purpose of this research, we highlight several characteristics that we consider key and distinguish DTs from, e.g., traditional simulation frameworks used in OM (Boschert and Rosen 2016) and RL. The reader is referred to Jones et al. (2020) or Cimino et al. (2019) for comprehensive reviews.

First and foremost, DTs represent a close to the real-world representation of the physical system at a granular level using high-fidelity interconnected physical models of system components (Glaessgen and Stargel 2012). Note that this does not necessarily require identical accuracy but can also mean partial accuracy if this is sufficient for the DT to fulfill its intended use (van der Valk et al. 2020). Second, a DT is data-driven, meaning that it primarily relies on real-world operational data input acquired directly from sensors of the device and the intended application environment (van der Valk et al. 2020). It is often not the case that raw data on all required parameters is available. In such cases, synthetic data from statistical models or simulation frameworks can be used to supplement the data requirements (Sierla et al. 2018). Third, there is eventual synchronization of the DT and the physical asset. This can be achieved via one-directional (physical world to DT) or bi-directional data flows (Tao et al. 2018). DTs have been successfully applied in the use phase of operational systems (Jones et al. 2020). Examples include asset status and health monitoring (Glaessgen and Stargel 2012), asset optimization, maintenance planning (Cimino et al. 2019), and staff training (Choi et al. 2022).

As argued previously, a core feature of DTs is their reliance on real-world sensor data for asset and process representation in the virtual world. That does not, however, disqualify DTs from the use in design/configuration applications of assets and operational systems that are yet to be built. Indeed, high-granularity data on many aspects of the intended application environment as well as on typical machine/process behavior is likely to already be available. This means that pre-use-phase DTs can be fed with live as well as realistic historic data (Boschert and Rosen 2016).

In this work, we leverage the DT concept to create a high-resolution, data-driven simulation environment that resembles real-world conditions as closely as possible. We show that DTs, in combination with scalable RL algorithms, can circumvent many of the model simplifications

and problem reduction measures unavoidable in traditional mathematical programming or brute-force search strategies.

4.3 Model

We now describe our model. A conceptual overview of its core elements and their interactions is shown in Figure 4.1. In this Section we describe the setup of the model starting with a definition of the planning problem including state boundaries, action space and objective followed by a description of the environment simulator (DT) and the SAC RL framework.




Figure 4.1: Overview of Core Model Elements, Inputs and Interactions

4.3.1 Defining the EVCH Planning Scope

We define an EVCH as an EV charging-capable parking lot, depot, or garage that is typically attached to an existing building with a given baseload. Both the building and the EVCH receive power from the same grid connection point, which is constrained to the capacity of the on-site substation. The integrated facility may have additional on-site behind-the-meter generation (e.g., photo-voltaic (PV)) and storage (e.g., Lithium-Ion battery). Crucially, charging docks can have multiple connectors that afford parallel charging of vehicles from a single charging dock. A simplified view of the EVCH components and system boundary is depicted in Figure 4.2.

We formulate the EVCH configuration challenge as a feasibility problem that aims to satisfy all or a specified amount of total charging demand in the most resource-efficient manner while considering any exogenous rate, space, and total capacity constraints. The problem then becomes a cost minimization planning with the objective to jointly minimize investment costs (C^{Φ}) and operations cost (C^{Ω}) over all stages $h \in \mathcal{H}$ contained in the planning horizon while ensuring a pre-defined service level η^h (typically 100% in the following simulations and benchmarks). As




Figure 4.2: EVCH Service System Layout and Asset Components in Planning Stage h

we lay out the variables and parameters of our model please refer to Table 4.1 for an overview of nomenclature used throughout this paper.

Formally, the objective function $f(\Gamma)$ (where Γ is the system configuration) can be expressed as follows:

$$Min_{\Gamma}[C^{\Phi}(x_{k,h}^{i,n}, \delta_h^{Trafo}, \kappa_h^{PV}, \epsilon_h^{Bat}) + C^{\Omega}(\omega_{k,j,h}, \psi_{k,j,h,t}, \beta_{h,t}^{Charge}, \beta_{h,t}^{Discharge}, e_{h,t}^G)] \quad (4.1)$$

Table 4.1: Nomenclature

Symbol	Description	Unit
Sets		
\mathcal{H}	Set of planning stages in planning horizon with index h	set
\mathcal{I}	Set of charging dock types with index i ($\mathcal{I} = \{AC, DC\}$)	set
\mathcal{J}_h	Set of unique EVs entering the EVCH during the planning period h with index j	set
\mathcal{K}	Set of charging dock candidate points (i.e., parking spots) with index k	set
\mathcal{N}	Set of charging dock connector options $\mathcal{N} = \{1, 2, 4\}$ with index n	set
\mathcal{T}	Set of time periods per each stage in planning horizon with index t	set
Ξ	Set of decision variables	set
Γ	Full configuration of EVCH system	set
Parameters		

$A_{j,s}$	Arrival time of vehicle j in stage h	period t
β^{max} , β^{min}	Maximum charge and maximum discharge rate of energy storage	kW
$c_h^{i,n}$	Cost per EV charging dock of type i with n connectors in stage h	USD
c_h^{Trafo}	Cost per kW of grid connection (i.e., transformer) in stage h	USD/kW
c_h^{PV}	Cost per kWp of PV in stage h	USD/kW
c_h^{Bat}	Cost per kWh of energy storage (battery) in stage h	USD/kWh
δ_j	Duration of stay of vehicle j	hours
Δ_t	Duration of a single planning period t	hours
$D_{j,h}$	Departure time of vehicle j in stage h	period t
e_j^d	Total energy requested by vehicle j over duration of stay	kWh
$\eta^{(dis)charge}$	Charge/discharge efficiency of energy storage	ratio
η^{Inv}	AC-DC inversion efficiency	ratio
η_h^{Serv}	Target service level expressed as ratio of fulfilled vs. actual demand	ratio
$f_{h,t}^{PV}$	Avg. PV load factor in period t of stage h	ratio
κ^i	Maximum power per charging dock of type i	kW
κ_0^{Grid}	Existing facility substation capacity	kW
κ^{Trafo}	Standard size of transformer that can be installed	kW
lax_j	Laxity of vehicle j	hours
$l_{h,t}$	Base load of attached facility during period t in stage h	kW
l_h^*	Maximum expected base load of attached facility in stage h	kW
M	big-M constraint (for linearization)	kW
μ^{Maint}	Cost ratio for maintenance (as share in total capital stock)	ratio
R	Maximum installable PV capacity (space constraint)	kWp
SoC^{max}	Maximum energy storage level	%
SoC^{min}	Minimum energy storage level	%
L	Space limitation in number of parking spots	count
T_h^p	Cost of induced power peak per accounting period (i.e., demand charge) in stage h	USD/kW
$T_{h,t}^e$	Cost of energy in period t of stage h as per TOU tariff	USD/kWh
$U_{j,h,t}$	Indicator of whether vehicle j is present during period t in stage h	boolean
Variables		
$\beta_{h,t}^{Charge}$	Charge rate of EVCH battery storage	kW
$\beta_{h,t}^{Discharge}$	Discharge rate of EVCH battery storage	kW
$\beta_{h,t}^{Direction}$	Indicator of whether the battery is charging or discharging	boolean
δ_h^{Trafo}	Number of additional transformers installed in stage h	integer
C^Φ	Total normalized investment cost for the EVCH over planning horizon	USD

C^Ω	Total cost of operating the EVCH over the planning horizon	USD
C^Ψ	Penalty for not serving charging demand	USD
$e_{j,h,t}^S$	Net energy supplied to vehicle j during period t of stage h	kWh
$e_{h,t}^{Grid}$	Net energy supplied from grid during period t of stage h	kWh
ϵ_h^{Bat}	Installed energy storage capacity in stage h	kWh
κ_h^{PV}	Installed PV capacity in stage h	kW
p_h^*	Induced max peak attributable to EVCH operations during stage h	kW
$\psi_{k,j,h,t}$	Charge rate of vehicle j connected to charging dock k during period t of stage h	kW
$SoC_{h,t}$	State variable that tracks state of charge of energy storage	kWh
$w_{k,j,h}$	Indicator for whether a vehicle j is connected to charging dock k in stage h	boolean
$x_{k,h}^{i,n}$	Indicator whether dock (type i , n connectors) is installed at k in stage h	boolean

The EVCH infrastructure decision space determining C^Φ extends over a large set of decision variables, which we briefly describe here. First, decisions on the charging infrastructure configuration and scale-up over the investment horizon \mathcal{H} are required. We allow full flexibility regarding the type of EV charging docks (22kW AC or 50kW DC docks) and the number of connectors per dock (ranging from single-connector setups to up to four connectors per dock). Crucially, for charging docks with multiple connectors, we allow for simultaneous charging of EVs, meaning the rated power per dock can be shared dynamically and flexibly by all connected vehicles. This is different from the more prevalent single-server docks, which either possess just a single connector or multiple connectors that may only be operated sequentially. A multi-server setup has the advantage of higher utilization (vehicles that have completed their charging cycle do not block charging docks) (Ferguson et al. 2018). We capture EV charging infrastructure decisions via a set of binary indicator variables of form $x_{k,h}^{i,n}$, indicating whether a dock of type $i \in \{22kW, 50kW\}$ with number of connectors $n \in \{1, 2, 4\}$ is to be installed at candidate point $k \in K$ during planning stage $h \in \mathcal{H}$. The total number of docks and connectors is naturally bounded by the size of the facility (i.e., number of parking spaces) L . Second, the initial size and expansion pathway of possible on-site generation (PV) and/or storage assets (Li-Ion battery) must be defined. We assume that PV generation κ_s^{PV} can be scaled close-to continuously across all stages h in the planning horizon and that it is limited only by local facility space constraints R (e.g., roof space). In terms of on-site storage, we consider Li-Ion battery technology whose energy capacity κ_s^{Bat} (in kWh) can be scaled continuously over \mathcal{H} . Finally, a decision is required on whether, by how much and by when the existing substation capacity should be extended to accommodate the desired level of charging service. Note that substations can be purchased in standard sizes (κ^{Trafo}) only. Consequently, the grid connection can only be scaled

step-wise in multiples δ_h^{Trafo} of κ^{Trafo} , where δ_h^{Trafo} is an integer value denoting the number of transformer modules to be added to the facility's substation in state h ⁵. Given the physical size of substations and the fact that local grid conditions may not allow for an unconstrained scale-up of the existing grid connection, we impose a maximum G on the final size of the substation. In sum, total investment cost over the planning horizon is determined as follows: $C^\Phi = \sum_{h \in \mathcal{H}} (c_h^{Trafo} \delta_h^{Trafo} + \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{I}} \sum_{n \in \mathcal{N}} c_h^{i,n} x_{k,h}^{i,n} + c_h^{PV} \kappa_h^{PV} + c_h^{Bat} \epsilon_h^{Bat}) (1 + (|\mathcal{H}| - h) \mu^{Maint})$. Note that this includes the maintenance costs incurred over the planning horizon, which is captured by the factor $(|\mathcal{H}| - h) \mu^{Maint}$. The parameters c_h^{Trafo} , $c_h^{i,n}$, c_h^{PV} and c_h^{Bat} are stage-dependent cost parameters that take into account expected technology cost trajectories over the planning horizon.

Underlying these higher-level infrastructure choices are smaller operational decisions, which determine the operational cost C^Ω . Naturally, the operational scope is constrained by the installed infrastructure highlighting the two-way interdependencies between both sets of decisions. Operations decisions focus on the assignment of a vehicle j to a connector k upon arrival (captured by $\omega_{k,j,h}$) and the periodic charging decisions over the duration of stay ($\psi_{k,j,h,t}$). Finally, the on-site battery state is controlled via $\beta_{h,t}^{Charge}$ and $\beta_{h,t}^{Discharge}$, two booleans that control the rate of charge/discharge. We consider PV generation and building baseload to be exogenous parameters that cannot be actively controlled by the EVCH operator. Given the different sources of power (battery, PV, grid), the operator also needs to decide on the power mix per each period t . This involves setting the desired energy drawn from the grid per period in each state $e_{h,t}^{Grid}$. Note that $e_{h,t}^{Grid}$ is typically accounted for based on a two-part tariff including a time-of-use-dependent energy charge $T_{h,t}^e$ and a monthly demand charge T_h^p that is a function of the maximum induced power p_h^* in that month. Operational costs are formally defined as follows: $C^\Omega = \sum_{h \in \mathcal{H}} (\sum_{t \in \mathcal{T}} T_{h,t}^e e_{h,t}^{Grid} + T_h^p p_h^*)$.

4.3.2 Parameterizing the Digital Twin Simulator

We now set up the environment with which the RL agent will interact. The goal is to create a close-to-exact digital representation (i.e., a DT) of the EVCH, which is made up of three core components: (1) physical assets, (2) operational policies that define how the physical assets are operated, and (3) preference/demand characteristics, i.e., the external requests and usage patterns that the service system needs to fulfill. In this section, we lay out these three components. Note that the EVCH DT operates on a discrete-time basis. To reduce any issues/inconsistencies related to discretization, we use period lengths of just one minute. Wherever practical the DT is fed with real-world sensor data to achieve a highly accurate representation of the physical

⁵We consider any interaction effects with the upstream electrical distribution grid to be out of scope for this problem. Specifically, we assume that the distribution grid is unconstrained and able to accommodate any additional load from the EVCH, provided sufficient substation capacity (i.e., transformers) is installed for voltage regulation.

world. Simulation is used in areas where data are not available. For an overview of data sources and digitalization approaches per DT component refer to Table 4.5 in Appendix D.

Digital EVCH Asset Models

Figure 4.2 provides an overview of the physical EVCH asset classes that are to be represented digitally in the DT environment. We draw on asset spec sheets along with real-world machine data to represent the physical EVCH components and the context they operate in as accurately as possible.

Local Substation We model the local substation as an integrated system consisting of transformers, circuit breakers, and other peripheral equipment that connect the site to the higher voltage levels of the distribution grid. The substation capacity is determined by the sum of rated transformer capacities. Although typically very low, we account for transformer losses using an efficiency factor of $1-\eta^{Trafo}=2\%$.

On-site Electricity Generation Assets (PV Panels) For electricity generation, we assume photovoltaic (PV) modules, a natural supplement to EV charging hubs due to their production patterns that are highly correlated with occupancy profiles (and thus charging demand) of most parking lots. PV generation is non-dispatchable, i.e., it cannot be actively controlled. PV power is therefore consumed on-site (by EVs, battery storage, building, etc.), or fed back into the grid. Note that PV installations require DC-AC conversion via inverters. The efficiency losses of DC-AC conversion are accounted for via an inversion efficiency factor of $1-\eta^{Inv}=4\%$. We use real-world PV load factors ($f_{h,t}^{PV}$) to model PV production from the regions corresponding to the intended EVCH facility locations. Load factors are a measure of real PV panel power output as a ratio of installed capacity (κ_h^{PV}) and depend on local solar irradiation conditions⁶. PV production at time t (excluding DC-AC conversion losses) is then given by $f_{h,t}^{PV} \kappa_h^{PV}$.

Electricity Storage Assets We model electricity storage as a lithium-ion battery with instantaneous ramp time. To avoid excessive battery degradation, we allow the state of charge to vary over the interval of $[5\%, 95\%]$, thus avoiding deep discharging and over-charging that are particularly strenuous for battery hardware. Setting upper and lower energy content boundaries is a common approach in storage management (Ghiassi-Farrokhfal et al. 2016). We also assume symmetric charge and discharge efficiency of $\eta^{charge} = \eta^{discharge} = 95\%$. Battery operations are simulated using what is sometimes referred to as a C/C/C model⁷. In addition, we implement typical battery constraints related to a maximum charge and discharge rate κ^{Bat}

⁶This data is available via local transmission system operators (TSOs) and generally comes in 15-minute intervals.

⁷C/C/C models assume constant battery charge/discharge efficiencies, constant energy content upper and lower bounds, and constant voltage (Kazhamiaka et al. 2019)

(symmetric). κ^{Bat} is dependent on the size of the battery and is set such that the battery can be charged/discharged to/from full charge within one hour.

Peripheral Building We use real-world building consumption data to model site baseload ($l_{s,t}$) that is served by the same grid connection, thus influencing total available grid capacity at any given period t (see Figure 4.2). Contrary to EV loads, we assume $l_{s,t}$ to be exogenous, i.e., it cannot be dynamically managed or even curtailed. Given the absence of smart energy management hardware in most existing building stock, this is a reasonable assumption. Note that granular consumption data is widely available for commercial buildings above a certain consumption threshold since these consumer classes are typically exposed to time-of-use tariffs as well as demand charges for induced peak load. Our 1-year dataset records peak building loads and consumption at a 15-minute resolution.

EV Charging Docks and Connectors We model two different types of charging docks that mainly differ in terms of maximum charging rates κ . Specifically, we allow for AC fast chargers with maximum charging capacity $\kappa^{i=AC}=22\text{kW}$ and DC super-fast chargers with maximum charging capacity $\kappa^{i=DC}=50\text{kW}$. For each charger type $i \in [AC, DC]$ we allow for different connector configurations with $n \in [1, 2, 4]$ connectors per dock. We assume that κ can be shared dynamically and flexibly between all connectors per dock. This means that connected EVs can be served both sequentially and simultaneously via the same dock. Losses related to AC-DC conversion are modeled using an efficiency factor of $1-\eta^{Inv}=4\%$ ⁸.

EVCH Operational Policies

We also implement a range of realistic operational policies that simulate real-world operations in the DT environment. These policies are inspired by standard operational practices currently used in EV charging operations as well as recent algorithms proposed in the EV charging literature (e.g., Lee et al. 2019, Ferguson et al. 2018). Given that we allow for multi-connector charging docks with simultaneous charging capability, the initial assignment of vehicles to charging stations becomes important due to heterogeneous energy demand and flexibility characteristics. Clearly, since EVs cannot be readily relocated while parked, the initial assignment to a connector influences future available charging capacities for the EV in scope as well as for current and future arrivals that are to be served by the same charge dock.

Vehicle Routing Algorithms Vehicles $j \in \mathcal{J}$ are routed/assigned to a connector k upon entry into the EVCH (captured by $\omega_{k,j,h}$). We implement two heuristic routing algorithms of varying levels of sophistication and varying information requirements.

⁸For AC charging, the AC-DC inverter is integrated with the vehicle charger unit, whereas for DC charging the inverter sits inside the charging dock

- **Lowest-utilization-first (LUF)**: This strategy operates on a sorting basis. At each new arrival, the algorithm sorts all available docks based on free capacity. New arrivals are routed to docks with low utilization first. In the case of a tie, the algorithm selects randomly between the charging docks it is indifferent between. An advantage of the lowest-utilization-first approach over other sorting methods (such as lowest-occupancy-first) is that it considers the different charging capacities of AC vs. DC docks in the assignment process.
- **Lowest-laxity-to-highest-capacity matching (LLHC)**: This strategy not only considers the state of individual charging docks but also that of the vehicle that is to be assigned. Specifically, it sorts arriving vehicles into baskets of low, medium, and high laxity (using bins obtained from historical data). Low-laxity vehicles are then matched with charging docks that have a high free capacity and vice versa, thus implicitly provisioning for future arrivals.

Vehicle Charging Algorithms Vehicle-level charging schedules are re-computed in an online manner every five minutes allowing for updating of pre-computed schedules as new information becomes available. We implement a selection of sorting-based algorithms and optimization-based approaches to periodically determine the charge rate per connector k , vehicle j , and time t $\psi_{k,j,h,t}$:

- **First-come-first-served (FCFS)**: Charging requests are served on a first-come-first-served basis at full charging dock capacity until the available power capacity (on-site generation, storage, and grid) is exhausted. This algorithm is largely consistent with standard off-the-shelf load management tools available in the market today.
- **Least-laxity-first (LLF)**: Equivalent to a first-come-first-served algorithm but using a least-laxity-first priority rule meaning that least flexible vehicles are charged first. The algorithm, therefore, explicitly considers the current state of a vehicle in the charging decision.
- **Optimal**: Optimal operations uses mathematical optimization to periodically (re-)compute cost-optimal charging schedules that satisfy charging demand for the planning period δt in scope. We implement a standard cost-optimal charging framework (e.g., Ferguson et al. 2018). In our simulations, we plan $\delta t = 12$ periods ahead. We also implement a smoothing constraint by limiting the maximum charging ramp rate and considering the parallel use of docks. Since future arrivals and their preference vectors are unknown at the time of planning, we use a carefully tuned safety margin to be able to accommodate these in future periods. Further details are provided in 4.7.1.

Electricity Storage Operational Algorithms The third and final system component requiring active operational management is the on-site energy storage system. We implement a heuristic approach that has been demonstrated to perform well under real-world conditions (Gust et al. 2021). We update battery-specific decision variables $\beta_{s,t}^{Charge}$ and $\beta_{s,t}^{Discharge}$ every five minutes⁹:

- **Temporal arbitrage (TA):** The algorithms exploit the structure of the underlying time-of-use electricity tariff. During off-peak hours the battery is charged at a constant rate until the upper bound of the allowable energy content is reached. During on-peak periods, the battery is discharged at a constant rate to reduce the amount of electricity purchased at the higher on-peak price.

EVCH Preference Data

We populate the simulation’s base architecture consisting of the above-described asset models and operational policies with an EVCH-specific high-resolution model of parking and charging preferences. EVCH user preferences (of an individual j) are described by the three-dimensional vector $v_j = (A_j, \delta_j, e_j^d)$ where A_j is the time of arrival, δ_j the duration of stay and e_j^d the requested energy. δ_j and e_j^d define what is referred to as laxity ($lax_j = \delta_j - \frac{e_j^d}{\kappa^i}$) (Lee et al. 2019). $lax_j = 0$ means that a vehicle j needs to charge at the maximum available rate κ^i for the entirety of its stay, while higher laxity values indicate more room for active charging management. Time of entry A_j determines the earliest planning period by which a certain charging event needs to be initiated. We start by building a model of current archetypical parking patterns. We do this in order to understand what typical parker types exist and how the user base composition can vary across facilities. A taxonomy of parker types can also be useful for building synthetic user population datasets based on assumed parker type shares wherever real-world data is not available. We leverage a unique, large-scale transaction-level parking dataset that was provided by a major European real-estate investor and includes transactions from seven large-scale parking garages¹⁰. We use a full year of data to capture daily, weekly and yearly seasonality. 2019 is chosen as a reference year to filter out pandemic-related effects. In total, our data comprises 3.84M parking events. We cluster parking events j based on A_j and δ_j , the two core parameters of interest at this modeling stage. All details related to data pre-processing, selection of clustering algorithms, and robustness tests are provided in e-companion 4.7.2. In Table 4.2 we summarize our results. The largest proportion of parking events in our dataset is made up of three short-term parker types (*Morning Short*, *Afternoon Short* and *Evening Short*). These users enter a parking lot in the morning, afternoon, or evening respectively and typically stay for periods of 1-2 hours. We also observe a *Business* cluster, which comprises parking events that commence in the

⁹Note that battery decisions are made after the previously described charging decisions are made and that the available battery capacity at the start of each planning period is available for EV charging.

¹⁰Each row in this dataset represents a single parking event j with corresponding arrival and departure preference information. For privacy reasons, individual users cannot be identified




Figure 4.3: Occupancy profiles per parker type (three archetypal parking facilities)

early morning (7:26 am on average) and last for an average of 8 hours. Two additional segments comprise longer-term parking events. These are *Overnight* parkers, which enter the parking lot in the late afternoon and stay until the next morning (typically 15.8 hours on average), and *Long-term* parkers that stay for periods longer than 24h on average.

We then look at the distribution of parker types across the different facilities in our dataset. Three archetypal facilities can be identified: The first facility type is a typical workplace facility that caters mostly to Business parkers. The second facility type is a destination facility. Apart from a small proportion of business users, such facilities mostly host short-term parkers. Finally, we also observe facilities with less conclusive usage patterns experiencing strong demand from all segments. We term these mixed-use facility¹¹. Typical occupancy profiles for each of the three facility types are shown in Figure 4.3.

Finally, we focus on the third required preference input variable: the requested energy per vehicle e_j^d . We employ a recently published real-world dataset by Lee et al. (2019) containing >25,000 charging transactions for the year 2019. Per each charging transaction the full preference vector $v_j = (A_j, \delta_j, e_j^d)$ is available. We blend the charging data (which only contains served sessions that are constrained by the available infrastructure) with our parking dataset (which contains all parking requests per facility) using techniques from collaborative filtering. We train a prediction model on the labeled Lee et al. (2019) dataset and use the resulting model to predict charging demand in the parking dataset. We obtain an exponentially distributed charging demand across the entire population of EVs with an average demand of 26.46 kWh ($\sigma = 17.20$ kWh) per parking session¹². The distributional shape of charging demand is consistent

¹¹The example shown in Figure 4.3 (right panel) is a large-scale inner-city parking facility that caters to workers, visitors, and residents.

¹²We also apply some limited post-processing by limiting e_j^d to a realistic maximum bounded by the typical size of batteries (100 kWh) and feasible energy transfer over the duration of stay assuming 50kW maximum charge rate.

with the one seen in other empirical EV charging settings (e.g., Ferguson et al. 2018). Crucially, however, Table 4.2 highlights important implications for charge management resulting from the different compositions of parker types in a facility. As can be seen, the average laxity varies significantly across parker types. Thus, parking facilities with a high proportion of high-laxity parkers (e.g., Business, Long-term) benefit from considerably higher flexibility characteristics with higher scope for optimization through intelligent charge management and parallel charging at lower rates.

Table 4.2: Preference characteristics per parker type

k	Name	Cluster Size		Characteristics (avg & std. in parentheses)		
		N	share	A_j	δ_j	lax_j ¹³
1	<i>Business</i>	671,384	17.47%	7:26am (1.43h)	7.92h (2.73h)	6.29h (2.84h)
2	<i>Morning Short</i>	1,279,646	33.30%	11:10am (1.33h)	2.12h (1.90h)	1.31h (1.75h)
3	<i>Afternoon Short</i>	985,710	25.65%	3:03pm (1.00h)	1.73h (1.35h)	1.08h (1.30h)
4	<i>Evening Short</i>	744,753	19.38%	6:17pm (1.84h)	1.47h (1.30h)	1.11h (1.35h)
5	<i>Overnight</i>	129,273	3.36%	5:22 pm (4.01h)	15.84h (4.02h)	12.65h (4.40h)
6	<i>Long-term</i>	32,241	0.84%	2:28pm (4.94h)	37.04h (6.70h)	35.80h (6.98h)

Combining Asset Models, Operational Algorithms and Preferences into an EVCH Digital Twin

Combining asset models, operational algorithms, and charging demand preferences yields a high-resolution DT of the envisioned EVCH system. Figure 4.4 and 4.5 highlight the internal mechanics of the DT environment¹⁴. Figure 4.4 represents the demand side and visualizes load curves for the various load sinks in the EVCH (EV charging, building baseload, battery storage charging) over the simulation horizon. Building load curves follow a recurring daily pattern ramping up during the day and down again during the night. Loads are slightly lower on the weekend (especially on Sunday). Note also the battery storage load, which reflects the temporal arbitrage strategy.

The power requested by the above-described load sinks is supplied by the grid, the on-site generation unit (PV), or the on-site electricity storage. The behavior of the supply side is shown in Figure 4.5. Note how the exact supply mix heavily depends on the time of day (e.g., no PV generation after sunset, battery disc) and even weather conditions (note the considerably higher PV output in the middle of the week).

¹³assuming 22kW max. charge rate

¹⁴Shown here for a random week in a mixed-use facility using Lowest-laxity-to-highest-capacity routing, optimal charging and temporal arbitrage storage operations




Figure 4.4: Power supply by load sink over one-week simulation horizon (Monday to Sunday, Mixed-use facility)




Figure 4.5: Power supply by load source over one-week simulation horizon (Monday to Sunday, Mixed-use facility)

4.3.3 Setting up the SAC RL Framework

Defining the design objective and the DT environment now allows us to address the system design challenge over multiple stages in an effort to obtain a (near-optimal¹⁵) solution.

For the case at hand and given that planning decisions are made over multiple stages h in planning horizon \mathcal{H} , the problem can be framed as a stochastic sequential decision-making problem. This decision process can be cast as a Markov Decision Process (MDP). An episodic task emerges, which starts at the beginning of the planning horizon ($h = 0$) and runs through the last investment stage $h = |\mathcal{H}|$, with the epochs being the individual decision stages (e.g., beginning of each year).

The MDP is formulated as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, where the elements represent the state space, action space, transition probability, reward function, and the discount factor respectively. The state $s = (h, s^{chargers}, s^{grid}, s^{PV}, s^{storage})$ is defined as a vector of the time (i.e., the planning stage h) and the infrastructure which is accumulated over all previous stages. Actions are described by the vector $a = (a^{chargers}, a^{grid}, a^{PV}, a^{storage})$ and comprise planning decisions, such as the number of fast/slow docks with a specified number of plugs, the number of transformers, the PV capacity and storage to be installed. In line with the planning objective defined in Step 1, we define the reward for moving from state s_h to s_{h+1} as $r = r(s, a) = -(C^\Phi + C^\Omega + C^\Psi)$, which includes the investment cost, the operational costs, and C^Ψ which represents a penalty related to unserved charging demand. Note that the reward is the negative value of costs. Despite a deterministic state transfer function, exact reward functions are stochastic and unknown. In other words, given a decision, the next state is known, while the expected reward is unknown unless the state is evaluated in the EVCH DT environment. Therefore, we employ model-free reinforcement learning to find an optimal EVCH configuration policy π^* . The facility investment policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ maps the state of the environment to an investment decision for each planning time step. Note that the output of the policy is standardized for all action components to be between 0 and 1. Furthermore, each action component is individually scaled based on the associated lower and upper bounds.

To find the optimal policy, two classes of RL algorithms have been proposed in the literature (Sutton and Barto 2018a): (1) value-based algorithms which learn the state/action-state values by interactions and shape the optimal policy using the learned values, and (2) policy-based algorithms which directly evaluate and improve the current policy until converging to near-optimal solutions.

As mentioned, value-based models such as Q-learning can run into issues of tractability in large state-action space environments like EVCH sizing. Therefore, we opt for a soft actor-critic (SAC) model which combines value-based and policy-based concepts. SAC is an off-policy actor-critic deep RL algorithm that works based on the maximum entropy learning framework (see Appendix C.1 for the differences between traditional actor-critic and SAC). In other words, the

¹⁵Note that global optimality cannot be guaranteed for most learning or search methods.

actor of SAC aims to learn a policy $\pi(a_h|s_h)$ that maximizes expected reward (negative values of costs) while also maximizing entropy to improve the exploration which is vital for large-scale problems. Therefore, the objective of learning the policy is defined as follows:

$$J(\pi) = \sum_{h=0}^{\mathcal{H}} \mathbb{E}_{s_h, a_h \sim \rho_\pi} [r(s_h, a_h) + \alpha \mathcal{H}(\pi(\cdot|s_h))] \quad (4.2)$$

Where ρ_π denotes the marginal of the trajectory distribution induced by policy $\pi(a_h|s_h)$. The temperature parameter α adjusts the importance of the entropy term against the reward, and thus controls the stochasticity of the optimal policy. A key strength of SAC, setting it apart from other RL algorithms, is its powerful exploration capability. First, the policy is inherently stochastic, adding randomness to action selection. Second, the algorithm includes an entropy term in its objective function, promoting the exploration of less-visited policies. By adjusting the weight of this entropy term, SAC effectively balances exploration and exploitation. Additionally, random noise is introduced to the actions to prevent the model from getting stuck in local optima. For a detailed explanation of the SAC algorithm, we refer to Haarnoja et al. (2018).

To overcome the curse of dimensionality we use deep neural networks to represent both critic (value network) and actor (policy network). The value network evaluates the value of the current policy through interaction with the environment and the policy network makes decisions based on the current state of the system. Each network is fully connected and includes multiple (4) hidden layers with different number of nodes (256, 512, 512, 256) (See Appendix 4.7.3 for details). We train both networks from the agent’s past experience using temporal-difference algorithms. This means that each experience in buffer contains one interaction with the environment, including state, action, immediate reward, and next state. To improve stability of the critic network, we also use a target network that gets updated less frequently and to increase the chance of more comprehensive exploration we add extra noise to the output of our policy network in the training phase.

In problems with multi-dimensional action spaces, using continuous-to-discrete mapping significantly enhances the scalability of the model compared to discrete action space models (e.g., Christodoulou 2019), which must account for all possible action combinations. This scalability is crucial for addressing large-scale problems in real-world scenarios. We will illustrate this advantage by conducting a scalability analysis of our proposed model and comparing it with traditional discrete action space models. Our MDP contains discrete integer actions which may be better suited for discrete RL models at first glance. However, even state-of-art deep learning function approximation approaches do not easily scale to the number of action combinations (5^8 alternatives for each decision step) encountered in this problem. Alternatively, making use of integer relaxation, we can employ a continuous SAC model which is considerably more scalable than discrete (value-based) alternatives. Similar to Dulac-Arnold et al. (2015), our model first takes actions within a continuous space and then maps them to a discrete action set. We will

define:

$$f_{\theta^\pi} : \mathcal{S} \rightarrow \mathbb{R}^D, f_{\theta^\pi}(\mathbf{s}) = \hat{\mathbf{a}} \quad (4.3)$$

f_{θ^π} is a function parameterized by θ^π (the policy network parameters), mapping from the state space \mathcal{S} to the continuous action space \mathbb{R}^D , where D is the dimensionality of the action space. The output of this function ($\hat{\mathbf{a}}$) is likely not a valid set of actions for the environment as it will contain non-integer values that might violate the physical constraints of the environment. Therefore we define a mapping function as follows:

$$g : \mathbb{R}^D \rightarrow \mathcal{A} \quad (4.4)$$

$$g(\mathbf{s}, \hat{\mathbf{a}}) = \underset{D}{\operatorname{argmin}} \sum (a_d - \hat{a}_d)^2 \quad (4.5)$$

$$a_d \in \mathcal{Z}^+ \cup \{0\} \quad \forall d \in D \quad (4.6)$$

$$s_d + a_d \leq u_d \quad \forall d \in D \quad (4.7)$$

g is a mapping function from a continuous space to a discrete space, constrained by the physical capacities of the environment.

Based on findings of Dulac-Arnold et al. (2015), although the architecture of our policy is not fully differentiable, we can nevertheless train our policy by following the policy gradient of f_{θ^π} . To do so we define a simpler policy $\pi_\theta = g \circ f_{\theta^\pi}$. In this initial case we can consider that the policy is f_{θ^π} and that the effects of g are a deterministic aspect of the environment. This allows us to adopt a standard policy gradient approach to train f_{θ^π} on its output $\hat{\mathbf{a}}$, effectively interpreting the effects of g as environmental dynamics. In addition to this, we include the physical restriction of each action component using constraint (7), whereby unfeasible actions are prevented. As an example of this, the number of installed PVs cannot exceed the area capacity of the EVCH. The addition of the continuous-to-discrete layer in the agent’s policy is the primary modification compared to traditional SAC models. In order to ensure that the integer relaxation does not cause our model to converge to a local optimum, we benchmark the SAC model results with a deep Q-network agent which uses a discrete action space as well as with a perfect-information mathematical model. The DQN model does not rely on integer relaxation while the mathematical model represents the optimal solution under perfect information. The details of this comparison are provided in Section 4 and offer reassuring evidence that divergence to a local optimum is unlikely an issue in our experiments.

4.4 Performance Evaluation and Benchmark

We now evaluate our SAC-based RL model. We first present the convergence and optimality characteristics of our approach in comparison to traditional optimization methods, as well as




Figure 4.6: Convergence and optimality characteristics of our proposed model as compared to DQN and mathematical programming - Values resampled to minimum out of 10 consecutive episodes

popular reinforcement learning techniques, including the value-based Deep Q-Network (DQN) and the actor-critic Deep Deterministic Policy Gradient (DDPG) algorithms. Subsequently, we analyze the scalability properties of our SAC-based reinforcement learning model. All experiments are performed on a research workstation with an AMD Ryzen Threadripper 3970X 32-Core processor and 256GB of RAM. To achieve tractability of the mathematical programming benchmark, we introduce several simplifications across the three modeling approaches to ensure a fair comparison. First, we limit the problem size to 200 parking spots. Second, facility operations are assumed to be optimal (optimal vehicle placement and charging) and stochasticity is neglected (i.e., assuming perfect foresight). Third, we reduce the operational detail of the EVCH simulation by adopting an hourly temporal resolution.

4.4.1 Evaluating Convergence and Optimality Properties

In Figure 4.6 we present convergence and optimality characteristics of our main model, the SAC-based RL framework, as compared against the optimal solution derived using mathematical programming. We also include two additional RL-based benchmark models in the comparison. First, we implement a value-based DQN reinforcement learner based on the algorithm presented in Van Hasselt et al. (2016) and used extensively in extant OM research. Second, we implement an alternative actor-critic method, the DDPG algorithm, based on Lillicrap (2015). This approach is more closely aligned with our proposed model, which utilizes the SAC algorithm. Details on all four models are presented in Appendix 4.7.3.

In our experiments, SAC achieves a near-optimal solution in just 400 episodes, significantly outpacing DQN and DDPG, which converge after approximately 700 and 550 episodes, respectively. The optimality gap of the SAC model is significantly lower than that of the DQN and DDPG approaches (10% for SAC vs. 19% for DQN, and 22% for DDPG). We attribute these superior convergence and optimality gaps characteristics of SAC compared to DQN to the large state-action space of the given problem (8 discrete decision with 5 options each). Scalability to large state-action spaces is a well-known drawback of value-based approaches such as DQN which require more and more exploration to train larger and larger deep-Q networks (number of output nodes equals the number of actions that can be taken (Dulac-Arnold et al. 2015)). Indeed, this is confirmed in our experiments with larger/real-sized problem instances, where the DQN fails to converge in a reasonable amount of time (48h). SAC can handle highly-dimensional problem spaces significantly better than DQN. This is due to the core “actor” component learning the policy function directly without the need to evaluate all possible actions per state. Instead, the actor returns actions directly in a continuous space, resulting in a smaller network. While DDPG leverages both actor and critic networks, it struggles with the trade-off between exploration and exploitation. The DDPG algorithm employs a deterministic policy and relies only on action noise added to the policy output for exploration. Our results demonstrate that in complex environments with high-dimensional action spaces, the SAC algorithm outperforms traditional actor-critic models like DDPG (Haarnoja et al. 2018).

4.4.2 Evaluating Scalability to Real-Sized Problems

We now run several additional experiments aimed at understanding scaling performance of the SAC-based planning approach and its performance against benchmark algorithms for smaller and larger problem sizes.

First, we explore whether SAC has advantages when it comes to scaling to practical problem sizes, i.e., EVCHs significantly larger than the previously explored 200 parking spots. To this end we iteratively increase the problem size from an initial facility size of just 20 parking spots to up to 1000 over the course of six experimental runs. 1000 parking spots is a size that is representative of a large parking lot (see Section 4.3.2). Results are shown in Table 4.3. Solution time as a function of EVCH facility size increases exponentially for the mathematical programming framework. The model requires just 4s to solve a 20 parking spot EVCH planning problem to optimality, but 5,937s to reach an optimal solution for a facility 10 times the size (200 parking spots). No convergence is achieved for problem sizes significantly larger than that (e.g., 500 parking spots and upward) meaning mathematical programming is not a practical option for many real-sized EVCH planning scenarios. Our SAC framework, on the other hand, exhibits significantly more favorable scaling and tractability properties. It achieves solution times 64% lower than mathematical programming for a 200 parking spot facility. It also scales to problem sizes of 1,000 parking spots for which convergence is reached in just under 12h. SAC converges

very closely to the global optimum reaching optimality gaps between 4% to 15% across our experiments. Notably this is also reflected in the specific planning decision derived by SAC, which closely resemble optimal planning decisions. Please refer to Appendix E for a deeper analysis of individual planning decisions for all benchmark models.

Our results also highlight the scalability advantages of our proposed model (SAC) compared to value-based (DQN) and traditional actor-critic (DDPG) algorithms. While both DQN and DDPG perform close to optimal solutions for small facility sizes, their performance significantly deteriorates for sizes exceeding 200, and neither converges even after 1000 episodes. For DQN, convergence is slower even for small problems, and for larger facilities (e.g., 500 and 1000), the algorithm struggles to find good solutions due to the curse of dimensionality, as the action space becomes excessively large. Similarly, DDPG requires more time to converge compared to SAC, and the gap widens for large facilities. In such cases, DDPG fails to converge, primarily due to its reliance on insufficient exploration mechanisms, which is a limitation of traditional actor-critic models like DDPG (Colas et al. 2018).

Table 4.3: Evaluation of scalability and optimality

EVCH facility size	20		50		100		200		500		1000	
	t	gap	t	gap	t	gap	t	gap	t	gap	t	gap
Mathematical program	4	0.00	248	0.00	842	0.00	5,937	0.00	no solution		no solution	
DQN	932	0.05	1,950	0.35	5,420	0.14	18,582	0.19	no solution		no solution	
DDPG	343	0.05	1,330	0.13	1,780	0.12	2,474	0.22	no solution		no solution	
SAC	174	0.04	396	0.15	990	0.11	2,115	0.10	15,919	-	42,714	-

Notes: Problem size given by number of parking spots in facility, t indicates the solution/convergence time in seconds (excluding time needed for hyperparameter tuning); gap indicates the optimality gap; DQN: Deep Double Q-Networks RL model, DDPG: Deep Deterministic Policy Gradient RL model, SAC: Soft Actor-Critic RL model

As noted above, for reasons of comparability, these benchmark results are for an abstracted version of the EVCH sizing problem (perfect foresight, low temporal granularity and simplified operational detail). An additional benefit of adopting SAC versus mathematical optimization lies in the fact that operational and temporal detail do not increase the problem size and thus do not significantly impact solution time. This is because the modeling of operations is relegated to the DT simulation. Consequently, simulating real-sized EVCH systems in close to full operational detail, real-time (e.g., 1 minute discretization) and over large sets of sensor data (e.g., months or even years) become possible. SAC-derived solutions can therefore be expected to generalize better to real-world stochastic conditions compared to the optimization-derived solutions. Note that stochastic and/or robust optimization approaches have not been explored in this work due to their significantly higher computational requirements compared to deterministic approaches which exacerbate scalability concerns.

4.5 Scenario Analyses

Finally, we use our SAC model in conjunction with the DT simulator to run extensive sensitivity testing under close-to-real-world conditions. Aspects of particular interest here are: user preference scenarios and operational policy choices and their impact on planning decisions. For illustrative purposes, we report configurations achieved in the final planning state ($h=9$) only, i.e., we neglect the scale-up pathway until that state is reached. Details on parameterization are provided in Appendix 4.7.4.

4.5.1 Impact of Variations in User Preferences

In Figure 4.7 we compare SAC-derived infrastructure investment decisions for three archetypical EVCH facilities (Destination, Mixed-use, Workplace). The results reveal significant sensitivity of optimal physical layout decisions to user preferences. Recall that Destination facilities are primarily used for short-term parking (large proportion of *Morning Short*, *Afternoon Short* and *Evening Short* parkers), Mixed-use facilities exhibit a heterogeneous user pool, while Workplace facilities are primarily used by commuters with long stays (see Table 4.2 and Figure 4.3). Consequently, average laxity characteristics vary considerably across facility types. For example, the average laxity of parkers in a Destination facility is considerably lower than that of a Workplace facility where users remain parked for prolonged periods. We see these preference differences reflected in the derived EVCH configurations across the three facilities. Specifically, for the facility with the lowest average laxity (Destination facility), the RL algorithm chooses to provision primarily single-connector AC docks along with a small number of single-connector fast chargers. A single-connector setup ensures that the full charging power is available at all times but comes at the risk of vehicles blocking an entire dock even after completing a charging cycle. The latter issue seems to be less problematic in a destination parking setting, where users do not stay long on average. At the other end, the infrastructure setup for a Workplace facility tends to favor multi-connector docks, particularly AC double-connector docks, thus taking advantage of the higher laxity of the underlying user population that affords longer charging cycles at lower rates. The third facility type (Mixed-use) falls somewhat in the middle between the previous two extremes. This is consistent with its laxity profile that lies between that of the Destination and Workplace facilities. It is noteworthy that DC charging plays a minor role in any scenario. EVCH charging use cases can mostly be satisfied at lower charging rates.

In terms of power supply infrastructure, the Destination facility requires an additional 200kW transformer versus the other scenarios to achieve the target service level and satisfy low-laxity charging requests at higher charging rates. Some PV and battery storage is installed in all scenarios.




Figure 4.7: RL-derived system configuration decisions for three archetypical EVCH facilities




Figure 4.8: RL-derived system configuration and performance against objective for different routing policies and a mixed-use facility




Figure 4.9: RL-derived system configuration and performance against objective for different charging policies and a mixed-use facility

4.5.2 Impact of Operational Policy Choices

Next, we leverage the flexibility characteristics of the DT approach by investigating the impact of different operational policies on the sizing decisions and the system’s cost performance. As mentioned, such analyses would mean major model reformulation for optimization frameworks but are easily implemented in a DT-based model. In Figure 4.8, we explore the impact of routing decisions for a Mixed-use facility, assuming FCFS charging operations.

We find that the planning outcome is sensitive to the choice of routing strategy, as is the cost performance of the derived system. The more sophisticated routing strategy (LLHC) relies on more multi-connector docks and requires fewer alternative power sources (PV and storage) in an optimal setup compared to the same facility operated with LUF routing. Total system cost savings amount to 8.4%.

We perform a similar sensitivity analysis for the choice of charging strategy. Figure 4.9 displays RL-derived infrastructure decisions and system cost performance for the same Mixed-use facility, assuming LLHC routing. We observe a largely similar picture here. Planning decisions are sensitive to the choice of charging strategies (e.g., more multi-dock chargers, more PV, and more battery with optimal strategy), as is cost performance (i.e., the best performance of the system with optimal charging with savings of 1.7 to 3.2% against the alternative strategies).

In sum, we show that the physical EVCH configuration (II) is highly sensitive to Ω , the EVCH operational policy. In general, the more sophisticated the operational policies, the lower the total infrastructure requirements and the better the overall cost performance of the system. Thus, in order to obtain optimal planning decisions, operations managers need alignment on how they intend to operate the service system. Different operational strategies require different infrastructure layouts to achieve optimal performance and result in different total system costs.

In sum, we show that the physical EVCH configuration (II) is highly sensitive to Ω , the operational policies that the EVCH operates on. In general, the more sophisticated the operational policies, the lower the total infrastructure requirements and the better the overall cost

performance of the system. Thus, in order to obtain optimal planning decisions, operations managers need alignment on how they intend to operate the service system. Different operational strategies require different infrastructure layouts to achieve optimal performance and result in different total system costs.

4.6 Discussion

In this work we consider the problem of planning large-scale electric charging hubs (EVCHs). We define EVCHs as locally concentrated and centrally operated clusters of charging infrastructure that are typically integrated with on-site generation, storage and adjacent building infrastructure. Examples include workplace parking facilities, EV-enabled inner city parking garages or EV fleet depots. Planning these complex operational systems over a multi-year investment horizon represents a high-dimensional, dynamic and stochastic decision problem. Such planning problems have traditionally been approached by means of mathematical programming (e.g., Kazemi et al. 2016, Li et al. 2020). These frameworks are subject to computational challenges (e.g., NP-hardness) that can limit scalability to practical system sizes. As a result, simplifying assumptions related to, for example, temporal granularity, operational detail, system size, decision horizon or stochasticity are required to achieve tractability. This can come at the expense of generalizability to real-world conditions and does not take advantage of the wealth of granular operational data that has become increasingly abundant (Choi et al. 2022).

We develop and evaluate an alternative data-driven solution approach to the EVCH planning challenge, thus responding to calls from the scientific community and real-world sectors to develop methods that make use of and incorporate granular operational and preference data into OM frameworks (Qi and Shen 2018, Cohen 2018, Choi et al. 2022, Ketter et al. 2023).

The proposed solution – the core contribution of this work – leverages modern reinforcement learning (RL) (specifically soft actor-critic (SAC) RL) in combination with fine-grained data-driven simulation, also referred to in this work as Digital Twin (DT). SAC, a policy-based RL method, is better suited for the significant size of the action combinations in EVCH planning (e.g., wide variety of asset classes such as different EV charger types, PV, on-site battery, transformers, etc. each with large set of discrete options over multiple investment periods) compared to value-based deep learning function approximations. This is primarily due to the continuous nature of the action space in a SAC model. To adapt the continuous SAC model to the discrete action space of EVCH planning, our model first takes actions within a continuous space and then maps them to a discrete action set. We show that, for the case of EVCH, the proposed SAC-based model delivers on the key theoretical and practical benefits of RL: (1) scalability, (2) incorporation of operational detail and large-scale stochastic preference data, and (3) modeling flexibility. We also demonstrate that concerns around the lack of optimality guarantee are largely unfounded with our model converging closely to the global optimum across

all our experiments despite the integer relaxation adopted in our approach. We provide further details on these key results in the following.

In terms of scalability, we demonstrate experimentally that using soft actor-critic RL in combination with a data-driven simulation environment is scalable to real-world EVCH system sizes of 1,000 parking spots for which the model converges in under 12h. Optimization methods (similar to the one proposed in Li et al. (2020)) only scale to EVCH facility sizes of approx. 200 parking spots and fail to converge (within the 48h time constraint) for larger problems. An important caveat is the need for significant modeling simplifications, such as coarser temporal discretization, less operational detail and deterministic realization of normally uncertain parameters (e.g., arrival times, charging demand) to achieve tractability for these problem sizes with mathematical programming. Scalability (and optimality gap performance) of our SAC-based method also compares very favorably against alternative RL approaches, specifically the popular value-based DQN approach. For the 200 lot benchmark problem SAC converges within 400 episodes versus 700 episodes for DQN and achieves an optimality gap of just 10% versus 19% for DQN.

Another theoretical benefit of RL that we are able to exploit for this work on EVCHs and that distinguishes our method from extant optimization-based planning frameworks is the fact that RL scales almost independently of operational detail and data set size, due to the framework's reliance on a simulation rather than a mathematical model to capture system dynamics. This allows us to leverage data on preferences, operations and asset characteristics in great detail (e.g., 1 minute intervals) and over long periods of operational data (e.g., months). For example, we use real-time parking and charging data to develop a novel taxonomy of parker types along and their charging preferences. We show that parking events can be classified into one of six categories (e.g., business parkers, overnight parkers, etc.) and that archetypical facility types (e.g., a workplace facility) exhibit very distinct parker population patterns. These data-driven and very granular preference models power our EVCH simulation and allow us to align our simulation with real-world conditions as much as possible. We posit that this will result in better performance of the target system under real-world conditions. Indeed, using high-detailed simulation environments both in terms of temporal granularity and preference granularity provides significant and quantifiable value. Our experiments reveal better cost performance of the derived planning decision vs. models where we either use coarser time periods (approximately 20% cost increase vs. the benchmark case as we increase modeling granularity to 2h) or where we use distributional assumptions of preferences instead of real-world sensor data (significant 45% drop in service level compared to the benchmark case). While traditional stochastic or robust optimization approaches can be specified to account for uncertainty in future realizations of parameters, this comes at the cost of larger models and associated performance penalties rendering these alternatives impractical for the large-scale multi-stage stochastic EVCH planning problem.

As an added benefits of relying on a simulator rather than a mathematical model of the physical EVCH environment, the proposed method is extremely flexible. The key benefit of this modeling flexibility is the ability to conduct extensive scenario analyses regarding user preferences, operational policy choices, asset configurations and cost assumptions that have practical use for operations managers looking to make data-driven EVCH planning decisions. This can be achieved without requiring extensive model reformulation. We explore how infrastructure requirements change as asset operations become more sophisticated highlighting the value of such operational policies (see Section 4.5). We are also able to model complex interactions between a wide variety of different loads (building, EV charging, battery charging) and power sources (grid supply, PB, on-site battery storage), which sets this work apart from extant EVCH research (e.g., Kazemi et al. 2016, Li et al. 2020, Babic et al. 2022a). These scenario analyses yield several interesting findings that have practical implications for EVCH operations management. For example, we find that integrating generation (PV) and storage (on-site) assets into the EV charging hub is beneficial in most cases, but varies by facility type and the adopted charging and routing policies. We also demonstrate that the use of multi-server chargers is cost effective and can improve EVCH economics, especially if active vehicle routing and smart charging strategies are adopted and if the user population has high average laxity as would be the case in a typical workplace facility. Another interesting finding is that, in many scenarios, no significant investment in DC fast charging is required to achieve the desired service level. Opportunities to build on, expand and tailor these scenario analyses while leveraging our RL approach abound and we leave them for future work.

There are also several limitations of RL which we have explored extensively in this work and which we lay out here.

First, RL does not guarantee optimality (Sutton and Barto 2018a). This concern may be further exacerbated by the need for integer relaxation of several discrete decision variables in the proposed SAC framework. To provide evidence to the contrary, we run extensive simulation experiments on different EVCH system sizes to explore how closely to the global optimum our solution converges. We show that the solutions obtained in these experiments can be considered near-optimal (gaps between 4% and 15%).

It should also be considered that despite the highly detailed nature of the DT simulation, it is still an abstraction of reality that is subject to several limitations and simplifications. For example, we consider routing and charging decisions separately instead of jointly. We also do not allow for vehicle-to-grid operations and consider loads of the attached building loads to be exogenous, among other simplifications. Such limitations represent exciting avenues for follow-up work, and we leave them for future research.

Finally, our method is data hungry, meaning that it requires large amounts of granular operational data. We argue, that with the emergence of inexpensive sensor technology, ubiquitous computing, and mobile connectivity such data tends to be increasingly available (Choi et al.

2022). While the above-mentioned benefits of RL should warrant these data and implementation costs there is an additional argument in favor of the proposed method: as opposed to traditional OM planning models, the simulation environment (DT) used as part of the RL framework is not single-use. Indeed, the DT can be readily bridged-over into the use phase of the EVCH by simply replacing historical sensor data streams with real-time data flows (Boschert and Rosen 2016). In the system’s use phase, the DT then affords real-time system monitoring and optimization. We aim to exploit this multi-use characteristic in future work by leveraging the developed DT simulation in the development of novel high-performing learning algorithms for real-world EVCH operations.

In sum, our framework offers a novel practical method for OM practitioners to incorporate data-driven, high-fidelity simulators (i.e., DTs) combined with state-of-the-art reinforcement learning methods in the design phase of large-scale EVCH systems. The method yields near optimal planning solutions, scalability to real-world systems, the ability to incorporate and account for large-scale stochastic user and systems behavior as well as modeling flexibility.

In terms of scalability, we demonstrate experimentally that using soft actor-critic RL in combination with a data-driven simulation environment is scalable to real-world EVCH system sizes of 1,000 parking spots for which the model converges in under 12h. Optimization methods (similar to the one proposed in Li et al. (2020)) only scale to EVCH facility sizes of approx. 200 parking spots and fail to converge (within the 48h time constraint) for larger problems. An important caveat is the need for significant modeling simplifications, such as coarser temporal discretization, less operational detail and deterministic realization of normally uncertain parameters (e.g., arrival times, charging demand) to achieve tractability for these problem sizes with mathematical programming. Scalability (and optimality gap performance) of our SAC-based method also compares very favorably against alternative RL approaches, specifically the popular value-based DQN approach. For the 200 lot benchmark problem SAC converges within 400 episodes versus 700 episodes for DQN and achieves an optimality gap of just 10% versus 19% for DQN.

Another theoretical benefit of RL that we are able to exploit for this work on EVCHs and that distinguishes our method from extant optimization-based planning frameworks is the fact that RL scales almost independently of operational detail and data set size, due to the framework’s reliance on a simulation rather than a mathematical model to capture system dynamics. This allows us to model preferences, operations and asset characteristics in great detail (e.g., 1 minute intervals) and over long periods of operational data (e.g., months). For example, we use real-time parking and charging data to develop a novel taxonomy of parker types along and their charging preferences. We show that parking events can be classified into one of six categories (e.g., business parkers, overnight parkers, etc.) and that archetypical facility types (e.g., a workplace facility) exhibit very distinct parker population patterns. These data-driven and very granular preference models power our EVCH simulation and allow us to align our simulation with real-

world conditions as much as possible. We posit that this will result in better performance of the target system under real-world conditions. While traditional stochastic or robust optimization approaches can be specified to account for uncertainty in future realizations of parameters, this comes at the cost of larger models and associated performance penalties rendering these alternatives impractical for the large-scale multi-stage stochastic EVCH planning problem.

As an added benefits of relying on a simulator rather than a mathematical model of the physical EVCH environment, the proposed method is extremely flexible. The key benefit of this modeling flexibility is the ability to conduct extensive scenario analyses regarding user preferences, operational policy choices, asset configurations and cost assumptions that have practical use for operations managers looking to make data-driven EVCH planning decisions. This can be achieved without requiring extensive model reformulation. We explore how infrastructure requirements change as asset operations become more sophisticated highlighting the value of such operational policies (see Section 4.5). We are also able to model complex interactions between a wide variety of different loads (building, EV charging, battery charging) and power sources (grid supply, PB, on-site battery storage), which sets this work apart from extant EVCH research (e.g., Kazemi et al. 2016, Li et al. 2020, Babic et al. 2022a). These scenario analyses yield several interesting findings that have practical implications for EVCH operations management. For example, we find that integrating generation (PV) and storage (on-site) assets into the EV charging hub is beneficial in most cases but varies by facility type and the adopted charging and routing policies. We also demonstrate that the use of multi-server chargers is cost effective and can improve EVCH economics, especially if active vehicle routing and smart charging strategies are adopted and if the user population has high average laxity as would be the case in a typical workplace facility. Another interesting finding is that in many scenarios no significant investment in DC fast charging is required to achieve the desired service level. Opportunities to build on, expand and tailor these scenario analyses while leveraging our RL approach abound and we leave them for future work.

There are also several limitations of RL which we have explored extensively in this work and which we lay out here.

First and foremost, RL does not guarantee optimality (Sutton and Barto 2018a). This concern may be further exacerbated by the need for integer relaxation of several discrete decision variables in the proposed SAC framework. To provide evidence to the contrary, we run extensive simulation experiments on different EVCH system sizes to explore how closely to the global optimum our solution converges. We show that solutions obtained in these experiments can be considered near-optimal (gaps between 4% and 15%).

It should also be considered that despite the highly detailed nature of the DT simulation, it is still an abstraction of reality that is subject to several limitations and simplifications. For example, we consider routing and charging decisions separately instead of jointly. We also do not allow for vehicle-to-grid operations and consider loads of the attached building loads to be

exogenous, among other simplifications. Such limitations represent exciting avenues for follow-up work, and we leave them for future research.

Finally, our method is data hungry, meaning it requires large amounts of granular operational data. We argue, however, that with the emergence of inexpensive sensor technology, ubiquitous computing, and mobile connectivity such data tends to be increasingly available (Choi et al. 2022). While the above-mentioned benefits of RL should warrant these data and implementation these costs there is an additional argument in favor of the proposed method: as opposed to traditional OM planning models, the simulation environment (DT) used as part of the RL framework is not single-use. Indeed, the DT can be readily bridged-over into the use phase of the EVCH by simply replacing historical sensor data streams with real-time data flows (Boschert and Rosen 2016). In the system’s use phase, the DT then affords real-time system monitoring and optimization. We aim to exploit this multi-use characteristic in future work by leveraging the developed DT simulation in the development of novel high-performing learning algorithms for real-world EVCH operations.

In sum, our framework offers a novel practical method for OM practitioners to incorporate data-driven, high-fidelity simulators (i.e., DTs) combined with state-of-the-art reinforcement learning methods in the design phase of large-scale EVCH systems. The method yields near optimal planning solutions, scalability to real-world systems, the ability to incorporate and account for large-scale stochastic user and systems behavior as well as modeling flexibility.

4.7 Appendix

4.7.1 Operational Modeling and Algorithms

Optimal Charging Approach

Our charging model reconsiders the charging rate for connected vehicles at the beginning of each planning interval, thus following an online optimization paradigm.

The objective is to minimize the energy costs while meeting the charging demand of all connected vehicles (up to a predefined service level) over the look ahead planning window \mathcal{T}^Ω . As we consider an uncertain case where the perfect information of upcoming vehicles is not available to the decision model, we make this conservative assumption to ensure that the service level is fulfilled.

First, the planning horizon must be chosen carefully due to its significant effect on model performance. In our simulations, we limit it to 6 hours, which, given a planning interval Δ^Ω of 15 minutes, breaks down to 24 decision steps (i.e., charging rates are re-computed every 15 minutes of simulation time). We term the set of decision steps \mathcal{T}^Ω .

Second, we consider a flexibility margin μ_t to accommodate future, yet unknown demand. Although this model outputs a vector of charging rates for each vehicle, we only use the first charging rate and reconsider decisions in the next charging time step based on the updated system state including the new arrived vehicles.

$$\text{Min}_{\Xi} \sum_{t \in \mathcal{T}^\Omega} T_t^e e_t^{Grid} + T^p p^* \quad (4.8)$$

The grid energy consumption e_t^{Grid} accounts for the charging of vehicles (variables) as well as the storage (dis)charging, PV generation and building loads (storage rate is given as parameter before charging management). Constraint Eq. (10) guarantees that the grid energy consumption does not exceed the grid capacity minus the safety threshold we consider for the following time steps. We compute the induced peak in Eq. (11).

$$e_t^{Grid} = \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}} \Delta_t (\psi_{k,j,t} + \beta_t^{Charge} - \beta_t^{Discharge} - f_t^{PV} \sum_{\tau=0} \kappa_\tau^{PV} + l_t) \quad \forall t \in \mathcal{T}^\Omega \quad (4.9)$$

$$\frac{e_t^{Grid}}{\Delta_t} \leq p_t^{Grid} - \mu_t \quad \forall t \in \mathcal{T}^\Omega \quad (4.10)$$

$$p^* \geq \frac{e_t^{Grid}}{\Delta_t} - l^* \quad \forall t \in \mathcal{T}^\Omega \quad (4.11)$$

We also ensure that all vehicles receive at least η percentage of their charging demands (Eq. (12)). Constraint Eq. (13) ensures that vehicle can only charge when they are physically present in the EVCH. Finally, Constraint Eq. (14) restricts the parallel charging of vehicles that are connected to charging dock k to its charging capacity.

$$\sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}^\Omega} \psi_{k,j,t} \geq \eta e_j^D \quad \forall j \in \mathcal{J} \quad (4.12)$$

$$0 \leq \psi_{k,j,t} \leq U_{j,t} M \quad \forall k \in \mathcal{K}, j \in \mathcal{J}, t \in \mathcal{T}^\Omega \quad (4.13)$$

$$\sum_{j \in \mathcal{J}_k} \psi_{k,j,t} \leq \kappa^k \quad \forall k \in \mathcal{K}, t \in \mathcal{T}^\Omega \quad (4.14)$$

4.7.2 Preference Modeling

In this Section we provide details on the clustering routine and robustness test employed to identify parking archetypes. We cluster parking events j based on A_j and δ_j , the two core parameters of interest at this modeling stage. To account for the circular nature of arrival time A_j , which is not captured accurately by any distance-based clustering algorithm (for example, entries at 23:59h and 0:00h would be considered furthest apart despite their obvious proximity), we create two circular features $A_j^{sin} = \sin(2\pi(A_j/24))$ and $A_j^{cos} = \cos(2\pi(A_j/24))$. This yields the following vector of clustering variables $v_j^{clust} = (A_j^{sin}, A_j^{cos}, \delta_j)$, which we normalize.

Given the size of our dataset (3.84M observations) we limit our algorithm search to clustering algorithms that are sufficiently scalable. We run initial tests with three clustering algorithms: k-means++, a centroid-based algorithm, Gaussian Mixture Models (GMM) and BIRCH, a scalable density-based clustering algorithm. Overall, we find k-means++ to perform best in terms of runtime and stability. While GMM yields relatively similar results, BIRCH performs very poorly, yielding unstable and non-cohesive clusters suggesting that relative density may not be a good identifier of clusters for the given dataset. We thus focus on fine tuning k-means++. A major challenge in the application of k-means++ is to select the number of centroids (clusters) k that are to be initialized and optimized for. To identify good candidate choices for k , We initially test integer values over an interval of reasonable values $[0, 20]$ and compute Calinski-Harabasz scores per each clustering outcome (Calinski and Harabasz 1974). These analyses suggest $k = 5$ or $k = 6$ to be good choices. To validate and further narrow down our choice for k , we perform silhouette analyses for both candidate choices (Rousseeuw 1987). We obtain the highest average silhouette coefficient $\bar{\mathcal{H}}$ for $k = 6$ ($\bar{\mathcal{H}} = 0.420$). Finally, taking $k = 6$ as the best performing choice across the above described internal validity measures, we conduct extensive cross-validation to assess cluster outcome robustness. We iteratively perform 2-1 splits of the data and re-run k-means++ on the larger dataset, then use the fitted algorithm to predict the labels of the smaller (test) dataset. We find our clustering results to be stable with observations in the test set having the same label 99.14% ($\sigma = 0.51$ %, 100 replications) of the time. We run an additional set of robustness analyses, this time focusing on the amount of preference data that is required to identify parker types reliably. This analysis draws on Griffin and Hauser (1993) who looked at the question of how many customer interviews were required for reliable insights. Clearly, there is a benefit to prospective EVCH planners if the need for data was smaller than the full one




Figure 4.10: Distribution of clustering variables per each cluster

year period we have considered thus far. We run tests for 2 weeks, 4 weeks and 12 weeks of data per each facility and check the robustness of the clustering results as compared to the clusters obtained on the full one-year facility dataset. We find that high quality clustering results can be obtained with just three weeks of data (95.28%, $\sigma = 3.49$ %) accuracy vs. full one-year facility dataset). Beyond this threshold the value of additional data appears to diminish. At six weeks of data, for example, we obtain very similar accuracy (95.85% ($\sigma = 1.95$ %)), albeit slightly lower variance. In addition to internal validity and robustness of our clustering results we look at interpretability (or external validity). For this purpose, we presented our final clustering results (for $k = 6$) to a range of practitioners and discussed their implications. The clusters were deemed consistent with the domain experts' experience. In sum, we obtain six parker types that are supported both by internal criteria and real-world observation and can be readily identified with just three weeks of data. Figure 4.10 shows the distributions of the two clustering variables per each final cluster.

4.7.3 Solution Frameworks and Models

In the following, we present the specifications of the model architectures developed and used in this research.

The common hyperparameters between the RL models are specified in Table 4.4. Model-specific parameters are explained separately. For each EVCH facility size, we tune the hyperparameters using a brute-force grid search, focusing on learning rate, batch size, training frequency, and hidden layers. The goal of the grid search is to find the best set of hyperparameters that converges to the highest objective function after training. Here we show the final hyperparameters for the main configuration used in the model benchmarking (200 parking spaces).

Table 4.4: Hyperparameter Configuration for Reinforcement Learning Algorithms

Parameter	DQN	DDPG	SAC
Optimizer (learning rate)	Adam (0.001)	Adam (0.0001)	Adam (0.0001)
Loss function	MSE	MSE	MSE
Discount factor (γ)	0.99	0.99	0.99
Memory capacity	1,000,000	1,000,000	1,000,000
Steps prior to learning	1024	256	256
Training frequency	10	10	10
Batch size	64	256	256
Exploration strategy	Epsilon decay	Action noise	Entropy maximization
Target network update rate (τ)	0.01	0.01	0.05
Hidden layer activation function	ReLU	ReLU	ReLU
Number of hidden layers (nodes)	3 (256, 512, 256)	4 (256, 512, 512, 256)	4 (256, 512, 512, 256)

Main Model: Soft Actor-Critic Reinforcement Learner

Actor-Critic and Soft Actor-Critic models are reinforcement learning algorithms designed to identify optimal policies for sequential decision-making problems, but they differ significantly in their approaches and objectives. Generally, Actor-Critic models consist of two key components: the actor, which defines the policy function responsible for selecting actions, and the critic, which evaluates the actor’s actions using a value function (Konda and Tsitsiklis 1999). This framework aims to enhance policy performance by reducing variance in policy gradients through value-based feedback. SAC extends this model by including an entropy term in the objective function, encouraging the agent to balance exploration and exploitation by optimizing a trade-off between maximizing expected returns and policy entropy (Haarnoja et al. 2018). This objective improves training stability and sample efficiency, especially in environments with continuous or large action spaces. Additionally, SAC employs an off-policy approach, utilizing a replay buffer for more efficient data usage, whereas traditional Actor-Critic methods are often on-policy and require fresh data for updates (though not all Actor-Critic models are strictly on-policy). These characteristics make SAC particularly well-suited for complex, high-dimensional tasks requiring robust exploration and stability, such as the problem addressed in our study.

We utilize identical network architectures for both the actor and critic. We experiment with various hyperparameters for the actor and critic networks but found no significant impact on their performance, except for doubling the grid search size. The only contrast between the two is that we use a Tanh activation function in the final layer of the actor network, which yields a value between -1 and 1 for each sub-action and must be proportionally scaled based on the sub-action range. The SAC model utilizes Adam with a learning rate of 0.0001 to update the networks while implementing a batch size of 256. An automatic entropy tuning feature is also employed which is critical for SAC to minimize the impact of entropy in the objective function in order to achieve stable policies (Haarnoja et al. 2018). As suggested by existing literature

(Haarnoja et al. 2018), a constant additional noise is applied during the training phase to enable the model to sufficiently explore the environment and mitigate the risk of local optima. The noise function follows a normal distribution with a mean of zero and variance of 0.05. We use Kaiming uniform initialization (also known as He initialization), which is the default weight initialization for a fully connected (feedforward) neural network (FNN) layer. This initialization is well suited for layers with ReLU or similar activation functions, as Kaiming initialization helps to maintain the variance of activations across layers. In addition, our model set the initial values of the biases to zero.

Benchmark RL Model: DQN Reinforcement Learner

We use deep double Q-networks as the benchmark model in our study. For further details, please refer to Van Hasselt et al. (2016). The rationale behind Double Q-learning is to mitigate overestimation by breaking down the maximum operation in the target into action selection and evaluation. This results in the usage of two networks, one each for action selection and evaluation purposes. The action evaluation network is termed target network and undergoes lesser updates compared to the main network. To adjust the parameters of the target network, we adopt soft update, which is implemented as follows:

$$\theta' \leftarrow (1 - \tau)\theta' + \tau\theta \quad (4.15)$$

Where θ' represents the target network parameters, θ represents the main network parameters, and τ is the soft update weight, which ranges from zero to one. The hyperparameter grid search, as shown in Table 1, indicates that DQN requires a higher learning rate, lower batch size, and smaller hidden layers compared to SAC. Our exploration strategy employs an epsilon-decay algorithm, where random actions are chosen with decreasing probability of epsilon during the training process. In our model, epsilon begins at 0.3 and decreases to 0.01 after 600 episodes, and remains unchanged thereafter.

Benchmark RL Model: Deep Deterministic Policy Gradients

Deep Deterministic Policy Gradients (DDPG) is a model-free, off-policy reinforcement learning algorithm designed for continuous action spaces Lillicrap (2015). It combines the strengths of both Q-learning and policy gradient methods. Similar to SAC, DDPG utilizes an actor-critic architecture, where the actor network learns a deterministic policy to map states to actions, and the critic network evaluates the Q-value of the state-action pairs. Inspired by Deep Q-Networks (DQN), DDPG uses a replay buffer to store experiences and sample mini-batches for training, ensuring decorrelated updates and improved stability. Additionally, we employ target networks for both the actor and critic to stabilize training by providing a slowly updated, consistent set

of parameters. To encourage exploration in deterministic policies, we add noise to the actions during training. Although DDPG is designed for continuous action spaces, we adapt it to the integer action space of our problem using the same modifications applied to our proposed SAC model.

Upper Bound Benchmark: Mathematical Programming Model

A Mathematical Programming Model acts as upper bound in our benchmarks and is used to compute optimality gaps of RL-based models. We formulate the decision challenge as a feasibility problem which aims to satisfy all or a specified amount of total charging demand most resource efficiently while considering rate, space, and total capacity constraints. In doing so we expand on and adapt extant EVCH planning models (e.g., Li et al. 2020).

In line with the planning objective, we frame the problem as a cost minimization planning with the goal to jointly minimize the investment cost (C^Φ) and the operations cost (C^Ω) of the EVCH while ensuring a certain service level η_h^{Serv} . Formally, the objective can be expressed as follows:

$$Min_{\Xi}[(C^\Phi(x_{k,h}^{i,n}, \delta_h^{Trafo}, \kappa_h^{PV}, \epsilon_h^{Bat}) + C^\Omega(\omega_{k,j,h}, \psi_{k,j,h,t}, \beta_{h,t}^{Charge}, \beta_{h,t}^{Discharge})] \quad (4.16)$$

Both cost items are defined as follows. The investment cost (C^Φ) is the sum of the grid expansion cost (if any), the cost of charging infrastructure plus any installed PV and battery capacity over the full investment horizon \mathcal{H} . The operations cost (C^Ω) is defined as the total sum of electricity costs over the investment horizon, where costs are only incurred on the electricity retrieved from the grid with $e_{h,t}^{Grid}$. Formally:

$$C^\Phi = \sum_{h \in \mathcal{H}} [(c_h^T \delta_h^{Trafo} + \sum_{k \in \mathcal{K}} c_h^{i,n} x_{k,j,h}^i + c_h^{PV} p_h^{PV} + c_h^{Bat} \epsilon_h^{Bat})(1 + (|\mathcal{H}| - h)\mu^{Maint})] \quad (4.17)$$

$$C^\Omega = \sum_{h \in \mathcal{H}} \sum_{t \in \mathcal{T}} T_{h,t}^e e_{h,t}^{Grid} + T_h^p p_h^* \quad (4.18)$$

Note that $e_{h,t}^{Grid}$ is accounted for on the basis of a two-part tariff charging for both the use of electricity from the grid (excl. PV generation and possible battery discharge $\beta_{h,t}^{Discharge}$) and demand charges arising from the induced peak load attributable to EVCH operations. Demand charges T_h^p are designed to incentivize efficient utilization of the grid (Gust et al. 2021) and are typically based on the monthly peak load induced by the facility. We therefore define p^* as the excess of the expected base facility peak load l^* (excl. EVCH operations) for state h (Eq. (20)).

$$e_{h,t}^{Grid} = \sum_{j \in \mathcal{J}_h} \sum_{k \in \mathcal{K}} \Delta_t (\psi_{k,j,h,t} + \beta_{h,t}^{Charge} - \beta_{h,t}^{Discharge} - f_{h,t}^{PV} \sum_{\tau=0}^h \kappa_\tau^{PV} + l_{h,t}) \quad (4.19)$$

$$p_h^* \geq \frac{e_{h,t}^{Grid}}{\Delta_t} - l_h^* \quad \forall h \in \mathcal{H}, t \in \mathcal{T} \quad (4.20)$$

The optimization is subject to additional operational and physical constraints.

$$\sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} \psi_{k,j,h,t} \geq \eta_h^{Serv} e_j^D \quad \forall h \in \mathcal{H}, j \in \mathcal{J}_h \quad (4.21)$$

$$x_{k,h}^{i,n}, \omega_{k,j,h} \in \{0, 1\} \quad \forall k \in \mathcal{K}, h \in \mathcal{H}, j \in \mathcal{J}_h, i \in \mathcal{I}, i \in \mathcal{N} \quad (4.22)$$

$$\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{I}} \sum_{n \in \mathcal{N}} \sum_{h \in \mathcal{H}} x_{k,h}^{i,n} n \leq L \quad (4.23)$$

$$\sum_{h \in \mathcal{H}} \sum_{i \in \mathcal{I}} \sum_{n \in \mathcal{N}} x_{k,h}^{i,n} \leq 1 \quad \forall k \in \mathcal{K} \quad (4.24)$$

$$\sum_{j \in \mathcal{J}_h} \omega_{k,j,h} U_{j,h,t} \leq \sum_{\tau=0}^h \sum_{i \in \mathcal{I}} x_{k,\tau}^{i,n} n \quad \forall k \in \mathcal{K}, h \in \mathcal{H}, t \in \mathcal{T} \quad (4.25)$$

$$\sum_{k \in \mathcal{K}} \omega_{k,j,h} \leq 1 \quad \forall h \in \mathcal{H}, j \in \mathcal{J}_h \quad (4.26)$$

$$0 \leq \psi_{k,j,h,t} \leq \omega_{k,j,h} U_{j,h,t} M \quad \forall k \in \mathcal{K}, h \in \mathcal{H}, j \in \mathcal{J}_h, t \in \mathcal{T} \quad (4.27)$$

$$\sum_{j \in \mathcal{J}_h} \psi_{k,j,h,t} \leq \sum_{\tau=0}^h \sum_{i \in \mathcal{I}} \sum_{n \in \mathcal{N}} x_{k,\tau}^{i,n} \kappa^i \quad \forall k \in \mathcal{K}, h \in \mathcal{H}, t \in \mathcal{T} \quad (4.28)$$

$$\sum_{h \in \mathcal{H}} \kappa_h^{PV} \leq R \quad (4.29)$$

$$SoC_{h,t} = SoC_{h,t-1} + (\beta_{h,t-1}^{Charge} - \beta_{h,t-1}^{Discharge}) \Delta t \quad \forall h \in \mathcal{H}, t \in \{1, 2, \dots, \mathcal{T}\} \quad (4.30)$$

$$SoC_{h,0} = \sum_{\tau=0}^h SoC^{\min} \epsilon_{\tau}^{Bat} \quad \forall h \in \mathcal{H} \quad (4.31)$$

$$\beta_{h,t}^{Charge} \leq \beta_{h,t}^{Direction} \beta^{max} \quad \forall h \in \mathcal{H}, t \in \mathcal{T} \quad (4.32)$$

$$\beta_{h,t}^{Discharge} \leq (1 - \beta_{h,t}^{Direction}) \beta^{max} \quad \forall h \in \mathcal{H}, t \in \mathcal{T} \quad (4.33)$$

$$\beta_{h,t}^{Charge} \Delta t \leq \sum_{\tau=0}^h SoC^{\max} \epsilon_{\tau}^{Bat} - SoC_{h,t-1} \quad \forall h \in \mathcal{H}, t \in \{1, 2, \dots, \mathcal{T}\} \quad (4.34)$$

$$\beta_{h,t}^{Discharge} \Delta t \leq SoC_{h,t-1} \quad \forall h \in \mathcal{H}, t \in \{1, 2, \dots, \mathcal{T}\} \quad (4.35)$$

$$e_{h,t}^{Grid} \leq \Delta t (\kappa_0^{Grid} + \sum_{\tau=0}^h \delta_{\tau}^{Trafo} \kappa^{Trafo}) \quad \forall h \in \mathcal{H}, t \in \mathcal{T} \quad (4.36)$$

First and foremost, service level is guaranteed in Eq. (21). Note that the summation is bounded by set \mathcal{T} , meaning that we consider the total supplied energy at the time of departure. This important constraint ensures that adequate infrastructure is provisioned despite the cost minimization objective.

EV charging infrastructure decisions and operations are controlled by means of decision variables $x_{k,h}^{i,n}$, $\omega_{k,j,h}$ (both binary indicators, see Eq. 22) and $\psi_{k,j,h,t}$. First, the number of charging docks and associated connectors is restricted by the space constraints L of the facility (Eq. (23)). Similarly, Eq. (24) ensures that candidate points can only be equipped with chargers

once and that this decision cannot be changed over the planning period, i.e., they cannot be removed once installed.

In terms of routing and charging operations, the model assigns vehicles to chargers upon arrival (one-off decision) and periodically adjust the charging power over the duration of their visit. Constraint Eq. (25) allocates vehicle j to spot k during stage h only if k is equipped with a charging dock and only if j is present in the EVCH (captured via $U_{j,h,t}$). Eq. (28) ensures that each vehicle connects to at most one charging dock. Constraint Eq. (27) guarantees that vehicle j receives non-negative energy (bounded by the maximum power of the specific dock in Eq. (28)) from charging dock k only if it is connected to k .

Battery and on-site generation constraints are set as follows. We assume PV generation to be non-controllable meaning no constraints are necessary to model their operations (in-feed is an exogenous parameter). We simply limit the maximum installable PV capacity $\sum_{h \in \mathcal{H}} \kappa^{PV}$ to the available on-site space (such as rooftop space) R (see Eq. (29)). Eq. (30) through (35) implement various battery-related constraints. Constraint Eq. (30) incrementally updates the battery state of charge $SoC_{h,t}$. Constraint Eq. (31) ensures that the battery SoC remains within a certain interval. Note that we neglect efficiency losses and assume battery depreciation to be independent of operations (Sharifi et al. 2020). We realize that these are simplifications, yet these are necessary to retain tractability of our model. We assume symmetric charge/discharge rate limits which are enforced through constraint Eq. (32) and (33), where $\beta^{max} \geq 0$. These constraints also ensure that the battery cannot be charged and discharged at the same time.

Our model ensures that the EVCH's base load as well as EV and battery charging loads cannot exceed the total grid capacity (existing and extension) plus current PV generation, which is enforced by Eq. (36). Note that if the battery was discharging (negative β_t^{Bat}) this would increase the available capacity.

4.7.4 Experimental Setup

In this Section we provide details on the experimental setup and parameterization of the Digital Twin (DT) simulation environment. Table 4.5 provides details on the key components of the DT environment and the digitalization approach adopted (real-world sensor data vs. simulation).

Note that we rely on real-world sensor data to model system dynamics wherever available and resort to simulations informed by research papers and/or asset specification sheets in all other cases (following Sierla et al. (2018)). In addition, we impose several physical constraints inherent to the various components of the EVCH system. These are summarized in Table 4.6.

Table 4.7 summarizes the core investment-related parameters (costs and space constraints) used in the benchmark experiments over the investment horizon (10 states $s \in S$). Energy costs are based on real-world electricity tariffs from the same region in which the charging data were gathered (i.e., California). Table 5.4 gives an overview of the tariff structure that is used throughout all experiments.

Table 4.5: Digital Twin (DT) components and associated datasets

DT Component	Type	Digitalization approach	Description	Source
Local Substation	Physical Asset	simulated	Transformation losses and physical limit	assumptions as detailed in Table 4.6
On-site Electricity Generation Assets (PV Panels)	Physical Asset	real-world sensor data	PV load factors (power output as percentage of installed capacity)	Open Power System Data provided by Neon Neue Energieökonomik and Technical University of Berlin and ETH Zürich and DIW Berlin (2024)
Electricity Storage Assets	Physical Asset	simulated	(dis-)charging efficiency curves, physical constraints (min/max state of charge)	Ghiassi-Farrokhfal et al. (2016); assumptions as detailed in Table 4.6
EV Charging Docks and Connectors	Physical Asset	simulated	AC-DC conversion losses, maximum charging capacity	assumptions as detailed in Table 4.6
Peripheral Building Electricity Consumption	Preference Pattern	real-world sensor data	Peak load in KW and consumption in kWh per 15-min interval	Unique real-world meter data provided by a major European real-estate investor; includes peak building loads and consumption at 15-minute resolution across thirteen facilities with different usage profiles
Parking Demand	Preference Pattern	real-world sensor data	Vehicle-level time of arrival and duration of stay obtained from parking garage sensors	Unique transaction-level parking dataset provided by a major European real-estate investor; includes transactions from seven large-scale parking garages catering to different parking use cases (office building, destination parking, mixed-use)
Charging Demand	Preference Pattern	real-world sensor data	Requested energy per charging session in kWh	Real-world dataset by Lee et al. (2019) containing >25,000 charging transactions for the year 2019

¹⁶All cost parameters include cost of installation and peripheral equipment (e.g., inverters for battery, PV and DC chargers)

¹⁷22kW, single connector

¹⁸50kW, single connector


¹⁹Assuming sales share equals penetration for given facility

²⁰Cost of transformer, cabling and contribution to upstream grid upgrades; assuming 5% yearly cost increase

²¹Assuming 500m² roof space and PV energy density of 0.2kWp/m²

Table 4.8: Time-of-use tariff and demand charge for large-scale EV charging customers (> 500 kW)

	Summer (Jun - Sep)	Winter (all other months)
Super Off-Peak (8am-4pm)	0.08 USD/kWh	0.06 USD/kWh
On-Peak (4pm to 9pm)	0.23 USD/kWh	0.23 USD/kWh
Off-Peak (9pm-8am)	0.08 USD/kWh	0.08 USD/kWh
Demand Charge (monthly)	15.48 USD/kW	

**Figure 4.11:** Investment decision derived by the upper limit mathematical programming model

4.7.5 Supplemental Results

Investment decision by time period

In this Section we take a closer look at the dynamics of the investment decisions derived by the four decision frameworks implemented in this work (optimal, DQN, DDPG, SAC). This supplements the benchmark results shown in Section 4. It allows us to analyze the differences in planning decisions between the different planning algorithms in much more detail. Figures 4.11 through 4.14, show the investment plans for charging infrastructure (left) and power supply infrastructure (right) derived via mathematical programming (optimal), DQN, DDPG and SAC, respectively. Note that in order to achieve comparability with the benchmark mathematical model, the experiments run in Section 4 do not consider on-site storage, hence storage is not being built in any planning period.

There are a few very interesting insights to be taken from this analysis. First, the planning decisions made by SAC are very close to the optimal decisions. Under both investment plans charging demand is primarily served through AC charging using 4 connectors per charger. Although the scale up path is slightly different for the supply side, both decision algorithms ultimately arrive at the same end state (100kW PV plus 200kW of grid extension). This further underlines the robust and near-optimal performance of SAC. DDPG achieves the same sup-




Figure 4.12: Investment decision derived by DQN




Figure 4.13: Investment decision derived by DDPG

ply decisions. However, it relies on significantly more charging infrastructure including a large amount of single-server AC chargers yielding suboptimal cost performance vs. the upper limit optimal investment plan and SAC.

DQN takes a notably different set of planning decisions compared to the other decision frameworks. This applies both to the number of charging docks that are being installed and to the investments in power supply. DQN installs the smallest number of docks and is the only framework to install DC chargers in significant numbers. Although this charging infrastructure may be able to serve the charging demand (primarily by providing higher average charge rates), this comes at the expense of significantly higher combined peak loads. As a result, the charging cluster configuration suggested by the DQN algorithm requires roughly twice the amount of grid updates to serve these loads compared to the other benchmark algorithms. This is suboptimal from a total cost perspective and highlights the value of multi-server charging docks and lower charging rates.

Algorithm: SAC




Figure 4.14: Investment decision derived by SAC

Chapter 5

A Pricing Decision Framework for Electric Vehicle Charging Hubs¹

5.1 Introduction

Transportation systems are a major contributor to climate change (Pal et al. 2023). A crucial step in reducing their environmental footprint is the widespread adoption of EVs (Nanaki and Koroneos 2016). Ensuring accessible and convenient charging at key locations such as workplaces and shopping centers is essential to accelerating EV adoption (Egbue and Long 2012). This not only reduces range anxiety for EV users, but also facilitates smart charging strategies that help balance grid demand and provide additional grid services (Kahlen et al. 2024). As a result, significant investment in large-scale non-residential charging infrastructure is needed, particularly in urban areas where many residents lack access to home charging (Lee et al. 2019). We define these high-density EV charging clusters as EV Charging Hubs (EVCHs).

Expanding charging infrastructure faces financial challenges due to high upfront costs (Engel et al. 2018). Pricing management can enhance profitability and attract investment (Lin et al. 2023). Utilities often use ToU pricing, where electricity rates vary by time of day, but this can shift demand to lower-priced periods, creating new peak loads and raising operational costs (Yang et al. 2021). Dynamic pricing offers a more effective solution by adjusting prices in real time based on supply and demand (Valogianni et al. 2020a). Unlike fixed or ToU rates, dynamic pricing provides precise price signals, encouraging users to optimize energy consumption, improving grid stability, and enabling more flexible load management. These benefits will become even more important as EV adoption and renewable energy integration increase, bringing new challenges to the system such as supply uncertainty and new significant loads.

¹This Chapter is currently under review at a leading peer-reviewed academic journal.

Parts of this Chapter have appeared in the following (non-copyrighted) peer-reviewed academic conference: Ahadi, R., Schroer, K., & Ketter, W. (2024). Managing electric vehicle charging hubs through dynamic capacity-based pricing. ECIS 2024 Proceedings, 38.

We aim to support EVCH managers in enhancing both economic outcomes and system efficiency (e.g., peak load reduction) through dynamic pricing management². EVCHs are profit-oriented entities, generating revenue by fulfilling charging requests while incurring energy costs and system-related expenses associated with grid stability disturbances and inefficient consumption patterns. Therefore, maximizing profits for EVCHs requires a balance between increasing energy sales at competitive prices and optimizing power consumption to reduce costs. Consequently, designing effective dynamic pricing strategies for EVCHs is challenging, given the complex interactions among EV users, charging operators, and the power grid, compounded by uncertainties on both demand (e.g., EVCH occupancy, and individual charging demand), and supply (e.g., local generation, electricity prices). To address this, we propose a machine learning-enhanced decision support system (DSS) that dynamically adapts to system conditions, helping managers optimize performance. Our model not only boosts EVCH profitability but also promotes sustainability by facilitating EV adoption and managing charging loads to prevent costly grid expansions. Additionally, we consider price caps to ensure affordability for users.

Our research lies within the field of energy informatics (Watson et al. 2010), which combines energy systems and information technology to improve the efficiency, sustainability, and reliability of energy systems through computational methods. In a branch of this field, researchers develop decision frameworks to determine pricing strategies for profit maximization and demand response management (e.g., Bichler et al. 2023, Valogianni et al. 2020a). In particular, in the area of EV charging management, researchers have developed pricing DSSs and optimization models as a means to reduce peak loads (Flath et al. 2014, Chen et al. 2024a), address uncertainties in energy demand and renewable energy production (Soares et al. 2017), and increase the profitability of charging services in parking lots (Zanvettor et al. 2022). Recent research has explored advanced pricing models for EV charging. For instance, Lu et al. (2023) show that menu-based pricing³ could surpass traditional pricing models for charging station management. We build on this research stream by proposing a machine learning-based DSS to optimally manage the charging demand of EVCHs. Our proposed model is one of the first pricing management solutions for charging services that not only maximizes profits by adjusting prices based on EV users' time-dependent demand preferences (e.g., price sensitivity), but also optimizes total energy consumption to achieve system-level goals such as peak shaving and reshaping consumption patterns. Accordingly, we define our research question as follows:

How can a machine learning-based DSS be designed to enhance the economic performance of EV charging facilities and simultaneously shape aggregated load curve?

The increasing availability of data, combined with advancements in computational power and the Internet of Things (IoT), enables seamless communication between EV charging hub (EVCH) components. This allows operators to leverage data-driven and machine learning tools

²Although our approach is applicable to public charging stations, this study focuses on EVCHs due to their pivotal role in the future energy landscape.

³Menu-based pricing allows users to choose lower-cost charging options by extending their stay.

to enhance decision-making (Provost and Fawcett 2013, Babic et al. 2022b). We propose a deep reinforcement learning (DRL) algorithm that learns efficient dynamic pricing policies by interacting with real-world and simulated environments, optimizing both the profitability and sustainability of EVCHs. Our results show that the proposed DSS, powered by machine learning, outperforms traditional pricing models for charging services. Beyond profitability, EVCHs must consider system-level sustainability factors, such as peak load reduction and grid integration. Neglecting these aspects can lead to inefficient charging behaviors, such as excessive electricity peaks, increasing costs for both the grid and charging providers. Our DSS mitigates these risks by using capacity-based dynamic pricing signals to shape electricity consumption patterns, avoiding high peaks and aligning demand with renewable energy production.

To solve the pricing management of EVCHs, traditional optimization models can become intractable due to the curse of dimensionality, even in deterministic settings. Dynamic pricing for EVCHs involves three key aspects: a) Sequential decision making - at each time step, the DSS must set price signals that influence the system and affect future decisions; b) Stochastic nature - the parameters of the problem, such as demand preferences and supply capacities, are uncertain and their distributions are largely unknown; and c) Large scale - EVCHs consist of numerous interconnected components (e.g., multiple charging vehicles, various operational processes) that must work in coordination. To address these challenges, we develop a DRL algorithm that combines reinforcement learning (RL) algorithms with deep learning techniques suitable for solving large-scale stochastic sequential decision problems (Sutton and Barto 2018b). Specifically, we use soft actor-critical (SAC) models, which do not require mathematical models of system dynamics (e.g., transition probabilities) and are capable of managing continuous action spaces (price signals) while converging to stable policies. The proposed DRL agent is trained using a simulated digital representation of an EVCH, including detailed operational and user behavioral characteristics, which allows the agent to interact with the environment and learn how to dynamically adjust price signals to maximize profits and system-level efficiency. We employ agent-based modeling (ABM) to capture the complex interactions between various components (e.g., EV users, EVCH, and the power grid interface) and enhance the simulation with unique real-world data for actionable insights.

We offer the following theoretical and practical contributions.

- We present a novel machine learning-based DSS to help EVCH operators maximize economic performance while effectively managing aggregated charging demand through interactions with EV users. Our DSS consists of two key components: a DRL algorithm that learns optimal decision-making policies, and a capacity-based pricing model that enables operators to better influence charging requests (Valogianni et al. 2024). In capacity-based model, costs are tied to charging power⁴, with higher power demands resulting in higher

⁴charging power refers to the rate at which energy is delivered to a device or system during charging, typically measured in watts (W) or kilowatts (kW)

prices. This approach is particularly beneficial for large-scale EVCHs, where uncontrolled charging from numerous EVs can cause significant power spikes.

- We develop a DRL algorithm to optimize pricing policies under realistic conditions, considering stochastic demand, large-scale charging facilities, grid and charging constraints, supply volatility from on-site renewable energy generation, and fluctuating building energy consumption⁵. Therefore, to optimize pricing decisions we need a model that is able to track a wide range of variables such as the state of EVCH, charging schedules and grid information. DRL has demonstrated strong performance in solving large-scale sequential and stochastic decision problems (Wang et al. 2023a). In our problem, several pricing parameters vary continuously, making it well-suited for a soft actor-critical (SAC) algorithm. SAC is model-free, with no need for prior knowledge of system dynamics, which is particularly useful given the complexity and high cost of estimating demand patterns.
- To train our DRL agent, we develop an advanced ABM of large-scale EVCHs, creating a high-fidelity simulation that captures the interactions between EVCH components, EV users, and the power grid. On the supply side, an operator agent manages pricing strategies, charging operations, and system constraints, including grid capacity, renewable energy production, and station availability. On the demand side, each EV is modeled as an individual agent with unique preferences. To provide actionable managerial insights, we calibrate our ABM with real-world data, accurately representing user arrival/departure patterns and energy demand. A realistic training environment is also essential for efficient agent deployment in real-world scenarios (Zhao et al. 2020).
- We assess our DSS by benchmarking it against an optimal pricing policy under perfect information (serving as an upper-bound), traditional dynamic pricing models, and a ToU pricing model. Our results indicate that despite incorporating non-optimal charging management—reflecting real-world charging behaviors—our learned policies achieve near-optimal performance and significantly outperform traditional dynamic and ToU pricing strategies. These findings demonstrate the effectiveness of our DRL-enhanced DSS for demand shaping and revenue management, particularly for EVCHs operating under peak electricity penalties or high renewable energy integration.
- In practice, our proposed DSS enables EVCH operators to maximize profitability, which is critical to overcoming financial barriers to charging station development and ensuring widespread EV adoption. In addition, because our DSS is highly effective at peak shaving, it benefits utilities by significantly reducing the peak consumption of EV demand, thereby helping to avoid the extremely high investment costs associated with grid expansion.

⁵EVCHs are often integrated with workplaces or commercial centers with variable energy usage.

In the forthcoming sections, we contextualize our paper within the existing literature and conduct a review of related studies. We then provide a concise description of the problem and our proposed model. Finally, we present numerical examples and thoroughly discuss our results.

5.2 Literature Review

This study contributes to two streams of literature: (i) Information Systems (IS) and sustainable electric mobility systems, and (ii) electric vehicle charging facility operations management. We demonstrate how our proposed DSS contributes to sustainability and solving stochastic decision-making problems by optimally using available information and learning hidden information of its environment. We also examine the management of operations at large-scale charging stations with an emphasis on service pricing.

5.2.1 Information Systems and Sustainable Electric Mobility Systems

Companies significantly contribute to climate change and must balance sustainability with profitability to minimize their environmental impact while maintaining economic viability (Böttcher et al. 2024). Emerging startups demonstrate that environmental responsibility and financial success can coexist. For example, Octopus Energy⁶ utilizes a digital platform to make sustainable energy more affordable. Our proposed IS-enabled DSS, integrated with digital technologies, assists large-scale charging platforms in mitigating unsustainable outcomes within their business models. Traditionally, IS research has examined economic and ecological transformations separately, and in this paper we aim to couple them together.

From an economic standpoint, several studies in the IS domain have investigated the relationship between information systems and profitability. Mithas et al. (2016) emphasize the role of user satisfaction in optimizing IT investments, demonstrating their contribution to increased profitability. Haki et al. (2024) explore how the complexity of market opportunities and the design of platform mechanisms influence partners' engagement with B2B innovation platforms, ultimately affecting the profitability of innovations for incumbents and their partners. Regarding pricing strategies, Aloysius et al. (2013) analyze how technological advancements enable sellers to implement sequential pricing models, facilitating price discrimination based on customers' revealed purchasing preferences.

Research on ecological sustainability within the IS domain, often referred to as Green IS, is relatively recent (Watson et al. 2010, Malhotra et al. 2013, Melville 2010b). Green IS aims to harness the potential of IS to promote an environmentally sustainable society (vom Brocke et al. 2013). There are two main orientations of Green IS research: one follows a solution-oriented or design science framework, while the other adopts a behavioral science perspective. Our work falls into the first category, focusing on the design of an IS artifact to help EVCHs improve

⁶<https://octopusenergy.com>

their efficiency by identifying and implementing improved decision strategies. In this context, Seidel et al. (2013a) outlines key design principles for sustainability transformations through a sensemaking support system, while Hilpert et al. (2013) develops a system for promoting sustainable logistics practices by tracking gas emissions.

More specifically, our research aligns with the field of energy informatics (Watson et al. 2010), which leverages IS expertise to improve energy efficiency. Energy informatics focuses on optimizing energy supply and demand by analyzing and designing energy systems, collecting and processing energy data to enhance energy distribution networks. For instance, Brandt et al. (2018) present an IS framework that identifies synergies between EVs and renewable energy sources, improving grid stability and sustainability. Similarly, Ketter et al. (2016) introduce a simulation platform that addresses societal challenges, including environmental sustainability and smart grid stabilization, through competitive benchmarking. Beyond the energy sector, transportation systems also require significant improvements to ensure more sustainable urban mobility. The call to action by Ketter et al. (2023) underscores the necessity for greater IS community involvement in the development of intelligent, sustainable transportation systems. Our research examines a socio-technical challenge at the intersection of sustainable energy and smart mobility, requiring both technical algorithmic methodologies and large-scale behavioral insights provided by IS approaches (Sarker et al. 2019b). Relevant studies include decision support systems for demand response activation and load shifting in electricity markets (Fridgen et al. 2016), data-driven microgrid operations (Gust et al. 2021), and decentralized management of EV charging processes (Valogianni et al. 2020a).

Our work contributes to the Green IS field by focusing on EV integration and adoption, a critical driver of sustainability in the transportation sector. At the same time, we address pressing challenges such as the extreme peak load increase in power grids caused by the uncontrolled charging demand of a growing EV fleet. Specifically, we propose a DSS for optimal pricing of capacity-based charging services at EVCHs (e.g., workplaces, shopping malls, depots), taking into account heterogeneous user behavior. Our research falls within the computational and optimization domain of IS, as outlined by Rai (2017). The integration of data analytics and artificial intelligence techniques also links our study to the broader field of Information Technology (IT) engineering (Jenkin et al. 2011, Valogianni et al. 2024). Indeed, we position our research at the intersection of Green IS and IT engineering, developing a DSS that leverages IT solutions to promote sustainability by facilitating the widespread adoption of EVs. Our approach employs machine learning techniques and ABM, both of which have recently been applied in IS to address complex socio-technical problems (Haki et al. 2020). We provide a detailed agent-interactive simulation of EVCHs, calibrate our ABM using empirical data analysis, and model service pricing decisions as a dynamic learning problem. Our study bridges Green IS and Economic IS by optimizing both economic and environmental objectives, and introduces a novel pricing DSS for EVCHs.

5.2.2 Management of Electric Vehicle Charging Hubs

EVCHs exhibit several unique features that distinguish them from other charging use cases. First, EVCHs typically represent large, locally concentrated loads that may require significant local electricity grid extensions and load management (Lee et al. 2019). Second, integration with building loads and generation units (renewable energy production, storage) may be possible/desirable (Nunes et al. 2016) to reduce induced peak loads. Third, they experience different user (i.e., charging) behavior compared to other charging use cases. User behaviors can vary substantially depending on the use case of the attached facility (workplace, mall, etc.).

We review state-of-the-art operations management approaches in the realms of (1) smart charging and (2) pricing EVCHs. In terms of EVCH operations, we acknowledge the extensive work on EV smart charging (see e.g., Mukherjee and Gupta (2015) for a review) that most EVCH operations-focused research is based on. A key differentiator from traditional smart charging literature is the inclusion of building/cluster-level constraints and optimization opportunities. Early examples include the development of a mixed-integer optimization framework for workplace charging strategies that takes into account different eligibility levels (Huang and Zhou 2015), and coordinated charging management models with solar energy production (Lee et al. 2019).

EV charging pricing is being studied in depth to integrate EV loads into the grid. Auction mechanisms have been used to coordinate EV charging (Hou et al. 2019). Pricing mechanisms are attracting more attention because they are easier to implement and are preferred by customers (Valogianni et al. 2020a). (Cui et al. 2021) propose a charging price optimization to coordinate the demand between multiple EVCHs. To integrate the uncertainty of the charging demand, Luo et al. (2017) propose a stochastic dynamic pricing that also deals with the volatility of the renewable energy generation. Mao et al. (2017) propose a vehicle-to-grid pricing regulation to use the EV batteries as storage systems in EVCHs. Lu et al. (2018) consider the competition among multiple EVCHs while designing a pricing scheme. Also, due to high uncertainty and computational complexity, many dynamic pricing related works employ DRL algorithms to enable the implementation of their works in large-scale problems (Lee and Choi 2021).

In more similar works, researchers design more advanced pricing models as a function of service quality. For example, Lin et al. (2023) include the waiting time to dynamically determine the price for fast charging services in public charging stations, optimizing the queue for limited resources. Valogianni et al. (2024) develop a capacity-based pricing model from the grid operator's perspective with the goal of reducing the additional peak from EV loads. Some closer related works study the menu-based EV charging services to utilize the flexibility of EV charging demands. For example, Moradipari and Alizadeh (2019) investigate optimal pricing mechanism to assign users with higher priority to charging stations with lower waiting time in order to maximize social welfare. Lu et al. (2023) design a deadline differentiated (i.e., users get a discount if they park longer) dynamic pricing model to reveal the real departure time of

EV users. To extend on the existing literature, we have developed interrelated dynamic pricing models for capacity-based charging services. In the case of EVCHs, where users are engaged in other activities (e.g., work and shopping), it is unlikely that they will wait for charging services or alter their plans (departure time) based on marginally varying prices Daina et al. (2017), Lee et al. (2019). We assume that EV users would rather adjust their energy demand according to the price signals based on their preferences. Our work is one of the first studies to design dynamic capacity-based pricing for charging services when the price is a function of energy request and dwell time. We use DRL algorithms that do not require prior information about EV users' utility functions (price responsiveness) while providing scalability and computational advantages for large-scale problems. This allows us to analyze heterogeneous utility functions for EV users as the model learns the optimal pricing policies by interacting with the EVCH environment.

5.3 Model

We develop a model to optimize EVCH performance through dynamic pricing management, focusing on profit maximization from the operator's perspective—service revenue minus operating costs. Our model contains two main decision problems: first, the EV user's single-period energy demand adjustment, and second, the EVCH operator's multi-period price management. In general, the overall problem can be framed as a single-leader, multi-follower optimization problem. At a higher level, the EVCH operator sets the pricing parameters at each time step, while at a lower level, EV users arriving at the charging station within the current time window determine how much to charge based on the service price and their actual demand. In the next section, we show that since the EV users' decision problem can be solved analytically, the closed-form solution for the follower level can be applied. This allows us to transform the bilevel optimization model into a single-level problem, making it easier to solve.

Like other EV charging providers, EVCHs purchase electricity from the distribution grid and sell charging services to EV users. Consistent with the literature, we model EV users as price-sensitive, adjusting their energy requests based on the service price and user characteristics (Lin et al. 2023, Lehmann et al. 2022). On the supply side, we consider real-world scenarios (e.g., current California rates⁷) where electricity tariffs for EVCHs include per-unit energy costs and additional sustainability fees (e.g., costs for high peak consumption). These fees aim to discourage undesirable consumption behaviors, like high peak demand. Alignment with preferred consumption patterns is expected to grow more critical as EVs impose additional grid loads and time-dependent energy sources (e.g., renewable energy sources) increase in use (Gilleran et al. 2021, Alam et al. 2020). This incentivize grid operators to design tariffs for large charging stations to discourage unwanted consumption, such as significant local peak loads (Ansarin et al.

⁷<https://www.cpuc.ca.gov/industries-and-topics/electrical-energy/infrastructure/transportation-electrification/electricity-rates-and-cost-of-fueling>

2022). As a result, the pricing model must increase revenues and avoid high energy consumption peaks to maximize profitability.

On the demand side, we assume that EV users are rational price takers who adjust their energy demand based on the published pricing parameters of the EVCH, as commonly assumed in the literature (e.g., Valogianni et al. 2020a, Huang and Kockelman 2020, Lu et al. 2023). Like many other studies (e.g., Limmer 2019, Abdalrahman and Zhuang 2020), we also assume that EV users who charge their vehicles at EVCHs, such as those located at workplaces or shopping malls, may adjust their charging demand based on the prices set by EVCHs, influenced by their preferences. All parameters and decision variables related to the EV user decision problem and EVCH pricing optimization are summarized in Appendix 5.8.

5.3.1 EV User Problem

At the arrival time, users receive the price signals and adjust their energy request accordingly. Each user has an objective function and aims to minimize all costs of charging and inconvenience (Eq. (5.1)). This is a single-period decision problem where users make their decisions after arriving at the EVCH and receiving the price parameters. Once the charging price is set at the time of arrival (e.g., \$0.7/kWh), it remains fixed for that user, even if the price parameters change during their stay. We define the utility function for EV users, including both charging costs and dissatisfaction terms, following standard formulations in the academic literature on EV charging (Limmer 2019). Charging costs are determined by the unit price of energy $p(x, \delta)$, which depends on the energy requested x and the length of stay δ for the user. Building on the charging cost considerations discussed in Han et al. (2012) and Limmer (2019), we propose a quadratic model to represent user dissatisfaction. This model uses $(x - D)^2$ to quantify the squared difference between the requested energy x and the user's maximum energy demand D , ensuring a non-negative value⁸. This implies that any energy request below or above the maximum demand will result in dissatisfaction. In addition, the requested energy must be greater than or equal to zero and cannot be greater than the maximum energy demand (Constraint 5.2).

$$\min f(x) = p(x, \delta)x + \beta(x - D)^2 \quad (5.1)$$

$$0 \leq x \leq D \quad (5.2)$$

EV users exhibit different charging behaviors characterized by their duration of stay (δ), willingness to charge (β), and maximum energy demand required to fully charge their battery (D). Users are assumed to behave rationally, seeking to minimize their charging costs while considering a penalty for any unmet energy demand upon departure. This creates a trade-off between cost optimization and demand satisfaction. Users with higher charging willingness (β)

⁸While this utility function is commonly used in the literature, in practice users may respond differently. Our proposed model-free approach is able to capture these different behaviors.

are willing to pay more to satisfy their full energy demand, while users with lower charging willingness adjust their energy demand to minimize costs.

5.3.2 Pricing Problem with Perfect Information

The primary goal of EVCH price management is to maximize profits, a complex task due to the interactions between multiple decision-makers and uncertainties on both the supply and demand sides. To formulate a mathematical model, we initially make the following assumptions: (a) deterministic parameters—where user information (arrival time, willingness to pay, and dwell time), renewable generation patterns, and building loads are known, (b) homogeneous chargers with optimal load allocation by the operator, and (c) no grid capacity constraints. These assumptions provide a foundational framework for model development. However, in our proposed machine learning-based DSS, we relax these constraints to more accurately capture real-world dynamics.

Employing Capacity-Based Pricing

Regarding the energy price function, we adopt a capacity-based model (Valogianni et al. 2024), defined as $Price = p^0 + \alpha \frac{x}{\delta}$, with two key parameters: (1) a fixed price per energy unit (p^0) and (2) a positive coefficient (α) that adjusts the price based on the average power of the charging request. Valogianni et al. (2024) demonstrate the effectiveness of this model for reshaping private EV charging loads. We apply a modified version of this pricing model to EVCHs to optimize profits while mitigating extreme consumption peaks.

Substituting the capacity-based pricing model into the objective function of the EV user's problem (Eq. (5.1)), we obtain the following:

$$\min f(x) = (p^0 + \alpha \frac{x}{\delta})x + \beta(x - D)^2 \quad (5.3)$$

$$0 \leq x \leq D \quad (5.4)$$

To find the optimal solution, it is not necessary to constrain the requested energy to less than the maximum demand (D), since such a constraint would not yield an optimal result (see Appendix 5.7.2). We analytically find the closed-form optimal solution for the requested energy, which is $x^* = \max(\frac{(2\beta D - p^0)\delta}{2(\beta\delta + \alpha)}, 0)$ (see Appendix 5.7.1). Finding the closed-form solution to the EV user's one-shot energy adjustment decision problem allows us to separate it from the pricing problem.

Mathematical Model for Perfect Information Model

Using the closed-form solution for the EV user's decision-making problem, we can express the mathematical formulation of the pricing problem as follows.

$$\text{Max } g(p_t^0, \alpha_t, z_i, y_{i,t}) = \sum_{i,t} (p_t^0 + \alpha_t \frac{x_{i,t}}{\delta_i}) x_{i,t} z_i - C^\Omega \quad (5.5)$$

Revenues include the summation of fulfilled requested energy ($x_{i,t}$) of all vehicles ($i \in \mathcal{I}$) at all time steps ($t \in \mathcal{T}$) multiplied by the charging price corresponded to each individual vehicle. Note that the charging price for each vehicle is based on its requested power ($\frac{x_i}{\delta_i}$) and, once determined, will remain constant for the duration of the user's stay. The costs (C^Ω) consist of electricity purchasing costs from the power grid and peak consumption penalties further defined in the below constraints.

Revenues include the sum of the fulfilled requested energy ($x_{i,t}$) of all vehicles ($i \in \mathcal{I}$) at all time steps ($t \in \mathcal{T}$) multiplied by the charging price corresponding to each individual vehicle. Note that the charge price for each vehicle is based on its requested power ($\frac{x_i}{\delta_i}$) and, once determined, remains constant for the duration of the user's stay. The cost (C^Ω) consists of the cost of buying electricity from the grid and the penalty for peak consumption, as defined in the constraints below. Without loss of generalizability, we use a realistic electricity tariff for large EV charging stations from California. See Appendix 5.9.3 for details.

$$x_{i,t} = \max\left(\frac{(2\beta_i D_i - p_t^0)\delta_i}{2(\beta_i \delta_i + \alpha_t)} A_{i,t}, 0\right) \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \quad (5.6)$$

$$C^\Omega = \sum_{i,t} e_t^{\text{Grid}} C_t^e + C^p p^* \quad (5.7)$$

$$e_t^{\text{Grid}} = \sum_i \Delta_t (y_{i,t} - f_t^{\text{PV}} \kappa^{\text{PV}}) \quad \forall t \in \mathcal{T} \quad (5.8)$$

$$y_{i,t} \leq z_i U_{i,t} R \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \quad (5.9)$$

$$\sum_t y_{i,t} \Delta_t \geq \sum_t x_{i,t} \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \quad (5.10)$$

$$\sum_i z_i U_{i,t} \leq \kappa \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \quad (5.11)$$

$$p^* \geq \frac{e_t^{\text{Grid}}}{\Delta_t} - l^* \quad \forall t \in \mathcal{T} \quad (5.12)$$

Constraint (6) guarantees that the requested energy of vehicle i at time t can only be greater than zero if the vehicle arrives at the charging station at this time slot ($A_{i,t} = 1$). The parameter $A_{i,t}$ is a boolean variable indicating whether the vehicle i arrives at time t . The requested energy ($x_{i,t}$) at the arrival time of vehicle i corresponds to the closed-form solution of the EV user's decision problem derived in the previous section⁹. The cost consists of the energy consumption

⁹The maximum function can be linearized using the Big-M trick.

from the grid multiplied by the temporal electricity prices (C_t^e), along with the penalties for exceeding the peak threshold (Constraint (7)). In Constraint (8), the grid consumption at time t (e_t^{Grid}) is defined as the total power consumed for all charging processes ($y_{i,t}$) minus the on-site generation power at each time step ($f_t^{PV} \kappa^{PV}$, where κ^{PV} is the maximum PV capacity and f_t^{PV} is the generation efficiency at time t). Constraint (9) allows charging only if the vehicle is assigned to a charger and is present at the EVCH. z_i is a binary variable that takes the value 1 if vehicle i is assigned to a charger, and 0 otherwise. $U_{i,t}$ is a boolean parameter indicating the presence of vehicle i at time t in the EVCH. It also ensures that the charging power does not exceed the maximum charging rate of the chargers (R)¹⁰. Constraint (10) ensures that each vehicle receives at least as much energy as it requested, and to avoid infeasible solutions, we assume that there is no grid power constraint¹¹. The maximum number of vehicles connected to chargers is set to the capacity of the EVCH (κ), which is considered in constraint (11). We do not allow vehicles to be moved after they arrive at the station—if there is no charger available when vehicle i arrives, the request will be missed. Finally, Constraint (12) calculates the exceeded peak (p^*) from the expected peak threshold (l^*).

5.3.3 Proposed Method: Learning Near-Optimal Dynamic Capacity-Based Pricing Policies

The actual dynamic pricing management of EVCHs is very complex with multiple stochastic parameters, such as EV user preferences and renewable energy production patterns. Therefore, the previous mathematical model for the perfect scenario is only built to jointly formulate the pricing and user decision problems, and later used as an upper-bound to validate our proposed solutions. Here, we show how to solve the real stochastic problem for large-scale EVCHs without prior knowledge of the dynamics of the environment. To achieve this, we design a DSS that uses machine learning to solve the pricing problem for EVCHs. These models account for expected future demand and other uncertainties, and are scalable to large-scale EVCHs.

Our proposed approach is illustrated in Figure 1. The goal is to develop a DSS capable of determining near-optimal dynamic capacity-based pricing strategies, where prices vary based on charging power, for EVCHs under real-world stochastic conditions. At the core of our model is a DRL agent that learns effective pricing policies by issuing capacity-based price signals and observing feedback (profits) from the environment. This feedback depends on EV users' responses to pricing, as well as dynamic factors such as grid costs and physical constraints.

During the training phase, before real-world deployment, the DRL agent learns by interacting with a simulation model of the EVCH, receiving feedback on each action. The DSS is initialized with key inputs, including the objective function, physical constraints, demand patterns, and facility configurations, derived from observational data and operator expertise. This setup allows

¹⁰We consider EV chargers to be homogeneous with the same maximum charging power for all

¹¹this assumption will be relaxed later in our proposed model




Figure 5.1: A Decision Framework: Learning Capacity-Based Pricing Policies for EVCHs

EVCH operators to interact with the DSS, verifying how pricing models influence profitability, demand, and facility configurations (e.g., on-site electricity generation capacity).

Once trained, the pricing agent can operate as a fully automated system, making real-time dynamic pricing decisions in simulations and real-world scenarios while continuously adapting through environmental interactions. Additionally, our DSS offers offline pricing policies that function independently of the agent, providing high-level managerial insights. For instance, EVCH operators can evaluate the impact of capacity-based pricing under different conditions. Finally, we analyze the learned pricing policies, comparing them to existing pricing models and conducting sensitivity analyses to assess the effects of user price elasticity and EVCH configurations.

The dynamic pricing decision for charging services is a sequential stochastic decision problem. Each decision has an impact on the next system state and depends on previous actions (pricing parameters) and expected load demands, which are affected by the service price. In addition to stochastic EV user arrival times and demand preferences, there are other sources of uncertainty such as renewable energy production patterns. The outcome of pricing decisions also depends on other operational strategies, including load management of charging vehicles. These complexities make it too difficult to track the dynamics of the environment and calculate the expected outcome. Therefore, we formulate the problem as a Markov Decision Process (MDP) and use model-free algorithms to determine near-optimal pricing policies.

Markov Decision Process

As the pricing decisions are made centrally by the EVCH operator, the entire environment is visible to the learning agent, which guarantees a fully observable MDP definition. The pricing problem is also an episodic problem where the agent makes pricing decisions at each decision time step until the end of the operational horizon (T). The problem is symmetric with equal time windows (e.g., 15 minutes) and the agent receives immediate feedback before taking the next action. The components of the MDP are as follows.

State: To represent the state of the EVCH comprehensively, we define the state at time t as $s_t = \{t, C_t^e, p_t^*, f_t^{PV} \kappa^{PV}, \kappa_t^{Grid}, \bar{E}_t, \bar{P}_t\}$. This includes the time step t , the purchase price of electricity from the grid (C_t^e), the current peak demand threshold (p_t^*), the current renewable electricity production ($f_t^{PV} \kappa^{PV}$), the current grid capacity (κ_t^{Grid}), the current average energy (\bar{E}_t) and power demand (\bar{P}_t)¹².

Time is one of the most critical state components because it helps the DRL agent anticipate demand and supply values, which significantly affects the payoff of the policy. To interact with the grid, the agent needs to track the current price of electricity, since the EVCH purchases electricity under a ToU contract. Another key factor affecting grid costs is the peak threshold: when electricity consumption exceeds this threshold, a peak charge is applied. On the supply side, a dynamic component is current on-site generation, which varies over time and affects grid costs and capacity constraints. The facility is also constrained by a given grid capacity, meaning that total energy consumption, including EV charging and building load, must not exceed this capacity. The agent should take this into account to effectively optimize price signals. Similarly, current demand (both energy and power) affects the availability of supply resources and influences grid costs, which are also included in the state vector¹³.

Action: The action is a vector of the capacity-based term and the power-free term of the pricing function $\mathbf{a} = (p^0, \alpha)$. The action space is continuous and bounded between zero and the cap for each component. Multidimensional continuous action space MDPs present significant challenges due to the infinite combination of possible actions (Manna et al. 2022, Khetarpal et al. 2022). Exploring this vast space and achieving stable policies requires advanced DRL algorithms and careful hyperparameter tuning, making these problems particularly complex.

Reward: We consider the same objective function as defined in the previous section (Eq. (5.5)), which is the maximization of EVCH profits. Therefore, for consistency, the reward function is formulated as Eq. (5.13), which represents the profit earned between two successive decision time steps. The profit of choosing action \mathbf{a} in state s is the revenue from selling

¹²The operator has access to all connected EV information, including their unserved energy demand and departure time. However, to avoid overfitting problems, we use aggregated information from charging vehicles to define the state of our pricing agent.

¹³we normalize the state space to prevent any single state component from exerting a biased influence and to account for the periodic nature of the time component.

electricity to EVs, which depends on the published price, minus grid costs, including the cost of buying electricity from the grid and additional peak charges (if applicable).

$$r_t(a, s) = \sum_i (p_t^0 + \alpha_t \frac{x_{i,t}}{\delta_i}) x_{i,t} z_i - C_t^\Omega - C_t^M \quad (5.13)$$

As discussed, we design the reward function to minimize delayed feedback and thereby improve the learning process. Consequently, the reward for the time window t includes the revenue from EVs that arrive during this period and connect to a charger (i.e., users pay upon arrival). Note that only vehicles that request charging and find an available charger are accommodated, otherwise their entire charging demand is not satisfied ($z_i = 0$). A shortage of chargers is not the only reason for unmet charging demand; other constraints, such as grid capacity, also play a role. Therefore, some charging demand may remain unmet when vehicles depart. To account for unmet demand, we impose a penalty cost, $C_t^M = \zeta_i \sum_i \sum_{\tau=0}^t (x_i - y_{i,t}) De_{i,t}$. Thus, in the time window t , we penalize the missed energy requests for vehicles departing during this period ($De_{i,t} = 1$) and ensure that the penalty rate exceeds the service price for vehicle i ($\zeta_i > p_t^0 + \alpha_t \frac{x_{i,t}}{\delta_i}$). Finally, the charging cost C_t^Ω is simply the cost of electricity to meet the charging demand from the grid and any peak load costs.

Transition Function: A transition from one state to another, along with the receipt of a given reward, is inherently a stochastic function in our dynamic pricing problem. In other words, when the pricing agent takes action a in state s , the reward and the subsequent state are highly dependent on the unknown arrival of EVs and their responses to published price parameters and other operational decisions, such as EV charging management. Estimating these transition probabilities is very complex, so we use a simulation to monitor the state of the system. This simulation approach allows us to model the randomness and uncertainty associated with EV behavior and system dynamics, providing a practical means to analyze and optimize the pricing strategy.

Solution: Deep Reinforcement Learning

Optimizing a pricing policy within an MDP framework is feasible using policy or value iteration when the dynamics of the environment are known and can be mathematically modeled. However, in dynamic capacity-based pricing for EVCHs, estimating user behavior distributions and transition probabilities is both costly and highly dynamic, making it impractical to rely on predefined models. To overcome this challenge, we adopt a model-free DRL approach that allows the agent to learn optimal policies through direct interaction with the environment without requiring prior knowledge of stochastic elements.

Given the large state space of our problem, traditional tabular RL methods are inadequate. Instead, we apply function approximation techniques to generalize the state space, enabling efficient learning in large-scale settings. In RL, function approximation uses parameterized

models, such as neural networks, to estimate value functions or policies in high-dimensional or continuous action spaces where tabular methods fail. Neural networks help approximate the action-value function for complex decision problems.

A widely used value-based RL algorithm is the deep Q-network (DQN), which learns state-action values while exploring the environment. However, DQN is limited to discrete action spaces. In continuous action spaces, where actions must be selected from an infinite range, more sophisticated methods are required. These methods must approximate the best action for any given state, which increases computational complexity. In such cases, the agent must represent policies over a continuous domain, typically using Gaussian distributions, which require optimization during training –introducing additional complexity and potential instability.

We use actor-critic models, which are well suited for multidimensional continuous action space problems (Grondman et al. 2012). Unlike actor-only (policy-based) methods, which directly optimize a parameterized policy but suffer from high variance in gradient estimates, actor-critic methods combine the advantages of both policy-based and value-based approaches. The actor efficiently computes continuous actions without requiring direct value function optimization, while the critic provides low-variance feedback that stabilizes and accelerates learning. This approach increases training efficiency and ensures robust pricing strategies for EVCHs.

Traditional actor-critic models may struggle with multidimensional problems like dynamic pricing, often leading to inadequate exploration of the state-action space. This can result in limited system information discovery, slowing down learning and reducing the effectiveness of feedback for complex decision-making tasks. Therefore, we adopt the Soft Actor-Critic (SAC) algorithm that enhances exploration-exploitation balance through entropy regularization.

The SAC agent is responsible for determining pricing decisions based on the current state of the environment. It consists of two main components: the actor (policy) network and the critic (value) networks. The actor network, a neural network, maps states to actions and outputs a stochastic policy $\pi(a|s)$, representing a probability distribution over possible actions for a given state. A key feature of SAC is its entropy-based policy, which promotes exploration by encouraging diverse action selection. The policy is parameterized by θ . To evaluate action values, SAC employs two critic networks that estimate the expected return of taking an action in a given state, represented as $Q(s, a)$. These networks are parameterized by ϕ_1 and ϕ_2 . Using two critics helps reduce overestimation bias by selecting the minimum value predicted, leading to more stable learning.

Entropy regularization is a crucial part of SAC, encouraging the policy to maintain a certain level of randomness. This is controlled by an entropy term α^{RL} in the objective function, which balances exploration (high entropy) and exploitation (low entropy). The entropy term is included in the loss functions of both the actor and the critic networks.

The main objective function is to maximize the expected reward while also maximizing entropy. This can be expressed as:

$$J(\pi) = \sum_{t=0}^T \mathbb{E}_{(s_t, a_t) \sim \rho_\pi} [r(s_t, a_t) + \alpha^{RL} \mathcal{H}(\pi(\cdot|s_t))] \quad (5.14)$$

where ρ_π represents the state-action distribution under the policy π , $r(s_t, a_t)$ is the reward function, $\mathcal{H}(\pi(\cdot|s_t)) = -\mathbb{E}_{a \sim \pi}[\log \pi(a|s_t)]$ is the entropy of the policy at state s_t , and α^{RL} is the temperature parameter that determines the relative importance of the entropy term.

Actor Loss The actor network is updated to maximize both the expected return and the entropy of the policy. The policy objective can be rewritten using the soft Q-function, leading to the following loss function for the actor:

$$J_\pi(\theta) = \mathbb{E}_{s_t \sim D^{RL}} [\alpha^{RL} \log \pi_\theta(a_t|s_t) - Q_\phi(s_t, a_t)] \quad (5.15)$$

where D^{RL} is the replay buffer (a container to save the experiences), α^{RL} is the entropy coefficient, and Q_ϕ is the minimum value from the two critic networks. This loss encourages the policy to maximize the expected return while maintaining high entropy.

Critic Loss The critic networks are updated to minimize the Temporal Difference (TD) error. The soft Q-value can be defined as:

$$Q^\pi(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \rho_\pi} [V^\pi(s_{t+1})] \quad (5.16)$$

where $V^\pi(s) = \mathbb{E}_{a \sim \pi} [Q^\pi(s, a) - \alpha^{RL} \log \pi(a|s)]$ is the soft state value function. The loss function for each critic is:

$$J_Q(\phi_i) = \mathbb{E}_{(s_t, a_t, r_t, s_{t+1}) \sim D} [(Q_{\phi_i}(s_t, a_t) - (r_t + \gamma \mathbb{E}_{s_{t+1} \sim \rho_\pi} [V_{\bar{\phi}}(s_{t+1})]))^2] \quad (5.17)$$

where γ is the discount factor, ρ_π is the state distribution, and $V_{\bar{\phi}}$ is the target value function. The discount factor links the state value at the next time step to the current time and is always a value between zero and one, indicating that the immediate reward is relatively more important than the value of the next state. Critic target networks are used to stabilize training. They are periodically updated using a moving average of the weights of the critic networks. Finally, the replay buffer stores transitions (s_t, a_t, r_t, s_{t+1}) that the agent has experienced. This allows the agent to learn from past experience and break the correlation between successive samples.

By leveraging the SAC algorithm, we can develop robust, adaptive policies for dynamic pricing in EVCHs, ensuring effective real-time decision-making in complex and variable environments.

5.3.4 Simulation Environment

As previously mentioned, we need to build a digital simulation of EVCHs to train the DRL pricing agent and evaluate the performance of the proposed DSS. Below, we will first outline the input data and information required to simulate realistic scenarios, followed by the development of our ABM.

5.3.5 Input and Data Preparation

The three main data inputs are EV user information (arrival time, duration of stay, and demand preferences), building energy consumption patterns, and on-site solar photovoltaic (PV) production.

Parking and Charging Behavior

To realistically generate demand characteristics in our scenarios, we need to model parking and charging behavior of users in EVCH based on historical data. User preferences (of an individual i) in an EVCH context are described by the four-dimensional vector $v_i = (A_i, \delta_i, D_i, \beta_i)$. The four individual components are: (1) time of arrival (A_i), (2) duration of stay δ_i , (3) maximum energy demand (D_i), and (4) willingness to charge (β_i).

To model the demand characteristics of our simulation, we take advantage of a unique observational parking data set provided by a major European real estate investor, which includes transactions from seven large parking garages (capacities ranging from 275 to 2200 spaces). A mix of workplace, downtown, and shopping center facilities are available. Each row in this dataset represents a single parking event (user i) with corresponding arrival and departure preference information. For privacy reasons, individual users cannot be identified. We use a full year of data to capture daily, weekly, and annual seasonality. The year 2019 is chosen to filter out the effects of the pandemic. In total, our data includes 3.84M parking events. For details, see the appendix 5.9.1.

In Figure 5.2, we present key population characteristics for a representative facility over the course of a week. These plots generally support our hypothesis that dynamic pricing could be beneficial for EVCHs. The distribution of vehicle arrival times indicates significant variability in arrival frequencies throughout the day. Arrival rates peak at the beginning of the workday, gradually decline until early evening, and significantly drop after late evening. The other two graphs show that users arriving at different times of day may have different demand preferences. For example, users arriving early in the morning or late in the evening tend to have higher

energy demands but also longer dwell times. These patterns create a highly uneven demand distribution over time, highlighting that users arriving at different times may respond differently to pricing signals.




Figure 5.2: Parking and Charging Characteristics of Electric Vehicle Users

Base-Load Energy Consumption Patterns

Since EVCHs are often integrated with buildings like workplaces or shopping malls, considering the building’s base load is essential. Without a separate grid connection for charging stations, the building’s energy consumption directly impacts real-time grid capacity and peak demand. We use real-world building consumption data to model the base load of EVCHs. Unlike EV loads, we treat building load as exogenous, assuming it cannot be actively managed or curtailed, a reasonable assumption given the lack of smart energy management in most existing buildings.

Electricity (PV) Production Patterns

Given the synergy between EV energy storage and renewable energy sources, EVCHs are likely to incorporate on-site renewable generation. Solar PV is particularly well-suited for urban areas, though its output varies with weather conditions. To model PV generation, we use real-world PV load factors from facility locations. These load factors represent the actual power output of PV panels as a ratio of installed capacity, influenced by local solar irradiation. We calculate them by dividing real PV infeed by the total installed PV capacity at each site. This data, provided by local transmission system operators.

5.3.6 Agent-based Modeling of EVCHs

We present the ABM of EVCHs’ digital representation, incorporating supply and demand specifics.

Supply Side

We describe the EVCH equipment and the operational processes including charging dock allocation and power management. An EVCH is a parking lot, depot, or garage equipped with EV charging stations, typically sharing a grid connection with an adjacent building, which may also have on-site PV generation and battery storage. The EVCH can feature different types of charging docks (22kW or 50kW) with single or multi-connector setups, allowing simultaneous charging for greater flexibility and utilization. On-site storage using Li-Ion batteries can further optimize energy use, with PV generation constrained by local space and weather-dependent availability. For details see Appendix 5.9.2.

Charging Hub Operational Decisions

We implement operational decision policies that model real-world processes in our EVCH simulation environment. There are three different operational decision modes: a) pricing, b) charger allocation, and c) load control. Service pricing is a critical component of EVCH operational management, as price signals can indirectly adjust charging demand and significantly impact profits. Our system supports multi-connector charging docks with simultaneous charging capability, making the initial assignment of vehicles to charging stations an important factor. This allocation affects not only the available charging capacity for the EV, but also for current and future arrivals at the same dock. Since EVs cannot be moved while connected to a charging dock, well-informed initial allocation decisions are essential. Managing the charging of connected vehicles impacts demand satisfaction and grid consumption patterns, which directly impacts costs. Typically, routing and charging decisions are made separately for simplicity, as noted by Ferguson et al. (2018), and we follow this approach as routing is not the focus of our paper. Finally, when a storage system is present, charging and discharging decisions for purposes such as peak shaving and electricity arbitrage play an important role in the revenue management of EVCHs. In the following sections, we describe algorithms for each of these operational decisions.

Pricing Scheme The EVCH operator needs to communicate the service price to EV users, which can be done in several ways. The simplest strategy is a fixed price regardless of time and service rate, or a ToU tariff where the charge price is published in advance for a period of time (e.g. daily). A more complicated scheme is dynamic pricing, where the service price is published at each updating time step and users do not know the price for the next time steps. To account for the service rate (i.e., the charging power), there are several options: a) offering discrete charging powers with different prices, known in the literature as menu-based pricing, where e.g. standard and fast charging have different prices, and b) a capacity-based pricing function, where the price is a continuous function of the charging power (see section 5.3.2 for details). We combine both dynamic and capacity-based strategies, adjusting the power term of the pricing function at each decision time step. By implementing dynamic, capacity-based pricing, the operator can achieve various objectives, such as maximizing profits and reducing peak usage. We assume that the

operator wants to maximize profits. We must consider that there is a trade-off between prices and demand, where higher prices reduce demand. In addition, it is worth noting that EVCH has a ToU electricity contract, which includes peak charges. Therefore, determining optimal prices for different charging levels is a complex problem due to stochastic charging demand patterns, including arrival times, length of stay and price sensitivity, as well as time-based electricity costs.

Charging Station Assignment Algorithms Vehicles are assigned to a connector when they enter the EVCH. We evaluated several heuristic routing algorithms with different levels of sophistication, including fill-one-after-the-other and lowest-occupancy-first. However, the lowest-occupancy-to-highest-laxity matching algorithm provided the best service level. This strategy takes into account the condition of both the loading docks and the vehicle prior to allocation. It categorizes incoming vehicles as low, medium, or high laxity (using bins derived from historical data) and then matches low laxity vehicles to docks with high remaining capacity and vice versa. In this way, the strategy implicitly prepares for future arrivals that may require more or less charging capacity.

Charging Adjustment Algorithms The charging loads of vehicles being serviced at the EVCH need to be controlled, as there are power capacity constraints and economic incentives to adjust charging rates. In our simulation model, the charging operator updates the charging rate for all connected vehicles every certain time step (e.g., five minutes) as new information becomes available. There are a variety of sorting-based and optimization-based algorithms to periodically determine the charging rate for each vehicle. Due to the focus of our work, which is on pricing models, and for the sake of scalability and realistic modeling, we use sorting-based algorithms. We use two methods: a) equal-sharing: in this model, the charging power of each vehicle is the average requested power of that vehicle ($\frac{p_i}{T_i}$), and in cases where the cumulative consumption load is greater than the grid capacity, this model reduces the charging power of all vehicles equally, b) least-laxity-first (LLF): this model uses the least-laxity-first priority rule, which means that the least flexible vehicles are charged first if the charging station has a grid usage restriction. Thus, the algorithm explicitly considers the current state of a vehicle (remaining energy demand and departure time) in the charging decision.

Demand Side

In the simulation, we model EVs entering and exiting the EVCH according to a schedule driven by real-world sensor data. Upon arrival, EV users are presented with the pricing function (including parameters) and prompted to estimate their intended duration of stay and the desired amount of energy (if any) to charge during that period. Note that we assume that users do not change their departure time based on the charging prices and that their duration of stay depends only on their activity (e.g., working and charging), but they might adjust their requested energy based on their needs and the service price. These are common assumptions in the literature (Lee et al. 2019, 2018). In most studies, users are assumed to stay for the duration they indicated

at the beginning, although this assumption is not necessary for our model. Upon arrival, and only if they wish to charge, the operator connects the EVs to the designated receptacle. It is important to note that EVs cannot be moved or relocated during their stay, and will occupy both the parking space and the connector assigned to them for the entire period, even if the charging process has already been completed. Once EVs reach their scheduled stay, any charging in progress is terminated, the connector is released, and the parking space is vacated and becomes available for the next period. See Section 3.1 for details on EV user charging demand decisions.

5.4 Simulation Experiments

We perform benchmark and sensitivity analyses to illustrate how our DSS helps EV charging providers improve their profitability and system-level impacts (e.g., peak load reduction). In addition, we evaluate the benefits of integrating capacity-based and dynamic pricing strategies for EVCHs. To validate the effectiveness of our proposed model, we compare its performance to a globally optimal upper-bound. This upper-bound represents a mathematical programming model with perfect information, as derived in Subsection 5.3.2. The goal of this validation analysis is to ensure that our DRL approach yields near-optimal solutions. In further evaluation of robustness, we compare our proposed model to the upper-bound under various scenarios, considering different facility sizes, user price sensitivity levels, and electricity tariffs. Finally, we conduct a comparative benchmark analysis to demonstrate that our model outperforms existing pricing models in the literature. The benchmark pricing policies include traditional dynamic pricing (power-free), dynamic menu-based pricing, and ToU pricing models.

We select a commercial facility with attached parking, for which we have access to both building load data and transaction-level parking and estimated charging demand. We consider a facility with 200 parking spaces, a 65% EV adoption rate, and 100 single-plug DC fast chargers (50 kW). We use lowest-occupancy-highest-laxity equal sharing for load control of charging requests (defined in subsection 5.3.6). Energy costs are calculated using California electricity tariffs, the same region where the charging data was collected. Table 5.4 provides an overview of the tariff structure used in all experiments. This is a ToU tariff where the cost of electricity is higher during peak hours (\$0.23/kWh) and lower during off-peak hours (\$0.08/kWh). According to this tariff, the charging operator will be penalized if the peak consumption exceeds a certain threshold. We consider a peak consumption threshold of 500 kW, above which the EVCH operator must pay an additional charge for excess power consumption.

Regarding the user price sensitivity characteristics we use uniform random distributions to parameterize the utility function of EV users (see Eq. (5.1)). This is a common assumption in the literature as EV users have heterogeneous preferences for the prices of charging services with different charging speeds (Babic et al. 2022a). The willingness to charge to the maximum demand (β) is uniformly generated between 0.1 and 0.3.

5.4.1 Hyperparameters of SAC Model

For finding the closest-to-optimal dynamic pricing policies, we use a grid search to set the hyperparameters of our DRL algorithm. For the above-mentioned simulation the grid search leads to a learning rate of 10^{-3} using Adam optimizer, 512 batch size, and a fully-connected neural network with 2 hidden layers of (512, 256, 512) number of nodes for both actor and critic networks. We exclude the replay buffer size and soft update parameter from the grid search and set them to 10^5 , and 10^{-2} as suggested by the pre-trained agents.

5.4.2 Upper-Bound Benchmarking Experiments

DRL algorithms do not guarantee optimality. Therefore, to properly evaluate our proposed machine learning-based pricing model we compare the results to a theoretical upper-bound generated using the mathematical programming model from Section 5.3.2.

Upper-Bound: Perfect Information Model

To compute the upper-bound using a mathematical programming model, we make some assumptions. Since we seek an optimal solution, we assume that all parameters (e.g., demand information) are deterministic and known to the model (i.e., perfect information scenario). To achieve the computational intractability of the perfect information model, we limit the size of the problem to 200 parking spaces and assume that there are enough charging stations to serve all EV users. With this assumption, we can remove the assigning variable (z_i), since all vehicles receive a charger, and convert the model to simple nonlinear optimization. Finally, instead of optimal charging management, we use equal sharing load control in the mathematical model (this model is explained in Section 3.6.2). This leads to a better comparison since we use the same load management models for both machine learning-based and perfect-information models. The modified mathematical model is shown below:

$$\text{Max } g(p_t^0, \alpha_t) = \sum_{i,t} (p_t^0 + \alpha_t \frac{x_{i,t}}{\delta_i}) x_{i,t} - C^\Omega \quad (5.18)$$

$$x_{i,t} = \max(\frac{(2\beta_i D_i) \delta_i - p_t^0}{2(\beta_i \delta_i + \alpha_t)} A_{i,t}, 0) \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \quad (5.19)$$

$$y_{i,t} = \frac{x_{i,t}}{T_i} U_{i,t} \quad \forall i \in \mathcal{I}, t \in \mathcal{T} \quad (5.20)$$

$$C^\Omega = \sum_{i,t} e_t^{Grid} C_t^e + C^p p^* \quad (5.21)$$

$$e_t^{Grid} = \sum_i \Delta_t (y_{i,t} - f_t^{PV} \kappa^{PV}) \quad \forall t \in \mathcal{T} \quad (5.22)$$

$$p^* \geq \frac{e_t^{Grid}}{\Delta_t} - l^* \quad \forall t \in \mathcal{T} \quad (5.23)$$

Evaluation Results

To assess the optimality gap of our proposed decision model, we compare its performance against the upper-bound case with perfect information, as described earlier. To ensure the robustness of the optimality gap across various numerical settings (i.e., the stability of our model across different scenarios), we evaluate its performance under different parameter combinations, including price sensitivity, facility size, and peak penalty in the electricity tariff. We consider three levels of EV user price sensitivity, where the values for low, medium, and high sensitivity are randomly drawn from the intervals $[0.03, 0.1]$, $[0.05, 0.2]$, and $[0.1, 0.3]$, respectively. Additionally, we incorporate two facility sizes (100 and 200 parking spaces) and two peak penalty cost levels (\$15.89/kW and \$47.67/kW).

Table 5.1 shows the results for all possible combinations. As can be seen, the DRL performs very well, coming quite close to the optimal perfect information scenario with an average optimality gap of only 16%. Given the complexity of the problem and the uncertainties in supply and demand, this is a relatively small optimality gap compared to a perfect information model. A second observation from the results is that for higher peak costs, the gap becomes slightly larger, which could be caused by the high negative reward for exceeding the peak threshold, which creates extreme gradients and makes the training process more difficult for the agent. In addition, as the price sensitivity of users decreases, they tend to accept higher prices, which increases the profits for EVCHs and slightly reduces the optimal gap. Also, increasing the facility size does not significantly affect the gap, demonstrating that our proposed machine learning-based model is not only applicable to large problem sizes, but also ensures high performance for large charging hubs.

Profits (Optimally Gap)	Small Facility		Large Facility	
	Low Peak Cost	High Peak Cost	Low Peak Cost	High Peak Cost
High Price Sensitivity	\$1172 (17%)	\$1139 (19%)	\$2198 (19%)	\$2213 (19%)
Medium Price Sensitivity	\$1945 (16%)	\$1906 (19%)	\$3566 (17%)	\$3450 (19%)
Low Price Sensitivity	\$4470 (12%)	\$4482 (12%)	\$7612 (13%)	\$6950 (15%)

Table 5.1: Optimally gap between the perfect information and the proposed model

5.4.3 Benchmark Pricing Policies

We compare its performance to two benchmark policies: a traditional dynamic pricing policy and a menu-based dynamic pricing policy. To see the performance of these models under more realistic conditions, we include building load and PV generation. The building baseload, derived from historical data, is scaled based on the size of the parking facility and is limited to between 75 kW and 250 kW. We consider a maximum of 500 kW of installed PV as an onsite energy generation source that varies at different times. Finally, to make these models comparable, we

limit the maximum charging price to \$1.5. This also ensures that the model takes into account the preferences of EV users and does not significantly increase the cost of charging.

Traditional Dynamic Pricing

To demonstrate the impact of capacity-based pricing, we first use a traditional dynamic pricing model as a benchmark policy (e.g., Cui et al. 2023, Zhao and Lee 2021). By traditional dynamic pricing, we mean that the pricing function depends only on the energy unit parameter (p_0), which can be changed in time. For this, we need to modify and solve the EV user decision problem by excluding the capacity-based from the pricing function, which leads to the following problem:

$$\min f(x) = P_0x + \beta(x - D)^2 \quad (5.24)$$

The optimal solution for this problem is $x^* = \frac{2\beta D - p^0}{2\beta}$. To train a pricing agent with the traditional dynamic model, we update the EV users' decisions in the simulation. Due to the high complexity of the dynamic decision problem, and also for the sake of model comparability, we use a modified version of the DRL algorithm to find near-optimal traditional dynamic pricing policies. We consider the same state space and reward function, but the action space is different, $a_t = (p_t^0)$, and contains only the power-free pricing parameter.

Dynamic Menu-based Pricing

In recent advances in pricing management for charging services (e.g., Abdalrahman and Zhuang 2020, Lu et al. 2022), researchers optimize the price of differentiated charging services. For example, Abdalrahman and Zhuang (2020) considers different prices for different charging classes (i.e., charging speed) to avoid over-utilization of scarce facilities, and Lu et al. (2022) considers a deadline differentiated dynamic price menu that offers multiple choice-pairs of deadlines and charging prices. Similarly, we also consider a menu-based pricing model where the operator offers two options of charging power (11 and 50 kW) with different prices¹⁴. To solve this, we use a similar DRL model, but with a different action space consisting of two different elements $a_t = (a_t^{11}, a_t^{50})$.

To solve the decision problem regarding the energy requested by EV users, we adopt the same objective function as in the traditional dynamic pricing model (Eq. (5.24)). Since users make discrete choices, it is not possible to solve the EV decision optimization problem analytically. Therefore, we assume that EV users choose the option that minimizes their individual costs. Based on this assumption, we compute costs for four different energy demand levels: $x = 0$, $x = 11\delta$, $x = 55\delta$, and $x = D$. These values correspond to the amount of energy requested by a user based on their choices, where δ represents the length of stay and D represents the user's maximum demand. For example, a user may decline all charging options ($x = 0$) or choose a

¹⁴Finding the optimal number of options and charging speed for each of them is a challenge for this model, but we choose these options similarly to practice where charging facilities might offer slow and fast charging options.

slow charge and request $x = 11\delta$ units of energy. This discrete approach ensures that users do not request more energy than their maximum demand, as this would increase their costs and dissatisfaction. Additionally, if the price is too high and charging is not urgent, users may forgo charging altogether and request zero energy.

Time of Use Pricing

First, we consider a ToU pricing model that follows the pattern of electricity costs. Similar to other benchmark policies, the maximum price for the ToU strategy is set at \$1.50. Thus, during periods of high peak electricity costs, the charge price is \$1.50 and decreases during periods of lower electricity costs.

Benchmark Results

Figure 5.3 shows the learning curves of the proposed model compared to the benchmark algorithms. As shown, all three RL-based dynamic algorithms reach convergence after several hundred episodes. Notably, our proposed DSS achieves the highest objective function after training, providing a profit increase of approximately 33% and 86% compared to other dynamic pricing models and the ToU model, respectively. While menu-based pricing models are expected to outperform traditional dynamic pricing by better accommodating heterogeneous user demand through multiple service offerings, they do not in this case due to the penalty costs incurred. Finally, our results are consistent with previous research (Lin et al. 2023) indicating that static pricing schemes, such as the ToU model, may not provide sufficient profitability for charging services.




Figure 5.3: Benchmark of Dynamic Capacity-Based Pricing against State-of-Art Models

To demonstrate that our DSS can meet sustainability goals while ensuring profitability, we examine the energy consumption patterns of the proposed and benchmark models. Figure 5.4 shows the energy supply from different sources. During the day, the EVCH benefits from on-site energy production, which peaks around noon. All three dynamic models were successful in attracting high charging demand during these hours, likely by lowering prices. However, the total load for the ToU model remains lower than the other models during these periods because the cost of electricity does not match the solar energy production curves. In terms of peak control, our proposed dynamic capacity pricing model is the only one that keeps grid usage below the 500 kW threshold set by the electricity tariff. This is achieved by the capacity term in the pricing function, which encourages users to reduce their energy demands during periods of high demand. The other models show nearly double the peak consumption compared to the capacity-based model, underscoring the advantage of our approach in managing peak demand and coordinating with renewable energy production.




Figure 5.4: Electric Supply Sources Over Simulation Horizon

5.5 Discussion and Conclusion

We present a machine learning-based DSS designed to optimize the operational management of large-scale EVCHs through dynamic capacity-based pricing. Our goal is to promote sustainable

mobility by providing EVCH operators with a DSS that enhances profitability while improving system-level performance, such as alleviating grid stress by reshaping aggregated charging loads. To achieve this, we integrate DRL with a dynamic capacity-based pricing model to derive time-varying pricing policies. Our approach accounts for diverse user preferences and attributes, including price sensitivity, and the stochastic nature of arrival times and energy demands. By incorporating demand uncertainty and user heterogeneity, our pricing model increases EVCH profitability while maintaining operational efficiency.

To optimally determine dynamic capacity-based pricing policies, we build a DRL agent that is able to learn from the stochastic environment. To train and evaluate our agent, we build a realistic simulation that closely mimics the real world. An important feature of the real world is stochasticity, which we model using an agent-based model calibrated with empirical data. Pricing decisions have time dependencies, as each decision affects the demand for charges at the next decision step and consequently affects the pricing decision. We define the problem as a Markov decision process. Since obtaining information about the dynamics of the environment is impractical (e.g., estimating the behavior of EV users is costly and varies over time), we use a model-free approach (DRL) to approximate the optimal pricing policies. By interacting with the environment, these models learn the policies over multiple episodes. Therefore, we create a digital representation of the EVCH environment using ABM and calibrate the charging demand using real-world parking and charging observations. A realistic and detailed simulation is essential for deriving actionable insights and mitigating the challenges associated with transferring RL agents from a simulated environment to real-world applications.

We demonstrate that our proposed decision support system (DSS) outperforms all benchmark policies, including traditional dynamic pricing, dynamic menu-based pricing, and time-of-use (ToU) pricing, achieving approximately 33% higher profits. Beyond profitability, our model enhances the overall performance of EV charging hubs (EVCHs) by effectively reducing peak electricity consumption and optimizing the utilization of fluctuating renewable energy sources throughout the day. This approach is particularly beneficial in future scenarios where high EV adoption and increased reliance on renewables could strain the power grid, leading to new peak loads and greater supply uncertainty. By incorporating a power-dependent pricing function and a powerful DRL agent that can learn near-optimal policies in complex environments, our model allows EVCH operators to exert greater control over aggregated charging demand, mitigating peak loads while maximizing the use of green energy more effectively than traditional pricing strategies.

Our study also shows that EVCH profits are significantly affected by user characteristics such as the cost of their alternative charging options and their willingness to fully charge. Although accurate measurements of user behavior were not obtained in our research, our model-free approach eliminates the need for user input in real-world implementation and instead learns

through environmental interactions. Therefore, the discussed assumption about user behavior serves only to analyze our model and offer managerial insights for EVCH and grid operators.

Finally, our findings offer insights for EVCH operators and energy providers. A well-designed dynamic pricing model can incentivize investment in EVCHs, which are crucial for widespread EV adoption and integration into power grid systems—without requiring extensive grid expansions or major system modifications. Additionally, our decision support platform serves as a powerful tool, enabling EVCH operators to determine near-optimal pricing strategies for various configurations. It also allows them to assess the value of potential investments, such as on-site renewable generation and large-scale energy storage, across different user segments and operational settings.

Our research has certain limitations. Methodologically, we decouple the load scheduling and pricing problems, addressing them separately. This separation allows us to focus on pricing management, which is the primary objective of this study. Moreover, pricing decisions operate on a different scale and complexity compared to load scheduling, further justifying this distinction. Future research could explore the potential benefits of training both scheduling and pricing agents simultaneously, allowing them to interact dynamically. Additionally, we plan to validate our assumptions regarding user price sensitivities through survey experiments. We anticipate that this validation will enhance the robustness of our sensitivity analyses.

5.6 Appendix

This appendix presents proofs for the analytical solution of EV user decision-making and provides additional details on the development of an agent-based model (ABM) for EV charging hubs (EVCHs).

5.7 Analytical Solutions for EV User Decision-Making Problem

Since the mathematical programming model for EV user decision-making involves a nonlinear objective function with simplex constraints, a closed-form solution for the optimal requested energy can be derived using Lagrangian methods. Additionally, we demonstrate that it is unnecessary to impose an explicit upper bound on the requested energy, as the optimal solution inherently remains within the user's maximum energy demand.

5.7.1 Solving the optimization problem using the Lagrangian method

We introduce a Lagrange multiplier λ for this inequality constraint. The Lagrangian function becomes:

$$\mathcal{L}(x, \lambda) = \left(p_0 + \alpha \frac{x}{\delta}\right) x + \beta(x - D)^2 - \lambda x$$

We write the first partial derivative of the Lagrangian function

- Derivative with respect to x :

$$\frac{\partial \mathcal{L}}{\partial x} = p_0 + 2\alpha \frac{x}{\delta} + 2\beta(x - D) - \lambda = 0$$

- Derivative with respect to λ :

$$\frac{\partial \mathcal{L}}{\partial \lambda} = -x = 0$$

This implies $x \geq 0$, and when $x = 0$, the constraint is active. Afterwards, we solve the system of equations.

- If the constraint $x \geq 0$ is inactive (meaning $x > 0$), then $\lambda = 0$.
- If $x = 0$, the constraint is active, so $\lambda \geq 0$.

To find the critical points we employ the first derivative condition.

$$p_0 + 2\alpha \frac{x}{\delta} + 2\beta x - 2\beta D = 0$$

$$x = \frac{(2\beta D - p_0)\delta}{2\alpha + 2\beta\delta}$$

If the critical point is positive and within the allowable range $x \geq 0$, it is a candidate for the maximum. If the critical point results in $x < 0$, then the maximum occurs at $x = 0$.

We also need to mention that the utility function always has a minimum, since the second derivative is always positive.

$$f''(x) = \frac{dx}{d} \left(p_0 + 2\alpha \frac{x}{\delta} + 2\beta(x - D) \right)$$

$$f''(x) = \frac{2\alpha}{\delta} + 2\beta$$

5.7.2 The optimal requested energy can not be greater than the raw demand

In the following we proof that the optimal requested energy cannot be larger than the raw demand (D):

$$\frac{(2\beta D - p_0)\delta}{2(\beta\delta + \alpha)} \leq D$$

We first multiply both sides by $2(\beta\delta + \alpha)$. Since $2(\beta\delta + \alpha)$ is positive, multiplying both sides by this term gives:

$$2\beta D\delta - p_0\delta \leq 2\beta D\delta + 2\alpha D$$

We then simplify the inequality by removing the common term $2\beta D\delta$ from both sides:

$$-p_0\delta \leq 2\alpha D$$

We divide the inequality by $-\delta$. Since $\delta > 0$, dividing both sides by $-\delta$ reverses the inequality:

$$p_0 \geq -\frac{2\alpha D}{\delta}$$

We bring all parameters to one side of the inequality:

$$-\frac{2\alpha D}{\delta} - p_0 \leq 0$$

Note that this inequality always holds since all parameters are non-negative.

5.8 Mathematical Programming Notations

To enhance the readability of our mathematical models, we present all the indexes used in this paper. We define the sets, variables, and parameters of the mathematical model in Table ??.

Symbol	Description
Sets	
\mathcal{T}	Set of decision-making time windows
\mathcal{I}	Set of EVs parking in the EVCHs over the time horizon T
Variables	
p_t^0	Fixed component of the pricing function at time t
α_t	Power-dependent component of the pricing function at time t
$y_{i,t}$	Continuous variable for the charging rate of vehicle i at time t
z_i	Boolean variable indicating whether vehicle i is assigned to a charger
$x_{i,t}$	Continuous variable representing the energy requested by vehicle i at time t
p^*	Continuous variable representing the excess peak electricity consumption
e_t^{Grid}	Continuous variable for the energy drawn from the grid at time t
\bar{E}_t	Average energy demand of EVs in the EVCH at time t
\bar{P}_t	Average power demand of EVs in the EVCH at time t
Parameters	
C_t^e	Electricity cost at time t
C^p	Peak demand charge
Δ_t	Duration of each time window
l^*	Peak threshold set by the electricity tariff
R	Maximum charging power of all stations
κ	Number of charging docks in the EVCH
$f_t^{PV} \kappa^{PV}$	Power generated from PV at time t
$U_{i,t}$	Boolean parameter indicating the presence of vehicle i at time t in the EVCH
$A_{i,t}$	Boolean parameter indicating if vehicle i arrives at time t
$De_{i,t}$	Boolean parameter indicating if vehicle i departs at time t
D_i	Maximum energy demand of user i
β_i	User i 's willingness to charge up to maximum demand
δ_i	Duration of stay of user i
ζ_i	Penalty for unmet energy demand of user i

Table 5.2: Notation Used in the Mathematical Programming Model

5.9 Additional Information for Agent-Based Modeling of EVCHs

In this section, we provide some details of input prepreation and information about the physical equipments of EVCHs.

5.9.1 Estimating charging demand

For estimating the requested energy per vehicle e_i^d . We employ a recently published real-world dataset by Lee et al. (2019) containing >25,000 charging transactions for the year 2019. Per each charging transaction the full preference vector $v_i = (A_i, \delta_i, e_i^d)$ is available. We combine the charging data (which only contains served sessions that are constrained by the available infrastructure) with our parking dataset. We train a prediction model on the labeled dataset (Lee et al. 2019) and use the resulting model to predict charging demand in the parking dataset.

Symbol	Description
π	Policy of the RL agent
r_t	Reward of the RL agent at time t
s_t	State of the RL agent at time t
a_t	Action of the RL agent at time t
C_t^M	Cost of missed demand at time t
J	Objective function of the RL agent
ρ_π	State-action distribution under policy π
\mathcal{H}	Entropy of the RL agent's policy
α^{RL}	Entropy temperature parameter of the RL agent
D^{RL}	Replay buffer of the RL agent
θ	Parameter of the actor network in the RL agent
ϕ	Parameter of the critic network in the RL agent
Q	Action-state value function of the RL agent
V	State-value function of the RL agent

Table 5.3: Notation Used in the Reinforcement Learning Model

Specifically, we train a kNN-model on the charging transaction dataset using the set of clustering variables from before as predictors and the requested energy in kWh as outcome variable. Cross-validation reveals $k=12$ neighbors to be a good value. We use the fitted kNN-model to predict charging demand per transaction in our unlabeled parking dataset. Using this approach, we obtain an exponentially distributed charging demand across the entire population of EVs with average demand of 26.46 kWh ($\sigma = 17.20$ kWh) per parking session. The distributional shape of charging demand is consistent with the one seen in other empirical EV charging settings (e.g., Ferguson et al. 2018).

5.9.2 EVCH physical equipments

We define an EVCH as an EV charging-capable parking lot, depot or garage that will typically be attached to an existing building. Both the building and the EVCH receive power from the same grid connection point, which is constrained to a certain capacity. The integrated facility may have additional on-site behind-the-meter generation (e.g., PV), and battery storage.

The EVCH could be with different types of EV charging docks (22kW AC or 50kW DC docks) and the number of connectors per dock (ranging from single-connector setups to up to four connectors per dock). Crucially, for charging docks with multiple connectors, we allow for simultaneous charging of EVs meaning the rated power per dock can be shared dynamically and flexibly by all connected vehicles. This is different from the more prevalent single-server docks which either possess just a single connector, or multiple connectors that can only be used sequentially (i.e., one after the other). A multi-connector setup offers important theoretical advantages over single connector and/or sequential dock architectures. These include higher potential utilization (vehicles that have completed their charging cycle do not block charging

docks) (Ferguson et al. 2018) and more flexibility of charge cycles over the full parking duration of each connected vehicle. We assume that PV generation can be scaled and is limited by local facility space constraints, such as roof space. Naturally, the actual available PV capacity at any given time will depend on local weather and solar irradiation conditions, which we capture by means of a time-dependent load factor. The agent-based simulation could incorporate on-site storage utilizing Li-Ion battery technology and regulate the storage system's charge and discharge processes.

5.9.3 Electricity Tariff for EVCH

	Summer (Jun - Sep)	Winter (all other months)
Super Off-Peak (8am-4pm)	0.08 USD/kWh	0.06 USD/kWh
On-Peak (4pm to 9pm)	0.23 USD/kWh	0.23 USD/kWh
Off-Peak (9pm-8am)	0.08 USD/kWh	0.08 USD/kWh
Peak Charge (monthly)	15.48 USD/kW	

Table 5.4: Time-of-use Tariff and Demand Charge for Large-scale EV Charging Customers (> 300 kW)

References

- A. Abbasi, F. Zahedi, D. Zeng, Y. Chen, H. Chen, and J. F. Nunamaker Jr. Enhancing predictive analytics for anti-phishing by exploiting website genre information. *Journal of Management Information Systems*, 31(4):109–157, 2015.
- A. Abdalrahman and W. Zhuang. Dynamic pricing for differentiated pev charging services using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 23(2):1415–1427, 2020.
- A. Abdelwahed, P. L. van den Berg, T. Brandt, J. Collins, and W. Ketter. Evaluating and optimizing opportunity fast-charging schedules in transit battery electric bus networks. *Transportation Science*, 54(6):1601–1615, 2020. ISSN 15265447.
- H. Abouee-Mehrizi, O. Baron, O. Berman, and D. Chen. Adoption of Electric Vehicles in Car Sharing Market. *Production and Operations Management*, 30(1):190–209, 2021. ISSN 19375956.
- M. Adam, K. Roethke, and A. Benlian. Human versus automated sales agents: How and why customer responses shift across sales stages. *Information Systems Research*, 0(0), 2022.
- I. E. Agency. Transport: Improving the sustainability of passenger and freight transport, 2022. URL <https://www.iea.org/topics/transport>.
- R. Ahadi, W. Ketter, J. Collins, and N. Daina. Siting and sizing of charging infrastructure for shared autonomous electric fleets. *In Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021.
- R. Ahadi, W. Ketter, J. Collins, and N. Daina. Cooperative Learning for Smart Charging of Shared Autonomous Vehicle Fleets. *Transportation Science*, 57(3):613–630, may 2023. ISSN 0041-1655.
- L. Al-Kanj, J. Nascimento, and W. B. Powell. Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Journal of Operational Research*, 284(3):1088–1106, 2020a. ISSN 03772217.
- L. Al-Kanj, J. Nascimento, and W. B. Powell. Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Journal of Operational Research*, 284(3):1088–1106, 2020b.
- M. S. Alam, F. S. Al-Ismael, A. Salem, and M. A. Abido. High-level penetration of renewable energy sources into grid utility: Challenges and solutions. *IEEE access*, 8:190277–190299, 2020.
- V. Albino, U. Berardi, and R. M. Dangelico. Smart cities: Definitions, dimensions, performance, and initiatives. *Journal of Urban Technology*, 22(1):3–21, 2015.

- J. Aloysius, C. Deck, and A. Farmer. Sequential pricing of multiple products: Leveraging revealed preferences of retail customers online and with auto-id technologies. *Information Systems Research*, 24(2):372–393, 2013.
- M. Ansarin, Y. Ghiassi-Farrokhfal, W. Ketter, and J. Collins. Economic inefficiencies of pricing distributed generation under novel tariff designs. *Applied Energy*, 313:118839, 2022.
- D. Ao, Z. Lai, and S. Li. Control of dynamic ride-hailing networks with a mixed fleet of autonomous vehicles and for-hire human drivers. *Transportation Research Part E: Logistics and Transportation Review*, 189:103680, 2024.
- M. Attaran and B. G. Celik. Digital twin: Benefits, use cases, challenges, and opportunities. *Decision Analytics Journal*, 6:100165, 2023. ISSN 2772-6622.
- Y. Babar and G. Burtch. Examining the heterogeneous impact of ride-hailing services on public transit use. *Information Systems Research*, 31:820–834, 9 2020.
- Y. Babar and G. Burtch. Recharging retail: Estimating consumer demand spillovers from electric vehicle charging stations. *Manufacturing & Service Operations Management*, 26(3):797–813, 2024.
- J. Babic, A. Carvalho, W. Ketter, and V. Podobnik. A data-driven approach to managing electric vehicle charging infrastructure in parking lots. *Transportation Research Part D: Transport and Environment*, 105:103198, 5 2022a. ISSN 13619209.
- J. Babic, A. Carvalho, W. Ketter, and V. Podobnik. A data-driven approach to managing electric vehicle charging infrastructure in parking lots. *Transportation Research Part D: Transport and Environment*, 105:103198, 2022b.
- J. Bai, K. C. So, C. S. Tang, X. Chen, and H. Wang. Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing & Service Operations Management*, 21(3):556–570, 2019.
- S. Banerjee, D. Freund, and T. Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. *Operations Research*, 70(3):1783–1805, 2022.
- R. Bapna, P. Goes, A. Gupta, and G. Karuga. Predicting bidders’ willingness to pay in online multiunit ascending auctions: Analytical and empirical insights. *INFORMS Journal on Computing*, 20(3):345–355, 2008.
- A. Barfar and B. Padmanabhan. Predicting presidential election outcomes from what people watch. *Big Data*, 5(1):32–41, 2017.
- M. Batty, K. W. Axhausen, F. Giannotti, A. Pozdnoukhov, A. Bazzani, M. Wachowicz, G. Ouzounis, and Y. Portugali. Smart cities of the future. *The European Physical Journal Special Topics*, 214(1):481–518, 2012.
- M. Ben-Akiva and S. R. Lerman. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, Massachusetts, 1985.
- C. Benevolo, R. P. Dameri, and B. D’auria. Smart mobility in smart city: Action taxonomy, ict intensity and public benefits. In *Empowering organizations: Enabling platforms and artefacts*, pages 13–28. Springer, 2016.
- S. Benjaafar, Z. Wang, and X. Yang. Human in the loop automation: Ride-hailing with remote (tele-) drivers. *Management Science*, 2024.

- A. Bensoussan, R. Mookerjee, V. Mookerjee, and W. T. Yue. Maintaining diagnostic knowledge-based systems: a control-theoretic approach. *Management Science*, 55(2):294–310, 2009.
- H. K. Bhargava and A. N. Mishra. Electronic medical records and physician productivity: Evidence from panel data analysis. *Management Science*, 60(10):2543–2562, 2014.
- M. Bichler, A. Gupta, and W. Ketter. Designing smart markets. *Information Systems Research*, 21(4):688–699, 2010.
- M. Bichler, J. Knörr, and F. Maldonado. Pricing in nonconvex markets: How to price electricity in the presence of demand response. *Information Systems Research*, 34(2):652–675, 2023.
- C. M. Bonina, K. M. Koskinen, B. D. Eaton, and A. Gawer. Digital platforms for development: Foundations and research agenda. *Information Systems Journal*, 31:869–902, 2021.
- P. M. Bösch, F. Becker, H. Becker, and K. W. Axhausen. Cost-based analysis of autonomous mobility services. *Transport Policy*, 64:76–91, 2018.
- S. Boschert and R. Rosen. Digital Twin—The Simulation Aspect. In *Mechatronic Futures*, pages 59–74. Springer International Publishing, Cham, 2016. ISBN 9783319321561.
- T. P. Böttcher, S. Empelmann, J. Weking, A. Hein, and H. Krcmar. Digital sustainable business models: Using digital technology to integrate ecological sustainability into the core of business models. *Information Systems Journal*, 34(3):736–761, 2024.
- T. Brandt, S. Feuerriegel, and D. Neumann. Modeling interferences in information systems design for cyberphysical systems: Insights from a smart grid application. *European Journal of Information Systems*, 27(2):207–220, 2018.
- E. Bruun and M. Givoni. Sustainable mobility: Six research routes to steer transport policy. *Nature*, 523(7558):29–31, 2015.
- G. Burtch, S. Carnahan, and B. N. Greenwood. Can You Gig It? An Empirical Examination of the Gig Economy and Entrepreneurial Activity. *Management Science*, 64(12):5497–5520, dec 2018.
- G. P. Cachon and M. Fisher. Supply chain inventory management and the value of shared information. *Management science*, 46(8):1032–1048, 2000.
- T. Calinski and J. Harabasz. A dendrite method for cluster analysis. *Communications in Statistics*, 3(1):1–27, 1974.
- S. Caussade, J. d. D. Ortúzar, L. I. Rizzi, and D. A. Hensher. Assessing the influence of design dimensions on stated choice experiment estimates. *Transportation Research Part B: Methodological*, 39(7):621–640, 2005.
- L. Chen, L. He, and Y. Zhou. An exponential cone programming approach for managing electric vehicle charging. *Operations Research*, 72(5):2215–2240, 2024a.
- T. D. Chen, K. M. Kockelman, and J. P. Hanna. Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle & charging infrastructure decisions. *Transportation Research Part A: Policy and Practice*, 94:243–254, 2016. ISSN 09658564.
- Y. Chen and Y. Liu. Integrated optimization of planning and operations for shared autonomous electric vehicle systems. *Transportation Science*, 0(0), 8 2022. ISSN 0041-1655.
- Z. Chen, Y. Miao, J. Ke, and Q.-C. He. Operations and regulations for a ride-sourcing market with a mixed fleet of human drivers and autonomous vehicles. *Transportation Research Part C: Emerging Technologies*, 160:104519, 2024b.

- Z. Cheng, M. S. Pang, and P. A. Pavlou. Mitigating traffic congestion: The role of intelligent transportation systems. *Information Systems Research*, 31(3):653–674, 2020.
- T. M. Choi, S. Kumar, X. Yue, and H. L. Chan. Disruptive Technologies and Operations Management in the Industry 4.0 Era and Beyond. *Production and Operations Management*, 31(1):9–31, 2022. ISSN 19375956.
- P. Christodoulou. Soft actor-critic for discrete action settings. *arXiv preprint arXiv:1910.07207*, 2019.
- C. Cimino, E. Negri, and L. Fumagalli. Review of digital twin applications in manufacturing. *Computers in Industry*, 113:103130, 2019. ISSN 01663615.
- M. C. Cohen. Big Data and Service Operations. *Production and Operations Management*, 27(9):1709–1723, 2018. ISSN 19375956.
- C. Colas, O. Sigaud, and P.-Y. Oudeyer. Gep-pg: Decoupling exploration and exploitation in deep reinforcement learning algorithms. In *International conference on machine learning*, pages 1039–1048. PMLR, 2018.
- J. Corbett. Designing and using carbon management systems to promote ecologically responsible behaviors. *Journal of the Association for Information Systems*, 14(7):339–378, 2013.
- P. Cramton, R. Geddes, and A. Ockenfels. Set road charges in real time to ease traffic. *Nature*, 560(2 August 2018):23–25, 2018.
- J. Cranefield and J. Pries-Heje. Boundary management strategies for leading digital transformation in smart cities. *MIS Quarterly Executive*, 22(2), 2023.
- L. Cui, Q. Wang, H. Qu, M. Wang, Y. Wu, and L. Ge. Dynamic pricing for fast charging stations with deep reinforcement learning. *Applied Energy*, 346:121334, 2023.
- Y. Cui, Z. Hu, and X. Duan. Optimal pricing of public electric vehicle charging stations considering operations of coupled transportation and power systems. *IEEE Transactions on Smart Grid*, 12(4):3278–3288, 2021.
- R. Curtale, F. Liao, and E. Rebalski. Transitional behavioral intention to use autonomous electric car-sharing services: Evidence from four european countries. *Transportation Research Part C: Emerging Technologies*, 135:103516, 2022.
- N. Daina, A. Sivakumar, and J. W. Polak. Electric vehicle charging choices: Modelling and implications for smart charging services. *Transportation Research Part C: Emerging Technologies*, 81:36–56, 2017. ISSN 0968090X.
- F. D. Davis. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, pages 319–340, 1989.
- M. de Reuver, C. Sørensen, and R. C. Basole. The digital platform: A research agenda. *Journal of Information Technology*, 33(2):124–135, 2018.
- J. Dedrick. Green IS: Concepts and issues for information systems research. *Communications of the Association for Information Systems*, 27(1):173–184, 2010. ISSN 15293181.
- M. Demircan, R. Ahadi, and W. Ketter. Sustainability vs. price: Analysis of electric multi-modal vehicle sharing systems under substitution effects. *ECIS 2022 Research Papers*, 162, 2022.
- E. M. Demirezen, S. Kumar, and B. Shetty. Managing co-creation in information technology projects: A differential games approach. *Information Systems Research*, 27(3):517–537, 2016.

- S. Devlin and D. Kudenko. Theoretical considerations of potential-based reward shaping for multi-agent systems. In *The 10th International Conference on Autonomous Agents and Multiagent Systems*, pages 225–232. ACM, 2011.
- S. Dhanorkar and G. Burtch. The heterogeneous effects of p2p ride-hailing on traffic: Evidence from uber’s entry in california. *Transportation Science*, 56:750–774, 5 2022.
- B. Dietvorst, J. Simmons, and C. Massey. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of experimental psychology. General*, 144, 11 2014.
- O. Dlugosch, T. Brandt, and D. Neumann. Combining analytics and simulation methods to assess the impact of shared, autonomous electric vehicles on sustainable urban mobility. *Information & Management*, page 103285, 2020.
- O. Dlugosch, T. Brandt, and D. Neumann. Combining analytics and simulation methods to assess the impact of shared, autonomous electric vehicles on sustainable urban mobility. *Information & Management*, 59(5):103285, 2022.
- Y. Dong, R. D. Koster, D. Roy, and Y. Yu. Dynamic vehicle allocation policies for shared autonomous electric fleets. *Transportation Science*, 56:1238–1258, 9 2022. ISSN 0041-1655.
- G. Dulac-Arnold, R. Evans, H. van Hasselt, P. Sunehag, T. Lillicrap, J. Hunt, T. Mann, T. Weber, T. Degris, and B. Coppin. Deep reinforcement learning in large discrete action spaces. *arXiv preprint arXiv:1512.07679*, 2015.
- F. D’Acunto, N. Prabhala, and A. G. Rossi. The promises and pitfalls of robo-advising. *The Review of Financial Studies*, 32(5):1983–2020, 2019.
- O. Egbue and S. Long. Barriers to widespread adoption of electric vehicles: An analysis of consumer attitudes and perceptions. *Energy policy*, 48:717–729, 2012.
- H. Engel, R. Hensley, S. Knupfer, and S. Sahdev. Charging ahead: Electric-vehicle infrastructure demand. *McKinsey Center for Future Mobility*, 8, 2018.
- E. Eren and V. E. Uz. A review on bike-sharing: The factors affecting bike-sharing demand. *Sustainable cities and society*, 54:101882, 2020.
- D. J. Fagnant and K. M. Kockelman. The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies*, 40:1–13, 2014. ISSN 0968090X.
- D. J. Fagnant and K. M. Kockelman. Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in austin, texas. *Transportation*, 45(1):143–158, 2018.
- D. J. Fagnant, K. M. Kockelman, and P. Bansal. Operations of shared autonomous vehicle fleet for austin, texas, market. *Transportation Research Record*, 2563(1):98–106, 2015.
- B. Ferguson, V. Nagaraj, E. C. Kara, and M. Alizadeh. Optimal Planning of Workplace Electric Vehicle Charging Infrastructure with Smart Charging Opportunities. *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2018-Novem:1149–1154, 2018.
- C. M. Flath, J. P. Ilg, S. Gottwalt, H. Schmeck, and C. Weinhardt. Improving electric vehicle charging coordination through area pricing. *Transportation Science*, 48(4):619–634, 2014.
- A. M. Florio, N. Absi, and D. Feillet. Routing electric vehicles on congested street networks. *Transportation Science*, 55(1):238–256, 2021. ISSN 15265447.

- G. Fridgen, L. Häfner, C. König, and T. Sachs. Providing utility to utilities: the value of information systems enabled flexibility in electricity consumption. *Journal of the Association for Information Systems*, 17(8):1, 2016.
- A. Froger, O. Jabali, J. E. Mendoza, and G. Laporte. The Electric Vehicle Routing Problem with Capacitated Charging Stations. *Transportation Science*, 56(2):460–482, 2022. ISSN 0041-1655.
- C. Fu, C. Fu, and M. Michael. *Handbook of simulation optimization*. Springer, 2015.
- Fuel economy guide. <https://www.fueleconomy.gov/feg/pdfs/guides/FEG2020.pdf>, 2020.
- M. Gansterer, R. F. Hartl, and M. Tzur. Transportation in the sharing economy. *Transportation Science*, 56:567–570, 5 2022.
- Z. Gao, Y. Gao, Y. Hu, Z. Jiang, and J. Su. Application of Deep Q-Network in Portfolio Management. In *2020 5th IEEE International Conference on Big Data Analytics (ICBDA)*, pages 268–275. IEEE, may 2020.
- R. Ge, Z. E. Zheng, X. Tian, and L. Liao. Human–robot interaction: When investors adjust the usage of robo-advisors in peer-to-peer lending. *Information Systems Research*, 32(3):774–785, 2021.
- Y. Ghiassi-Farrokhfal, C. Rosenberg, S. Keshav, and M. B. Adjaho. Joint Optimal Design and Operation of Hybrid Energy Storage Systems. *IEEE Journal on Selected Areas in Communications*, 34(3):639–650, 2016. ISSN 07338716.
- J. Gijbrecchts, R. N. Boute, J. A. Van Mieghem, and D. J. Zhang. Can Deep Reinforcement Learning Improve Inventory Management? Performance on Lost Sales, Dual-Sourcing, and Multi-Echelon Problems. *Manufacturing & Service Operations Management*, 24(3):1349–1368, may 2022. ISSN 1523-4614.
- M. Gilleran, E. Bonnema, J. Woods, P. Mishra, I. Doebber, C. Hunter, M. Mitchell, and M. Mann. Impact of electric vehicle charging on the power demand of retail buildings. *Advances in Applied Energy*, 4:100062, 2021.
- E. Glaessgen and D. Stargel. The digital twin paradigm for future nasa and us air force vehicles. In *53rd AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics and materials conference 20th AIAA/ASME/AHS adaptive structures conference 14th AIAA*, page 1818, 2012.
- E. Glikson and A. W. Woolley. Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2):627–660, 2020.
- Global petrol prices. https://www.globalpetrolprices.com/Germany/gasoline_prices, 2021.
- U. Gnewuch, S. Morana, O. Hinz, R. Kellner, and A. Maedche. More than a bot? the impact of disclosing human involvement on customer interactions with hybrid service agents. *Information Systems Research*, 2023. ISSN 1047-7047 1526-5536.
- F. Golbabaee, T. Yigitcanlar, and J. Bunker. The role of shared autonomous vehicle systems in delivering smart urban mobility: A systematic review of the literature. *International Journal of Sustainable Transportation*, 15(10):731–748, 2021.
- W. H. Greene and D. A. Hensher. A latent class model for discrete choice analysis: contrasts with mixed logit. *Transportation Research Part B: Methodological*, 37(8):681–698, 2003.
- M. Grieves and J. Vickers. *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*, pages 85–113. Springer International Publishing, Cham, 2017. ISBN 978-3-319-38756-7.

- A. Griffin and J. R. Hauser. The Voice of the Customer. *Marketing Science*, 12(1):1–27, feb 1993. ISSN 0732-2399.
- I. Grondman, L. Busoniu, G. A. Lopes, and R. Babuska. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, part C (applications and reviews)*, 42(6):1291–1307, 2012.
- H. Guda and U. Subramanian. Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science*, 65(5):1995–2014, 2019.
- M. Guillet, G. Hiermann, A. Kröller, and M. Schiffer. Electric vehicle charging station search in stochastic environments. *Transportation Science*, 56(2):483–500, 2022. ISSN 0041-1655.
- Z. Guo, J. Li, and R. Ramesh. Optimal management of virtual infrastructures under flexible cloud service agreements. *Information Systems Research*, 30(4):1424–1446, 2019.
- G. Gust, T. Brandt, S. Mashayekh, M. Heleno, N. DeForest, M. Stadler, and D. Neumann. Strategies for microgrid operation under real-world conditions. *European Journal of Operational Research*, 292(1):339–352, 2021. ISSN 03772217.
- T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.
- C. J. Haboucha, R. Ishaq, and Y. Shifan. User preferences regarding autonomous vehicles. *Transportation Research Part C: Emerging Technologies*, 78:37–49, 2017.
- K. Haki, J. Beese, S. Aier, and R. Winter. The evolution of information systems architecture: An agent-based simulation model. *MIS Quarterly*, 44(1), 2020.
- K. Haki, H. Tanriverdi, D. Safaei, M. Schmid, S. Aier, and R. Winter. Generativity and profitability on b2b innovation platforms: A simulation-based theory development. *MIS Quarterly*, 48(2), 2024.
- Y. Han, Y. Chen, F. Han, and K. R. Liu. An optimal dynamic pricing and schedule approach in v2g. In *Proceedings of The 2012 Asia Pacific Signal and Information Processing Association Annual Summit and Conference*, pages 1–8. IEEE, 2012.
- L. A. Hannah. Stochastic Optimization. In *International Encyclopedia of the Social & Behavioral Sciences*, volume 23, pages 473–481. Elsevier, second edi edition, 2015.
- L. He, H. Y. Mak, Y. Rong, and Z. J. M. Shen. Service region design for urban electric vehicle sharing systems. *Manufacturing and Service Operations Management*, 19(2):309–327, 2017. ISSN 15265498.
- L. He, G. Ma, W. Qi, and X. Wang. Charging an electric vehicle-sharing fleet. *Manufacturing & Service Operations Management*, 23(2):471–487, 2021.
- K. Heineke, B. Kloss, T. Möller, and C. Wiemuth. Shared mobility: Where it stands and where it’s going. *McKinsey Center for Future Mobility*, 2021.
- K. Heineke, R. Heuss, P. Kampshoff, A. Kelkar, and M. Kellner. The road to affordable autonomous mobility. Report, McKinsey, 2022.
- D. Hensher, J. Rose, and W. Greene. *Applied Choice Analysis*. Cambridge University Press, Cambridge, Massachusetts, 01 2005.
- D. A. Hensher and W. H. Greene. The mixed logit model: the state of practice. *Transportation*, 30: 133–176, 2003.

- A. R. Hevner, S. T. March, J. Park, and S. Ram. Design science in information systems research. *MIS quarterly*, pages 75–105, 2004.
- H. Hilpert, J. Kranz, and M. Schumann. Leveraging green is in logistics: Developing an artifact for greenhouse gas emission tracking. *Business & Information Systems Engineering*, 5:315–325, 2013.
- S. C. Ho, W. Y. Szeto, Y.-H. Kuo, J. M. Leung, M. Petering, and T. W. Tou. A survey of dial-a-ride problems: Literature review and recent developments. *Transportation Research Part B: Methodological*, 111:395–421, 2018.
- S. J. Hong, J. M. Bauer, K. Lee, and N. F. Granados. Drivers of supplier participation in ride-hailing platforms. *Journal of Management Information Systems*, 37(3):602–630, 2020.
- Z. Hoover, E. Polymeneas, and S. Sahdev. How to integrate electric-vehicle charging into existing building and grid infrastructure. Technical report, McKinsey & Company, 2021.
- L. Hou, C. Wang, and J. Yan. Bidding for preferred timing: An auction design for electric vehicle charging station scheduling. *IEEE Transactions on Intelligent Transportation Systems*, 21(8):3332–3343, 2019.
- Y. Huang and K. M. Kockelman. Electric vehicle charging station locations: Elastic demand, station congestion, and network equilibrium. *Transportation Research Part D: Transport and Environment*, 78:102179, 2020.
- Y. Huang and Y. Zhou. An optimization framework for workplace charging strategies. *Transportation Research Part C: Emerging Technologies*, 52:144–155, 2015. ISSN 0968090X.
- M. Hyland and H. S. Mahmassani. Dynamic autonomous vehicle fleet operations: Optimization-based strategies to assign AVs to immediate traveler demand requests. *Transportation Research Part C: Emerging Technologies*, 92:278–297, 2018. ISSN 0968090X.
- R. Iacobucci, B. McLellan, and T. Tezuka. Optimization of shared autonomous electric vehicles operations with charge scheduling and vehicle-to-grid. *Transportation Research Part C: Emerging Technologies*, 100:34–52, 2019. ISSN 0968090X.
- E. Ismagilova, L. Hughes, Y. K. Dwivedi, and K. R. Raman. Smart cities: Advances in research—an information systems perspective. *International journal of information management*, 47:88–100, 2019.
- P. Jabbari, J. Auld, and D. MacKenzie. How do perceptions of safety and car ownership importance affect autonomous vehicle adoption? *Travel Behaviour and Society*, 28:128–140, 2022.
- T. A. Jenkin, J. Webster, and L. McShane. An agenda for ‘green’information technology and systems research. *Information and organization*, 21(1):17–40, 2011.
- L. Jian, Y. Zheng, and Z. Shao. High efficient valley-filling strategy for centralized coordinated charging of large-scale electric vehicles. *Applied Energy*, 186:46–55, 2017. ISSN 03062619.
- P. Jing, H. Hu, F. Zhan, Y. Chen, and Y. Shi. Agent-based simulation of autonomous vehicles: A systematic literature review. *IEEE Access*, 8:79089–79103, 2020.
- D. Jones, C. Snider, A. Nassehi, J. Yon, and B. Hicks. Characterising the Digital Twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29:36–52, 2020. ISSN 17555817.
- M. Jun and A. Meintz. Workplace Charge Management with Aggregated Building Loads. In *2018 IEEE Transportation and Electrification Conference and Expo, ITEC 2018*, pages 519–524. IEEE, 2018.

- M. Kahlen, K. Schroer, W. Ketter, and A. Gupta. Smart markets for real-time allocation of multiproduct resources: the case of shared electric vehicles. *Information Systems Research*, 35(2):871–889, 2024.
- M. T. Kahlen, W. Ketter, and J. van Dalen. Electric Vehicle Virtual Power Plant Dilemma: Grid Balancing Versus Customer Mobility. *Production and Operations Management*, 27(11):2054–2070, nov 2018a. ISSN 10591478.
- M. T. Kahlen, W. Ketter, and J. van Dalen. Electric Vehicle Virtual Power Plant Dilemma: Grid Balancing Versus Customer Mobility. *Production and Operations Management*, 27(11):2054–2070, 2018b. ISSN 19375956.
- M. A. Kazemi, M. Sedighizadeh, M. J. Mirzaei, and O. Homaei. Optimal siting and sizing of distribution system operator owned EV parking lots. *Applied Energy*, 179:1176–1184, 2016. ISSN 03062619.
- F. Kazhamiaka, C. Rosenberg, and S. Keshav. Tractable lithium-ion storage models for optimizing energy systems. *Energy Informatics*, 2(1), 2019. ISSN 25208942.
- W. Ketter, M. Peters, J. Collins, and A. Gupta. Competitive benchmarking: An is research approach to address wicked problems with big data and analytics. *MIS Quarterly*, 40(4):1057–1080, 2016.
- W. Ketter, J. Collins, M. Saar-Tsechansky, and O. Marom. Information Systems for a Smart Electricity Grid. *ACM Transactions on Management Information Systems*, 9(3):1–22, 2018a. ISSN 2158656X.
- W. Ketter, J. Collins, M. Saar-Tsechansky, and O. Marom. Information systems for a smart electricity grid: Emerging challenges and opportunities. *ACM Transactions on Management Information Systems (TMIS)*, 9(3):1–22, 2018b.
- W. Ketter, K. Schroer, and K. Valogianni. Information systems research for smart sustainable mobility: A framework and call for action. *Information Systems Research*, 34(3):1045–1065, 2023.
- K. Khetarpal, M. Riemer, I. Rish, and n. Precup. Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research*, 75:1401–1476, 2022.
- S. Y. X. Komiak and I. Benbasat. The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS Quarterly*, 30(4):941–960, 2006.
- V. Konda and J. Tsitsiklis. Actor-critic algorithms. *Advances in neural information processing systems*, 12, 1999.
- T. D. Kulkarni, K. Narasimhan, A. Saeedi, and J. Tenenbaum. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. *Advances in neural information processing systems*, 29, 2016.
- N. D. Kullman, M. Cousineau, J. C. Goodson, and J. E. Mendoza. Dynamic ride-hailing with electric vehicles. *Transportation Science*, 56(3):775–794, 2021a.
- N. D. Kullman, J. C. Goodson, and J. E. Mendoza. Electric vehicle routing with public charging stations. *Transportation Science*, 55(3):637–659, 2021b. ISSN 15265447.
- A. Kumar, A. Mehra, and S. Kumar. Why do stores drive online sales? evidence of underlying mechanisms from a multichannel retailer. *Information Systems Research*, 30(1):319–338, 2019.
- S. Kumar, V. Mookerjee, and A. Shubham. Research in operations management and information systems interface. *Production and Operations Management*, 27(11):1893–1905, 2018.
- J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1):50–80, 2004.

- J. H. Lee, D. Chakraborty, S. J. Hardman, and G. Tal. Exploring electric vehicle charging patterns: Mixed usage of charging infrastructure. *Transportation Research Part D: Transport and Environment*, 79: 102249, feb 2020. ISSN 13619209.
- S. Lee and D.-H. Choi. Dynamic pricing and energy management for profit maximization in multiple smart electric vehicle charging stations: A privacy-preserving deep reinforcement learning approach. *Applied Energy*, 304:117754, 2021.
- Z. J. Lee, D. Chang, C. Jin, G. S. Lee, R. Lee, T. Lee, and S. H. Low. Large-scale adaptive electric vehicle charging. *2018 IEEE Global Conference on Signal and Information Processing, GlobalSIP 2018 - Proceedings*, pages 863–864, 2018.
- Z. J. Lee, T. Li, and S. H. Low. ACN-Data: Analysis and applications of an open EV charging dataset. *e-Energy 2019 - Proceedings of the 10th ACM International Conference on Future Energy Systems*, pages 139–149, 2019.
- C. A. Lehmann, C. B. Haubitz, A. Fügener, and U. W. Thonemann. The risk of algorithm transparency: How algorithm complexity drives the effects on the use of advice. *Production and Operations Management*, Advance online publication, 2022.
- S. Leitham, R. W. McQuaid, and J. D. Nelson. The influence of transport on industrial location choice: a stated preference experiment. *Transportation Research Part A*, 34:515–535, 2000.
- M. W. Levin, M. Odell, S. Samarasena, and A. Schwartz. A linear program for optimal integration of shared autonomous vehicles with public transit. *Transportation Research Part C: Emerging Technologies*, 109:267–288, 2019. ISSN 0968090X.
- M. Li, Z. Qin, Y. Jiao, Y. Yang, J. Wang, C. Wang, G. Wu, and J. Ye. Efficient ridesharing order dispatching with mean field multi-agent reinforcement learning. In *The world wide web conference*, pages 983–994, 2019.
- S. Li, F. Xie, Y. Huang, Z. Lin, and C. Liu. Optimizing workplace charging facility deployment and smart charging strategies. *Transportation Research Part D: Transport and Environment*, 87:102481, 2020. ISSN 13619209.
- W. Li and M. Kamargianni. An integrated choice and latent variable model to explore the influence of attitudinal and perceptual factors on shared mobility choices and their value of time estimation. *Transportation Science*, 54(1):62–83, 2020.
- Y. Liang, Z. Ding, T. Ding, and W.-J. Lee. Mobility-aware charging scheduling for shared on-demand electric vehicle fleet using deep reinforcement learning. *IEEE Transactions on Smart Grid*, 12(2): 1380–1393, 2020.
- T. Lillicrap. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- S. Limmer. Dynamic pricing for electric vehicle charging—a literature review. *Energies*, 12(18):3574, 2019.
- C.-A. Lin, K. Shang, and P. Sun. Wait time-based pricing for queues with customer-chosen service times. *Management Science*, 69(4):2127–2146, 2023.
- M. L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.

- N. Liu, S. R. Finkelstein, M. E. Kruk, and D. Rosenthal. When waiting to see a doctor is less irritating: Understanding patient preferences and choice behavior in appointment scheduling. *Management Science*, 64(5):1975–1996, 2018.
- B. Loeb and K. M. Kockelman. Fleet performance and cost evaluation of a shared autonomous electric vehicle (saev) fleet: A case study for austin, texas. *Transportation Research Part A: Policy and Practice*, 121:374–385, 2019.
- B. Loeb, K. M. Kockelman, and J. Liu. Shared autonomous electric vehicle (SAEV) operations across the Austin, Texas network with charging infrastructure decisions. *Transportation Research Part C: Emerging Technologies*, 89:222–233, 2018. ISSN 0968090X.
- M. Lokhandwala and H. Cai. Dynamic ride sharing using traditional taxis and shared autonomous taxis: A case study of NYC. *Transportation Research Part C: Emerging Technologies*, 97:45–60, 2018. ISSN 0968090X.
- M. Lokhandwala and H. Cai. Siting charging stations for electric vehicle adoption in shared autonomous fleets. *Transportation Research Part D: Transport and Environment*, 80:102231, 2020. ISSN 13619209.
- C. Lu, J. Wu, J. Cui, Y. Xu, C. Wu, and M. C. Gonzalez. Deadline differentiated dynamic ev charging price menu design. *IEEE Transactions on Smart Grid*, 14(1):502–516, 2022.
- C. Lu, J. Wu, J. Cui, Y. Xu, C. Wu, and M. C. Gonzalez. Deadline differentiated dynamic ev charging price menu design. *IEEE Transactions on Smart Grid*, 14:502–516, 1 2023. ISSN 19493061.
- M. Lu, Z. Chen, and S. Shen. Optimizing the Profitability and Quality of Service in Carshare Systems Under Demand Uncertainty. *Manufacturing & Service Operations Management*, 20(2):162–180, May 2018. URL <http://pubsonline.informs.org/removed/10.1287/msom.2017.0644>.
- C. Luo, Y.-F. Huang, and V. Gupta. Stochastic dynamic pricing for ev charging stations with renewable integration and energy storage. *IEEE Transactions on Smart Grid*, 9(2):1494–1505, 2017.
- Lyft. Form s-1. www.sec.gov/Archives/edgar/data/1759509/000119312519059849/d633517ds1.htm, 2019.
- T.-Y. Ma and S. Xie. Optimal fast charging station locations for electric ridesharing with vehicle-charging station assignment. *Transportation Research Part D: Transport and Environment*, 90:102682, 2021.
- A. Mahdavian, A. Shojaei, S. McCormick, T. Papandreou, N. Eluru, and A. A. Oloufa. Drivers and barriers to implementation of connected, automated, shared, and electric vehicles: An agenda for future research. *IEEE Access*, 9:22195–22213, 2021.
- H. S. Mahmassani. 50th Anniversary invited article autonomous vehicles and connected vehicle systems: Flow and operations considerations. *Transportation Science*, 50(4):1140–1162, 2016. ISSN 15265447.
- A. Malhotra, N. P. Melville, and R. T. Watson. Spurring impactful research on information systems for environmental sustainability. *MIS Quarterly*, 37(4):1265–1274, 2013.
- S. Manna, T. D. Loeffler, R. Batra, S. Banik, H. Chan, B. Varughese, K. Sasikumar, M. Sternberg, T. Peterka, M. J. Cherukara, et al. Learning in continuous action space for developing high dimensional potential energy models. *Nature communications*, 13(1):368, 2022.
- T. Mao, W.-h. Lau, C. Shum, H. S.-H. Chung, K.-F. Tsang, and N. C.-F. Tse. A regulation policy of ev discharging price for demand scheduling. *IEEE Transactions on Power Systems*, 33(2):1275–1288, 2017.

- D. McFadden. Conditional logit analysis of qualitative choice behavior. 1972.
- N. Melville. Information Systems Innovation for Environmental Sustainability. *MIS Quarterly*, 34(1): 1–21, jun 2010a.
- N. P. Melville. Information Systems Innovation for Environmental Sustainability. *MIS Quarterly*, 34(1): 1–21, 2010b.
- G. Meyer, G. Adomavicius, P. E. Johnson, M. Elidrisi, W. A. Rush, J. M. Sperl-Hillen, and P. J. O’Connor. A machine learning approach to improving dynamic decision making. *Information Systems Research*, 25(2):239–263, 2014.
- J. H. Miller and S. E. Page. *Complex adaptive systems: an introduction to computational models of social life: an introduction to computational models of social life*. Princeton university press, 2009.
- S. Mithas, M. S. Krishnan, and C. Fornell. Research note—information technology, customer satisfaction, and profit: Theory and evidence. *Information Systems Research*, 27(1):166–181, 2016.
- V. Mnih. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. ISSN 14764687.
- A. Moradipari and M. Alizadeh. Pricing and routing mechanisms for differentiated services in an electric vehicle public charging station network. *IEEE Transactions on smart grid*, 11(2):1489–1499, 2019.
- J. C. Mukherjee and A. Gupta. A Review of Charge Scheduling of Electric Vehicles in Smart Grid. *IEEE Systems Journal*, 9(4):1541–1553, 2015. ISSN 19379234.
- E. A. Nanaki and C. J. Koroneos. Climate change mitigation and deployment of electric vehicles in urban areas. *Renewable energy*, 99:1153–1160, 2016.
- S. Narayanan, E. Chaniotakis, and C. Antoniou. Shared autonomous vehicle services: A comprehensive review. *Transportation Research Part C: Emerging Technologies*, 111:255–293, 2020.
- U. Nations. World urbanization prospects: The 2018 revision, 2018. URL <https://population.un.org/wup/>.
- Neon Neue Energieökonomik and Technical University of Berlin and ETH Zürich and DIW Berlin. Open power system data, 2024. URL <https://open-power-system-data.org/>.
- P. Nunes, R. Figueiredo, and M. C. Brito. The use of parking lots to solar-charge electric vehicles. *Renewable and Sustainable Energy Reviews*, 66:679–693, 2016. ISSN 18790690.
- G. Oestreicher-Singer and L. Zalmanson. Content or community? a digital business strategy for content providers in the social age. *MIS Quarterly*, pages 591–616, 2013.
- Open charge map. <https://www.openchargemap.org/>, 2020.
- A. Oroojlooyjadid, M. Nazari, L. V. Snyder, and M. Takáč. A deep q-network for the beer game: Deep reinforcement learning for inventory optimization. *Manufacturing & Service Operations Management*, 24(1):285–304, 2022.
- A. Osterwalder, Y. Pigneur, and C. Tucci. Clarifying business models: origins, present, and future of the concept. *Communications of the Association for Information Systems*, 15(May):1–43, 2005.
- P. Pal, P. Gopal, and M. Ramkumar. Impact of transportation on climate change: An ecological modernization theoretical perspective. *Transport Policy*, 130:167–183, 2023.

- T. P. Pantelidis, L. Li, T.-Y. Ma, J. Y. Chow, and S. E. G. Jabari. A node-charge graph-based online carshare rebalancing policy with capacitated electric charging. *Transportation Science*, 56(3):654–676, 2022.
- M. Panzer and B. Bender. Deep reinforcement learning in production systems: a systematic literature review. *International Journal of Production Research*, 60(13):4316–4341, 2022. ISSN 1366588X.
- A. Papu Carrone, V. M. Hoening, A. F. Jensen, S. E. Mabit, and J. Rich. Understanding car sharing preferences and mode substitution patterns: A stated preference experiment. *Transport Policy*, 98:139–147, 2020.
- E. Peer, L. Brandimarte, S. Samat, and A. Acquisti. Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163, 2017.
- D. Pevec, J. Babic, A. Carvalho, Y. Ghiassi-Farrokhfal, W. Ketter, and V. Podobnik. A survey-based assessment of how existing and potential electric vehicle owners perceive range anxiety. *Journal of cleaner Production*, 276:122779, 2020.
- W. B. Powell. Energy and uncertainty: Models and algorithms for complex energy systems. *AI Magazine*, 35(3):8–21, 2014. ISSN 07384602.
- F. Provost and T. Fawcett. Data science and its relationship to big data and data-driven decision making. *Big data*, 1(1):51–59, 2013.
- Q. Qi and F. Tao. Digital twin and big data towards smart manufacturing and industry 4.0: 360-degree comparison. *IEEE Access*, 6:3585–3593, 2018.
- W. Qi and Z.-J. M. Shen. A Smart-City Scope of Operations Management. *Production and Operations Management*, 0(0):1–14, 2018. ISSN 10591478.
- W. Qi, M. Sha, and S. Li. When shared autonomous electric vehicles meet microgrids: Citywide energy-mobility orchestration. *Manufacturing & Service Operations Management*, 24:2389–2406, 9 2022. ISSN 1523-4614.
- Z. Qin, X. Tang, Y. Jiao, F. Zhang, Z. Xu, H. Zhu, and J. Ye. Ride-hailing order dispatching at didi via reinforcement learning. *INFORMS Journal on Applied Analytics*, 50(5):272–286, 2020.
- M. Rahman, M. Rashed, and K. Alam. Challenges in implementing smart mobility solutions in smart cities. *Smart Cities Journal*, 4(3):459–477, 2021.
- A. Rai. Diversity of Design Science Research. *MIS Quarterly*, 41(1):iii–xviii, 2017.
- A. Rai, P. Constantinides, and S. Sarker. Next generation digital platforms : toward human-ai hybrids. *Management Information Systems Quarterly*, 43, 2019.
- Y. Ren and R. E. Kraut. Agent-based modeling to inform online community design: Impact of topical breadth, message volume, and discussion moderation on member commitment and contribution. *Human-Computer Interaction*, 29(4):351–389, 2014.
- K. S. Rhee, J. Zheng, Y. Wang, and Y. Tan. Value of information sharing via ride-hailing apps: An empirical analysis. *Information Systems Research*, 2022.
- S. A. Rijdsdijk and E. J. Hultink. “honey, have you seen our hamster?” consumer evaluations of autonomous domestic products. *Journal of Production and Innovation Management*, 20:204–216, 2003.

- H. Ritchie, V. Samborska, and M. Roser. The world population is moving to cities. why is urbanization happening and what are the consequences? Our World in Data, 2024. URL <https://ourworldindata.org/urbanization>. Accessed: January 30, 2025.
- M. S. Roni, Z. Yi, and J. G. Smart. Optimal charging management and infrastructure planning for free-floating shared electric vehicles. *Transportation Research Part D: Transport and Environment*, 76:155–175, nov 2019. ISSN 13619209.
- F. Rossi, R. Zhang, Y. Hindy, and M. Pavone. Routing autonomous vehicles in congested transportation networks: structural properties and coordination algorithms. *Autonomous Robots*, 42(7):1427–1442, 2018. ISSN 15737527.
- P. J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20(C):53–65, 1987. ISSN 03770427.
- R. G. Sargent. Verification and validation of simulation models. In *Proceedings of the 2010 winter simulation conference*, pages 166–183. IEEE, 2010.
- S. Sarker, S. Chatterjee, X. Xiao, and A. Elbanna. The Sociotechnical Axis of Cohesion for the IS Discipline: Its Historical Legacy and its Continued Relevance. *MIS Quarterly*, 43(3):695–719, jan 2019a.
- S. Sarker, S. Chatterjee, X. Xiao, and A. Elbanna. The sociotechnical axis of cohesion for the is discipline: Its historical legacy and its continued relevance. *MIS quarterly*, 43(3):695–720, 2019b.
- B. Schleich, N. Anwer, L. Mathieu, and S. Wartzack. Shaping the digital twin for design and production engineering. *CIRP Annals*, 66(1):141–144, 2017. ISSN 00078506.
- C. Schlereth and B. Skiera. Two new features in discrete choice experiments to improve willingness-to-pay estimation that result in sdr and sadr: Separated (adaptive) dual response. *Management Science*, 63(3):829–842, 2017.
- K. Schroer, W. Ketter, T. Y. Lee, A. Gupta, and M. Kahlen. Data-driven competitor-aware positioning in on-demand vehicle rental networks. *Transportation Science*, 56(1):182–200, 2022.
- S. Seidel, J. Recker, and J. v. Brocke. Sensemaking and Sustainable Practicing: Functional Affordances of Information Systems in Green Transformations. *MIS Quarterly*, 37:1275–1299, Dec. 2013a.
- S. Seidel, J. Recker, and J. vom Brocke. Sensemaking and Sustainable Practicing: Functional Affordances of Information Systems in Green Transformations. *MIS Quarterly*, 37(4):1275–1299, 2013b. ISSN 02767783.
- P. Setia and P. C. Patel. How information systems help create om capabilities: Consequents and antecedents of operational absorptive capacity. *Journal of operations management*, 31(6):409–431, 2013.
- S. Shaheen, A. Cohen, I. Zohdy, et al. Shared mobility: current practices and guiding principles. Technical report, United States. Federal Highway Administration, 2016.
- S. Shaheen, A. Cohen, N. Chan, and A. Bansal. Sharing strategies: carsharing, shared micromobility (bikesharing and scooter sharing), transportation network companies, microtransit, and other innovative mobility modes. In *Transportation, land use, and environmental planning*, pages 237–262. Elsevier, 2020.
- A. Shariff, J. F. Bonnefon, and I. Rahwan. Psychological roadblocks to the adoption of self-driving vehicles. *Nature Human Behaviour*, 1(10):694–696, 2017.

- P. Sharifi, A. Banerjee, and M. J. Feizollahi. Leveraging owners' flexibility in smart charge/discharge scheduling of electric vehicles to support renewable energy integration. *Computers and Industrial Engineering*, 149(July):106762, 2020. ISSN 03608352.
- T. L. Sheldon and R. Dua. Consumer preferences for ride-hailing: Barriers to an autonomous, shared, and electric future. *Journal of Cleaner Production*, 434:140251, 2024.
- J. Shi, Y. Gao, W. Wang, N. Yu, and P. A. Ioannou. Operating Electric Vehicle Fleet for Ride-Hailing Services with Reinforcement Learning. *IEEE Transactions on Intelligent Transportation Systems*, 21(11):4822–4834, 2020. ISSN 15580016.
- J. F. Shortle, J. M. Thompson, D. Gross, and C. M. Harris. *Fundamentals of Queueing Theory*. John Wiley & Sons, Apr. 2018.
- Z. Shou and X. Di. Reward design for driver repositioning using multi-agent reinforcement learning. *Transportation research part C: emerging technologies*, 119:102738, 2020.
- A. Siddiq and T. A. Taylor. Ride-hailing platforms: Competition and autonomous vehicles. *Manufacturing & Service Operations Management*, 24(3):1511–1528, 2022.
- S. Sierla, V. Kyrki, P. Aarnio, and V. Vyatkin. Automatic assembly planning based on digital product descriptions. *Computers in Industry*, 97:34–46, 2018. ISSN 01663615.
- D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. ISSN 14764687.
- M. D. Simoni, K. M. Kockelman, K. M. Gurusurthy, and J. Bischoff. Congestion pricing in a world of self-driving vehicles: An analysis of different strategies in alternative future scenarios. *Transportation Research Part C: Emerging Technologies*, 98:167–185, 2019.
- J. Soares, M. A. F. Ghazvini, N. Borges, and Z. Vale. Dynamic electricity pricing for electric vehicles using stochastic programming. *Energy*, 122:111–127, 2017.
- D. Sperling. *Three revolutions: steering automated, shared, and electric vehicles to a better future*. Island Press, Washington, D.C., USA, 2018a.
- D. Sperling. *Three revolutions: Steering automated, shared, and electric vehicles to a better future*. Island Press, 2018b.
- R. Sutton. *The Bitter Lesson*, 2019.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018a.
- R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018b.
- R. S. Sutton, r. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- T. M. Sweda, I. S. Dolinskaya, and D. Klabjan. Optimal recharging policies for electric vehicles. *Transportation Science*, 51(2):457–459, 2017. ISSN 15265447.
- A. Talebian and S. Mishra. Predicting the adoption of connected autonomous vehicles: A new approach based on the theory of diffusion of innovations. *Transportation Research Part C: Emerging Technologies*, 95:363–380, 2018.

- F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, and F. Sui. Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology*, 94(9-12):3563–3576, feb 2018. ISSN 0268-3768.
- Taxi costs in Berlin. <http://taxihowmuch.com/location/berlin-de>, 2022.
- T. A. Taylor. On-demand service platforms. *Manufacturing & Service Operations Management*, 20(4):704–720, 2018.
- M. Tran, D. Banister, J. D. Bishop, and M. D. McCulloch. Realizing the electric-vehicle revolution. *Nature climate change*, 2(5):328–333, 2012.
- Transportation & Environment. How clean are electric cars? <https://www.transportenvironment.org/discover/how-clean-are-electric-cars>, 2022.
- Uber. Form s-1. www.sec.gov/Archives/edgar/data/1543151/000119312519103850/d647752ds1.htm, 2019.
- K. Valogianni, W. Ketter, and J. Collins. A multiagent approach to variable-rate electric vehicle charging coordination. In *Proceedings of the 2015 international conference on autonomous agents and multiagent systems*, pages 1131–1139, 2015.
- K. Valogianni, W. Ketter, J. Collins, and D. Zhdanov. Sustainable Electric Vehicle Charging using Adaptive Pricing. *Production and Operations Management*, 29(6):1550–1572, 2020a. ISSN 19375956.
- K. Valogianni, W. Ketter, J. Collins, and D. Zhdanov. Sustainable electric vehicle charging using adaptive pricing. *Production and Operations Management*, 29(6):1550–1572, 2020b.
- K. Valogianni, W. Ketter, J. Collins, and G. Adomavicius. Toward sustainable electricity markets: Capacity-based pricing for electric vehicle smart charging. *Available at SSRN 5054307*, 2024.
- H. van der Valk, H. Haße, F. Möller, M. Arbter, J. L. Henning, and B. Otto. A Taxonomy of Digital Twins. *26th Americas Conference on Information Systems, AMCIS 2020*, pages 1–10, 2020.
- H. Van Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- L. van Hezewijk, N. Dellaert, T. Van Woensel, and N. Gademann. Using the proximal policy optimisation algorithm for solving the stochastic capacitated lot sizing problem. *International Journal of Production Research*, 2022. ISSN 1366588X. URL https://no_doi.org/10.1080/00207543.2022.2056540.
- V. Venkatesh and F. D. Davis. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, 46:186–204, 2000.
- V. Venkatesh, J. Thong, and X. Xu. Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *MIS Quarterly*, 36:157–178, 03 2012.
- J. vom Brocke, R. T. Watson, C. Dwyer, S. Elliot, and N. Melville. Green Information Systems: Directives for the IS Discipline. *Communications of the Association for Information Systems*, 33:509–520, 2013. ISSN 15293181.
- K. W. Axhausen, A. Horni, and K. Nagel, editors. *The Multi-Agent Transport Simulation MATSim*. Ubiquity Press, 2016.
- B. Wali, P. Santi, and C. Ratti. Modeling consumer affinity towards adopting partially and fully automated vehicles – the role of preference heterogeneity at different geographic levels. *Transportation Research Part C: Emerging Technologies*, 129:103276, 2021.

- E. Wang, R. Ding, Z. Yang, H. Jin, C. Miao, L. Su, F. Zhang, C. Qiao, and X. Wang. Joint charging and relocation recommendation for e-taxi drivers via multi-agent mean field hierarchical reinforcement learning. *IEEE Transactions on Mobile Computing*, 2020.
- W. Wang, L. Qiu, D. Kim, and I. Benbasat. Effects of rational and social appeals of online recommendation agents on cognition- and affect-based trust. *Decision Support Systems*, 86:48–60, 2016.
- W. Wang, B. Li, X. Luo, and X. Wang. Deep reinforcement learning for sequential targeting. *Management Science*, 69(9):5439–5460, 2023a.
- X. Wang, Y. Li, and W. Zhang. Artificial intelligence in smart mobility: Opportunities and challenges. *Transportation Research Part C: Emerging Technologies*, 151:104129, 2023b.
- R. T. Watson, M.-C. Boudreau, and A. J. Chen. Information Systems and Environmentally Sustainable Development: Energy Informatics and New Directions for the IS Community. *MIS Quarterly*, 24(4):665–694, 2010. ISSN 02767783.
- K. Wei, V. Vaze, and A. Jacquillat. Transit planning optimization under ride-hailing competition and traffic congestion. *Transportation Science*, 56(3):725–749, 2022.
- C. Whittle, L. Whitmarsh, P. Haggard, P. Morgan, and G. Parkhurst. User decision-making in transitions to electrified, autonomous, shared or reduced mobility. *Transportation Research Part D: Transport and Environment*, 71:302–319, 2019.
- World Health Organization. GLOBAL STATUS REPORT ON ROAD SAFETY. Technical Report 4, World Health Organization, 2018.
- D. Wu, H. Zeng, C. Lu, and B. Boulet. Two-Stage Energy Management for Office Buildings with Workplace EV Charging and Renewable Energy. *IEEE Transactions on Transportation Electrification*, 3(1):225–237, 2017. ISSN 23327782.
- O. Q. Wu, Şafak Yücel, and Y. H. Zhou. Smart charging of electric vehicles: An innovative business model for utility firms. *Manufacturing & Service Operations Management*, 24:2481–2499, 9 2022. ISSN 1523-4614.
- J. Xie, Y. Liu, and N. Chen. Two-Sided Deep Reinforcement Learning for Dynamic Mobility-on-Demand Management with Mixed Autonomy. *Transportation Science*, 57(4):1019–1046, jul 2023. ISSN 0041-1655.
- R. E. Yagoda and D. J. Gillan. You want me to trust a robot? the development of a human–robot interaction trust scale. *International Journal of Social Robotics*, 4(3):235–248, 2012.
- M. Yang, L. Zhang, Z. Zhao, and L. Wang. Comprehensive benefits analysis of electric vehicle charging station integrated photovoltaic and energy storage. *Journal of Cleaner Production*, 302:126967, 2021.
- Y. Yang, R. Luo, M. Li, M. Zhou, W. Zhang, and J. Wang. Mean field multi-agent reinforcement learning. In *International conference on machine learning*, pages 5571–5580. PMLR, 2018.
- F. Yao, J. Zhu, J. Yu, C. Chen, and X. M. Chen. Hybrid operations of human driving vehicles and automated vehicles with data-driven agent-based simulation. *Transportation Research Part D: Transport and Environment*, 86:102469, 2020.
- N. Yaraghi, A. Y. Du, R. Sharman, R. D. Gopal, and R. Ramesh. Health information exchange as a multi-sided platform: Adoption, usage, and practice involvement in service co-production. *Information Systems Research*, 26(1):1–18, 2015.

- Y. Yoo, O. Henfridsson, and K. Lyytinen. Research commentary: The new organizing logic of digital innovation: An agenda for information systems research. *Information Systems Research*, 21(4):724–735, 2010a.
- Y. Yoo, O. Henfridsson, and K. Lyytinen. Research commentary—the new organizing logic of digital innovation: an agenda for information systems research. *Information systems research*, 21(4):724–735, 2010b.
- E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE access*, 8:58443–58469, 2020.
- G. G. Zanvettor, M. Casini, R. S. Smith, and A. Vicino. Stochastic energy pricing of an electric vehicle parking lot. *IEEE Transactions on smart grid*, 13(4):3069–3081, 2022.
- T. Zeng, S. Bae, B. Travacca, and S. Moura. Inducing human behavior to maximize operation performance at pev charging station. *IEEE Transactions on Smart Grid*, 12(4):3353–3363, 2021.
- H. Zhang, X. Zhao, X. Fang, and B. Chen. Proactive resource request for disaster response: A deep learning-based optimization model. *Information Systems Research*, 35(2):528–550, 2024.
- J. Zhang, G. Adomavicius, A. Gupta, and W. Ketter. Consumption and Performance: Understanding Longitudinal Dynamics of Recommender Systems via an Agent-Based Simulation Framework. *Information Systems Research*, 31(1):76–101, Mar. 2020a.
- J. Zhang, Y. Chen, and Q. Li. The impact of shared autonomous vehicles on urban mobility: A simulation-based study. *Transportation Research Part A: Policy and Practice*, 163:302–316, 2022.
- T. Z. Zhang and T. D. Chen. Smart charging management for shared autonomous electric vehicle fleets: A Puget Sound case study. *Transportation Research Part D: Transport and Environment*, 78:102184, 2020. ISSN 13619209.
- Y. Zhang, B. Li, and R. Krishnan. Learning individual behavior using sensor data: The case of global positioning system traces and taxi drivers. *Information Systems Research*, 31(4):1301–1321, 2020b.
- W. Zhao, J. P. Queralta, and T. Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE symposium series on computational intelligence (SSCI)*, pages 737–744. IEEE, 2020.
- Z. Zhao and C. K. Lee. Dynamic pricing for ev charging stations: A deep reinforcement learning approach. *IEEE Transactions on Transportation Electrification*, 8(2):2456–2468, 2021.
- Y. Zheng, S. Rajasegarar, and C. Leckie. Smart traffic management systems: A survey of models and techniques. *Journal of Big Data Analytics in Transportation*, 2(1):1–22, 2020.
- Z. Zhu, S. Lambbotharan, W. H. Chin, and Z. Fan. A mean field game theoretic approach to electric vehicles charging. *IEEE Access*, 4:3501–3510, 2016.
- C. Zhuge and C. Wang. Integrated modelling of autonomous electric vehicle diffusion: From review to conceptual design. *Transportation Research Part D: Transport and Environment*, 91, 2021.
- F. Zwick, N. Kuehnel, and S. Hörl. Shifts in perspective: Operational aspects in (non-) autonomous ride-pooling simulations. *Transportation Research Part A: Policy and Practice*, 165:300–320, 2022.
- P. Ówiakowski, M. Giergiczny, and M. Krawczk. Pirates in the lab: Using incentivized choice experiments to explore preference for (un)authorized content. *MIS Quarterly*, 40(3):709–716, 2016.

Eidesstattliche Erklärung nach § 8 Abs. 3 der Promotionsordnung vom 17.02.2015

Hiermit versichere ich an Eides Statt, dass ich die vorgelegte Arbeit selbstständig und ohne die Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe. Die aus anderen Quellen direkt oder indirekt übernommenen Aussagen, Daten und Konzepte sind unter Angabe der Quelle gekennzeichnet.

Weitere Personen, neben den ggf. in der Einleitung der Arbeit aufgeführten Koautorinnen und Koautoren, waren an der inhaltlich-materiellen Erstellung der vorliegenden Arbeit nicht beteiligt. Insbesondere habe ich hierfür nicht die entgeltliche Hilfe von Vermittlungs- bzw. Beratungsdiensten in Anspruch genommen. Niemand hat von mir unmittelbar oder mittelbar geldwerte Leistungen für Arbeiten erhalten, die im Zusammenhang mit dem Inhalt der vorgelegten Dissertation stehen.

Die Arbeit wurde bisher weder im In- noch im Ausland in gleicher oder ähnlicher Form einer anderen Prüfungsbehörde vorgelegt.

Ich versichere, dass ich nach bestem Wissen die reine Wahrheit gesagt und nichts verschwiegen habe.

Ich versichere, dass die eingereichte elektronische Fassung der eingereichten Druckfassung vollständig entspricht.

Die Strafbarkeit einer falschen eidesstattlichen Versicherung ist mir bekannt, namentlich die Strafandrohung gemäß § 156 StGB bis zu drei Jahren Freiheitsstrafe oder Geldstrafe bei vorsätzlicher Begehung der Tat bzw. gemäß § 161 Abs. 1 StGB bis zu einem Jahr Freiheitsstrafe oder Geldstrafe bei fahrlässiger Begehung.

Köln, den 04. März, 2025

A handwritten signature in blue ink, appearing to read 'Ramin Ahadi', is written over a light blue horizontal line. The signature is stylized and cursive.

Ramin Ahadi

About the Author



Ramin Ahadi, born in 1944 in Ghochan, Iran, is a Ph.D. candidate in Operations Management (OM) and Information Systems (IS) at the University of Cologne and a Visiting Ph.D. Scholar at IE University Business School. He started his Ph.D. in October 2019 under the supervision of Prof. Dr. Wolfgang Ketter, with funding from the Institute of Energy Economics at the University of Cologne (EWI). He has been working at the Chair of Information Systems for Sustainable Society at the University of Cologne throughout his doctoral studies.

Ramin holds a Master's degree in Industrial Engineering and a Bachelor's degree in Mechanical Engineering from the Ferdowsi University of Mashhad. Prior to embarking on his academic career, he gained extensive teaching and research experience as a research assistant at the Ferdowsi University of Mashhad.

His research focuses on applying data-driven and machine learning methods to understand socio-technical systems and enhance efficiency through advanced decision-making models, primarily in the domains of sustainable energy and mobility. In his work, he leverages agent-based modeling, machine learning, mathematical programming, and reinforcement learning to design and develop mathematical decision-making algorithms for planning and managing digital energy and mobility systems. These algorithms aim to balance objectives such as environmental impact, cost, and service levels. Additionally, his research seeks to understand and predict the impact of user behaviors on the adoption of new technologies, such as autonomous vehicles.

Ramin's research has been presented at leading OM, IS, and computer science conferences and is published in highly ranked journals. His notable contributions include developing cooperative routing and scheduling models for the charging management of shared autonomous electric vehicles through hierarchical multi-agent reinforcement learning, analyzing the impacts of user behavior in electric multi-modal vehicle sharing platforms, and managing large-scale electric vehicle charging hubs through AI-enabled pricing models.

Beyond his academic activities, Ramin enjoys advancing his calisthenics skills, hiking, and playing video games. He is also keen on staying updated with emerging technologies and is currently exploring deep learning architectures, including transformer models and large language models.

Research Portfolio

Publications in Peer-Reviewed Academic Journals

- **Ahadi, R.**, Ketter, W., Collins, J., & Daina, N. (2023). *Cooperative Learning for Smart Charging of Shared Autonomous Vehicle Fleets*. *Transportation Science* 57(3):613-630. <https://doi.org/10.1287/trsc.2022.1187>
- Khalilzadeh, M., Neghabi, H., & **Ahadi, R.** (2021). *An application of approximate dynamic programming in multi-period multi-product advertising budgeting*. *Journal of Industrial & Management Optimization*, 19(1). <https://doi.org/10.3934/jimo.2021202>

Working Papers

- Schroer, K., **Ahadi, R.**, Lee, T. Y., & Ketter, W. (2024). Data-driven Planning of Large-Scale Electric Vehicle Charging Hubs using Deep Reinforcement Learning. (2nd-round)
- **Ahadi, R.**, Taudien, A., Ketter, W., & Gupta, A. (2024). Adoption of Autonomous Vehicles in Ride-Hailing Services: The Role of User Preferences. (1st-round)
- **Ahadi, R.**, Valogianni, K., Schroer, & K., Ketter., W. (x 2025). Dynamic Capacity-Based Pricing of Charging Hubs with Supply Constraints. (1st-round)

Publications & Presentations in Peer-Reviewed Academic Conference Proceedings

- Sund, L., Muires, J., **Ahadi, R.**, & Ketter. (2024). A Charging Capacity Market for Competing Fleets. Workshop on Information Technologies and Systems (WITS).
- **Ahadi, R.**, Schroer, K., & Ketter, W. (2024). Managing Electric Vehicle Charging Hubs Through Dynamic Capacity-Based Pricing. European Conference on Information Systems.
- **Ahadi, R.**, Taudian, A., & Ketter, W. (2023). Human versus automated agents: how user preferences affect future mobility systems. European Conference on Information Systems.
- Demircan, M., **Ahadi, R.**, & Ketter, W. (2022). Sustainability vs. Price: Analysis of Electric Multi-Modal Vehicle Sharing Systems under Substitution Effects. European Conference on Information Systems.
- Schroer, K., **Ahadi, R.**, Lee, Y.T., Ketter, W. (2021). *Preference-aware Planning and Operations of Electric Vehicle Charging Clusters: A Data-Driven Prescriptive Framework*. In Proceedings of the SIG GREEN Workshop 2021.
- Schroer, K., **Ahadi, R.**, Lee, Y.T., Ketter, W. (2021). *Preference-Aware Planning and Operations of Electric Vehicle Charging Clusters: A Prescriptive Framework*. In Workshop on Information Systems and Technology (WITS) 2021.
- **Ahadi, R.**, Ketter, W., Collins, J., & Daina, N. (2021). Siting and Sizing of Charging Infrastructure for Shared Autonomous Electric Fleets. In Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems.
- Mozafari, M., Rezaee, B., & **Ahadi, R.** (2019). Optimal Charging Planning for Electric-Vehicle in highways. 12th International Conference of Operations Research Society, Mazandaran University of science and technology 2019.
- Bashtani, F., Rezaee, B., & **Ahadi, R.** (2019). Charge scheduling of electric vehicles for smart parking lot with consideration the satisfaction of car owners. 12th International Conference of Operations Research Society, Mazandaran University of science and technology 2019.
- Yazdi, L., **Ahadi, R.**, & Rezaee, B. (2018). Optimal electric vehicle charging station placing with integration of renewable energy. 15th Iran International Industrial Engineering Conference, Yazd University 2018.
- **Ahadi, R.**, & Rezaee, B. (2018). Charge scheduling and planning for electric vehicles in highways. 4th International conference on industrial and system engineering, Ferdowsi University of Mashhad 2018.

Teaching Portfolio

Lecturing & Tutoring

As a Ph.D. candidate at the University of Cologne, I had the privilege to gain extensive experience in (co-)lecturing and tutoring at the Bachelor's, Master's, and MBA levels. I contributed to the development of a new version of an advanced machine learning curriculum designed for information systems and computer science students. This involved strategic curriculum planning, the creation of lecture content, and the development of entirely new programming tutorials. My teaching responsibilities spanned various topics, and I was actively involved in leading weekly programming tutorials (in Python), preparing lecture materials, co-lecturing, and organizing, conducting, and assessing examinations. A detailed list of my lecturing activities is provided below.

Bachelor's

Bachelorseminar Next Generation Information Systems – Seminar Tutor & Organizer

- I lead the kick-off session for the Bachelor's seminar "Next Generation Information Systems," where I teach students the fundamentals of writing scientific essays, with a specific focus on conducting literature reviews. During this session, I also introduce the chair, present the seminar topics, and handle organizational tasks, including topic integration, student allocation, and facilitating student presentation sessions.
- In addition, I supervise several students throughout the seminar, providing guidance on their chosen topics, assisting them in structuring and writing their essays, and evaluating both their written submissions and presentations.

Master's

Advanced Analytics and Applications – Co-Lecturer and Tutor

- I co-lectured and tutored for the Master's and Ph.D. course "Advanced Analytics and Applications." The curriculum, which I co-developed, spans the entire advanced machine learning life cycle, including classification (e.g., kernel SVM, & deep learning), soft clustering (e.g.,

Gaussian Mixture Models), Markov decision process, and reinforcement learning algorithms, model building, and evaluation.

- In the tutorials, I teach data processing, data analysis, and advanced machine learning techniques using Python, emphasizing practical applications of theoretical concepts. I also prepare and teach numerical examples that complement the theoretical parts of the lecture. Furthermore, I mentor student groups on team assignments that involve solving real-world data science and prediction problems, typically based on large-scale mobility or energy datasets.

MBA

Elective on Operations Management for Future Mobility Systems – Guest Lecturer
[Cologne-Rotterdam EMBA]

- As a Guest Lecturer in the annual Cologne-Rotterdam Executive MBA elective on the Future of Mobility, I delivered two-hour interactive lecture sessions designed to explore core technology trends and their implications for data-driven operations management. These sessions were entirely self-developed and included real-world case studies from my research, such as cooperative charging management of shared autonomous electric vehicles. Additionally, I incorporated insights from the latest research in Operations Management (OM) and Information Systems (IS) to illustrate how technology can address critical societal challenges.
- To enhance student engagement, I developed and moderated in-class group case studies. These activities required students to analyze scenarios, prepare solutions, and present their findings for discussion, fostering a collaborative and practical learning environment.

Thesis Supervision (Total Count: 16)

In addition to teaching, I took primary responsibility for supervising and mentoring 15 Master's, MBA, and Bachelor's theses. These projects were closely aligned with my research interests in data-driven and machine learning-based decision-making, as well as the empirical analysis of emerging energy and mobility systems. My supervision responsibilities included guiding students in formulating their research questions, providing relevant datasets, conducting regular feedback and discussion sessions, and evaluating the final theses. A comprehensive list of supervised thesis titles is provided below.

Master's & MBA Theses (Count: 6)

- *Analysis of the current EV charging infrastructure and optimal placement of new EV charging station* (2025)
- *Rebalancing Multi-Modal Vehicle Sharing Fleets using Centralized Reinforcement Learning* (2024)
- *Dynamic pricing policies for multi-modal shared mobility fleets* (2023)
- *Dynamic Decision Making for Free-floating Shared Electric Fleets to Enable Demand-side Rebalancing: Applications of Reinforcement Learning for Pricing and Vehicle Assignment Policies* (2022)
- *Evaluating the impact of strategic decisions on a multi-mode shared autonomous electric fleet using agent-based modeling* (2021)
- *Siting and Sizing of Shared Charging Infrastructure for Shared Electric Vehicle Fleets* (2021)

Bachelor's Theses (Count: 10)

- *A review of the use of reinforcement learning as a new methodology in information systems* (2024)
- *A Review of User Behavior Analysis in Hybrid Autonomous Shared Mobility Platforms* (2024)
- *Evaluating the impact of shared charging facilities on the performance of electric shared mobility platforms* (2024)
- *Analysis of On-Demand Multimodal Shared Mobility Platforms: an Application of Demand Estimation and Pattern Recognition* (2023)
- *A Review of Smart Charging Management and Pricing of Electric Vehicle Charging Facilities* (2023)
- *Charging Load Prediction of Shared Electric Vehicles: A Data-driven Approach* (2022)
- *A literature review of users' patterns, selection behavior and impactful factors on demand for the shared electric urban mobility* (2021)
- *Recommending Charging Station Destinations for Electric Vehicle Users - A Simulation-Based Approach to Evaluate Rule-based Strategies* (2021)

- *Building a Simulation Environment for Charging Infrastructure Planning* (2021)
- *Understanding Parking Behavior as a basis for charging station design* (2021)

RAMIN AHADI

Email: ramin.ahadi1994@gmail.com ◊ Mobile-number: (+49) (0) 176 3267 9149

Address: Bernhard-letterhaus-str.49, 50670 Cologne, Germany ◊ Date of birth: 01.04.1994

[HomePage](#) ◊ [LinkedIn](#) ◊ [GitHub](#) ◊ [Google Scholar](#)

I am a last-year **Ph.D. candidate** at the Department of **Information Systems** at the **University of Cologne/IE Business School**, specializing in practical **operations management and data science** work-streams. My focus is on the development of **quantitative and data-driven** decision support platforms, leveraging **agent-based simulation**, advanced **machine learning, optimization**, and **deep reinforcement learning** algorithms. My core competencies lie in **smart mobility**, encompassing shared autonomous and electric vehicles, as well as **energy systems** and their impact on **climate change**. I have experience in a range project stages, from high-level collaborative problem identification and task structuring to the atomic independent development of detailed solutions and presentations. In addition to applying cutting-edge analytical models to real-world problems (with advanced **Python** proficiency), I **teach machine learning theory** to master's and Ph.D. students. My aim is to leverage technological advancements to help build a more sustainable future.

EDUCATION

University of Cologne, Germany Ph.D. in Information Systems & Operations Management Advisor: Prof. Wolfgang Ketter	<i>2025 (exp)</i>
IE Business School Madrid, Spain Visiting Ph.D. Scholar in Information Systems & Technology Advisor: Prof. Konstantina Valogianni	<i>2024</i>
Ferdowsi University of Mashhad, Iran M.Sc in Industrial Engineering & Operations Management Grade: First Class (top 5% of cohort)	<i>2018</i>
Ferdowsi University of Mashhad, Iran B.Sc in Mechanical Engineering Grade: Distinction (top 10% of cohort)	<i>2016</i>

EXPERIENCE

University of Cologne & IE Business School <i>Ph.D. Researcher & Lecturer</i>	Oct '19 - present <i>Cologne, Germany; Madrid, Spain</i>
<ul style="list-style-type: none">· Developing data-driven decision making (artificial intelligence, econometrics, machine learning, optimization, computer simulation, reinforcement learning) with a focus on future energy and mobility systems (EV charging, (autonomous) shared fleet management).· Designing advanced analytical algorithms/solutions (e.g., heuristic (genetic algorithm), online optimization (model-predictive control), and dynamic (reinforcement learning)) for EV smart charging/pricing management and car-sharing/ride-hailing fleet management.· Multiple real-world (agent-based) simulation environments, calibrated with observational data (creating Digital Twins), for shared (autonomous) vehicles, multi-modal mobility platforms, and EV charging hubs.· Real-world data collection and data management in the domain of shared mobility platforms for multiple large EU cities (e.g., Berlin, Paris).· Implementing large-scale numerical analysis using novel technologies: advanced Python packages (Tensorflow and Pytorch for predictive/deep-learning and stochastic decision-making tasks, Simpy for modeling complex simulation environments, and Pyomo for optimization problems), R-studio, SQL for statistics and descriptive analysis, and CPLEX studio, GAMS and JULIA for mathematical programming problems.	

- Publishing papers at leading international information systems (ICIS, ECIS, WITS), operations management, and computer science (AAMAS) conferences and high-impact journals (e.g., Transportation Science, INFORMS)
- Teaching "*Advanced Analytic and Applications/ Advanced Machine Learning*" (Master and Ph.D.), and annual guest lecture on "*Smart Sustainable Mobility*" (MBA/EMBA); supervising theses at Bachelor's, Master's and MBA level
- Co-leading EU-level grant proposals and managing active company research collaborations; e.g., facilitating parking lots with charging equipment and electrifying delivery fleets.

Institute of Energy Economics at the University of Cologne (EWI) Oct '19 - Sep '22
Scholarship Holder and Affiliated Researcher *Cologne, Germany*

- Investigating operations management and smart load scheduling of charging facilities (results presented in group research seminars).
- Researching the interactions between the next generation of energy and mobility sectors with respect to the emerging technologies: high penetration of renewable energy sources and electric vehicles.

Service and Production Operations Management Lab at Ferdowsi University of Mashhad Sep '18 - Sep '19
Volunteer Associate *Mashhad, Iran*

- Co-led research strategy managing & collaborating with manufacturing companies (inventory and supply chain management)
- Developing customer relationship management (CRM) strategies for large restaurants to improve customer service and increase profitability (sales)
- Tutoring programming languages (Python, R, and CPLEX studio), and supervising Bachelor and Master students.

Creative Mind Educational Institute Mar '14 - May '18
Tutor/Teacher *Mashhad, Iran*

- Teaching mathematics and physics in privately managed high schools and educational institutes.

HONOURS & AWARDS

Research Visit Scholarship 2024, C-SEB at the University of Cologne

Fully funded research stay scholarship from the Center for Social and Economic Behavior (C-SEB) for conducting a research visit in IE Business school.

AIS Doctoral Consortium Member 2022, Association for Information Systems (AIS)

One of only 20 Ph.D. students globally to be invited to the prestigious ECIS Doctoral Consortium 2022 in Romania; competitive peer review-based selection process; fully funded.

The most Innovative Research Award, One of my practical projects on demand-side management of electric multi-modal shared vehicles won the most innovative award in ECIS 2022, the largest EU conference in Information Systems

Ph.D. Scholarship, Institute of Energy Economics at the University of Cologne (EWI)

Full scholarship (research, equipment and traveling) to pursue my doctoral study in the field of Operations Management and Information Systems for future energy and mobility systems (Oct 2019 - Oct 2022)

Conference Participation Scholarship, Scholarship from International Conference on Autonomous Agents and Multiagent Systems for Attendance and Paper Registration (2021)

SKILLS

Programming APIs & Technologies	Python, R, MATLAB, Julia, GAMS, CPLEX-studio, SQL, LaTeX, Git, Machine/reinforcement learning (scikit-learn, TensorFlow, Pytorch), optimization (CPLEX, Gurobi), large-scale simulation (SimPy), causal inference/econometrics (R and Python statistics packages), GIS (geopandas, h3, kepler.gl), big data analysis, cloud computing (Microsoft Azure), advanced visualisation (matplotlib and seaborn)
Languages	English (fluent), German (B1+), Persian (native)

SELECTED PUBLICATIONS

Journals

- **Ahadi, R.**, Ketter, W., Collins, J., & Daina, N. (2023). *Cooperative Learning for Smart Charging of Shared Autonomous Vehicle Fleets*. *Transportation Science* 57(3):613-630. <https://doi.org/10.1287/trsc.2022.1187>
- Khalilzadeh, M., Neghabi, H., & **Ahadi, R.** (2021). *An application of approximate dynamic programming in multi-period multi-product advertising budgeting*. *Journal of Industrial & Management Optimization*, 19(1). <https://doi.org/10.3934/jimo.2021202>

Working Papers

- Schroer, K., **Ahadi, R.**, Lee, T. Y., & Ketter, W. (2024). Data-driven Planning of Large-Scale Electric Vehicle Charging Hubs using Deep Reinforcement Learning. (2nd-round)
- **Ahadi, R.**, Taudien, A., Ketter, W., & Gupta, A. (2024). Adoption of Autonomous Vehicles in Ride-Hailing Services: The Role of User Preferences. (1st-round)
- **Ahadi, R.**, Valogianni, K., Schroer, & K., Ketter., W. (x 2025). Dynamic Capacity-Based Pricing of Charging Hubs with Supply Constraints. (1st-round)

Conferences

- Sund, L., Muire, J., **Ahadi, R.**, & Ketter. (2024). A Charging Capacity Market for Competing Fleets. Workshop on Information Technologies and Systems (WITS).
- **Ahadi, R.**, Schroer, K., & Ketter, W. (2024). Managing Electric Vehicle Charging Hubs Through Dynamic Capacity-Based Pricing. European Conference on Information Systems.
- **Ahadi, R.**, Taudian, A., & Ketter, W. (2023). Human versus automated agents: how user preferences affect future mobility systems. European Conference on Information Systems.
- Demircan, M., **Ahadi, R.**, & Ketter, W. (2022). Sustainability vs. Price: Analysis of Electric Multi-Modal Vehicle Sharing Systems under Substitution Effects. European Conference on Information Systems.
- Schroer, K., **Ahadi, R.**, Lee, Y.T., Ketter, W. (2021). *Preference-Aware Planning and Operations of Electric Vehicle Charging Clusters: A Prescriptive Framework*. In Workshop on Information Systems and Technology (WITS) 2021.
- **Ahadi, R.**, Ketter, W., Collins, J., & Daina, N. (2021). Siting and Sizing of Charging Infrastructure for Shared Autonomous Electric Fleets. In Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems.
- Yazdi, L., **Ahadi, R.**, & Rezaee, B. (2018). Optimal electric vehicle charging station placing with integration of renewable energy. 15th Iran International Industrial Engineering Conference, Yazd University 2018.

Cologne, January 2025